

**UNIVERSITÀ DEGLI STUDI DI NAPOLI
FEDERICO II**



SCUOLA POLITECNICA E DELLE SCIENZE DI BASE

Department of Industrial Engineering

Doctoral Programme in Industrial Engineering

**Statistical Methodologies of Functional Data Analysis
for Industrial Applications**

Doctoral Dissertation of:
Fabio Centofanti

Thesis Advisors:

Prof. Ing. Biagio Palumbo

Prof. Ing. Simone Vantini

The Chair of the Doctoral Program:

Prof. Ing. Michele Grassi

Year 2021 - XXXIII Cycle

Summary

This thesis stands as one of the first attempt to connect the statistical object oriented data analysis (OODA) methodologies with the industry field. Indeed, the aim of this thesis is to develop statistical methods to tackle industrial problems through the paradigm of the OODA. The new framework of Industry 4.0 requires factories that are equipped with sensor and advanced acquisition systems that acquire data with a high degree of complexity. OODA can be particularly suitable to deal with this increasing complexity as it considers each statistical unit as an atom or a data object assumed to be a point in a well-defined mathematical space. This idea allows one to deal with complex data structure by changing the resolution of the analysis. Indeed, from standard methods where the atom is represented by vector of numbers, the focus now is on methodologies where the objects of the analysis are whole complex objects. In particular, this thesis focuses on functional data analysis (FDA), a branch of OODA that considers as the atom of the analysis functions defined on compact domains.

The cross-fertilization of FDA methods to industrial applications is developed into three parts in this dissertation. The first part presents methodologies developed to solve specific applicative problems. In particular, a first consistent portion of this part is focused on *profile monitoring* methods applied to ship CO₂ emissions. A second portion deals with the problem of predicting the mechanical properties of an additively manufactured artifact given the particle size distribution of the powder used for its production. And, a third portion copes with the cluster analysis for the quality assessment of metal sheet spot welds in the automotive industry based on observations of dynamic resistance curve.

Stimulated by these challenges, the second part of this dissertation turns towards a more methodological line that addresses the notion of *interpretability* for functional data. In particular, two new interpretable estimators of the coefficient function of the function-on-function linear regression model are proposed, which are named S-LASSO and AdaSS, respectively. Moreover, a new method, referred to as SaS-Funclust, is presented for sparse clustering of functional data that aims to classify a sample of curves into homogeneous groups while jointly detecting the most informative portions of domain.

In the last part, two ongoing researches on FDA methods for industrial application are presented. In particular, the first one regards the definition of a new robust nonparametric functional ANOVA method (Ro-FANOVA) to test differences among group functional means by being robust against the presence of outliers with an application to additive manufacturing. The second one sketches a new methodological framework for the real-time profile monitoring.

Contents

Contents	ii
Introduction	1
Part I	9
1 A Functional Data Analysis Approach for the Monitoring of Ship CO₂ Emissions	11
1.1 Introduction	11
1.2 Methodology	13
1.3 A Real-case Study in the Shipping Industry	16
1.4 Conclusions	19
2 Functional Regression Control Chart	21
2.1 Introduction	21
2.2 The Functional Regression Control Chart Framework	23
2.3 Performance Analysis	26
2.4 Real-case Study: Fuel Consumption Monitoring in the Shipping Industry . .	34
2.5 Conclusions	36
2.6 Supplementary Materials	37
3 Functional Regression Control Chart for the Monitoring of Ship CO₂ Emissions	45
3.1 Introduction	45
3.2 Technological Background and Data Structure	47
3.3 Methodology	50
3.4 Results and Discussion	53
3.5 Conclusions	63
4 Function-on-Function Regression for Assessing Production Quality in Industrial Manufacturing	65
4.1 Introduction	65
4.2 Methodology	67
4.3 A Real-case Study in Laser Powder Bed Fusion Additive Manufacturing . . .	70
4.4 Conclusions	74
4.5 Appendix	75

5	Functional Clustering Methods for Resistance Spot Welding Process Data in the Automotive Industry	77
5.1	Introduction	77
5.2	Technological Background and Data Structure	79
5.3	Functional Data Clustering Approaches for Dynamic Resistance Curves . . .	81
5.4	Results and Discussion	86
5.5	Conclusions	90
5.6	Appendix	91
Part II		95
6	Smooth LASSO Estimator for the Function-on-Function Linear Regression Model	97
6.1	Introduction	97
6.2	Methodology	100
6.3	Theoretical Properties of the S-LASSO Estimator	105
6.4	Simulation Study	108
6.5	Real-Data Examples	111
6.6	Conclusions	115
6.7	Supplementary Materials	116
7	Adaptive Smoothing Spline Estimator for the Function-on-Function Linear Regression Model	135
7.1	Introduction	135
7.2	The Adaptive Smoothing Spline Estimator	138
7.3	Simulation Study	142
7.4	Real-data Examples	147
7.5	Conclusions	150
7.6	Appendix	151
8	Sparse and Smooth Functional Data Clustering	153
8.1	Introduction	153
8.2	The SaS-Funcclust Method for Functional Clustering	156
8.3	Simulation Study	163
8.4	Real-data Examples	166
8.5	Conclusions and Discussions	169
8.6	Supplementary Materials	170
Part III		173
9	Robust Functional ANOVA with Application to Additive Manufacturing	175
9.1	Introduction	175
9.2	The Robust Functional Analysis of Variance	180
9.3	Simulation Study	185
9.4	Real Case-study: Analysis of Variance of Applied to the Analysis of Spatter Behaviour in Laser Powder Bed Fusion	192
9.5	Conclusions	198

CONTENTS

9.6 Supplementary Materials	199
10 Functional Real-time Monitoring Control Chart	203
Conclusion and Future Developments	207
Bibliography	209

Introduction

Nowadays, new technologies allow gathering massive amount of data as never before. Terms as *big data* and *artificial intelligence* have become of common use, and much more interest has been drawn by academic disciplines as computer science, statistic and machine learning. In this scenario, emphasis is usually placed on the size of the data. Indeed, the term big data is now mainstream (Secchi, 2018; Dunson, 2018). However, parallelly data are also becoming more complex. The notion of *complexity* has appeared only recently in the statistical literature, where most of the approaches are designed for datasets where each observation is either a number or a vector. On the contrary, nowadays, datasets are composed of very complex objects such as curves, images and shapes. The increasing data complexity has motivated the birth and development of a branch of statistics which is referred to as Object Oriented Data Analysis (OODA) (Wang et al., 2007; Marron and Alonso, 2014). The idea of OODA is to consider each statistical unit as an *atom* or a *data object* assumed to be a point in a well-defined mathematical space. This idea allows dealing with complex objects by changing the resolution of the analysis.

The industry is not immune from this revolution. The new paradigm of Industry 4.0 requires that factories are equipped with sensor and advanced acquisition systems to increase automation, improve communication and monitor the processes (Xu et al., 2018). However, the acquired datasets are often complex and require the use of new advanced statistical methodologies, as OODA methods, to provide valuable results. Indeed, in industrial applications, the standard practice to deal with complex datasets is to restore the vector representation of the data through some pre-processing operations, as averaging or discretization. With this approach the risk is discarding valuable information and, thus, reducing the effectiveness of the analysis.

This thesis stands as one of the first attempt to connect the statistical OODA methodologies with the industry. The aim of this thesis is to develop statistical methods to tackle modern industrial problems through the new OODA paradigm. In particular, this thesis focuses on *functional data analysis* (FDA), a branch of OODA that considers as the atom of the analysis functions defined on compact domains (Ramsay and Silverman, 2005; Kokoszka and Reimherr, 2017). Most of the datasets collected in the modern industry bring information about curves or surfaces that are apt to be modelled as functional data. FDA has been obtaining increasing attention in the literature since the first appearance of the seminal book by Ramsay and Silverman (2005). The term FDA was coined by Ramsay (1982) and Ramsay and Dalzell (1991), even though the history of this area is much older and dates back to Grenander (1950) and Rao (1958). The greatest issue of FDA is that functions, differently from vectors, are intrinsically infinite dimensional, i.e., infinite parameters are needed to represent them. This poses challenges under both the theoretical and computational points of view. On the other hand, the intrinsic functional data complexity, if utilised properly,

could be the basis for more effective and successful data analyses. In this scenario, FDA extends and adapts multivariate methods, specifically designed for either vectors or numbers, to the case where the atoms of the analysis are functions. For instance, the principal component analysis (Jolliffe, 2011), as dimension reduction tool for multivariate data, has been extended to functional data and termed functional principal component analysis (Shang, 2014); functional regression stands as the counterpart of the multivariate regression analysis where either the response responses or the covariates or both are functional objects (Morris, 2015); clustering analysis has also been proposed for functional data with the aim of grouping functions in such a way that objects in the same group are more similar to each other than to those in other groups (Jacques and Preda, 2014).

In this thesis, the cross-fertilization of FDA methods with industrial applications is developed in three parts. The first part presents methodologies developed to solve specific applicative problems. In particular, a consistent portion is focused on *statistical process control* (SPC) (Montgomery, 2007) methods for functional data which are also known as *profile monitoring* (Noorossana et al., 2012). The aim of SPC is to monitor and control the stability over time of a quality characteristic and to trigger a signal when assignable sources of variations (i.e., special causes) act on the process. When this happens, the process is said to be out-of-control (OC). On the contrary, the process is said to be in-control (IC) when only normal sources of variation (i.e., common causes) apply. We here develop a very transparent set of steps for monitoring profiles in real-world case studies, based on three main steps: (i) data smoothing, where the raw observations are converted to functional data; (ii) functional principal component analysis, where the infinite dimensional problem is reduced to a finite dimensional one by means of an optimal functional data approximation; (iii) monitoring procedure, where the principal component scores are used as input to build appropriate control charts. The effectiveness of this procedure is shown through a real-case study from the maritime field to monitor CO₂ emissions during the navigation phase of a roll-on/roll-off passenger (Ro-Pax) cruise ship, i.e., a ship designed to carry both passengers and wheeled vehicles. In practice, there are situations where the quality characteristic is influenced by one or more functional covariates. In this scenario, if one of these covariates manifests itself with an extreme realization, the quality characteristic may wrongly be judged to be OC. Otherwise, there may be situations where the covariates are not extreme and the quality characteristic may wrongly appear IC. As well as, the quality characteristic may wrongly appear IC because the variance explained by the covariates is overlooked. To address this issue, the functional regression control chart (FRCC) framework is presented. The basic idea behind this chart is to consider the quality characteristic after being adjusted for the effects of the covariates, that is, monitoring the residuals of the regression of the quality characteristic on the covariates. Also in this case, the application is in the shipping industry, where the FRCC is used to identify reductions of cumulative fuel consumption, and thus CO₂ emissions (which are stoichiometrically related to it) after that an energy efficiency initiative was performed off-line on a Ro-Pax ship. Moreover, the potential of the FRCC is also assessed to monitor ship CO₂ emissions throughout each voyage in order to identify special causes, at given values of the functional covariates.

As already introduced before, functional regression is particularly relevant in FDA (Morris, 2015), where the aim is to model the relationship between a quality characteristic (referred to as response) and one or more independent variables (referred to as covariates), in which at least one of them is apt to be modelled as a function. Particularly significant is the setting where both the response and covariate are functions and the relationship between them is assumed to be linear, that is usually referred to as *function-on-function (FoF) linear*

regression model (Ramsay and Silverman, 2005). In order to show the practical applicability and to highlight the benefits of it, the FoF linear regression model is applied to a real-life case study in the additive manufacturing (AM) field. The impact of many characteristics of raw materials on the final properties of the produced parts is in fact still an open research issue and hampers the industrialization of additive technologies (Khajavi et al., 2018; Mani et al., 2017), which are, on the other hand, also very costly. Specifically, the aim is to predict the mechanical properties, summarized into stress-strain curves, of an additively manufactured artefact, given the particle size distribution (PSD) of the powder used for its production. In this setting, practitioners classically settle the analyses by extracting scalar features from both the sample PSD and the observed stress-strain curve. The application of the FoF linear regression allows instead the estimation and prediction of the full stress-strain curve of produced part (i.e., the functional response) given the powder PSD (i.e., the functional covariate).

In industry, another paramount issue constantly faced by practitioners is related to the need of identifying homogeneous groups of observations in a specific dataset. This problem is referred to as cluster analysis in the statistical literature (Everitt et al., 2011). Also in this case, the most common practice is to base cluster analysis on one or a few scalar features extracted from acquired dataset, even though feature extraction is known to be problem-specific, arbitrary, and with the risk to compress useful information. On the contrary, FDA approaches stands as a possible way of minimizing the loss of information. This idea could be successfully applied in the modern automotive Industry 4.0 framework, where automatic acquisition systems allow a control through the continuous record of a large volume of process parameters. In this setting, resistance spot welding (RSW) is the most common technique employed in joining metal sheets during body-in-white assembly of automobiles (Ighodaro et al., 2016) and their quality is routinely controlled in order to guarantee the structural integrity and solidity of welded assemblies in each vehicle. In particular, the so-called dynamic resistance curve (DRC) is the most important process parameter acquired on-line that is recognized as the full technological signature of the metallurgical development of a spot weld (Dickinson et al., 1980). In this scenario, it is important to identify clusters of spot welds based on DRC observations, with a convenient interpretation in terms of mechanical and metallurgical properties especially since the difficulty in matching each DRC with a quantitative measure of the final quality of the corresponding spot welds.

Part II of this dissertation turns towards a more methodological line. Sometimes, literature FDA methods are in fact not able to straightforwardly address problems that arise in industrial applications. The increasing complexity, brought by the Industry 4.0 revolution, stimulates new challenges and calls for totally new statistical techniques. In addition, most of industrial applications require methods that are able to enhance the understanding of the process under study. This idea is related to the broad notion of *interpretability*. There is no mathematical definition of interpretability. A (non-mathematical) definition by Miller (2019) is: *Interpretability is the degree to which a human can understand the cause of a decision*. The higher the interpretability, the easier it is for someone to comprehend why certain decisions have been made. This idea translates directly to FDA, where new methods could be developed to improve the degree to which a human can understand the process under study.

In the FoF linear regression model, the functional predictor contributes linearly to the response through the coefficient function, which works as a continuous weight function. Sparse coefficient functions have a great practical interest because allow identifying domain regions where the predictor has not influence on the response. Methods that are able to

recover the sparse nature of the coefficient function are referred to as sparse or interpretable ones (James et al., 2009), because they facilitate the interpretation of the relationship between the covariate and the response, and, thus, improve interpretability. In the literature, no effort has been made to obtain an interpretable estimator for the FoF linear regression model. We here introduce the sparse S-LASSO (Smooth plus LASSO) estimator of the coefficient function of the FoF linear regression model. The property of sparseness of the S-LASSO estimator is provided by a the functional LASSO penalty, which is the functional generalization of the classical Least Absolute Shrinkage and Selection Operator (LASSO) (Tibshirani, 1996). Sparse estimators are not the only ones able to enhance interpretability in the FoF linear regression setting. Indeed, in this setting, a higher degree of interpretability is reflected by a higher accuracy in the estimation of the coefficient function. If the coefficient function is sparse then sparse estimators are needed, otherwise the sparseness property is not essential. Most of the estimation methods for the FoF linear regression model rely on penalty terms that reduce the complexity by constraining the coefficient function to own a certain degree of smoothness (Ramsay and Silverman, 2005). However, such penalties fail to accommodate the local behaviour of coefficient function because they uniformly apply penalization throughout the domain. To solve this problem, we here present the adaptive smoothing spline (AdaSS) estimator that considers two adaptive roughness penalties that are able to produce different amount of penalty over the domain. The idea of adaptive penalties is not new in the statistical literature (Ruppert and Carroll, 2000), indeed, it is very popular and well established in the field of nonparametric regression (Wahba, 1990). On the contrary, in the FoF linear regression setting, the AdaSS estimator is the first attempt to improve the estimation accuracy of the coefficient function through the use of adaptive penalties.

Cluster analysis is another branch of statistics, where interpretability could have a prominent role. Also in this case, interpretability could be reflected in the notion of sparsity. Sparse clustering methods are capable of selecting informative features, i.e., the features in which respect the clusters differ the most, and eliminating the remaining ones, referred to as noninformative. In the multivariate context, sparse clustering methods have received increasing attention in the statistical literature (Pan and Shen, 2007), however only recently, the notion of sparsity has been translated into a functional data clustering framework (Floriello and Vitelli, 2017). Specifically, sparse functional clustering methods aim to cluster a bunch of curves while jointly detecting the most informative portion of domain to the clustering in order to improve both the accuracy and the interpretability of the analysis. In this dissertation, we introduce the sparse and smooth functional clustering (SaS-Funclust) method, where the basic idea is to provide accurate and interpretable cluster analysis. Differently from the methods already present in the literature (Floriello and Vitelli, 2017; Vitelli, 2019), the SaS-Funclust method is able to detect portions of domain that are noninformative pairwise, i.e., for at least a specific cluster pair, and it owns the great flexibility of the model-based procedures.

In the Part III of this dissertation, two ongoing researches on FDA methods for industrial application are presented. The first one tackles the problem of identifying the presence of significant differences, in terms of functional mean, among groups of a functional data, is treated. This problem is the functional extension of the classical analysis of variance, that is, the functional analysis of variance (FANOVA) (Ramsay and Silverman, 2005). However, when dealing with real data, the functional sample under study is usually contaminated by some outliers, which can strongly bias the analysis results (Hubert et al., 2015). In order to take into account the possible presence of anomalous functional observations, a new robust nonparametric functional ANOVA method (Ro-FANOVA) is presented that is able to test

differences among group functional means by reducing the influence of anomalous observations. The last work presented in this dissertation sketches the main elements of a new method for real-time monitoring of functional data. Most of the profile monitoring methods aim to assess the stability of the functional quality characteristic completely observed (Centofanti et al., 2020b; Capezza et al., 2020). However, in some applications, the interest relies in understanding if the process is working properly before its completion. For this reason, a new method, referred to as functional real-time monitoring control chart (FRTMCC), is going to be developed for the real-time monitoring of a quality characteristic profile. It is based on the idea of real-time alignment and simultaneously monitoring of phase and amplitude variations.

The dissertation is developed in seven chapters, whose structure is briefly summarized hereinafter.

Part I

Chapter 1: *A Functional Data Analysis Approach for the Monitoring of Ship CO₂ Emissions*

The objective of the present chapter is to provide the reader with a very transparent set of steps for the SPM of functional data in real-world case studies: i) identifying a finite dimensional model for the functional data, based on functional principal component analysis; ii) estimating the unknown parameters; iii) designing control charts on the estimated parameters, in a nonparametric framework. The proposed SPM procedure is applied to a real-case study from the maritime field in monitoring CO₂ emissions from real navigation data of a roll-on/roll-off passenger cruise ship, i.e., a ship designed to carry both passengers and wheeled vehicles that are driven on and off the ship on their own wheels.

The results here presented are illustrated in

- Capezza, C., Centofanti, F., Lepore, A., Palumbo, B., (2021), ‘A Functional Data Analysis Approach for the Monitoring of Ship CO₂ Emissions’. To appear in *Gestão & Produção*.

Chapter 2: *Functional Regression Control Chart*

In this chapter, we present a new framework for monitoring a functional quality characteristic when functional covariates are available. In particular, the Phase II monitoring is considered, where the unknown parameters and control limits are estimated using a set of IC data (Phase I). This framework is henceforth referred to as functional regression control chart (FRCC). Specifically, we consider the case when the model which links the functional response and functional covariates is linear and, and we monitor residuals by using a profile monitoring approach based on the simultaneous application of the Hotelling’s T^2 and the squared prediction error (SPE) control charts.

The results here presented are illustrated in

- Centofanti, F., Lepore, A., Menafoglio, A., Palumbo, B., Vantini, S. (2020). ‘Functional regression control chart’. *Technometrics*, 1-14.

Chapter 3: *Functional Regression Control Chart for Monitoring Ship CO₂ Emissions*

In this chapter, we show an application of the FRCC with the ultimate goal of answering, at the end of each ship voyage, the question: *given the value of the covariates, is the observed CO₂ emission profile as expected?*. To this aim, the FRCC

focuses on the monitoring of residuals obtained from a multivariate functional linear regression of the CO₂ emission profiles on the functional covariates. The applicability of the FRCC is demonstrated through a real-case study of a Ro-Pax ship operating in the Mediterranean Sea.

The results here presented are illustrated in

- Centofanti, F., Lepore, A., Menafoglio, A., Palumbo, B., Vantini, S. (2020). ‘Functional regression control chart for monitoring ship CO₂ emissions’. To appear in *Quality and Reliability Engineering International*.

Chapter 4: *Function-on-function regression for assessing production quality in industrial manufacturing*

In order to show its practical applicability and to highlight the benefits of a functional data approach, in this chapter, the function-on-function linear regression model is applied to a real-life case study in the additive manufacturing (AM) field. In particular, the potential of the proposed method is demonstrated by applying it to a real-case study in powder bed fusion AM for metals to predict the mechanical properties of an additively manufactured artefact given the particle size distribution of the powder used for its production.

The results here presented are illustrated in

- Palumbo, B., Centofanti, F., Del Re, F. (2020). ‘Function-on-function regression for assessing production quality in industrial manufacturing’. *Quality and Reliability Engineering International*, **36**(8), 2738-2753.

Chapter 5: *Functional clustering methods for resistance spot welding process data in the automotive industry*

The present chapter means to show the potentiality and the practical applicability of clustering methods for functional data, which avoid the need for arbitrary and often controversial feature extraction. The aim is to find out homogeneous groups of dynamic resistance curves (DRC), which is recognized as the full technological signature of the spot weld, produced in joining metal sheets in the automotive industry. Homogeneous groups of DRCs likely pertain to spot welds sharing common mechanical and metallurgical properties. Moreover, an essential hands-on overview of the most promising functional clustering methods is provided.

The results here presented are illustrated in

- Capezza, C., Centofanti, F., Lepore, A., Palumbo, B. (2021+). ‘Functional clustering methods for resistance spot welding process data in the automotive industry’. Submitted to *Applied Stochastic Models in Business and Industry* awaiting for final decision.

Part II

Chapter 6: *Smooth LASSO Estimator for the Function-on-Function Linear Regression Model*

In this chapter, a new estimator, named S-LASSO, is presented for the coefficient function of the function-on-function linear regression model. The S-LASSO estimator is shown to be able to increase the interpretability of the model, by better locating regions where the coefficient function is zero, and to smoothly estimate non-zero values of the coefficient function. The sparsity of the estimator is ensured by a functional LASSO penalty, which pointwise shrinks toward zero the coefficient function, while

the smoothness is provided by two roughness penalties that penalize the curvature of the final estimator.

The results here presented are illustrated in

- Centofanti, F., Fontana, M., Lepore, A., Vantini, S. (2020). ‘Smooth Lasso Estimator for the Function-on-Function Linear Regression Model’. *arXiv preprint arXiv:2007.00529*.

Chapter 7: *Adaptive Smoothing Spline Estimator for the Function-on-Function Linear Regression Model*

In this chapter, we present an adaptive smoothing spline (AdaSS) estimator for the function-on-function linear regression model where each value of the response, at any domain point, depends on the full trajectory of the predictor. The AdaSS estimator is obtained by the optimization of an objective function with two spatially adaptive penalties, based on initial estimates of the partial derivatives of the regression coefficient function. This allows the proposed estimator to adapt more easily to the true coefficient function over regions of large curvature and not to be undersmoothed over the remaining part of the domain.

The results here presented are illustrated in

- Centofanti, F., Lepore, A., Menafoglio, A., Palumbo, B., Vantini, S. (2020). ‘Adaptive Smoothing Spline Estimator for the Function-on-Function Linear Regression Model’. *arXiv preprint arXiv:2011.12036*.

Chapter 8: *Sparse and Smooth Functional Data Clustering*

In this chapter, a new model-based procedure is developed for sparse clustering of functional data that aims to classify a sample of curves into homogeneous groups while jointly detecting the most informative portions of domain. The proposed method is referred to as sparse and smooth functional clustering (SaS-Funclust) and relies on a general functional Gaussian mixture model whose parameters are estimated by maximizing a log-likelihood function penalized with a functional adaptive pairwise penalty and a roughness penalty. The former allows identifying the noninformative portion of domain by shrinking the means of separated clusters to some common values, whereas the latter improves the interpretability by imposing some degree of smoothing to the estimated cluster means.

The results here presented are illustrated in

- Centofanti, F., Lepore, A., Palumbo, B. (2021). ‘Sparse and Smooth Functional Data Clustering’. *arXiv preprint arXiv:2103.15224*.

Part III

Chapter 9: *Robust Functional ANOVA with Application to Additive Manufacturing*

In this chapter, we propose a robust nonparametric functional ANOVA method (Ro-FANOVA) that is able to test differences among group functional means by being robust against the presence of outliers. Indeed, when dealing with real data, it is common that the functional sample under study is contaminated by some anomalous observations, which can strongly bias the analysis. In order to take into account this possibility, the Ro-FANOVA method reduces the weights of outliers on the results of the analysis. The performance of the proposed approach is demonstrated in the framework of a motivating real case-study in additive manufacturing. The aim of the robust FANOVA consists of identifying significant effects of relevant process factors

and their interactions on by-products of the process, i.e., spatter ejections, measured through high-speed video imaging and translated into a functional form.

The results of this chapter are part of an on-going research in collaboration with Bianca Maria Colosimo (Politecnico di Milano, Milan, Italy), Marco Grasso (Politecnico di Milano, Milan, Italy), Alessandra Menafoglio (Politecnico di Milano, Milan, Italy), Biagio Palumbo (University of Naples, Naples, Italy), and Simone Vantini (Politecnico di Milano, Milan, Italy).

Chapter 10: *Functional Real-time Monitoring Control Chart*

In this chapter, a new method, referred to as functional real-time monitoring control chart (FRTMCC), to real-time monitor a functional quality characteristic is described. Indeed, most of the profile monitoring methods aim to assess the stability of the functional quality characteristic in its entirety (Centofanti et al., 2020b; Capezza et al., 2020). However, in some applications, the interest relies in understanding if the process is working properly before its completion. To tackle this issue, the FRTMCC combines an alignment step with a monitoring scheme of phase and amplitude variations.

The results of this chapter are part of an on-going research in collaboration with Max Spooner (DTU Compute, Technical University of Denmark, Kgs. Lyngby, Denmark), Murat Kulahci (DTU Compute, Technical University of Denmark, Kgs. Lyngby, Denmark), Biagio Palumbo (University of Naples, Naples, Italy), and Antonio Lepore (University of Naples, Naples, Italy).

Part I

Chapter 1

A Functional Data Analysis Approach for the Monitoring of Ship CO₂ Emissions

Abstract

Sensing networks provide nowadays massive amounts of data that in many applications vary over a continuum and thus, can be suitably modelled as functional data. Their proper modelling by means of functional data analysis approaches naturally addresses new challenges also in the statistical process monitoring (SPM) of functional data, also known as profile monitoring. The objective of the present motivation-applied chapter is to provide the reader with a very transparent set of steps for the SPM of functional data in real-world applications: i) identifying a finite dimensional model for the functional data, based on functional principal component analysis; ii) estimating the unknown parameters; iii) designing control charts on the estimated parameters, in a nonparametric framework. The profile monitoring strategy is applied to a real-case study from the maritime field in monitoring CO₂ emissions during the navigation phase of a roll-on/roll-off passenger cruise ship, whose data are courtesy of the owner Grimaldi Group. We show different scenarios highlighting clear and interpretable indications that can be extracted from the data set that support the detection of anomalous voyages.

1.1 Introduction

In many applications, the development of data-acquisition systems allows the gathering of massive amount of data that can be suitably modelled as functional data, that is as functions varying over a continuum. Functional data analysis (FDA) refers to the set of statistical methods where the observation units are functional data. Thorough overviews of FDA techniques are provided by Ramsay and Silverman (2005); Horváth and Kokoszka (2012); Kokoszka and Reimherr (2017) and Hsing and Eubank (2015); Bosq (2012) for more specific theoretical insight. Functional data are usually observed on a discrete grid of points. Thus, standard multivariate methods could be in principle applied. Unfortunately, often the number of observations is much less than the number of observed discrete points and, thus, issues of high dimensionality arise, and must be properly overcome when the aim of the analysis is the monitoring and controlling of the stability over time of quality characteristics

apt to be modelled as a functional data. Statistical process monitoring of functional data is known also as *profile monitoring*, where functional data are referred to as *profiles*. As in the classical univariate and multivariate setting, where data are represented by scalars or vectors, profile monitoring control charts have the task of continuously monitoring the quality characteristic and of triggering a signal when assignable sources of variations (i.e., special causes) act on it. When this happens, the process is said to be out of control (OC). Otherwise, when only normal sources of variation (i.e., common causes) apply, the process is said to be in-control (IC). As discussed by Woodall et al. (2004), all the approaches for profile monitoring share the following structure: i) identifying a finite dimensional model for the functional data; ii) estimating the unknown parameters; iii) designing control charts on the estimated parameters. In particular, the book of Noorossana et al. (2012) represents a comprehensive overview of profile monitoring methods. Pini et al. (2017) proposed a two-steps profile monitoring approach where, firstly, the informative parts of the functional data to be monitored are selected by means of the inferential interval-wise testing procedure (Pini and Vantini, 2017), and, then, the monitoring strategy is performed on the basis of the information that the functions contain in the selected domains. Menafoglio et al. (2018) introduced a new approach for monitoring probability density functions based on simplicial functional principal component analysis. Grasso et al. (2016) presented a novel approach for profile monitoring that combines the functional principal component analysis and the use of parametric warping functions. More recently, Capezza et al. (2019) extends classical multivariate techniques to the monitoring of multivariate functional data and a scalar quality characteristic related to them. Whereas, Centofanti et al. (2020b) expand the Mandel's regression control chart idea (Mandel, 1969) to the functional setting, that is a control chart elaborated on the functional residuals obtained from a function-on-function regression of the quality characteristic profile on concurrent functional covariates. Other relevant contributions in this field include the work of Jin and Shi (1999); Colosimo and Pacella (2007, 2010), and Grasso et al. (2017).

The objective of the present motivation-applied chapter is to provide the reader with a very transparent set of steps for monitoring profiles in real-world applications. In particular, the proposed method can be divided into three main steps. Firstly, the functional data are obtained from the raw data through a smoothing technique based on spline functions. Then, a functional principal component analysis (FPCA), that is the functional extension of the classical (non-functional) principal component analysis (PCA) (Jolliffe, 2011), is performed in order to extract the relevant principal component scores. Lastly, the retained principal component scores are used in a monitoring strategy that is based on the simultaneous application of the *Hotelling's* and the *squared prediction error* (SPE) control charts in a nonparametric framework.

A complete overview of smoothing techniques for functional data is provided by Ramsay and Silverman (2005), where methods based on least squares and roughness penalties are presented under a practical point of view. More generally, references on smoothing spline estimators for nonparametric regression are Wahba (1990); Green and Silverman (1993); Eubank (1999), and Gu (2013). A survey of FPCA, and its use in explanatory analysis, modeling and forecasting, and classification of functional data is provided by Shang (2014). The T^2 and SPE control charts are widely used for multivariate statistical process monitoring (Montgomery, 2007). See Lowry and Montgomery (1995) for a review on multivariate control charts.

Eventually, the proposed monitoring strategy is applied to a real-case study from the maritime field in monitoring CO₂ emissions during the navigation phase of a roll-on/roll-off

passenger (Ro-Pax) cruise ship, whose data are courtesy of the owner Grimaldi Group.

The chapter is structured as follows. Section 1.2 introduces the proposed procedure. In particular, in Section 1.2 we discuss how to obtain the functional data from the raw observations through data smoothing techniques; Section 1.2 describes the functional principal component analysis, and, Section 1.2 introduces the monitoring strategy based on the T^2 and SPE control charts. The real-case study in the shipping industry is presented in 1.3. Section 1.4 concludes the chapter. All computations and plots have been obtained by using the software environment R (R Core Team, 2020a).

1.2 Methodology

As stated before, the proposed methods for profile monitoring is composed of three main steps.

1. Data smoothing: the raw observations are converted to functional data.
2. Functional Principal Component Analysis: the infinite dimensional problem is translated into a finite dimensional one by means of an optimal functional data approximation.
3. Monitoring strategy: the principal component scores are used as input to build the T^2 and SPE control charts.

In the following Section 1.2, 1.2 and 1.2 these steps are illustrated, respectively.

Data Smoothing

Data are collected by devices in a discrete fashion, that is as n discrete observed curves $\{Y_i(t_j), j = 1, \dots, p\}_{i=1, \dots, n}$, where $\{t_j\}_{j=1, \dots, p}$ are the observation points in a given closed interval $\mathcal{T} \subseteq \mathbb{R}$. Hence, appropriate methods are required to convert discrete raw data $\{Y_i(t_j)\}$ into functional data $\{X_i(t)\}$ computable for any $t \in \mathcal{T}$, which are random realizations of a functional quality characteristic $X(t)$. If the discrete data are assumed without any measurement error, functional data can be theoretically drawn up by merely connecting the whole set of points $\{Y_i(t_j), j = 1, \dots, p\}_{i=1, \dots, n}$. However, this does not represent the ordinary situation. When measurement error is present, each discrete observation is expressed as

$$Y_i(t_j) = X_i(t_j) + \varepsilon_{ij}, \quad (1.1)$$

for $i = 1, \dots, n$ and $j = 1, \dots, p$, where ε_{ij} are zero mean random errors with equal variances. Trivially, note that Equation (1.1) degenerates in the previous case when the variance ε_{ij} tends to zero. From Equation (1.1), data smoothing techniques aim to recover the functional data by discarding exogenous perturbation due to error terms ε_{ij} . Functional data are intrinsically infinite dimensional, that is infinite features are needed to completely specify them, precisely, the values at each possible argument $t \in \mathcal{T}$. To this end, a common approach consists of representing, for $i = 1, \dots, n$, each functional datum $\{X_i(t)\}$ through a linear combination of K known basis functions $\mathbf{\Phi} = (\phi_1, \dots, \phi_K)^T$, as follows

$$X_i(t) = \sum_{l=1}^K c_{il} \phi_l(t) = \mathbf{c}_i^T \mathbf{\Phi}(t) \quad t \in \mathcal{T}, \quad (1.2)$$

1. A FUNCTIONAL DATA ANALYSIS APPROACH FOR THE MONITORING OF SHIP CO₂ EMISSIONS

where $\mathbf{c}_i = (c_{i1}, \dots, c_{ik})^T$ is the coefficient vector for each curve. Then, the problem of recovering the functional data $\{X_i(t)\}$ reduces to the estimation of the unknown coefficient vectors \mathbf{c}_i for each i . In particular, for every $i = 1, \dots, n$, the coefficient vector \mathbf{c}_i is estimated as $\hat{\mathbf{c}}_i$ by minimizing the following penalized sum of squares error

$$\hat{\mathbf{c}}_i = \underset{\mathbf{c} \in \mathbb{R}^K}{\operatorname{argmin}} \sum_{j=1}^p \left(Y_i(t_j) - \mathbf{c}^T \boldsymbol{\Phi}(t_j) \right)^2 + \lambda \mathbf{c}^T \mathbf{R} \mathbf{c}, \quad (1.3)$$

where $\lambda > 0$ is a smoothing parameter and \mathbf{R} is a matrix whose entries i, j are $\int_{\mathcal{T}} \phi_i^{(m)}(t) \phi_j^{(m)}(t) dt$, with $\phi^{(m)}$ the m -derivative of ϕ . Finally, the functional data we are interested in are as follows

$$\hat{X}_i(t) = \hat{\mathbf{c}}_i^T \boldsymbol{\Phi}(t) \quad t \in \mathcal{T}, \quad (1.4)$$

for $i = 1, \dots, n$.

Note that, to obtain the functional data as in Equation (1.4), some choices should be done and discussed in the following. As basis functions $\boldsymbol{\Phi}$, the B-spline basis system is the most common choice in case of non-periodic functional data owing good computational properties and great flexibility (Ramsay and Silverman, 2005). This implicitly assumes that the considered curves are well approximated by a spline function. Splines are optimal in the sense of being the smoothest possible functions interpolating the data (Green and Silverman, 1993). Spline functions divide the functional domain into subintervals, by means of break points. Over any subinterval, the spline is a polynomial of specific order q , with $q - 1$ non-zero derivatives and matching proper derivative constraints between adjacent polynomials (De Boor et al., 1978). The smoothing parameter λ in Equation (1.3) is chosen as that corresponding to the minimum value assumed by the generalized cross-validation (GCV) criterion, which is a well-known method to tradeoff between variance and bias. This criterion takes into account the degrees of freedom of the estimated curve that vary according to λ . See Ramsay and Silverman (2005) for further details. The penalty on the right-hand side of Equation (1.3) is computed by setting $m = 2$, i.e., by penalizing the function roughness. The value of K is not crucial (Cardot et al., 2003), unless it is sufficiently large to capture the local behaviour of the functional data.

Functional Principal Component Analysis

FPCA is a key method aimed at reducing the infinite dimensionality of the functional data, by retaining a finite number L of *principal component scores* or simply *scores* $\{\xi_{il}\}_{l=1, \dots, L}$, which explain the largest part of the sample variability, for each functional observation $\{X_i(t)\}$, obtained as described in Section 1.2 defined for $t \in \mathcal{T}$. Without loss of generality, in what follows, let us assume that $\{X_i(t)\}$ have zero mean or that, they are opportunely centered by subtracting the functional sample mean. Then, scores are defined as

$$\xi_{il} = \int_{\mathcal{T}} \psi_l(t) X_i(t) dt, \quad (1.5)$$

where $\{\psi_l\}_{l=1, \dots, L}$ are weight functions referred to as *principal components*. The principal components are subject to size restrictions of normalization and orthogonality. Respectively, they are subject to the constraints: $\int_{\mathcal{T}} \psi_l(t)^2 dt = 1$ and $\int_{\mathcal{T}} \psi_i(t) \psi_j(t) dt = 0$, for $i \neq j$. In this way, each weight function provides new information with respect to those brought by previous principal components.

Principal components are calculated defined by an opportune iterative algorithm which at each step finds the weight function that maximizes the following mean square of the scores, i.e., their sample variance,

$$\psi_l = \underset{\psi}{\operatorname{argmax}} \sum_{i=1}^n \xi_{il}^2 = \sum_{i=1}^n \left(\int_{\mathcal{T}} \psi(t) X_i(t) dt \right)^2, \quad l = 1 \dots, L, \quad (1.6)$$

under the normalization and orthogonality constraints. Moreover, the principal components correspond to eigenfunctions of the covariance function of the process X (Ramsay and Silverman, 2005). For each $i = 1, \dots, n$, let us consider the function \hat{X}_i^{PC} , $i = 1, \dots, n$ defined as

$$\hat{X}_i^{PC}(t) = \sum_{l=1}^L \xi_{il} \psi_l(t) \quad t \in \mathcal{T}, \quad (1.7)$$

that is the linear combination of the principal components and the scores. It can be demonstrated that for each i , \hat{X}_i^{PC} is the best L -dimensional approximation of X_i in terms of mean squared error, i.e., the quantity $E \left[\int_{\mathcal{T}} \left(\hat{X}_i^{PC}(t) - X_i(t) \right)^2 dt \right]$ is minimum over all the other L -dimensional linear combinations.

The choice of the number of retained components L depends on different necessities. Generally, the retained principal components are chosen such that the latter explain at least a given percentage of the total variability. However, more sophisticated method could be used as well (Jolliffe, 2011).

Monitoring Strategy

In this step, the information provided by FPCA is used to continuously monitor the functional quality characteristic X over time. To this aim, two functional control charts are introduced based on the following T^2 and the SPE statistics, defined for each $i = 1, \dots, n$, respectively. The T^2 statistic is as follows

$$T_i^2 = \sum_{l=1}^L \frac{\xi_{il}^2}{\lambda_l}, \quad (1.8)$$

where $\lambda_1, \dots, \lambda_L$ are the variances of $\xi_{i1}, \dots, \xi_{iL}$ and correspond to the eigenvalues of the covariance function of X . The statistic T^2 is the square distance of the projection of X from the origin of the space spanned by the principal components $\{\psi_l\}$. Changes along directions orthogonal to the latter space are monitored by means of the SPE_i statistic, defined for each i as

$$SPE_i = \int_{\mathcal{T}} \left(X_i(t) - \hat{X}_i^{PC}(t) \right)^2 dt, \quad (1.9)$$

where \hat{X}_i^{PC} are defined in Equation (1.7).

In this chapter, we focus on perspective (Phase II) monitoring. Thus, a set of IC data must be preliminarily set up in the design phase of the control charts (Phase I). Let assume that the functional observations $\{X_i\}$ are acquired under IC conditions, principal components $\{\psi_l\}$ and eigenvalues $\lambda_1, \dots, \lambda_L$ shall be estimated from the sample covariance function and denoted by $\{\hat{\psi}_l\}$ and $\hat{\lambda}_1, \dots, \hat{\lambda}_L$, respectively. The estimation of the control limits for both the T^2 and the SPE control charts are obtained by means of $(1 - \alpha)$ -quantiles of the empirical

1. A FUNCTIONAL DATA ANALYSIS APPROACH FOR THE MONITORING OF SHIP CO₂ EMISSIONS

distribution of the two statistics based on the estimated $\{T_i^2\}$ and $\{SPE_i\}$. The parameter α is chosen by using the Bonferroni correction $\alpha = \alpha^*/2$, where α^* is the overall type I error probability, in order to control the *family wise error rate* (FWER). Other corrections are also possible, such as the Sidák correction (Lehmann and Romano, 2006) $\alpha = 1 - (1 - \alpha^*)^{0.5}$. In Phase II, let $X^*(t)$, $t \in \mathcal{T}$, denote a new observation of the functional quality characteristic. Then, the new estimated scores are calculated as $\hat{\xi}_l^* = \int_{\mathcal{T}} \hat{\psi}_l(t) X^*(t) dt$, for $l = 1, \dots, L$, where $\{\hat{\psi}_l\}$ are the estimated principal components. The new realization of the T^2 and SPE statistics are calculated as

$$T^{2*} = \sum_{l=1}^L \frac{\hat{\xi}_l^{*2}}{\hat{\lambda}_l}, \quad (1.10)$$

and,

$$SPE^* = \int_{\mathcal{T}} \left(X^*(t) - \hat{X}^{PC*}(t) \right)^2 dt, \quad (1.11)$$

where $\hat{X}^{PC*}(t) = \sum_{l=1}^L \hat{\xi}_l^* \hat{\psi}_l(t)$, $t \in \mathcal{T}$. An out-of-control signal is issued if at least one of T^{2*} and SPE^* violates the control limits.

The more the portion of variability of the functional data explained by the first L principal components retained into the FPCA model, the more coherent is the following interpretation of T^{2*} and SPE^* statistics. In fact, as the former is based on the first scores, we expect that the larger the value, the larger the deviation in *magnitude* from the reference mean of the new functional observation. Accordingly, as the latter is based on the last principal components (see Equation (1.9)), we expect that a new profile with large SPE^* exhibits non-negligible deviation in the covariance structure, which roughly controls the *shape* of the current functional observation, from that estimated on the reference data set.

1.3 A Real-case Study in the Shipping Industry

We illustrate the proposed monitoring strategy by means of a real-case study from the maritime field in monitoring CO₂ emissions during the navigation phase of a roll-on/roll-off passenger (Ro-Pax) cruise ship. The data analysed in this chapter are courtesy of the owner Grimaldi Group. Information about ports, name of the ship and CO₂ emissions are omitted for confidentiality reasons. Two years' worth of data are available with five-minute frequency. In the proposed application, we focus only on one route sailed by the ship to link two ports. The available data set contains the discrete values for 194 voyages of CO₂ emissions due to propulsion, which is the functional quality characteristic to be monitored at the end of each voyage. The functional domain is the fraction of total distance travelled from the beginning of the voyage. The first 146 voyages are used as training data set to perform FPCA and estimate control chart limits. Then, each of the following 48 voyages is sequentially numbered using a voyage number (VN) and monitored as described in Section 1.2. Note that, since we focus on Phase II monitoring only, we do not report details about Phase I, which was devoted to overarching statistical and engineering considerations to filter out from the training data set data that do not reflect standard navigation conditions and thus may introduce bias in the model parameter and control chart limit estimation. For each voyage, functional data are drawn from discrete observations by means of 50 B-spline basis functions and equally spaced knots. Functional data are smoothed by penalizing the integrated squared second derivative and by choosing the smoothing parameter through GCV criterion, as discussed in




Figure 1.1. First four eigenfunctions obtained with FPCA. For each eigenfunction, the percentage of variability explained is reported.

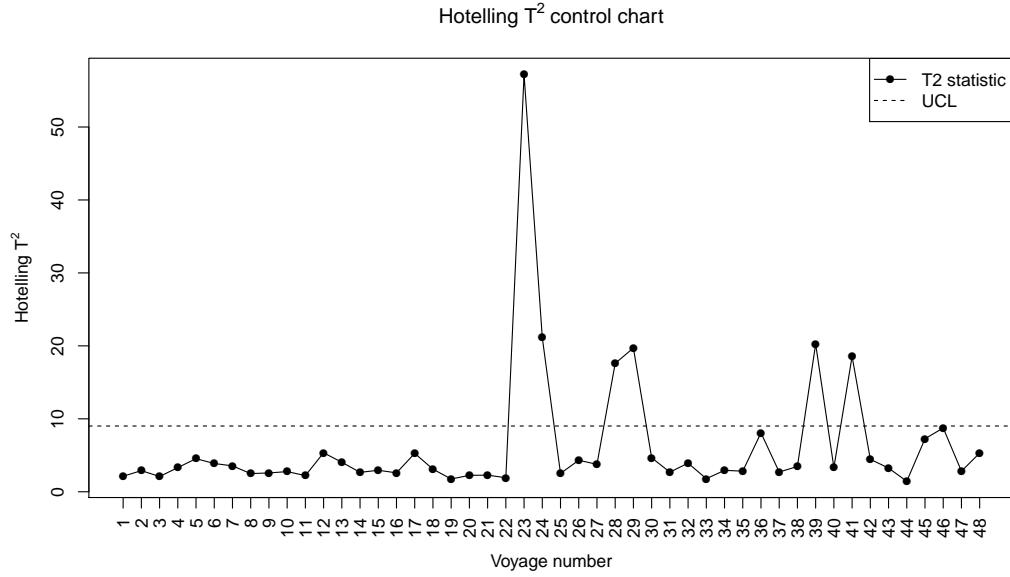
Section 1.2. Then, FPCA is applied on the training data set. Figure 1.1 reports the first four eigenfunctions and the percentage of variability explained by each principal component.

As an example, note that the first component explains 56.8% of the total variability in the data, which is mainly attributed to the beginning and the end parts of the voyage. This can be confirmed by a closer look to Figure 1.3, where the Phase I profiles used to perform FPCA are plotted in grey. The second component explains 14.5% of the variability, which is attributed to the average value of CO₂ emissions alongside the voyage, whereas the third component explains 12.1% of the variability and attributes the main weight to the end part of the voyage. Starting from the fourth component, the explained variability is less than 5% and interpretation becomes cumbersome. For the reasons discussed above, in this application it is convenient to retain the first $L = 3$ functional principal components, which explain together 83.4% of the total variability in the data, to approximate functional data and use the corresponding scores to calculate the statistic, while the residual functions can be used to calculate the SPE statistic as in Equation (1.9). Once control limits are estimated based on the T^2 and SPE statistics calculated on the training data set, it is possible to use control charts to monitor new voyages. Figure (1.9) shows the T^2 and SPE control charts used for Phase II monitoring.

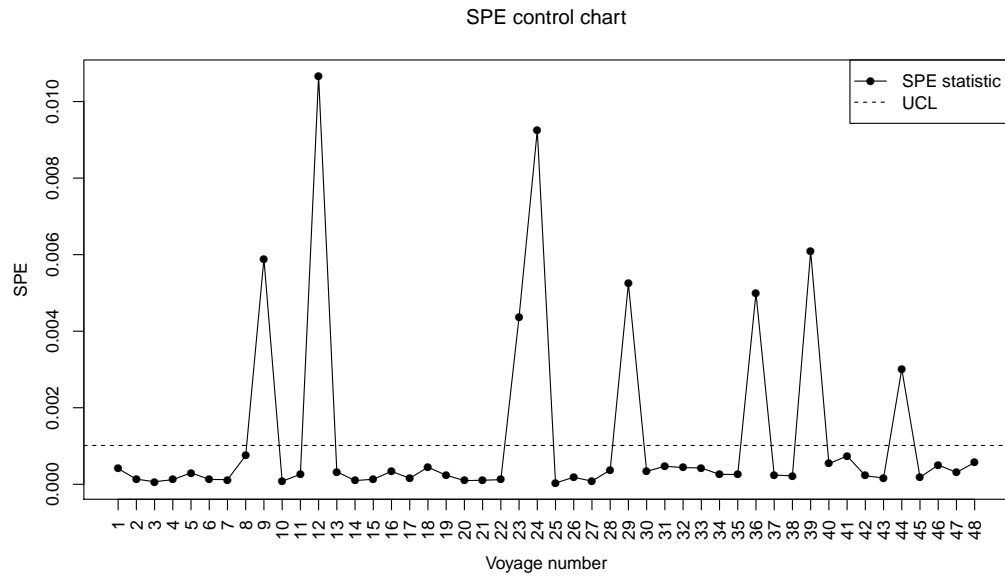
Several scenarios are possible, and it is interesting to notice the use of both control charts supports the interpretation of the type of anomalies encountered. In Figure 1.3 we report OC profiles against those of the training data set, which are plotted for ease of comparison as grey lines.

In particular, note that VN 28 and 41 are OC in the T^2 control chart only (Figure 1.3a), VN 9, 12, 36, and 44 are OC in the SPE control chart only (Figure 1.3b), whereas VN 23, 24, 29, and 39 are OC in both control charts (Figure 1.3c). In Figure 1.3a, profiles of VN 28 and 41 show a clear shift in magnitude only, that is the CO₂ emissions plot below the average. Strictly speaking, it is worth noting that lower CO₂ emissions, which are in fact desired, often are trivially associated to voyages sailed at lower-than-usual speed over ground that, in turn, imply other types of undesired costs for the shipping company due to arrival delay.

1. A FUNCTIONAL DATA ANALYSIS APPROACH FOR THE MONITORING OF SHIP CO₂ EMISSIONS



(a)



(b)

Figure 1.2. (a) T^2 and (b) SPE phase II control charts. In each control chart, points joint by solid line indicate monitoring statistic values at each voyage, while dashed lines indicate upper control limit (UCL), at $\alpha = 0.05$.

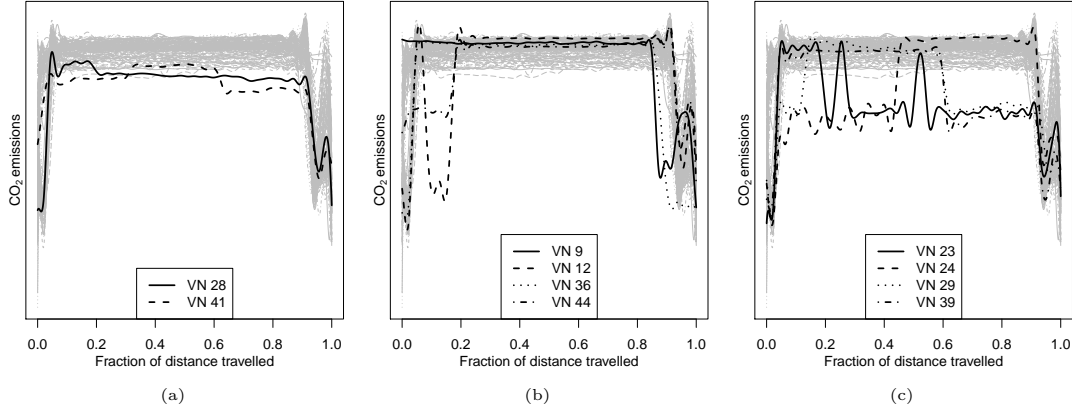


Figure 1.3. : OC CO₂ emission profiles (black lines) are superimposed on phase I reference ones (grey lines) and grouped by OC in (a) in T^2 control chart only, (b) SPE control chart only, (c) both T^2 and SPE control charts.

Therefore, it is crucial that the proposed control charting procedure can signal. In Figure 1.3b, VN 9, 12, 36, and 44 show that, during most of the voyage, the CO₂ emissions were not particularly different from the reference profiles in the training data set. However, some non-negligible slow down are highlighted in brief parts of these voyages. The most important shift occurs in fact at the beginning of VN 12, which shows the largest SPE in Figure 1.2b. The other voyages seem to postpone the acceleration phase at the beginning of the voyage, or to anticipate the slow down at the end of the voyage, then they show lower amounts of CO₂ emissions. More generally, these voyages show a different shape from standard Phase I profiles. Finally, In Figure 3c, with respect to the other voyages we discussed above, the CO₂ emissions for VN 23, 24, 29, and 39 show much larger deviations from the reference profiles in terms of both magnitude and shape. All these voyages have a lower-than-usual amount of CO₂ emissions for most of the voyage (deviation in magnitude), but also show sudden accelerations/slowdowns (deviation in shape). This is plausibly due to bad weather conditions, which has forced the ship to sail at an unusual navigation speed profile.

1.4 Conclusions

We presented a motivation-applied chapter to show benefits and practical applicability of the functional data analysis approach, with a very transparent set of steps. The two most important advantages in using functional data analysis over a discrete approach are the possibility (i) to analyze data theoretically defined over a continuum domain even when, over different replicates, the discrete observations are unequally spaced; (ii) to assume that the nearer the data points, the more similar, i.e., smoothness assumption. This is very reasonable in most of the practical cases and results in smooth parameters, such as eigenfunctions in functional principal component analysis, which supports model interpretability and fault diagnosis. The proposed functional control charting scheme is shown to be able to effectively monitor CO₂ emissions from real navigation data and to support the detection and interpretation of anomalous voyages. Different scenarios have in fact validated the capability of distinguishing the type of deviation associated to OC signals by the two proposed control

1. A FUNCTIONAL DATA ANALYSIS APPROACH FOR THE MONITORING OF SHIP CO₂ EMISSIONS

charts with respect to the reference profiles characterizing the normal operating conditions.

Chapter 2

Functional Regression Control Chart

Abstract

The modern development of data acquisition technologies in many industrial processes is facilitating the collection of quality characteristics that are apt to be modelled as functions, which are usually referred to as *profiles*. At the same time, measurements of concurrent variables, which are related to the quality characteristic profiles, are often available in a functional form as well, and usually referred to as *covariates*. In order to adjust the monitoring of the quality characteristic profiles by the effect of this additional information, a new functional control chart is elaborated on the residuals obtained from a function-on-function linear regression of the quality characteristic profile on the functional covariates. Furthermore, by means of a Monte Carlo simulation study, the performance of the proposed control chart are compared with those of other charts proposed in the literature. Eventually, a real-case study in the shipping industry is presented with the purpose of monitoring ship fuel consumption and thus, CO₂ emissions from a Ro-Pax ship, with particular regard to detecting CO₂ emission reduction after a specific energy efficiency initiative.

2.1 Introduction

In many industrial contexts, the development in data acquisition systems allow massive amounts of data to be recorded at high-rate and modelled as functions defined on multidimensional domains, i.e., *functional data* (Ramsay and Silverman, 2005; Ferraty and Vieu, 2006; Hsing and Eubank, 2015). In this scenario, new *statistical process control* (SPC) methods must be developed to monitor and control the stability over time of the quality characteristic, when functional data are available. In the classical SPC literature, functional data are more often referred to as *profiles* (Woodall et al., 2004). An overview of the main achievements in profile monitoring can be found in Noorossana et al. (2012). Other relevant contributions include the works of Jin and Shi (1999); Colosimo and Pacella (2010); Grasso et al. (2016, 2017); Menafiglio et al. (2018). As in the classical SPC (i.e., where data are scalars) profile monitoring control charts have the task of continuously monitoring the quality characteristic and of triggering a signal when assignable sources of variations (i.e. *special causes*) act on it. When this happens, the process is said to be out-of-control (OC). On the contrary, the process is said to be in-control (IC) when only normal sources of variation (i.e., *common causes*) apply.

In practice, there are situations where the quality characteristic is influenced by one or

more covariates. In this scenario, if one of these covariates manifests itself with an extreme realization, the quality characteristic may wrongly be judged to be OC. Otherwise, there may be situations where the covariates are not extreme and the quality characteristic may wrongly appear IC. As well as, the quality characteristic may wrongly appear IC, because the variance explained by the covariates is overlooked. In the multivariate SPC literature these issues have been addressed by means of the *regression control chart* (RCC) (Mandel, 1969). The basic idea behind this chart is to consider the quality characteristic after being adjusted for the effects of the covariates, that is monitoring the residuals of the regression of the quality characteristic on the covariates. Hawkins (1991) applied this idea to the multivariate setting by considering the regression of a variable on covariates that are assumed to be IC, and developed *Shewhart* and *cumulative sum* (CUSUM) control charts (Montgomery, 2007) based on the regression residuals. Hawkins (1993) applied the regression adjustment to particular kind of processes, said cascade processes in order to take advantage of the correlation between measures. Another application of the idea of regression adjustment appeared in Wade and Woodall (1993) where the so called cause-selecting control chart has been used for monitoring and control multistage processes. Other studies on the RCC include those of Shu et al. (2004), which studied the effects of the parameter estimation on the run length performance of the control chart, and those of Zhou and Goh (2016), where the influence of the regression model choice on the control chart performance was analysed.

In the literature of the RCC, the model used to describe the relation between the quality characteristic (hereinafter referred to also as *response variable*) and the covariates (hereinafter referred to also as *predictor variables*) is the linear regression model. In the functional context, functional linear regression models with one scalar response and one functional covariate have been deeply analysed in Cardot et al. (2003) and Hall et al. (2007); whereas, the extension to functional response was study by Ramsay and Silverman (2005) and Yao et al. (2005a). Functional linear models with functional response and multiple functional covariates have been far less studied. Matsui et al. (2009) developed estimation and evaluation methods based on regularized functional regression, whereas Fan et al. (2014) introduced the *functional response additive model estimation* (FRAME) and applied it to a case study in the online virtual stock markets. Chiou et al. (2016) proposed a functional liner model where both the response and the predictor variables are multivariate functional data, which relies on the *multivariate functional principal component* (MFPC) (or Karhunen–Loève) decomposition (Chiou et al., 2014; Happ and Greven, 2018).

In this chapter, we propose a new framework for monitoring a functional quality characteristic when functional covariates are available. This framework is henceforth referred to as *functional regression control chart* (FRCC) and can be regarded as an extension of the RCC to the functional context. In particular, we consider the case when the model which links the functional response and functional covariates is linear and, we monitor residuals by using the profile monitoring approach introduced by Woodall et al. (2004) and then used in Noorossana et al. (2012); Grasso et al. (2016); Pini et al. (2017), which is based on the simultaneous application of the *Hotelling's* T^2 and the squared prediction error (*SPE*) control charts.

A Monte Carlo simulation study is performed to quantify the FRCC *average run length* (*ARL*) (Montgomery, 2007), in identifying mean shifts in the functional response in presence or absence of drifts in the covariate means. This is done by comparing the proposed FRCC with other two control charts widely used in both the industrial context and the literature. In addition, a real-case study in the shipping industry is presented to illustrate the practical applicability of the proposed control chart. In particular, the FRCC is shown to adequately

identify reductions of cumulative fuel consumption, and thus CO₂ emissions (which are stoichiometrically related to it) after an energy efficiency initiative (EEI) was performed off-line on the considered Ro-Pax ship. In addition, a functional bootstrap procedure is developed and applied to evaluate uncertainty of the obtained results.

The chapter is structured as follows. Section 2.2 introduces the proposed FRCC. In Section 2.3 the performance of the FRCC is compared to that of other two popular control charts used for the same purpose. The real-case study is presented in Section 2.4. Section 2.5 concludes the chapter. All computations and plots have been obtained using the programming language R (R Core Team, 2020b).

2.2 The Functional Regression Control Chart Framework

The proposed FRCC can be regarded as a general framework for profile monitoring that can be divided into three main steps. Firstly, (i) define a functional regression model to be fitted

$$Y = g(\mathbf{X}) + \varepsilon, \quad (2.1)$$

where Y is the functional response variable and ε is a functional error term, both defined on the compact domain \mathcal{T} , g is a generic function of a vector \mathbf{X} of random functional covariates X_1, \dots, X_p , defined on the compact domain \mathcal{S} .

Secondly, (ii) define the estimation method of the chosen model, and, thirdly (iii) define the monitoring strategy of the functional residual defined as

$$e = Y - \hat{Y}, \quad (2.2)$$

where \hat{Y} is the fitted value of Y .

In what follows, after some preliminaries, we describe the FRCC when the following choices are made: (i) the *multivariate functional linear regression* (MFLR) model (Section 2.2) is set for the first step, (ii) an estimation method based on the Karhunen-Loève's decomposition (Section 2.2) is chosen for the second, and (iii) the *Hotelling's T²* and the *squared prediction error (SPE)* control charts (Section 2.2) are built in the third step.

Preliminaries

Assume that X_1, \dots, X_p and Y have smooth realizations in $L^2(\mathcal{S})$ and $L^2(\mathcal{T})$, i.e., the Hilbert spaces of square integrable functions defined on the compact sets \mathcal{S} and \mathcal{T} , respectively. Moreover, let us denote with $\mathbb{H}^X = (L^2(\mathcal{S}))^p$ the Hilbert space whose elements are vectors of functions in $L^2(\mathcal{S})$. Then, $\mathbf{X} = (X_1, \dots, X_p)^T$ is random vector of functions whose realizations are in \mathbb{H}^X . Accordingly, for a compact set \mathcal{Z} , the inner product of two functions f and g in $L^2(\mathcal{Z})$ is $\langle f, g \rangle = \int_{\mathcal{Z}} f(z)g(z)dz$, with dz the Lebesgue measure on \mathcal{Z} , and the norm is $\|\cdot\| = \sqrt{\langle \cdot, \cdot \rangle}$. The inner product of two function vectors $\mathbf{f} = (f_1, \dots, f_p)^T$ and $\mathbf{g} = (g_1, \dots, g_p)^T$ in \mathbb{H}^X is $\langle \mathbf{f}, \mathbf{g} \rangle_{\mathbb{H}^X} = \sum_{i=1}^p \langle f_i, g_i \rangle$ and the norm is $\|\cdot\|_{\mathbb{H}^X} = \sqrt{\langle \cdot, \cdot \rangle_{\mathbb{H}^X}}$. Further, assume that \mathbf{X} has mean function $\boldsymbol{\mu}^X = (\mu_1^X, \dots, \mu_p^X)^T$, with $\mu_i^X = E(X_i)$ and covariance function $\mathbf{C}^X = \{C_{i,j}^X\}_{1 \leq i,j \leq p}$, with $C_{i,j}^X(s_1, s_2) = \text{Cov}(X_i(s_1), X_j(s_2))$, for $s_1, s_2 \in \mathcal{S}$. Analogously, let $\mu^Y = E(Y)$ and $C^Y(t_1, t_2) = \text{Cov}(Y(t_1), Y(t_2))$, for $t_1, t_2 \in \mathcal{T}$, be the mean and the covariance function of the response variable $Y(t)$, respectively.

The transformation approach of Chiou et al. (2014) is here used, as covariates can exhibit different amount of variation. In what follows, all the operations between functions have

to be considered pointwise. Let $\mathbf{X}^{std} = (X_1^{std}, \dots, X_p^{std})^T = (\mathbf{V}^X)^{-1}(\mathbf{X} - \boldsymbol{\mu}^X)$, be the vector of the standardized covariates, with the matrix $\mathbf{V}^X = \text{diag}\left((v_1^X)^{1/2}, \dots, (v_p^X)^{1/2}\right)$ where $v_i^X(s) = C_{i,i}^X(s, s)$, for $s \in \mathcal{S}$. The response variable Y is also standardized as $Y^{std} = (v^Y)^{-1/2}(Y - \mu^Y)$, with $v^Y(t) = C^Y(t, t)$, for $t \in \mathcal{T}$. Let \mathbf{C}_{std}^X and C_{std}^Y be the covariance functions of the standardized covariate and response variables, respectively. Let us consider for \mathbf{C}_{std}^X the expansion $\mathbf{C}_{std}^X(s_1, s_2) = \sum_{i=1}^{\infty} \lambda_i^X \boldsymbol{\psi}_i^X(s_1) \boldsymbol{\psi}_i^X(s_2)^T$, for $s_1, s_2 \in \mathcal{S}$, where $\{\boldsymbol{\psi}_i^X\}$ are the orthonormal (i.e., $\langle \boldsymbol{\psi}_i^X, \boldsymbol{\psi}_j^X \rangle_{\mathbb{H}^X} = \delta_{ij}$, with δ_{ij} the Kronecker delta) multivariate eigenfunctions of \mathbf{C}_{std}^X corresponding to the eigenvalues $\{\lambda_i^X\}$ in descending order. Similarly, we consider for C_{std}^Y the expansion $C_{std}^Y(t_1, t_2) = \sum_{i=1}^{\infty} \lambda_i^Y \psi_i^Y(t_1) \psi_i^Y(t_2)$, for $t_1, t_2 \in \mathcal{T}$ where $\{\lambda_i^Y\}$ and $\{\psi_i^Y\}$ are defined in the same way. Note that both expansions are well-defined in virtue of the multivariate and univariate versions of the Mercer's Theorem (Happ and Greven, 2018), respectively.

The Model

For the MFLR model we assume that the covariates \mathbf{X}^{std} linearly influence the response Y^{std} as follows

$$Y^{std}(t) = \int_{\mathcal{S}} (\boldsymbol{\beta}(s, t))^T \mathbf{X}^{std}(s) ds + \varepsilon(t) \quad t \in \mathcal{T}, \quad (2.3)$$

that is a particular version of Equation (2.1). The regression coefficient vector $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$, is in $(L^2(\mathcal{S} \times \mathcal{T}))^p$, whose elements are vectors of bivariate functions in $L^2(\mathcal{S} \times \mathcal{T})$ (i.e., the space of square integrable function on the closed interval $\mathcal{S} \times \mathcal{T}$), and the random error function ε has $E(\varepsilon) = 0$ and $\text{Var}(\varepsilon) = v_\varepsilon^2$, and is independent of \mathbf{X}^{std} . Thus, the regression function is

$$E(Y^{std}(t) | \mathbf{X}^{std}) = \int_{\mathcal{S}} (\boldsymbol{\beta}(s, t))^T \mathbf{X}^{std}(s) ds \quad t \in \mathcal{T}. \quad (2.4)$$

The Estimation Method

From the multivariate and univariate Karhunen-Loève's Theorem (Happ and Greven, 2018), standardized covariate and response variables can be represented as follows

$$\mathbf{X}^{std} = \sum_{i=1}^{\infty} \xi_i^X \boldsymbol{\psi}_i^X \quad Y^{std} = \sum_{i=1}^{\infty} \xi_i^Y \psi_i^Y, \quad (2.5)$$

where $\xi_i^X = \langle \mathbf{X}^{std}, \boldsymbol{\psi}_i^X \rangle_{\mathbb{H}^X}$ and $\xi_i^Y = \langle Y^{std}, \psi_i^Y \rangle$ are random variables, said *principal component scores* or simply *scores*, such that $E(\xi_i^X) = 0$, $E(\xi_i^X \xi_j^X) = \lambda_i^X \delta_{ij}$ and $E(\xi_i^Y) = 0$, $E(\xi_i^Y \xi_j^Y) = \lambda_i^Y \delta_{ij}$, respectively. In this context, the eigenfunctions $\{\boldsymbol{\psi}_i^X\}$ and $\{\psi_i^Y\}$ (as defined in the preliminaries) are referred to as *principal components* as well. As demonstrated in Chiou et al. (2016), the regression coefficient vector is as follows

$$\boldsymbol{\beta}(s, t) = \sum_{i,j=1}^{\infty} \frac{E(\xi_i^X \xi_j^Y)}{\lambda_i^X} \boldsymbol{\psi}_i^X(s) \psi_j^Y(t) \quad s \in \mathcal{S}, t \in \mathcal{T}, \quad (2.6)$$

which is also the minimizer of the expected squared $L^2(\mathcal{T})$ distance between Y^{std} and $\int_{\mathcal{S}} (\mathbf{f}(s, \cdot))^T \mathbf{X}^{std}(s) ds$, with $\mathbf{f} \in (L^2(\mathcal{S} \times \mathcal{T}))^p$, i.e.

$$\beta = \underset{\mathbf{f} \in (L^2(\mathcal{S} \times \mathcal{T}))^p}{\operatorname{argmin}} \mathbb{E} \|Y^{std} - \int_{\mathcal{S}} (\mathbf{f}(s, \cdot))^T \mathbf{X}^{std}(s) ds\|^2. \quad (2.7)$$

By plugging Equation (2.5) and Equation (2.6) into Equation (2.4), and using the orthonormality of $\{\psi_i^X\}$, we obtain

$$\mathbb{E}(Y^{std}(t) | \mathbf{X}^{std}) = \sum_{i,j=1}^{\infty} b_{ij} \xi_i^X \psi_j^Y(t) \quad t \in \mathcal{T}, \quad (2.8)$$

where $b_{ij} = \mathbb{E}(\xi_i^X \xi_j^Y) / \lambda_i^X$. Therefore, the best least squares predictor of Y given \mathbf{X} is

$$\mathbb{E}(Y(t) | \mathbf{X}) = \mu^Y(t) + v^Y(t)^{1/2} \mathbb{E}(Y^{std}(t) | \mathbf{X}^{std}) \quad t \in \mathcal{T}. \quad (2.9)$$

An estimation method of the above unknown quantities is described in the Supplementary Materials. Broadly speaking, it is based on the truncated versions of Equation (2.5), namely

$$\mathbf{X}_L^{std} = \sum_{i=1}^L \xi_i^X \psi_i^X \quad Y_M^{std} = \sum_{i=1}^M \xi_i^Y \psi_i^Y, \quad (2.10)$$

where the number of retained scores L and M are chosen such that they explain at least given proportions δ_X and δ_Y of total variation respectively (Ramsay and Silverman, 2005). The estimation method provides estimators $\hat{\beta}_{LM}$ of β , in Equation (2.3), and \hat{Y}_{LM}^{std} of $\mathbb{E}(Y^{std} | \mathbf{X}^{std})$, in Equation (2.8), through the estimators \hat{C}_{std}^X and \hat{C}_{std}^Y of C_{std}^X and C_{std}^Y calculated using the sample mean and covariance functions (Hsing and Eubank, 2015).

The Monitoring Strategy

Upon using the estimator \hat{Y}_{LM}^{std} , the functional residual in Equation (2.2) particularizes as

$$e^{std} = Y^{std} - \hat{Y}_{LM}^{std}. \quad (2.11)$$

To monitor the residuals, we follow the strategy of Woodall et al. (2004); Noorossana et al. (2012); Grasso et al. (2016); Pini et al. (2017). In particular, the *Hotelling's* T^2 and the *SPE* control charts are applied on the coefficients obtained from the univariate functional principal component decomposition (Hsing and Eubank, 2015) of e^{std} , i.e.,

$$e^{std} = \sum_{i=1}^{\infty} \xi_i^e \psi_i^e, \quad (2.12)$$

where the scores $\xi_i^e = \langle e^{std}, \psi_i^e \rangle$ and the principal components $\{\psi_i^e\}$ are the eigenfunctions corresponding to the eigenvalues $\{\lambda_i^e\}$ in descending order of the covariance function $C^e(t_1, t_2) = \operatorname{Cov}(e^{std}(t_1), e^{std}(t_2))$, for $t_1, t_2 \in \mathcal{T}$. As a matter of fact, C^e is different from C_{std}^Y (and thus $\{\psi_i^Y\}$ from $\{\psi_i^e\}$), because the former refers to the distribution of

Y^{std} , whereas the latter is related to the conditional distribution of Y^{std} given \mathbf{X}^{std} . A straightforward approximation of e^{std} can be thus obtained as

$$e_K^{std} = \sum_{i=1}^K \xi_i^e \psi_i^e, \quad (2.13)$$

where K is the number of retained scores.

The *Hotelling's* statistic T^2 can be then particularized as follows

$$T_{e^{std}}^2 = \boldsymbol{\xi}^e \boldsymbol{\Sigma}_{\boldsymbol{\xi}^e}^{-1} \boldsymbol{\xi}^e, \quad (2.14)$$

where $\boldsymbol{\Sigma}_{\boldsymbol{\xi}^e} = \text{diag}(\lambda_1^e, \dots, \lambda_K^e)$ is the variance-covariance matrix of $\boldsymbol{\xi}^e = (\xi_1^e, \dots, \xi_K^e)^T$. Note that $T_{e^{std}}^2$ is the standardized square distance of the projection of e^{std} from the origin of the space spanned by the principal components $\{\psi_i^e\}$. Analogously, changes along directions orthogonal to the latter space are monitored by the statistic

$$SPE_{e^{std}} = \int_{\mathcal{T}} \left(e^{std}(t) - e_K^{std}(t) \right)^2 dt. \quad (2.15)$$

The design phase of the control charts (Phase I) can be performed by means of a set of n functional residuals e_i^{std} , $i = 1, \dots, n$, obtained by n independent observations (\mathbf{X}_i, Y_i) acquired under IC conditions. This phase involves the estimation of the MFLR model unknown parameters (Supplementary Materials), of the principal components $\{\psi_i^e\}$ and of the matrix $\boldsymbol{\Sigma}_{\boldsymbol{\xi}^e}$ (calculated by means of the sample covariance) as well as the estimation of the control limits for both the *Hotelling's* T^2 and the SPE control charts, which can be obtained by means of $(1 - \alpha^*)$ quantiles of the empirical distribution of the two statistics. Note that, to control the *family wise error rate* (FWER) in the strong sense (Lehmann and Romano, 2006), α^* is chosen by using the Šidák correction (Lehmann and Romano, 2006) $\alpha^* = 1 - (1 - \alpha)^{1/2}$, where α is the overall Type I error. In the monitoring phase (Phase II), functional residuals of a new observation (\mathbf{X}^*, Y^*) are calculated and an alarm signal is issued if at least one realization of the $T_{e^{std}}^2$ and $SPE_{e^{std}}$ statistics violates the control limits.

2.3 Performance Analysis

Data Generation

The overall performance of the proposed FRCC are evaluated by means of a Monte Carlo simulation. Profile patterns have been generated with signal and correlation structures similar to those in the real-case study presented in Section 2.4. Details about the data generation process are provided in the Supplementary Materials. The compact domains \mathcal{S} and \mathcal{T} are set, without loss of generality, equal to $[0, 1]$ and the number of covariates p is set equal to 3.

Moreover, in this section, the data are generated with $R^2 = \int_{[0,1]} \frac{\text{Var}(\mathbb{E}(Y^{std}(t)|\mathbf{X}^{std}))}{\text{Var}(Y^{std}(t))} dt$ set equal to 0.97 (Yao et al., 2005b). Additional analysis at different values of R^2 are provided in the Supplementary Materials.

The mean functions μ^X and μ^Y of the generated data are obtained through the following reference model

$$\mu(z) = P(z) + r \sum_{i=1}^I g_i(z; m_i, s_i) \quad z \in [0, 1], \quad (2.16)$$

Table 2.1. Different types of shift in μ^X and μ^Y .

Shift	δ_a	δ_b	δ_c
A	d	0	0
B	0	d	0
C	0	0	d
D	d	d	0

where

$$P(z) = az^2 + bz + c \quad z \in [0, 1], \quad (2.17)$$

a, b, c are real numbers, and the terms $g_i(z; m_i, s_i)$ are normal probability density functions having parameters m_i and s_i with values given in the Supplementary Materials. The right side term of Equation (2.16) is inspired by the data generation process proposed in Pini et al. (2017); Grasso et al. (2017).

The aim of the simulation is to assess the FRCC performance in identifying mean function shifts in the response in presence of

1. mean function shifts in Y conditional on \mathbf{X} , i.e. $E(Y|\mathbf{X})$, resulting from changes in μ^Y ;
2. mean function shifts in \mathbf{X} , i.e., $E(\mathbf{X}) = \mu^X$, and $E(Y|\mathbf{X})$.

The types of shift are consistent with those of Shu et al. (2004) and Wade and Woodall (1993). Note from Equations (2.4) and (2.9) that shifts in $E(Y|\mathbf{X})$ can result from changes in μ^Y and β . However, the latter, in addition, can affect variability of the functional regression residuals as well. Because we are interested in the FRCC performance in identifying mean function shifts in the response, given that the variability of the residuals are assumed constant, then, only shifts caused by changes in μ^Y are considered.

The functional patterns with shift in the mean function are generated using the model in Equation (2.16) with P defined as follows

$$P(z) = (a + \delta_a)z^2 + (b + \delta_b)z + (c + \delta_c) \quad z \in [0, 1], \quad (2.18)$$

where the real number δ_a, δ_b , and δ_c define the shift type. Without loss of generality δ_a, δ_b , and δ_c are set equal to a positive severity level d as reported in Table 2.1 where four different types of shift (namely A,B,C,D) in μ^X and μ^Y are considered. Shift A is representative of a change in the mean function curvature, whereas shift B and C represent slope modification and translation of the profile pattern, respectively. Shift D consists of both curvature and slope modifications of the mean function. These types of shift are indeed consistent with the real-case study of Section 2.4 and apt to model usual ways as ship performance increase or decrease in reality.

Simulation Results and Discussion

Three different profile monitoring methods are compared: (a) monitoring residuals by means of the FRCC, (b) monitoring model coefficients of the response variable Y via a *Hotelling's* T^2 and *SPE* control charts (hereinafter denoted by RESP -RESPonse- control chart), and

2. FUNCTIONAL REGRESSION CONTROL CHART

Table 2.2. Severity levels associated to each type of shift in μ^Y for Scenario 1 (a) and for Scenario 2 (b).

Shift			Severity			Shift			Severity		
μ^Y	A	$d \in \{0.5, 1.0, 1.5, 2.0\}$		μ^Y	A	$d \in \{0.5, 1.0, 1.5, 2.0\}$		μ_1^X μ_2^X μ_3^X	D	$d \in \{0.5\}$	
	B	$d \in \{0.5, 1.0, 1.5, 2.0\}$				$d \in \{0.5, 1.0, 1.5, 2.0\}$			A	$d \in \{0.5\}$	
	C	$d \in \{0.5, 1.0, 1.5, 2.0\}$				$d \in \{0.5, 1.0, 1.5, 2.0\}$			D	$d \in \{1.0\}$	
	D	$d \in \{0.5, 1.0, 1.5, 2.0\}$				$d \in \{0.5, 1.0, 1.5, 2.0\}$					

(a)

(b)

(a)

(b)

(c) monitoring the area under the response variable Y considered also in Pini et al. (2017) (hereinafter denoted by INBA -INdex BAsed- control chart).

Performance analysis of the FRCC is carried out by considering shifts in the conditional mean $E(Y|\mathbf{X})$, firstly, by means of changes in μ^Y only (Scenario 1), and secondly by means of changes in both μ^Y and μ^X (Scenario 2). The first scenario aims to analyse FRCC performance in absence of shift in the regressor mean μ^X ; whereas, the second aims to study the unwanted influence of shifts in μ^X on the FRCC performance. Severity levels and types of shift considered by Scenario 1 and Scenario 2 are listed in Table 2.2a and Table 2.2b, respectively. Note that in Scenario 2 we consider only shift type A (with the same severity levels of Scenario 1) for μ^Y , shift type D for μ_1^X , μ_3^X , and shift type A for μ_2^X . The latter three shifts are explored at only one severity level d as reported in Table 2.2b. For each shift type and severity level for Scenario 1 and for each severity level combination for Scenario 2, 100 simulation runs were performed. Each run consists of the following steps:

Phase I) A design set of 4000 IC patterns is randomly generated. In particular, $N_1 = 1000$ patterns from the training set are used to estimate MFLR model unknown quantities along with mean and covariance functions. The remaining $N_2 = 3000$ IC profiles are used as tuning set to estimate the empirical quantiles via the kernel density estimation (KDE) approach (Chou et al., 2001) with gaussian kernel, 2000 equally spaced points and bandwidth chosen by means of the Silverman's rule of thumb (Silverman, 1986). The number of retained scores L , M and K in Equation (2.10) and Equation (2.13) are chosen such that the retained principal components explain at least 95% of the total variability.

Phase II) A testing set of further 4000 OC patterns is randomly generated to carry out the monitoring phase and to evaluate the chart performance.

As is usually done in the literature (Montgomery, 2007), the FRCC and the competitor chart performance are compared by means of the average run length (ARL), that is referred to as ARL_0 in the case of no response mean shift ($d = 0$ for μ^Y), and as ARL_1 otherwise. For the sake of simplicity, we set $ARL_0 = 100$ and denote indistinctly by \widehat{ARL} the estimated ARL (regardless whether it is referred to ARL_0 or ARL_1).

For Scenario 1, Table 2.3 shows the \widehat{ARL} s along with 95% approximate confidence intervals based on the Student's t approximation. Graphical representation of the latter are in Figure 2.1, which shows that the FRCC outperforms the RESP and INBA control charts for all the considered shifts. The gain in efficiency is less evident for Shift C (i.e., in presence of translations of the profile pattern, only) at high severity level ($d = 2$). Whereas, in Scenario 2, Table 2.4 and Figure 2.2 point out that shifts in covariate mean functions strongly impact the ARL of the FRCC. Indeed, when the response variable is IC ($d = 0$), \widehat{ARL} s for the FRCC are usually lower than $ARL_0 = 100$. Table 2.4 shows that this issue

2.3. Performance Analysis

Table 2.3. Estimated ARLs (\widehat{ARL} s) and 95% confidence intervals (CI) for Scenario 1.

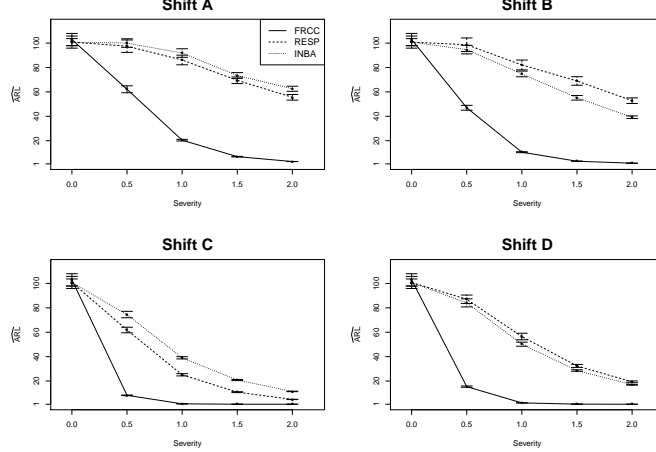
Shift	Severity	FRCC		RESP		INBA	
	d	\widehat{ARL}	CI	\widehat{ARL}	CI	\widehat{ARL}	CI
In-control	-	102.94	[97.95, 107.94]	100.81	[95.86, 105.76]	100.74	[97.81, 103.68]
	0.5	62.14	[59.29, 64.98]	97.47	[92.37, 102.57]	100.03	[96.52, 103.53]
	1.0	20.43	[19.76, 21.10]	86.17	[82.21, 90.14]	91.85	[88.36, 95.34]
	1.5	6.94	[6.73, 7.14]	69.50	[66.80, 72.20]	73.44	[71.07, 75.80]
	2.0	2.93	[2.87, 2.99]	55.57	[53.26, 57.88]	62.59	[60.51, 64.67]
A	0.5	47.00	[45.02, 48.99]	98.65	[93.08, 104.22]	94.68	[91.16, 98.21]
	1.0	10.58	[10.18, 10.98]	82.10	[78.07, 86.13]	74.71	[72.38, 77.04]
	1.5	3.19	[3.11, 3.27]	68.97	[65.48, 72.45]	55.19	[53.31, 57.07]
	2.0	1.56	[1.54, 1.58]	52.78	[50.49, 55.08]	39.25	[38.18, 40.31]
	0.5	8.30	[8.07, 8.54]	61.82	[59.56, 64.09]	74.51	[71.90, 77.12]
B	1.0	1.33	[1.32, 1.34]	25.13	[24.20, 26.06]	39.03	[38.12, 39.95]
	1.5	1.00	[1.00, 1.00]	10.89	[10.52, 11.26]	20.82	[20.36, 21.27]
	2.0	1.00	[1.00, 1.00]	4.93	[4.81, 5.06]	11.48	[11.29, 11.67]
	0.5	15.42	[14.89, 15.94]	87.00	[83.54, 90.47]	84.08	[80.77, 87.39]
	1.0	2.09	[2.06, 2.12]	56.50	[53.82, 59.18]	50.19	[48.53, 51.86]
C	1.5	1.07	[1.07, 1.08]	32.38	[31.16, 33.60]	28.61	[27.99, 29.23]
	2.0	1.00	[1.00, 1.00]	19.37	[18.65, 20.08]	16.93	[16.64, 17.23]

Table 2.4. Estimated ARL s (\widehat{ARL}) and 95% confidence intervals (CI) for the FRCC in Scenario 2 at different severity levels d of shifts in the response mean (μ^Y) (Table 2.2b) as a function of which and how many covariates are subject to mean shift (each at the severity level reported in Table 2.2b). Zeros and ones in the triplets (000, 100, 010, 001, 110, 101, 011, 111) indicate IC and OC covariates, respectively. For instance, the triplet 100 means that only the first covariate is OC, 111 means that all the covariates are OC, and so on.

μ^Y	d	shifted covariate combination	FRCC		RESP		INBA	
			\widehat{ARL}	CI	\widehat{ARL}	CI	\widehat{ARL}	CI
0	0	0 0 0	99.79	[95.42, 104.17]	105.07	[99.21, 110.93]	104.19	[100.32, 108.05]
	0	1 0 0	78.09	[74.65, 81.53]	89.28	[85.67, 92.90]	100.1	[95.73, 104.48]
	0	0 1 0	100.61	[95.74, 105.49]	88.02	[84.26, 91.77]	79.51	[76.60, 82.41]
	0	0 0 1	100.59	[96.22, 104.97]	99.58	[94.68, 104.48]	102.34	[98.62, 106.05]
	0	1 1 0	73.43	[70.03, 76.84]	66.59	[63.48, 69.71]	93.63	[90.52, 96.75]
	0	1 0 1	64.73	[62.41, 67.05]	96.74	[92.93, 100.56]	99.88	[96.30, 103.45]
	0	0 1 1	97.21	[92.56, 101.86]	89.87	[85.42, 94.33]	80.11	[77.25, 82.98]
	0	1 1 1	63.26	[60.72, 65.8]	80.75	[76.73, 84.77]	92.93	[89.57, 96.29]
0.5	0	0 0 0	63.65	[60.80, 66.50]	105.16	[99.96, 110.37]	101.95	[98.24, 105.66]
	0	1 0 0	44.65	[43.18, 46.11]	85.52	[81.62, 89.41]	85.17	[82.11, 88.24]
	0	0 1 0	62.50	[59.69, 65.31]	96.76	[91.19, 102.32]	98.52	[95.03, 102.00]
	0	0 0 1	57.61	[55.37, 59.84]	97.09	[92.40, 101.79]	100.44	[96.94, 103.93]
	0	1 1 0	45.68	[43.87, 47.49]	77.36	[73.17, 81.56]	100.75	[96.94, 104.56]
	0	1 0 1	39.19	[37.58, 40.80]	94.04	[89.26, 98.82]	89.21	[86.48, 91.94]
	0	0 1 1	59.47	[56.76, 62.17]	96.78	[92.13, 101.42]	91.13	[87.46, 94.80]
	0	1 1 1	38.89	[37.34, 40.44]	84.76	[80.41, 89.11]	102.87	[98.66, 107.08]
1.0	0	0 0 0	21.08	[20.26, 21.90]	85.43	[81.25, 89.61]	89.22	[86.16, 92.28]
	0	1 0 0	17.23	[16.73, 17.74]	77.09	[73.45, 80.73]	69.95	[67.71, 72.19]
	0	0 1 0	21.18	[20.32, 22.04]	99.23	[94.10, 104.36]	102.16	[98.34, 105.97]
	0	0 0 1	20.28	[19.45, 21.11]	80.57	[76.86, 84.28]	91.30	[88.00, 94.60]
	0	1 1 0	17.68	[17.08, 18.27]	80.68	[76.93, 84.43]	99.91	[96.41, 103.41]
	0	1 0 1	16.41	[15.89, 16.93]	81.56	[77.90, 85.22]	77.33	[74.69, 79.97]
	0	0 1 1	20.64	[19.82, 21.45]	96.69	[92.09, 101.29]	98.09	[94.38, 101.81]
	0	1 1 1	16.53	[15.92, 17.14]	89.58	[84.46, 94.71]	100.85	[96.47, 105.22]
1.5	0	0 0 0	6.90	[6.70, 7.09]	69.54	[66.41, 72.67]	76.22	[73.66, 78.78]
	0	1 0 0	6.31	[6.14, 6.48]	68.29	[64.69, 71.89]	59.58	[57.74, 61.42]
	0	0 1 0	7.15	[6.91, 7.40]	91.92	[87.36, 96.47]	102.86	[98.68, 107.03]
	0	0 0 1	6.78	[6.58, 6.99]	66.91	[64.08, 69.75]	78.76	[76.14, 81.37]
	0	1 1 0	6.69	[6.50, 6.87]	76.35	[72.85, 79.84]	92.75	[89.30, 96.20]
	0	1 0 1	6.05	[5.89, 6.22]	66.45	[63.77, 69.12]	61.33	[59.57, 63.10]
	0	0 1 1	7.04	[6.83, 7.25]	86.49	[82.15, 90.84]	103.80	[99.44, 108.16]
	0	1 1 1	6.32	[6.12, 6.53]	80.95	[77.09, 84.82]	96.19	[92.72, 99.66]
2.0	0	0 0 0	2.92	[2.86, 2.98]	54.72	[52.55, 56.89]	63.70	[61.61, 65.80]
	0	1 0 0	2.83	[2.77, 2.89]	53.80	[51.71, 55.90]	48.41	[47.12, 49.71]
	0	0 1 0	3.01	[2.94, 3.08]	73.53	[70.12, 76.94]	96.58	[92.72, 100.43]
	0	0 0 1	2.90	[2.84, 2.96]	55.65	[53.00, 58.30]	66.34	[64.11, 68.57]
	0	1 1 0	2.86	[2.80, 2.92]	66.80	[63.87, 69.73]	80.97	[78.51, 83.43]
	0	1 0 1	2.76	[2.71, 2.81]	53.47	[51.42, 55.52]	50.93	[49.48, 52.39]
	0	0 1 1	2.94	[2.88, 3.00]	76.18	[72.50, 79.86]	99.79	[95.80, 103.77]
	0	1 1 1	2.77	[2.71, 2.83]	68.11	[64.88, 71.35]	82.87	[80.14, 85.61]

2. FUNCTIONAL REGRESSION CONTROL CHART

Figure 2.1. Estimated ARLs (\widehat{ARL}) and 95% confidence intervals for different response mean shifts for Scenario 1 (Table 2.3).



occurs to both the RESP and INBA control charts as well. For the latter charts this is expected because they do not account for the correlation between covariates and the response variable. On the contrary, this is not intuitive for the FRCC because it is expected not to be sensitive to unusual covariate realizations. However, this is completely consistent with what happens in the multivariate case (Shu et al., 2004). Indeed, let us denote with Δ^Y , Δ^{X_1} , Δ^{X_2} and Δ^{X_3} the shift size of Y , X_1 , X_2 and X_3 means, respectively. Moreover, Y^{std} and \mathbf{X}^{std} (resp. Y_{Δ}^{std} and $\mathbf{X}_{\Delta}^{std}$) indicate the standardized response and predictors in absence (resp. presence) of shift. Then, when mean shifts in both response variable and covariates occur, the functional residual (Equation (2.11)) can be rewritten, for $t \in [0, 1]$, as

$$\begin{aligned}
 e_{\Delta}^{std}(t) &= Y_{\Delta}^{std}(t) - \hat{Y}_{LM\Delta}^{std}(t) \\
 &= Y^{std}(t) + \int_0^1 (\beta(s, t))^T \mathbf{V}^X(s)^{-1} \Delta^{\mathbf{X}}(s) ds + \frac{\Delta^Y(t)}{v^Y(t)^{1/2}} - \hat{Y}_{LM\Delta}^{std}(t) \\
 &\quad - \int_0^1 (\hat{\beta}_{LM}(s, t))^T \mathbf{V}^X(s)^{-1} \Delta^{\mathbf{X}}(s) ds \\
 &= e^{std}(t) + \frac{\Delta^Y(t)}{v^Y(t)^{1/2}} - \int_0^1 (\beta(s, t) - \hat{\beta}_{LM}(s, t))^T \mathbf{V}^X(s)^{-1} \Delta^{\mathbf{X}}(s) ds, \quad (2.19)
 \end{aligned}$$

where e^{std} is as in Equation (2.11), $\hat{\beta}_{LM}$ is the estimator of β given in the Supplementary Materials and \mathbf{V}^X is defined in Section 2.2 with $p = 3$, $\Delta^{\mathbf{X}} = (\Delta^{X_1}, \Delta^{X_2}, \Delta^{X_3})^T$ and

$$\hat{Y}_{LM\Delta}^{std}(t) = \int_0^1 (\hat{\beta}_{LM}(s, t))^T \mathbf{X}_{\Delta}^{std}(s) ds \quad t \in [0, 1]. \quad (2.20)$$

Then expected value of e_{Δ}^{std} conditioned on $\hat{\beta}_{LM}$ is


$$\mathbb{E}(e_{\Delta}^{std}(t)) = \frac{\Delta^Y(t)}{v^Y(t)^{1/2}} - \int_0^1 (\beta(s, t) - \hat{\beta}_{LM}(s, t))^T \mathbf{V}^X(s)^{-1} \Delta^{\mathbf{X}}(s) ds \quad t \in [0, 1]. \quad (2.21)$$

Table 2.5. Estimation error magnitude of the entries of the coefficient vector β

	X_1^{std}	X_2^{std}	X_3^{std}
Error	0.086	0.009	0.010

From Equation (2.19), it is clear that a mean shift in the residual, caused by covariate mean shifts ($\Delta^X \neq 0$), occurs when the difference $\beta - \hat{\beta}_{LM}$ is not negligible, even though the response variable is IC, i.e., $\Delta^Y = 0$.

Figure 2.2. Estimated $ARLs$ (\widehat{ARL} s) and 95% confidence intervals for the FRCC in Scenario 2 at different severity levels $d = \{0, 0.5, 1\}$ of shifts in the response mean (μ^Y) (Table 2.2b) as a function of which and how many covariates are subject to mean shift (each at the severity level reported in Table 2.2b). Zeros and ones in the triplets (000, 100, 010, 001, 110, 101, 011, 111) indicate IC and OC covariates, respectively. For instance, the triplet 100 means that only the first covariate is OC, 111 means that all the covariates are OC, and so on. The severity levels $d = \{1.5, 2\}$ are not reported because the ARL s are all very small. The Figure depicts part of the results reported in Table 2.4.



From Figure 2.2 and Table 2.4, it is clear that the magnitude of the deviations depends both on the number of covariates with shift in the mean function and the estimation error magnitude of the entries of the coefficient vector β in Table 2.5. The latter is measured as the L_2 distance between β and $\hat{\beta}_{LM}$, estimated by means of 50 random realization of $\hat{\beta}_{LM}$. Therefore, shifts in μ_1^X affect the performance of the FRCC stronger than mean shifts in the other covariates (viz. μ_2^X and μ_3^X). Obviously, this effect is more evident when the term $\frac{\Delta^Y}{(v^Y)^{1/2}}$ in Equation (2.19) is zero and do not cover the contribution to e_{Δ}^{std} of $\beta - \hat{\beta}_{LM}$. The problem of issuing an alarm only when a mean shift occurs in the response variable regardless of covariate mean shifts is addressed in Section 2.3 and solutions are proposed in this respect.

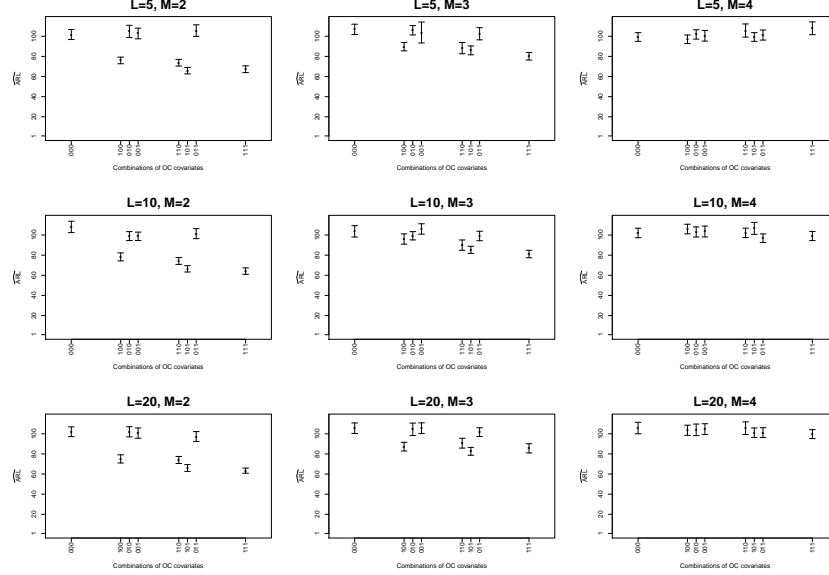
Remarks on the Use of the FRCC in Presence of Covariate Mean shifts

As stated before, the FRCC performance in identifying OC condition (of the response variable) can be affected by the number of covariates with shift in the mean function and the estimation error of the coefficient vector β (Table 2.4). In particular, when mean shifts occur in the covariates only, the interpretation of the FRCC becomes cumbersome, because a point falling outside the FRCC control limits is wrongly assigned to a shift in the response variable. In this section, we propose some solutions to enhance the FRCC performance in presence of covariate mean shifts.

As pointed out in Scenario 1 of the simulation study (Table 2.2a), which assumes no covariate mean shift, the FRCC performance in identifying response mean shifts is always better than that achieved by the competitor control charts (Figure 2.1). A straightforward solution is to verify the assumption that no covariate mean shift occurs by extending to the functional setting the control charts proposed by Wade and Woodall (1993) in the

2. FUNCTIONAL REGRESSION CONTROL CHART

Figure 2.3. Estimated \widehat{ARL} s (\widehat{ARL}) and 95% confidence intervals for the FRCC in Scenario 2 no shift ($d = 0$) in the response mean (μ^Y) at different values of truncation parameters L (increasing from top to bottom panel) and M (increasing from left to right panel) as a function of which and how many covariates are OC (each at the severity level reported in Table 2.2b). Zeros and ones in the triplets (000, 100, 010, 001, 110, 101, 011, 111) indicate IC and OC covariates, respectively. For instance, the triplet 100 means that the first covariate mean has shifted, 111 means that all the covariate mean are shifted, and so on.



multivariate case. In this regard, Capezza et al. (2019) suggest to monitor the covariates \mathbf{X} through the jointly use of the *Hotelling's* T^2 and *SPE* control charts built on the principal component decomposition of \mathbf{X}^{std} .

When the number of observations available in Phase I, N_1 , is large, an alternative solution, which does not the construction of additional control charts, can be based on the consistency of the estimator $\hat{\beta}_{LM}$ of the coefficient vector β (Chiou et al., 2016). Indeed, when truncation parameters L and M in Equation (2.10) go to infinity with N_1 , the impact of $\Delta^{\mathbf{X}}$ in Equation (2.19) and (2.21) fades out, as in (Yao et al., 2005b), even though they do not provide more detailed indications on the convergence rate. That is

$$\lim_{N_1 \rightarrow \infty} \int_S \int_{\mathcal{T}} \left[\beta(s, t) - \hat{\beta}_{LM}(s, t) \right]^2 ds dt = 0 \quad \text{in probability.} \quad (2.22)$$

This result indicates that, theoretically, for large N_1 , L and M can be increased in order to ensure convergence. In this perspective, we perform again the simulations in Scenario 2 (that assumes covariate mean shifts reported in Table 2.2b in the case of no response mean shifts ($d = 0$ for μ^Y)) with $L = 5, 10, 20$ and $M = 2, 3, 4$. The results are shown in Figure 2.3, where \widehat{ARL} s (which is in this case estimating $ARL_0 = 100$) in function of the number of covariates with shift in the mean function (according to Table 2.4) are reported by varying L (row-wise) and M (column-wise). Comparing these results with the top-left panel of Figure 2.2 ($d = 0$), it is clear that the negative effect (i.e., \widehat{ARL} not equal to $ARL_0 = 100$) caused by covariate mean shifts are attenuated by choosing L and M as large as possible (see the bottom-right panel of Figure 2.3). In general, the latter recommendation is expected to attenuate the effect on \widehat{ARL} of the covariate mean shifts, when there is a shift also in the response mean (depicted in panels $d = 0.5, 1.0, 1.5, 2.0$ of Figure 2.2).

Otherwise, when N_1 is small, we propose to use in the monitoring strategy (Section 2.2) the following scaled functional residual (hereinafter referred to as *studentized residual*)

$$e_{stu}^{std}(t) = \frac{Y^{std}(t) - \hat{Y}_{LM}^{std}(t)}{\left(\hat{\sigma}_\varepsilon^2(t) + \hat{\omega}_{LM}(t, t)\right)^{1/2}} \quad t \in \mathcal{T}, \quad (2.23)$$

instead of that in Equation (2.11). In Equation (2.23), $\hat{\sigma}_\varepsilon^2$ is an estimator of $\text{Var}(\varepsilon)$ and $\hat{\omega}_{LM}$ is defined as


$$\hat{\omega}_{LM}(s, t) = \text{Cov}\left(\hat{Y}_{LM}^{std}(t) | \mathbf{X}^{std}\right) = \left(\hat{\xi}_L^{\mathbf{X}}\right)^T \left(\hat{\Xi}_{\mathbf{X}}^T \hat{\Xi}_{\mathbf{X}}\right)^{-1} \hat{\xi}_L^{\mathbf{X}} \hat{\psi}_M^Y(s)^T \hat{\Sigma}_{\epsilon_M} \hat{\psi}_M^Y(t) \quad s \in \mathcal{S}, t \in \mathcal{T}, \quad (2.24)$$

where $\hat{\xi}_L^{\mathbf{X}}$ is the estimator of the score vector $\xi_L^{\mathbf{X}}$ of \mathbf{X}^{std} , $\hat{\Xi}_{\mathbf{X}}^T \hat{\Xi}_{\mathbf{X}}$ is the estimator of $N_1 \text{Cov}(\xi_L^{\mathbf{X}}, \xi_L^{\mathbf{X}})$, $\hat{\psi}_M^Y$ is the estimator of the vector of the first M eigenfunctions of Y^{std} , and $\hat{\Sigma}_{\epsilon_M}$ is the estimator of $\text{Cov}(\epsilon_M)$. Given e_{stu}^{std} in Equation (2.23), the *Hotelling's T^2* and *SPE* statistics, defined in Section 2.2, are particularized by the statistics $T_{e_{stu}^{std}}^2$ and $SPE_{e_{stu}^{std}}$ obtained by replacing e^{std} with e_{stu}^{std} , in Equation (2.14) and (2.15), respectively. This particular choice for the FRCC will be referred to as *studentized functional regression control chart* (sFRCC) which can be regarded as an extension to functional data of the regression control chart with prediction interval proposed by Wade and Woodall (1993). The studentized functional residual is the functional extension of the studentized residual that arises in the multivariate case (Woodall et al., 2004), with $\hat{\sigma}_\varepsilon^2(t) + \hat{\omega}_{LM}(t, t)$, for $t \in \mathcal{T}$, the variance function of e^{std} . The effect of $\hat{\sigma}_\varepsilon^2(t) + \hat{\omega}_{LM}(t, t)$, for $t \in \mathcal{T}$, on the latter is that of reducing the influence of covariate mean shifts on the residual mean. Indeed, the larger the term $\hat{\omega}_{LM}$, i.e., the more extreme realization of \mathbf{X}^{std} , the heavier the corresponding residual is rescaled and thus the higher the probability of the observation to be judged IC. And this is because the more the observations are far from the center of the sample cloud the larger the residual uncertainty. However, for a large N_1 , consistently with the dataset complexity, the use of the sFRCC leads to the same results of the FRCC defined before, because in this case, $\hat{\omega}_{LM}(t, t)$, for $t \in \mathcal{T}$, tends to zero. Therefore, all the residuals are equally rescaled regardless of the values achieved by the covariates.

To investigate the performance of the sFRCC, we carry out again simulation in Scenario 2 only in the case of no response mean shift ($d = 0$ for μ^Y) by setting N_1 equal to 100 and 50. Results are reported in Figure 2.4 and show that $\widehat{\text{ARL}}$ s achieved by the sFRCC are closer to the true value ($\text{ARL}_0 = 100$) than those obtained by means of the FRCC. In this simulation, the truncation parameters L , M and K in Equation (2.10) and Equation (2.13) shall be chosen large enough to avoid truncation bias, as addressed before, (in case of large N_1) but small enough to avoid overfitting problems due to small N_1 . In this simulation, we found appropriate to choose L , M and K such that the retained principal components explain the 99% of total variation. Too small values of L , M and K would not highlight the benefit of the term $\hat{\omega}_{LM}(t, t)$, for $t \in \mathcal{T}$, in the studentized residuals of Equation (2.23). Therefore, as in the multivariate case studied by Woodall et al. (2004), also in the functional setting, the sFRCC is able to control the Type I error (i.e., $\text{ARL}_0 = 100$) in case of covariate mean shifts. However, when the assumption of no covariate mean shift can be given as satisfied (e.g., for technological reasons), the FRCC is recommended since it results more sensitive than the sFRCC in detecting OC condition (of the functional response mean) as the latter has control limits wider than those of the FRCC.

2. FUNCTIONAL REGRESSION CONTROL CHART

Figure 2.4. Estimated $ARLs$ (\widehat{ARLs}) and 95% confidence intervals for the FRCC (left column panels) and sFRCC (right column panels) in Scenario 2 no shift ($d = 0$) in the response mean (μ^Y) at different sample sizes $N_1 = 100, 50$ as a function of which and how many covariates are subject to mean shift (each at the severity level reported in Table 2.2b). Zeros and ones in the triplets (000, 100, 010, 001, 110, 101, 011, 111) indicate IC and OC covariates, respectively. For instance, the triplet 100 means that only the first covariate is OC, 111 means that all the covariates are OC, and so on.



2.4 Real-case Study: Fuel Consumption Monitoring in the Shipping Industry

To demonstrate the potential and the applicability of the proposed control chart in practical situations, a real-case study in the shipping industry is presented henceforth. It addresses the issue of monitoring ship fuel consumption and, thus, CO_2 emissions, which, in view of the dramatic climate change, is of great interest in the maritime field in the very last years. Indeed, the new Regulation (EU 2015/757) of the European Union (EU) Council of 25 April 2015, coherently with the previous guidelines of the International Maritime Organization (IMO), compel operators having ships sailing in the Mediterranean Sea to monitor, report and verify (MRV) CO_2 emissions. In account of this, shipping companies are nowadays setting-up multi-sensor systems for massive high frequency recordings of operational data to be available. A large portion of these is suitably modelled as functional data.

In this study, data recorded during 2015, 2016, and 2017 from a Ro-Pax ship (owned by the shipping company *Grimaldi Group*) are analysed. The total number of voyages is 315. During each voyage the Data Acquisition (DAQ) device mounted on the ship have been collecting signals at five-minute frequency. The data refer only to the *navigation phase*, i.e., the time interval between the *finished with engine order* (when the ship leaves the departure port) and the *stand by engine order* (when the ship enters the arrival port). In particular, the percentage of travelled distance over the voyage is chosen as functional domain so that variables coming from different voyages are defined on interval of the same width. The *cumulative fuel consumption (CFC)* per each voyage is assumed as the functional response variable. It is the cumulative sum of the fuel consumption during the navigation phase. The following set of covariates are assumed as influencing the response: *sailing time* (T), measured in hours (h), which is the cumulative navigation time during the navigation phase; *speed over ground (SOG)*, measured in knots (kn), which is the ratio between the sailed distance over ground, i.e., the distance travelled by the vessel during the navigation, and the sailing time; *longitudinal and transverse wind components* (W_{lo} and W_{tr}), measured in knots (kn), which are functions of the true wind speed and the difference between the true wind angle (in the earth system) and the course over ground. Covariates are chosen on the basis of both engineering and statistical considerations. Additional information about the response and regressor variables can be found in Bocchetti et al. (2015); Erto et al. (2015). The 315 profiles observed for the response and covariates in the considered period are shown in the Supplementary Materials. Throughout February 2016, EEI (energy efficiency initiative) was performed on the considered ship, which mainly consisted in a silicone foul realising coat of

2.4. Real-case Study: Fuel Consumption Monitoring in the Shipping Industry

Table 2.6. Estimated ARL (\widehat{ARL}), mean \overline{ARL}^* of the empirical bootstrap ARL distribution, 95% confidence interval (CI) for the ARL statistic, and p -values of bootstrap test on the ARL mean differences for each chart combination.

	\widehat{ARL}	\overline{ARL}^*	CI	p -value			
				FRCC	sFRCC	RESP	INBA
FRCC	2.07	2.03	[1.78, 2.34]	-	-	-	-
sFRCC	2.11	2.21	[1.90, 2.56]	0.000999	-	-	-
RESP	9.46	10.09	[6.13, 17.92]	0.000999	0.000999	-	-
INBA	11.28	11.71	[7.81, 18.45]	0.000999	0.000999	0.000999	-

the hull. As guaranteed by the paint company, and described in Erto et al. (2015), this EEI plausibly produces a shift in the CFC mean. As shown in the Supplementary Materials, this is also confirmed by visual inspection of the mean function of the response before and after the EEI. In light of this, the 112 profiles relating to the period before dry-dock operations, are used to form the Phase I sample. Whereas, the remaining 203 profiles are used in Phase II to evaluate the proposed chart performance.

Implementation Details and Results

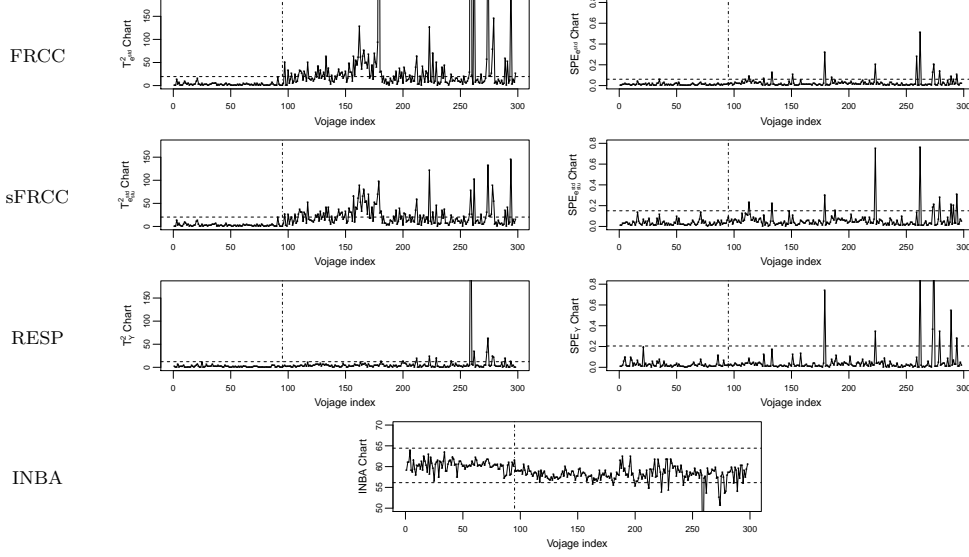
The Phase I sample consists of 112 profiles observed at five-minute frequency during each voyage. The functional observations are obtained by solving a regularization problem where profiles are approximated by means of a cubic B -spline basis expansion (i.e., of order 4) with 100 basis and 98 equispaced knots, and a smoothing parameter on the second derivative equal to 10^{-10} , chosen by means of generalized cross validation (Ramsay and Silverman, 2005). Then, the IC observations are identified by extending the approach proposed by Colosimo and Pacella (2007) to the multivariate functional case. As in Section 2.3, the appropriate values of L , M and K in Equation (2.10) and Equation (2.13) are found as those for which the retained principal components explain at least 95% of total variation. Then, the control limits are estimated using the empirical quantiles of the T_e^2 and SPE^e distribution estimated through the KDE procedure, as in the simulation study (Section 2.3), with the overall Type I error α equal to 0.0027 (which corresponds to $ARL_0 = 370$). The choice of α is prompted by the common industrial practice in analogy with the classic Shewhart control chart with 3-sigma limits (Montgomery, 2007). The remaining 203 profiles used in Phase II are obtained in the same way as in Phase I. Even if the use of the sFRCC has been recommended in all cases (Section 2.3), for the sake of completeness both the FRCC and the sFRCC have been applied. Each observation is plotted onto those two control charts and the two competitor ones (RESP and INBA) as shown in Figure 2.5.

By comparing the four charts, the responsiveness of the FRCC and the sFRCC is evidently higher than that of the INBA and the RESP control charts which signal a much lower number of OCs. In particular, in the FRCC and the sFRCC the change in the response mean is almost exclusively captured by the T^2 control chart, which means that dissimilarities between the Phase I and Phase II samples occur mostly in the space spanned by the retained principal components. As expected by remarks given in Section 2.3, Figure 2.5 shows that the sFRCC is less sensitive than the FRCC in detecting OC condition (of the functional response). However, the former chart should be used in this case because the assumption of no covariate mean shift cannot be given as satisfied. More precisely, by looking at the first column of Table 2.6, the estimated ARLs (\widehat{ARL}) achieved by FRCC and sFRCC are at least a fourth of those achieved by the RESP and INBA control charts.

To quantify the uncertainty of \widehat{ARL} s, a bootstrap analysis (Efron and Tibshirani, 1986)

2. FUNCTIONAL REGRESSION CONTROL CHART

Figure 2.5. T^2_{estd} and SPE_{estd} charts for the FRCC, the sFRCC, and, the RESP and INBA control charts. The vertical dotted line corresponds to the last voyage before the EEI.



was performed (see Supplementary Materials for more details).

Table 2.6 shows the mean \overline{ARL}^* of the empirical bootstrap ARL distribution and the bootstrap confidence intervals for each chart. Moreover, to test whether the mean of the empirical bootstrap ARL distributions differ significantly, bootstrap tests on the ARL mean differences (Efron and Tibshirani, 1986) for each chart combination were performed. The p -values are shown in Table 2.6 for each chart combination test. The bootstrap analysis, i.e., tests and confidence intervals, further confirms that both the FRCC and the sFRCC outperform the competitor control charts. Indeed, 95% confidence intervals are strictly below those of the RESP and INBA control charts and accordingly the tests reject the hypothesis of equal means for each chart combination.

Table 2.7 shows for the FRCC, sFRCC, and RESP control charts the estimated ARL s (\widehat{ARL}), the mean \overline{ARL}^* of the empirical bootstrap ARL distribution and the bootstrap 95% confidence intervals at different δ_Y and δ_X (i.e., percentages of variance explained by the retained scores in the response and covariates). The analysis for the INBA control charts is not influenced by different values of δ_Y and δ_X and therefore, the results are equal to those already reported in Table 2.6. Results in Table 2.7 show that both FRCC and sFRCC outperform the competitor ones for all δ_Y and δ_X values. However, as expected by remarks given in Section 2.3, the choice of the number L and M (Equation (2.10)) of the retained scores (related to δ_Y and δ_X) affects the performance of the FRCC and sFRCC.

2.5 Conclusions

In this chapter, we propose a new general framework for monitoring a functional quality characteristic when functional covariates are available, referred to as *functional regression control chart* (FRCC). In particular, the quality characteristic is adjusted for the effects of the covariates by means of multivariate functional linear regression model and then monitored by

Table 2.7. Estimated ARL (\widehat{ARL}), mean \overline{ARL}^* of the empirical bootstrap ARL distribution, and 95% confidence interval (CI) for the ARL statistic at different δ_Y , and δ_X , for the FRCC, sFRCC and RESP control chart.

δ_X	δ_Y	0.95 ($M = 2$)			0.98 ($M = 3$)			0.99 ($M = 5$)		
		\widehat{ARL}	\overline{ARL}^*	CI	\widehat{ARL}	\overline{ARL}^*	CI	\widehat{ARL}	\overline{ARL}^*	CI
0.90 ($L = 19$)	FRCC	2.99	3.06	[2.51, 3.77]	1.97	2.00	[1.77, 2.31]	1.96	2.00	[1.74, 2.29]
	sFRCC	2.89	2.95	[2.46, 3.62]	2.19	2.20	[1.92, 2.55]	2.25	2.28	[1.94, 2.67]
	RESP	9.46	9.99	[6.23, 17.87]	9.08	9.72	[5.90, 17.95]	5.72	6.50	[2.12, 12.92]
0.95 ($L = 36$)	FRCC	2.02	2.03	[1.76, 2.34]	1.66	1.67	[1.51, 1.87]	1.24	1.23	[1.17, 1.31]
	sFRCC	2.16	2.21	[1.90, 2.56]	1.71	1.72	[1.54, 1.94]	1.28	1.28	[1.21, 1.37]
	RESP	9.46	10.09	[6.13, 17.92]	9.08	9.45	[5.74, 16.55]	5.72	6.43	[2.41, 11.53]
0.97 ($L = 56$)	FRCC	1.36	1.36	[1.26, 1.47]	1.27	1.26	[1.19, 1.34]	1.62	1.63	[1.46, 1.82]
	sFRCC	1.46	1.48	[1.36, 1.62]	1.42	1.43	[1.31, 1.57]	1.84	1.86	[1.64, 2.13]
	RESP	9.46	10.24	[6.05, 17.14]	9.08	9.60	[5.75, 17.16]	5.72	6.41	[2.39, 12.19]

using jointly the *Hotelling's* T^2 and the *SPE* control charts built on its functional principal component decomposition. However the approach is very general, indeed the choice of the model, the estimation method as well as the monitoring strategy can be easily extended. To the best of the authors' knowledge, profile monitoring methods that are promptly able to enhance the monitoring by exploiting additional information on covariates (even possibly functional ones) are not present in the literature, whose attention is mainly focused on procedures that consider measurements of the functional quality characteristic only.

A Monte Carlo simulation is carried out with the aim of investigating the performance of the proposed control chart in identifying mean shifts in the response. The FRCC is then compared with other two control charts (named Response and INDEX-BASED control charts) that are widely used both in the literature and in real profile monitoring applications. The results showed that, firstly, the FRCC is far better than the competitor control charts in identifying response mean shifts, when no covariate mean shift occurs; secondly, the covariate mean shift implies estimation error of the coefficient vector and thus, strongly affects FRCC performance in terms of average run length. When the assumption of no covariate mean shift cannot be given as satisfied, some solutions are proposed in case of both large and small Phase I sample sizes. In the latter case, a studentized version of the FRCC (sFRCC) is proposed to take into account the different residual variance at different covariate values. Eventually, by means of a real-case study in the shipping industry, the FRCC and sFRCC is shown to outperform the competitor control charts in identifying CO₂ emission reduction after a specific energy efficiency initiative.

Future researches can be addressed on extending the FRCC framework to different types of regression models and to different residual monitoring strategies. Moreover, the effect on the FRCC performance in detecting shifts in the variance function of both the response and covariates deserve further investigations.

2.6 Supplementary Materials

Estimation of the Multivariate Functional Linear Regression Model

The estimation of the regression coefficient vector β in Equation (2.3), relies on the truncated principal components decomposition \mathbf{X}_L^{std} and Y_M^{std} of \mathbf{X}^{std} and Y^{std} , namely

$$\mathbf{X}_L^{std} = \sum_{i=1}^L \xi_i^X \psi_i^X, \quad Y_M^{std} = \sum_{i=1}^M \xi_i^Y \psi_i^Y. \quad (2.25)$$

This leads to truncated version of β in Equation (2.6), that is

$$\beta_{LM}(s, t) = \sum_{i=1}^L \sum_{j=1}^M b_{ij} \psi_i^X(s) \psi_j^Y(t) \quad s \in \mathcal{S}, t \in \mathcal{T}, \quad (2.26)$$

with $b_{ij} = E(\xi_i^X \xi_j^Y) / \lambda_i^X$. By considering the following expansion of ε in Equation (2.3),

$$\varepsilon = \sum_{i=1}^{\infty} \epsilon_i \psi_i^Y, \quad (2.27)$$

with $\epsilon_i = \langle \varepsilon, \psi_i^Y \rangle$, and the related truncated version ε_M of order M , then the truncated version of Equation (2.3) is

$$Y_M^{std}(t) = \int_{\mathcal{S}} (\beta_{LM}(s, t))^T \mathbf{X}_M^{std}(s) ds + \varepsilon_M(t) \quad t \in \mathcal{T}. \quad (2.28)$$

Plugging Equation (2.26) and Equation (2.27) in to Equation (2.28), due to the orthonormality of the principal components $\{\psi_i^X\}$ and $\{\psi_i^Y\}$, Equation (2.28) becomes

$$\xi_M^Y = (\mathbf{B}_{LM})^T \xi_L^X + \epsilon_M, \quad (2.29)$$

where $\xi_M^Y = (\xi_1^Y, \dots, \xi_M^Y)^T$, $\xi_L^X = (\xi_1^X, \dots, \xi_L^X)^T$, $\epsilon_M = (\epsilon_1, \dots, \epsilon_M)^T$ and $\mathbf{B}_{LM} = \{b_{ij}\}_{i=1, \dots, L, j=1, \dots, M}$. The least squares estimator of \mathbf{B}_{LM} is

$$\mathbf{B}_{LM}^{ls} = \text{Cov}(\xi_L^X, \xi_L^X)^{-1} \text{Cov}(\xi_L^X, \xi_M^Y). \quad (2.30)$$

Given n independent realizations (\mathbf{X}_i, Y_i) of (\mathbf{X}, Y) , an estimator $\hat{\beta}_{LM}$ of β_{LM} , - and thus of β , based on the least squares estimator \mathbf{B}_{LM}^{ls} of \mathbf{B}_{LM} - is obtained by means of the estimators $\hat{\mu}^X$ and $\hat{\mu}^Y$ of μ^X and μ^Y , and the estimators $\hat{\mathbf{C}}_{std}^X$ and $\hat{\mathbf{C}}_Z^Y$ of \mathbf{C}_{std}^X and \mathbf{C}_Z^Y . They are calculated using the sample mean and covariance functions (Hsing and Eubank, 2015). Estimators $\hat{\psi}_i^X$, $\hat{\lambda}_i^X$ and $\hat{\psi}_i^Y$, $\hat{\lambda}_i^Y$ of ψ_i^X , λ_i^X and ψ_i^Y , λ_i^Y are obtained through the method proposed by Happ and Greven (2018), based on the theory of integral equations. Then, the estimator $\hat{\beta}_{LM}$ of β can be calculated as

$$\hat{\beta}_{LM}(s, t) = \left(\hat{\psi}^Y(t) \right)^T \left(\hat{\mathbf{B}}_{LM}^{ls} \right)^T \hat{\Psi}^X(s) \quad s \in \mathcal{S}, t \in \mathcal{T}, \quad (2.31)$$

where $\hat{\psi}^Y = (\hat{\psi}_1^Y, \dots, \hat{\psi}_M^Y)^T$, $\hat{\Psi}^X = (\hat{\psi}_1^X, \dots, \hat{\psi}_L^X)^T$ and $\hat{\mathbf{B}}_{LM}^{ls} = \left(\hat{\Xi}_X^T \hat{\Xi}_X \right)^{-1} \hat{\Xi}_X^T \hat{\Xi}_Y$, with $\hat{\Xi}_X^T \hat{\Xi}_X$ and $\hat{\Xi}_X^T \hat{\Xi}_Y$ estimators of $n \text{Cov}(\xi_L^X, \xi_L^X)$ and $n \text{Cov}(\xi_L^X, \xi_M^Y)$. Finally, an estimator \hat{Y}_{LM} of the best prediction $E(Y|\mathbf{X})$ of Y in Equation (2.9) is

$$\hat{Y}_{LM}(t) = \hat{\mu}^Y(t) + \hat{v}^Y(t)^{1/2} \hat{Y}_{LM}^{std}(t) \quad t \in \mathcal{T}, \quad (2.32)$$

where \hat{v}^Y is an estimator of v^Y and $\hat{Y}_{LM}^{std} = \sum_{i=1}^L \sum_{j=1}^M \hat{b}_{ij} \hat{\xi}_i^X \hat{\psi}_j^Y$ is an estimator of $E(Y^{std}|\mathbf{X}^{std})$ in Equation (2.8), with $\{\hat{b}_{ij}\}$ the entries of $\hat{\mathbf{B}}_{LM}^{ls}$ and $\hat{\xi}_i^X = \langle \mathbf{X}^{std}, \hat{\psi}_i^X \rangle_{\mathbb{H}^X}$.

Table 2.8. Correlation functions and parameters for data generation in the simulation study.

Type		ρ	ν
Bessel	$J_\nu(z) = \left(\frac{ z }{2}\right)^\nu \sum_{j=0}^{\infty} \frac{(- z /\rho)^{2j}/4^j}{j!\Gamma(\nu+j+1)}$	0.25	0
Gaussian	$G(z) = \exp\left[-\left(\frac{ z }{\rho}\right)^2\right]$	1	-
Powered exponential	$P(z) = \exp\left[-\left(\frac{ z }{\rho}\right)^\nu\right]$	1	0.5

Details on Data Generation

The compact domains \mathcal{S} and \mathcal{T} are set, without loss of generality, equal to $[0, 1]$ and the number of covariates p is set equal to 3. The eigenfunctions sets $\{\psi_i^X\}$ and $\{\psi_i^Y\}$ are generated by the spectral decomposition of pre-specified correlation functions. In particular, the eigenfunction set $\{\psi_i^X\}$ is obtained considering the correlation function \mathbf{G}^X through the following steps.

1. Set the diagonal elements G_{ll}^X , $l = 1, 2, 3$ of \mathbf{G}^X as the *Bessel* correlation function of the first kind (Abramowitz and Stegun, 1964), the *gaussian* correlation function (Abrahamsen and Regnesentral, 1997) and the *powered exponential* correlation function (Stein, 1999). The general form of the correlation functions and parameters used in the simulation study are listed in Table 2.8. Then, calculate the eigenvalues $\{\eta_{lk}^X\}$ and the corresponding eigenfunctions $\{\vartheta_{lk}^X\}$, $k = 1, 2, \dots$, of G_{ll}^X , $l = 1, 2, 3$.
2. Obtain the cross-correlation function G_{lj}^X , $l, j = 1, 2, 3$ and $l \neq j$, by

$$G_{lj}^X(s_1, s_2) = \sum_{k=1}^{\infty} \tilde{\eta}_k^X \tilde{\vartheta}_{lk}^X(s_1) \tilde{\vartheta}_{jk}^X(s_2) \quad s_1, s_2 \in \mathcal{S}, \quad (2.33)$$

where $\tilde{\eta}_k^X = (1/3) \sum_{l=1}^3 \eta_{lk}^X$ and $\tilde{\vartheta}_{lk}^X = (1/\sqrt{3}) \vartheta_{lk}^X$.

3. Calculate the eigenvalues $\{\lambda_i^X\}$ and the corresponding eigenfunctions $\{\psi_i^X\}$ through the spectral decomposition of $\mathbf{G}^X = \{G_{lj}^X\}_{l,j=1,2,3}$, for $i = 1, \dots, L^*$.

The eigenvalues $\{\lambda_i^Y\}$ and the corresponding eigenfunctions $\{\psi_i^Y\}$, $i = 1, \dots, M^*$ are calculated by means of the spectral decomposition of G^Y set as the *Bessel correlation function* of the first kind with $\rho = 0.2$ and $\nu = 0$ (Abramowitz and Stegun, 1964). Further, set L^* and M^* equal to 50 and 10, respectively. Then, β is calculated as follows,

$$\beta(s, t) = \left(\psi^Y(t)\right)^T (\mathbf{B}_{L^*M^*})^T \Psi^X(s) \quad s, t \in [0, 1], \quad (2.34)$$

where the matrix $\mathbf{B}_{L^*M^*}$, is set as a partitioned matrix $[\mathbf{B}_{L^*M^*11} \mathbf{B}_{L^*M^*21}]^T$, where $\mathbf{B}_{L^*M^*11}$ is a diagonal matrix of dimension M^* and $\mathbf{B}_{L^*M^*21}$ is a $(L^* - M^*) \times M^*$ matrix of all zeros. Diagonal values of $\mathbf{B}_{L^*M^*11}$ are listed in Table 2.9 for three different settings, along with the corresponding R^2 values (Horváth and Kokoszka, 2012; Yao et al., 2005a), defined as

$$R^2 = \int_{[0,1]} \frac{\text{Var}\left(\mathbb{E}(Y^{std}(t) | \mathbf{X}^{std})\right)}{\text{Var}(Y^{std}(t))} dt. \quad (2.35)$$

2. FUNCTIONAL REGRESSION CONTROL CHART

Table 2.9. Diagonal values b_{ii} of \mathbf{B}_{LM11} and corresponding R^2 for three different settings.

R^2	b_{ii}
0.97	0.698, 0.838, 0.315, 0.0504, 0.002, 0.000, 0.000, 0.000, 0.000, 0.000
0.86	0.658, 0.795, 0.275, 0.010, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000
0.76	0.608, 0.745, 0.225, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000

The R^2 value measures globally the proportion of the variance in the response explained by the covariates. Then, in order to ensure the validity of the model in Equation (2.3), $\Sigma_{\epsilon_{M*}}$ is chosen such that the following relation holds

$$\Sigma_{\xi_{M*}^Y} = \Lambda^Y = (\mathbf{B}_{L*M*})^T \Lambda^X \mathbf{B}_{L*M*} + \Sigma_{\epsilon_{M*}}, \quad (2.36)$$

with $\Sigma_{\xi_{M*}^Y} = \text{Cov}(\xi_{M*}^Y)$, $\Lambda^Y = \text{diag}(\lambda_1^Y, \dots, \lambda_{M*}^Y)$, and $\Lambda^X = \text{diag}(\lambda_1^X, \dots, \lambda_{L*}^X)$. Realizations of \mathbf{X}^{std} are obtained through

$$\mathbf{X}^{std} = \sum_{i=1}^{L*} \xi_i^X \psi_i^X, \quad (2.37)$$

with $\xi_{L*}^X = (\xi_1^X, \dots, \xi_{L*}^X)^T$ generated by means of a multivariate normal distribution with covariance $\text{Cov}(\xi_{L*}^X) = \Sigma_{\xi_{L*}^X} = \Lambda^X$. In the same way, realizations of Y^{std} are generated by means of

$$Y^{std} = \sum_{i=1}^{M*} \xi_i^Y \psi_i^Y. \quad (2.38)$$

Realizations of the score vector $\xi_{M*}^Y = (\xi_1^Y, \dots, \xi_{M*}^Y)^T$ are obtained as

$$\xi_{M*}^Y = (\mathbf{B}_{L*M*})^T \xi_{L*}^X + \epsilon_{M*}, \quad (2.39)$$

with $\epsilon_{M*} = (\epsilon_1, \dots, \epsilon_{M*})^T$ generated by means of a multivariate normal distribution with covariance matrix $\text{Cov}(\epsilon_{M*}) = \Sigma_{\epsilon_{M*}}$ independent of ξ_{L*}^X . Further, the mean functions μ^X and μ^Y and the variance functions $\mathbf{v}^X = (v_1^X, v_2^X, v_3^X)^T$ and v^Y are generated through the following reference model

$$\mu(z) = P(z) + r \sum_{i=1}^I g_i(z; m_i, s_i) \quad z \in [0, 1], \quad (2.40)$$

where

$$P(z) = az^2 + bz + c \quad z \in [0, 1], \quad (2.41)$$

and the terms $g_i(t; m_i, s_i)$ are normal probability density functions with parameters m_i and s_i , and a, b, c are real numbers. The values of all unknown parameters are listed in Table 2.10 and Table 2.11. Then, given the mean functions μ^X and μ^Y and the variance functions \mathbf{v}^X and v^Y , realizations of \mathbf{X} and Y are easily obtained. Finally, \mathbf{X} and Y are assumed to be observed at 150 equally spaced time points $[0, 1]$ with measurement errors $\zeta_i^X \sim N(\mathbf{0}, \sigma_X^2)$ and $\zeta_i^Y \sim N(0, \sigma_Y^2)$ where $\sigma_X = (0.3, 0.05, 0.3)^T$ and $\sigma_Y = 0.3$. For illustrative purposes,

2.6. Supplementary Materials

Table 2.10. Values of m_i and s_i to generate \mathbf{X} and Y mean and variance functions in the simulation study.


	m_i	s_i
μ_1^X	0.075, 0.100, 0.250, 0.350, 0.500, 0.650, 0.850, 0.900, 0.950	0.050, 0.030, 0.050, 0.050, 0.100, 0.050, 0.100, 0.040, 0.035
μ_2^X	-	-
μ_3^X	0.075, 0.100, 0.150, 0.225, 0.400, 0.525, 0.550, 0.600, 0.625, 0.650, 0.850, 0.900, 0.925	0.050, 0.060, 0.050, 0.040, 0.050, 0.035, 0.045, 0.045, 0.040, 0.030, 0.015, 0.010, 0.015
μ^Y	-	-
v_1^X	0.075, 0.100, 0.125, 0.150, 0.400, 0.650, 0.850, 0.900, 0.925	0.050, 0.060, 0.075, 0.075, 0.075, 0.045, 0.045, 0.040
v_2^X	-	-
v_3^X	0.075, 0.100, 0.150, 0.225, 0.400, 0.525, 0.550, 0.600, 0.625, 0.650, 0.850, 0.900, 0.925	0.050, 0.060, 0.050, 0.040, 0.050, 0.035, 0.045, 0.045, 0.040, 0.030, 0.015, 0.010, 0.015
v^Y	-	-

Table 2.11. Values of a , b , c and r to generate \mathbf{X} and Y mean and variance functions in the simulation study.

	a	b	c	r
μ_1^X	-20	20	-20	0.05
μ_2^X	0	4	0	0
μ_3^X	-10	14	-8	0.05
μ^Y	0	30	0	0
v_1^X	0	0	1	0.1
v_2^X	0	0.02	1	0
v_3^X	-40	150	30	2
v^Y	0	8	1	0

a sample of 1000 randomly generated realizations of $\mathbf{X} = (X_1, X_2, X_3)^T$ and Y with their mean functions are shown in Figure 2.6 for $R^2 = 0.97$.

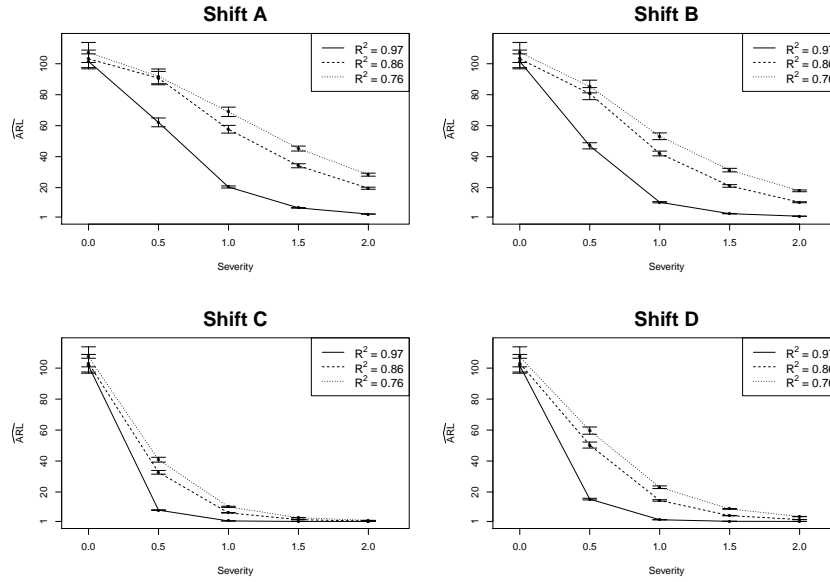
Figure 2.6. 1000 randomly generated realizations of the predictors and the response with their mean functions (black solid line).



Additional Simulations for Different Values of R^2

In the simulation study, the analysis were performed using data generated by considering diagonal values of $\mathbf{B}_{L \times M \times 11}$ corresponding to $R^2 = 0.97$. To assess the effect on the FRCC performance of changes in the proportion of the response variance explained by the covariates, the same analysis of Scenario 1 in Section 2.3 are performed by considering the three different settings in Table 2.9. As expected, Figure 2.7 and Table 2.12 show that the FRCC performance decrease with the R^2 values and that tend to be equal to those of the RESP control chart. This confirms the fact that when no linear relation hold between the response and the covariates the FRCC and the RESP control chart perform equivalently. However, already at $R^2 = 0.76$ the FRCC performs better than the competitors control charts.

Figure 2.7. Estimated $ARLs$ (\widehat{ARL}) and 95% confidence intervals for different R^2 values.



Bootstrap Analysis

Given the n observations in Phase II (\mathbf{X}_i, Y_i) , $i = 1, \dots, 203$, of the covariates and response, the bootstrap analysis can be summarized in the following steps.

1. Compute the standardized versions $(\mathbf{X}_i^{std}, Y_i^{std})$ of (\mathbf{X}_i, Y_i) , using the quantities estimated in Phase I.
2. Obtain B standard bootstrap samples of size n $(\mathbf{X}_{1b}^{std}, Y_{1b}^{std}), \dots, (\mathbf{X}_{nb}^{std}, Y_{nb}^{std})$, $b = 1, \dots, B$, resampling with replacement from the standardized observations $(\mathbf{X}_1^{std}, Y_1^{std}), \dots, (\mathbf{X}_n^{std}, Y_n^{std})$.
3. Use the B bootstrap samples $(\mathbf{X}_{1b}^{std}, Y_{1b}^{std}), \dots, (\mathbf{X}_{nb}^{std}, Y_{nb}^{std})$, $b = 1, \dots, B$, to compute B values of the statistic, ARL_1, \dots, ARL_B for each chart.

Table 2.12. Estimated ARLs (\widehat{ARL}) and 95% confidence intervals for the three different values of R^2 in Scenario 1.

Shift	Severity	R^2					
		0.97		0.86		0.76	
	d	\widehat{ARL}	CI	\widehat{ARL}	CI	\widehat{ARL}	CI
In-control	-	102.94	[97.95, 107.94]	103.42	[97.75, 109.08]	101.64	[96.20, 107.08]
A	0.5	62.14	[59.29, 64.98]	90.73	[86.43, 95.04]	91.85	[87.09, 96.61]
	1.0	20.43	[19.76, 21.10]	57.71	[55.23, 60.20]	68.96	[65.92, 72.00]
	1.5	6.94	[6.73, 7.14]	34.12	[32.82, 35.42]	45.23	[43.64, 46.83]
	2.0	2.93	[2.87, 2.99]	19.63	[18.96, 20.30]	28.38	[27.41, 29.36]
B	0.5	47.00	[45.02, 48.99]	80.70	[76.81, 84.60]	85.29	[81.20, 89.38]
	1.0	10.58	[10.18, 10.98]	42.06	[40.53, 43.58]	53.17	[50.99, 55.35]
	1.5	3.19	[3.11, 3.27]	21.10	[20.27, 21.94]	31.32	[30.18, 32.47]
	2.0	1.56	[1.54, 1.58]	10.56	[10.25, 10.88]	17.94	[17.37, 18.51]
C	0.5	8.30	[8.07, 8.54]	32.72	[31.53, 33.92]	40.87	[39.32, 42.43]
	1.0	1.33	[1.32, 1.34]	6.57	[6.40, 6.74]	10.43	[10.06, 10.81]
	1.5	1.00	[1.00, 1.00]	2.15	[2.11, 2.19]	3.43	[3.36, 3.50]
	2.0	1.00	[1.00, 1.00]	1.22	[1.21, 1.23]	1.67	[1.65, 1.70]
D	0.5	15.42	[14.89, 15.94]	50.31	[48.37, 52.25]	59.64	[57.32, 61.97]
	1.0	2.09	[2.06, 2.12]	14.56	[14.09, 15.02]	23.10	[22.24, 23.95]
	1.5	1.07	[1.07, 1.08]	4.76	[4.64, 4.87]	9.08	[8.84, 9.33]
	2.0	1.00	[1.00, 1.00]	2.20	[2.17, 2.24]	4.15	[4.05, 4.24]

4. Build the confidence interval with confidence level $1 - \alpha$ for the ARL statistics using the $\alpha/2$ and $1 - \alpha/2$ quantiles of the empirical bootstrapped ARL distribution and calculate the mean, \overline{ARL}^* , of the empirical bootstrapped ARL distribution for each control chart.

The number of bootstrap samples B is set equal to 500, and confidence intervals are built with $\alpha = 0.05$.

Additional Plots for the Real-case Study

Figures 2.8 shows the 315 profiles observed for the response and covariates in the real-case study of Section 4. The functional response is the *cumulative fuel consumption* (CFC) per each voyage. The scale on the ordinate axis is omitted for confidentiality reasons. The covariates are the *sailing time* (T), measured in hours (h) the *speed over ground* (SOG), measured in knots (kn), and the *longitudinal* and *transverse wind components* (W_{lo} and W_{tr}), measured in knots (kn).

2. FUNCTIONAL REGRESSION CONTROL CHART

Figure 2.8. Covariates and response in the real-case study.

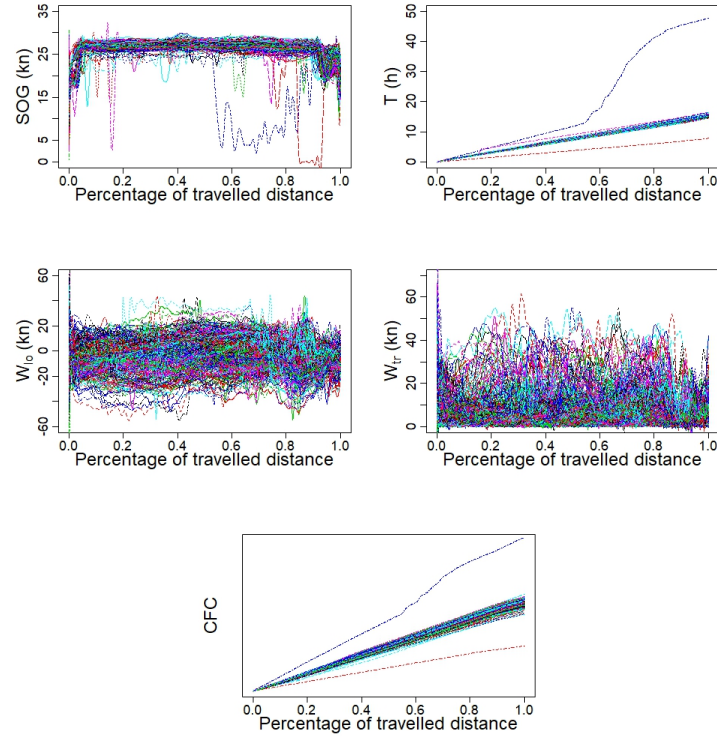



Figure 2.9 shows the mean function of the response before and after the EEI (energy efficiency initiative). By visual inspection, it is clear that a shift downward of the cumulative fuel consumption occurred.

Figure 2.9. *CFC* before (solid line) and after (dashed line) the EEI.



Chapter 3

Functional Regression Control Chart for the Monitoring of Ship CO₂ Emissions

Abstract

On the modern ships, the quick development in data acquisition technologies is producing data-rich environments where variable measurements are continuously streamed and stored during navigation and thus can be naturally modelled as functional data or profiles. Then, both the CO₂ emissions (i.e., the quality characteristic of interest) and the variable profiles that have an impact on them (i.e., the covariates) are called to be explored in the light of the new worldwide and European regulations on the monitoring, reporting and verification of CO₂ emissions. In this chapter, we show an application of the functional regression control chart (FRCC) with the ultimate goal of answering, at the end of each ship voyage, the question: *given the value of the covariates, is the observed CO₂ emission profile as expected?*. To this aim, the FRCC focuses on the monitoring of residuals obtained from a multivariate functional linear regression of the CO₂ emission profiles on the functional covariates. The applicability of the FRCC is demonstrated through a real-case study of a Ro-Pax ship operating in the Mediterranean Sea. The proposed FRCC is also compared with other alternatives available in the literature and its advantages are discussed over some practical examples.

3.1 Introduction

Nowadays, the quick development in the data acquisition (DAQ) technologies is producing data-rich industrial environments where massive amounts of data are available. In particular, a large portion of the ship observational data are complex measurement signals that, even though they consist of discrete values of the quantities of interest ordered in time or space, they may be envisaged as reflecting smooth variations of quantities generated by continuous functions defined on a infinite compact domain, i.e., as *functional data*. Functional data analysis (FDA) is a thriving area of statistics. For a comprehensive overview, the reader could refer to Ramsay and Silverman (2005), Horváth and Kokoszka (2012) and Kokoszka and Reimherr (2017). In this functional data setting, profile monitoring (Noorossana et al., 2012) is the name of a new branch of statistical process control (SPC) that provides a suite

3. FUNCTIONAL REGRESSION CONTROL CHART FOR THE MONITORING OF SHIP CO₂ EMISSIONS

of methods to continuously give a solution to the urging issue of evaluating the stability over time of functional quality characteristics. Recent contributions are Colosimo and Pacella (2010), Grasso et al. (2016), Menafoglio et al. (2018). As in the classical SPC, where data are scalars, profile control charts have the task of continuously monitoring the functional quality characteristic and of triggering a signal when assignable sources of variations (i.e., special causes) act on it. When this happens, the process is said to be out-of-control (OC). Otherwise, when only normal sources of variation (i.e., common causes) apply, the process is said to be in-control (IC).

Only recently, Centofanti et al. (2020b) introduced the functional regression control chart (FRCC) framework to deal with the situations where a functional quality characteristic is influenced by one or more functional covariates. In this scenario, if one of these covariates manifests itself with an extreme realization, the quality characteristic may wrongly be judged to be OC. Otherwise, there may be situations where the covariates are not extreme and the quality characteristic may wrongly appear IC. The FRCC framework is the functional extension of the basic Mandel's idea (Mandel, 1969), where the quality characteristic is monitored after being adjusted for the effect of covariates, i.e., the control variable is the residuals obtained from a regression of the quality characteristic on the covariates. In this way, the focus is on the residual variability not explained by the knowledge of the observed value of the covariates. In a more direct phrasing, the FRCC answers the question: *given the value of the covariates, is the quality characteristic as expected?* If the answer is no, then special causes may have occurred that are beyond the information brought by the covariates through the chosen regression model. In particular, in the FRCC framework, the quality characteristic and the covariates are linked through a multiple functional linear regression model (MFLR), where both the response and the explanatory variables can be functional data. Recent examples of MFLR model can be found in Palumbo et al. (2020), Centofanti et al. (2020a) and Chiou et al. (2014).

In recent years, profile monitoring has emerged as an effective technique in the field of maritime transport to tackle the issue of CO₂ emission monitoring (Capezza et al., 2019; Centofanti et al., 2020b; Lepore et al., 2018). Indeed, in view of climate change and global warming crises, the maritime transport industry is currently facing new challenges related to harmful emission control and reduction. Indeed, the Marine Environment Protection Committee of the International Maritime Organization (IMO, 2012a,b, 2014) has urged shipping companies to set up a framework for monitoring, reporting and verification of CO₂ emissions based on fuel consumption. In the face of these regulation, shipping companies are updating DAQ systems on their fleets, enabling large volumes of observational data to be automatically streamed and transferred to a remote server, bypassing human intervention. Indeed, a large proportion of these data can be modeled as functional data, thus representing a new challenge for FDA and related SPC methods in this area. The DAQ system installed on modern ships facilitates in fact the collection of functional data relating to CO₂ emissions as well as other functional variables affecting them. However, most of the approaches that have already appeared in the maritime literature (Lepore et al., 2017; Erto et al., 2015; Bocchetti et al., 2015) do not take advantage of the potential help to managerial decision-making represented by the modelling of the entire voyage profiles acquired and usually collapse information in one or more scalar features extracted from them. In this setting, a recurrent request posed by the maritime engineers is related to the CO₂ emissions corresponding to the given values of other recorded covariates. That is, they want, in particular, to assess if CO₂ emissions are coherent with the values of the covariates, in order to identify unexpected behaviours and take corrective measures. Engineers are less concerned in identifying CO₂

Features	Value
Gross Tonnage	32728
Length	203.9 m
Beam	25 m
Draft	6.8 m
Maximum power for propulsion	46080 kW
Maximum speed	28.9 knots

Table 3.1. Technical features of the considered Ro-Pax ship.

emission profiles that are extreme with respect to their marginal distribution if this can be explained well by some extreme value in the covariates. They are rather interested in CO₂ emission profiles that are not consistent with covariates affecting them because, for example, this can reveal anomalous ship performance.

The FRCC can be used to meet this engineering need. Therefore, in this chapter we propose to use the FRCC to monitor ship CO₂ emissions throughout each voyage in order to identify special causes given the values of some functional covariates. Specifically, we consider a particular instance of the FRCC framework where the functional quality characteristic, hereinafter referred to as response, and the functional covariates are related through the multivariate functional linear regression (MFLR) model, whose estimation is based on the multivariate functional principal components analysis (FPCA) (Happ, 2018; Chiou et al., 2014). Then, studentized residuals are monitored through the simultaneous application of the Hotelling's T^2 and the squared prediction error (SPE) control charts (Woodall et al., 2004; Noorossana et al., 2012; Grasso et al., 2016; Colosimo and Pacella, 2010). In particular, in this chapter the FRCC framework is used both *retrospectively*, as an aid to the practitioner to determine the IC state of a process and to identify an IC reference sample (Phase I), and, *prospectively* to monitor any departure from the IC state at future voyages (Phase II). The applicability of the FRCC is demonstrated through a real-case study of a Ro-Pax ship operating in the Mediterranean Sea, courtesy of the shipping company Grimaldi group.

The chapter is structured as follows. In Section 3.2, the structure of the data and technological details of the ship equipment, for the case study at hand are provided. In Section 3.3 the main materials and methods behind the particular instance of the FRCC framework are summarized. In Section 3.4, we apply the FRCC to the real-case study at hand to monitor CO₂ emissions, and, a comparison with competing methods, which do not take into account the information coming from the covariates, is illustrated as well. In Section 3.5, we draw conclusions. All computations and plots have been obtained using the software environment R (R Core Team, 2020b).

3.2 Technological Background and Data Structure

For confidentiality reasons, we omit the ship as well as port names. In Section 3.2, we provide the technical features of the ship. In Section 3.2, we describe all the variables and the data used for the analysis.

Technical Features

The main technical features of the ship are illustrated in Table 3.1. The ship is characterized

3. FUNCTIONAL REGRESSION CONTROL CHART FOR THE MONITORING OF SHIP CO₂ EMISSIONS

by two engine sets, each consisting of two main diesel engines for propulsion Wärtsilä, Type 16ZAV40S, four-stroke, with a maximum continuous rating of 11520 kW at 510 revolutions per minute (rpm) and by two variable pitch propellers and a shaft generator for electric power supply. The main engine power is used both for propulsion and electrical generation through the shaft generators, which are themselves keyed on a gearbox. The gearbox has two fast inlet shafts powered by the engine shaft, a slow outlet shaft for the propeller and a faster one to which the shaft alternator is connected. The gear ratio between the engine shaft and the propeller shaft is equal to 3.24. The gear ratio between the engine shaft and the shaft alternator is equal to 0.32. The main diesel engine of the considered Ro-Pax ship can be powered by three types of fuel with different percentage of sulphur (S) content compliant with the regulation in force on the geographical area to be sailed: heavy fuel oil, very low sulphur fuel oil ($\leq 0.5\%S$), ultra-low sulphur fuel oil or marine gas oil ($\leq 0.1\%S$). The electrical power supply of the ship consists of three diesel generators (1840 kVA, 690 V), two shaft generators (2100 kVA, 690 V) and one emergency diesel generator (480 kVA). The main engines can supply power in two different modes, at fixed rpm (constant mode) or at variable rpm (combined mode). In the constant mode, shaft generators can be used to supply electric power, even though the maximum speed cannot be reached because speed variations are only possible by changing the pitch of the propellers. Whereas, in the combined mode, the ship speed can be regulated by increasing both pitch propeller and engine rpm, but, vice-versa, the possibility to engage the shaft generator is lost.

Data Description

Data come from a DAQ system installed on the ship that are transmitted to the cloud with at different frequencies varying from 2 to 5 minutes. The data refer to a specific route, i.e. each observation in the data set corresponds to a voyage of the ship and all voyages have the same departure and arrival port. A period of 11 consecutive months following a dry-dock operation on the ship is considered. Observations in the first 9 months (i.e. from the beginning of February to the end of October 2020) are used in Phase I, i.e. to identify a reference data set, estimate the model and control chart limits. Observations in the last 2 months (i.e. from the end of October 2020 to the end of December 2020) are used in Phase II to show the performance of the FRCC on monitoring new voyages. We start with a data set of 190 voyages for the Phase I (Section 3.4) and 22 voyages for the Phase II (Section 3.4).

Note that the data refer to the navigation phase. More specifically, the navigation phase begins with the finished with engine order (when the ship leaves the departure port) and ends with the stand by engine order (when the ship enters the arrival port). Moreover, we need to identify an adequate functional domain for each voyage. Even if time is naturally suitable as a functional domain, total travel time can vary from voyage to voyage. Therefore, we prefer to use the fraction of distance traveled over the voyage as the common domain $(0, 1)$ of the data.

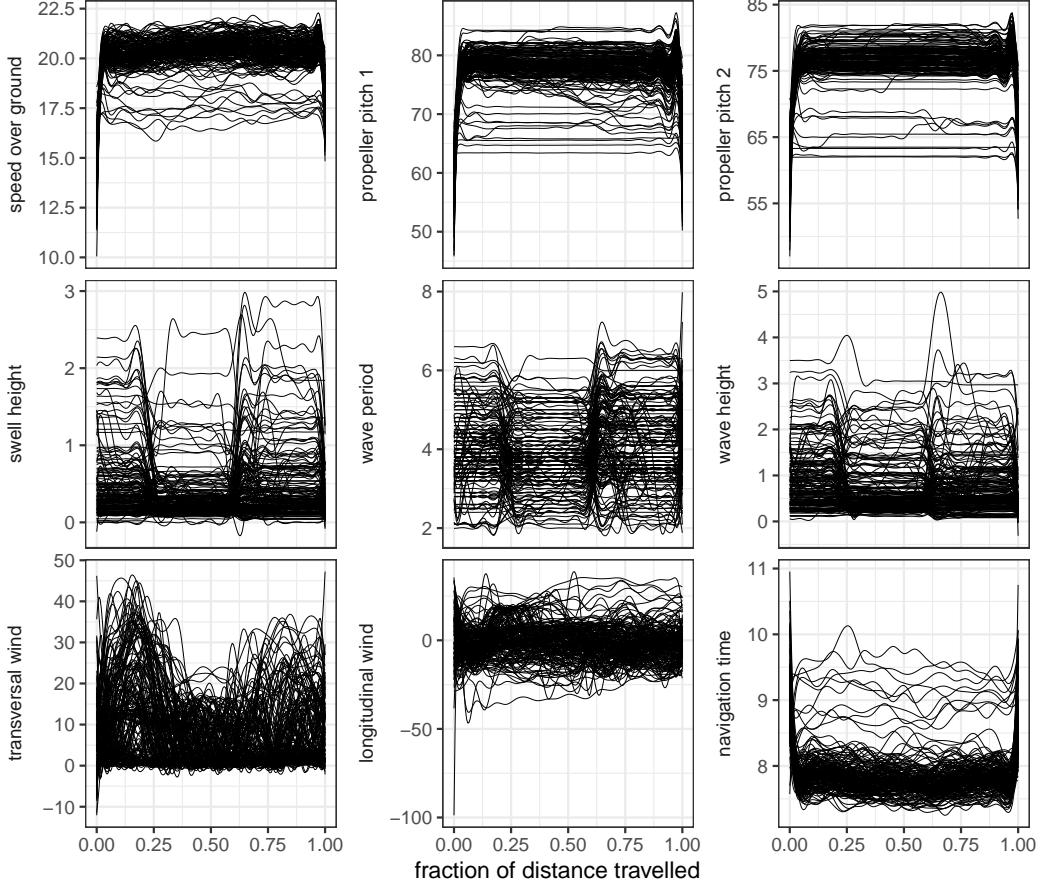
All the signals acquired by the DAQ system are summarized into several variables that we describe here to select later the functional covariates and response considered in the MFLR model. The ship is tracked by its global positioning system (GPS), which provides longitude and latitude coordinates. The course over ground (COG) is the actual direction of progress of a vessel, between two points, with respect to the surface of the earth, measured in degrees. The sailed distance over ground (SDOG) is the distance traveled by the vessel between two points, measured in nautic miles (NM), calculated from the GPS sensor through the Haversine formula. The speed over ground (SOG) is measured in knots (kn) and is

the ratio between SDOG and the sailing time, measured in hours. The propeller pitch (P) is measured in degrees and represents the angle between the intersection of the chord line of the blade section and a plane normal to the propeller axis. An anemometer sensor provides data about the true speed (V), measured in knots, and direction (ψ), measured in degrees, of the wind. The latter is obtained as the difference between the true wind angle in earth system and the COG. Additional information on the wind variables can be found on Bocchetti et al. (2015). From the two anemometer variables, the longitudinal component of the wind is calculated as $V \cos \psi$, while the transversal component is calculated as $|V \sin \psi|$. Note that a positive (respectively negative) longitudinal component of the wind means that the wind blows from stern (respectively bow). Moreover, a data fusion process also allows the integration of marine data into the data set, i.e. weather forecasts about the sea state furnished by private held weather service provider. The sea state is characterized by the provider through the typical parameters, viz. height and period, used to model waves that, in turn, are roughly divided into two components: wind-driven waves, or simply waves (generated by the immediate local wind) and swell (generated by distant weather systems and usually having larger period). In particular, the height, measured in meters, is defined as the vertical distance from wave crest to wave trough; whereas, the period, measured in seconds, represents the time between successive crests of a train of waves passing a fixed point in a ship, at a fixed angle of encounter (Lackenby, 1978). Regarding the CO₂ measurement, MRV regulations propose direct and indirect methods. The direct method determines the amount of CO₂ emitted measuring the flow of these emissions passing in exhaust gas funnels. Instead, the indirect method calculates the CO₂ emissions based on the fuel consumption. The direct method is based on the determination of CO₂ emitted that flow in exhaust gas stacks based on the measurement of the CO₂ in the exhaust gas and the measurement of the volume of the exhaust gas flow per unit of time. This method is very sensitive to the calibration and the uncertainty related to the measurement devices. Whereas, the class of indirect method determines calculates CO₂ emissions as a product of the whole amount of fuel consumption of the main and auxiliary engines, boilers, gas turbines and inert gas generators times the so called *emission factor*, which is calculated as the average emission rate of a GHG relative to the activity data of a source stream, assuming complete oxidation for combustion and complete conversion for all other chemical reactions. In this chapter, we use the indirect method and we focus on the main engines only.

In what follows, we list the functional variables chosen for the analysis in this work that are obtained from the signals acquired by the DAQ system. The functional response is the signal corresponding to the CO₂ emissions per hour along the entire voyage. In order to select functional covariates among the available signals, a very long preliminary investigation was carried out to identify the covariates that could better explain the CO₂ emissions. However, in practice, many signals that could have played the role of covariates were not able to be measured accurately. The intersection between the set of candidate and truly measurable covariates was finally identified after an intensive exchange of information and experience with marine engineers, shipping managers and operators. The following nine functional covariates have been identified, which are, thus, assumed as a characterization of the ship operational conditions. In particular, we include in the analysis the SOG, the left propeller pitch, the right propeller pitch, the transversal component of the wind, the longitudinal component of the wind, the wave height, the wave period, the swell height and the derivative of the cumulative navigation time. Figures 3.1 and 3.2 show the profiles of covariates and response, respectively, in the reference data set used for model building and control chart limits estimation. The functional data are obtained from discrete observations as described

3. FUNCTIONAL REGRESSION CONTROL CHART FOR THE MONITORING OF SHIP CO₂ EMISSIONS

Figure 3.1. Functional covariates in the reference data set.

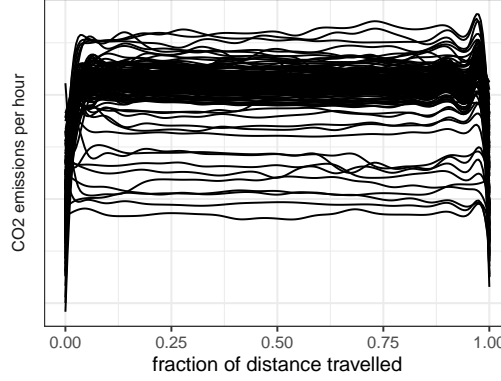


in Section 3.4.

3.3 Methodology

The FRCC is a general framework for profile monitoring that can be divided into three main steps. Firstly, *(i)* define a MFLR model which links the functional response variable \tilde{Y} , defined on the compact domain \mathcal{T} and a vector $\tilde{\mathbf{X}}$ of random functional covariates $\tilde{X}_1, \dots, \tilde{X}_p$, defined on the compact domain \mathcal{S} . Secondly, *(ii)* define the estimation method of the chosen model, and thirdly, *(iii)* define the monitoring strategy of the functional residual defined as the difference between the fitted value and the observed value of \tilde{Y} . In what follows, we assume that $\tilde{X}_1, \dots, \tilde{X}_p$ and \tilde{Y} have smooth realizations in $L^2(\mathcal{S})$ and $L^2(\mathcal{T})$, i.e., the Hilbert spaces of square integrable functions defined on $\mathcal{S}, \mathcal{T} \subset \mathbb{R}$.

To obtain a specific implementation of the FRCC, we assume that the vector of covariates

Figure 3.2. Functional response in the reference data set. For confidentiality reasons, y -axis labels are omitted.

linearly influence the response through the MFLR model, that is

$$Y(t) = \int_{\mathcal{S}} (\boldsymbol{\beta}(s, t))^T \mathbf{X}(s) ds + \varepsilon(t) \quad t \in \mathcal{T}, \quad (3.1)$$

where Y and \mathbf{X} are the standardized versions of \tilde{Y} and $\tilde{\mathbf{X}}$, obtained through the transformation approach of Chiou et al. (2014). The regression coefficient $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$, is a vector where β_i 's are square integrable bivariate functions defined on the closed interval $\mathcal{S} \times \mathcal{T}$, and the random error function ε has zero mean and variance function v_ε^2 , and is independent of \mathbf{X} .

Following Centofanti et al. (2020b), we propose to use an estimation method based on the multivariate Karhunen-Loève's Theorem Happ and Greven (2018). In particular, we assume that the standardized covariate and response variables can be represented as follows

$$\mathbf{X} = \sum_{i=1}^{\infty} \xi_i^X \boldsymbol{\psi}_i^X \quad Y = \sum_{i=1}^{\infty} \xi_i^Y \psi_i^Y, \quad (3.2)$$

where $\boldsymbol{\psi}_i^X = (\psi_{i1}^X, \dots, \psi_{ip}^X)^T$ and ψ_i^Y are *principal components* (PCs), i.e., the eigenfunctions of the covariance operator of the standardized covariates and response variable corresponding to the eigenvalues λ_i^X and λ_i^Y in descending order, respectively, and $\xi_i^X = \sum_{j=1}^p \int_{\mathcal{S}} X_j(s) \psi_{ij}^X(s) ds$ and $\xi_i^Y = \int_{\mathcal{T}} Y(t) \psi_i^Y(t) dt$ are the *scores*, such that $E(\xi_i^X) = 0$, $E(\xi_i^X \xi_j^X) = \lambda_i^X \delta_{ij}$ and $E(\xi_i^Y) = 0$, $E(\xi_i^Y \xi_j^Y) = \lambda_i^Y \delta_{ij}$, with δ_{ij} the Kronecker delta. As demonstrated in Chiou et al. (2016), the regression coefficient can be expressed as follows

$$\boldsymbol{\beta}(s, t) = \sum_{i,j=1}^{\infty} \frac{E(\xi_i^X \xi_j^Y)}{\lambda_i^X} \boldsymbol{\psi}_i^X(s) \psi_j^Y(t) \quad s \in \mathcal{S}, t \in \mathcal{T}. \quad (3.3)$$

An estimator of the regression coefficient is readily obtained by considering the truncated version of Equation (3.3), i.e.,

$$\boldsymbol{\beta}_{LM}(s, t) = \sum_{i=1}^L \sum_{j=1}^M b_{ij} \boldsymbol{\psi}_i^X(s) \psi_j^Y(t) \quad s \in \mathcal{S}, t \in \mathcal{T}, \quad (3.4)$$

3. FUNCTIONAL REGRESSION CONTROL CHART FOR THE MONITORING OF SHIP CO₂ EMISSIONS

with $b_{ij} = E(\xi_i^X \xi_j^Y) / \lambda_i^X$, and $L, M < \infty$. Plugging Equation (3.4) into Equation (3.1), due to the orthonormality of the PCs ψ_i^X and ψ_i^Y , Equation (3.1) becomes

$$\boldsymbol{\xi}_M^Y = (\mathbf{B}_{LM})^T \boldsymbol{\xi}_L^X + \boldsymbol{\epsilon}_M, \quad (3.5)$$

where $\boldsymbol{\xi}_M^Y = (\xi_1^Y, \dots, \xi_M^Y)^T$, $\boldsymbol{\xi}_L^X = (\xi_1^X, \dots, \xi_L^X)^T$, $\boldsymbol{\epsilon}_M = (\epsilon_1, \dots, \epsilon_M)^T$ and $\mathbf{B}_{LM} = \{b_{ij}\}_{i=1, \dots, L, j=1, \dots, M}$, with $\epsilon_i = \int_{\mathcal{T}} \varepsilon(t) \psi_i^Y(t) dt$. Therefore, the problem of estimating $\boldsymbol{\beta}$ reduces to estimate the matrix \mathbf{B}_{LM} that can be obtained through least squares given n independent realizations $(\tilde{\mathbf{X}}_i, \tilde{Y}_i)$ of (\mathbf{X}, Y) . Then, given the least squares estimator $\hat{\mathbf{B}}_{LM}$ of \mathbf{B}_{LM} , the estimator $\hat{\boldsymbol{\beta}}_{LM}$ of $\boldsymbol{\beta}$ can be calculated as

$$\hat{\boldsymbol{\beta}}_{LM}(s, t) = \left(\hat{\boldsymbol{\psi}}^Y(t) \right)^T \left(\hat{\mathbf{B}}_{LM} \right)^T \hat{\boldsymbol{\Psi}}^X(s) \quad s \in \mathcal{S}, t \in \mathcal{T}, \quad (3.6)$$

where $\hat{\boldsymbol{\psi}}^Y = \left(\hat{\psi}_1^Y, \dots, \hat{\psi}_M^Y \right)^T$, $\hat{\boldsymbol{\Psi}}^X = \left(\hat{\psi}_1^X, \dots, \hat{\psi}_L^X \right)^T$, with $\hat{\psi}_i^Y$ and $\hat{\psi}_i^X = \left(\hat{\psi}_{i1}^X, \dots, \hat{\psi}_{ip}^X \right)^T$ estimators of ψ_i^Y and ψ_i^X , respectively. Finally, an estimator \hat{Y}_{LM} of Y is

$$\hat{Y}_{LM} = \sum_{i=1}^L \sum_{j=1}^M \hat{b}_{ij} \hat{\xi}_i^X \hat{\psi}_j^Y, \quad (3.7)$$

with \hat{b}_{ij} the entries of $\hat{\mathbf{B}}_{LM}$ and $\hat{\xi}_i^X = \sum_{j=1}^p \int_{\mathcal{S}} X_j(s) \hat{\psi}_{ij}^X(s) ds$.

The raw functional residual is defined as

$$e(t) = Y(t) - \hat{Y}_{LM}(t) \quad t \in \mathcal{T}. \quad (3.8)$$

Instead of considering the raw residual in Equation (3.8), following the remarks in Centofanti et al. (2020b), we considered a scaled version of it. In particular, we consider the *studentized functional residual*, defined as

$$e_{stu}(t) = \frac{Y(t) - \hat{Y}_{LM}(t)}{\text{Cov}_Y \left(Y - \hat{Y}_{LM} \right)^{1/2}(t)} \quad t \in \mathcal{T}. \quad (3.9)$$

The residual variance function is estimated as $\widehat{\text{Cov}}_Y \left(Y - \hat{Y}_{LM} \right)(t) = \hat{v}_\varepsilon^2(t) + \hat{\omega}_{LM}(t, t)$, for $t \in \mathcal{T}$, where \hat{v}_ε^2 is an estimator of v_ε^2 and $\hat{\omega}_{LM}$ is defined as

$$\hat{\omega}_{LM}(s, t) = \text{Cov} \left(\hat{Y}_{LM} \right)(s, t) = \left(\hat{\boldsymbol{\xi}}_L^X \right)^T \left(\hat{\boldsymbol{\Xi}}_X^T \hat{\boldsymbol{\Xi}}_X \right)^{-1} \hat{\boldsymbol{\xi}}_L^X \hat{\boldsymbol{\psi}}_M^Y(s)^T \hat{\boldsymbol{\Sigma}}_{\epsilon_M} \hat{\boldsymbol{\psi}}_M^Y(t) \quad s \in \mathcal{S}, t \in \mathcal{T}, \quad (3.10)$$

where $\hat{\boldsymbol{\xi}}_L^X$ is the estimator of the score vector $\boldsymbol{\xi}_L^X$ of \mathbf{X} , $\hat{\boldsymbol{\Xi}}_X^T \hat{\boldsymbol{\Xi}}_X$ is the estimator of $n \text{Cov}(\boldsymbol{\xi}_L^X, \boldsymbol{\xi}_L^X)$, $\hat{\boldsymbol{\psi}}_M^Y$ is the estimator of the vector of the first M eigenfunctions of Y , and $\hat{\boldsymbol{\Sigma}}_{\epsilon_M}$ is the estimator of $\text{Cov}(\boldsymbol{\epsilon}_M)$. As described in Centofanti et al. (2020b), the mean of the studentized functional residual is less influenced by covariate mean shifts with respect to the raw residual. Indeed, the aim of $\text{Cov}_Y \left(Y - \hat{Y}_{LM} \right)^{1/2}$ is to weight the raw residual

on the basis of its uncertainty, such that for an extreme realization of \mathbf{X} , the residual is heavily scaled.

We use a monitoring strategy based on the Hotelling's T^2 and the SPE control charts (Woodall et al., 2004; Noorossana et al., 2012; Grasso et al., 2016; Pini et al., 2017) applied to e_{stu} . In particular, the studentized functional residual e_{stu} is approximated as

$$e_{stu,K} = \sum_{i=1}^K \xi_i^e \psi_i^e, \quad (3.11)$$

where the scores $\xi_i^e = \int_{\mathcal{T}} e(t) \psi_i^e(t) dt$ and the PCs ψ_i^e are the eigenfunctions corresponding to the eigenvalues λ_i^e in descending order of the covariance function of e_{stu} . The Hotelling's statistic T^2 is obtained as follows

$$T_e^2 = \boldsymbol{\xi}^e \boldsymbol{\Sigma}_{\boldsymbol{\xi}^e}^{-1} \boldsymbol{\xi}^e, \quad (3.12)$$

where $\boldsymbol{\Sigma}_{\boldsymbol{\xi}^e} = \text{diag}(\lambda_1^e, \dots, \lambda_K^e)$ is the variance-covariance matrix of $\boldsymbol{\xi}^e = (\xi_1^e, \dots, \xi_K^e)^T$. Note that T_e^2 is the squared distance of the projection of e_{stu} from the origin of the space spanned by the PCs standardized for the score variances. Analogously, changes along directions orthogonal to the latter space are monitored by the statistic

$$SPE_e = \int_{\mathcal{T}} (e_{stu}(t) - e_{stu,K}(t))^2 dt. \quad (3.13)$$

The control charts are designed in Phase I by means of a set of n functional studentized residuals $e_{stu,i}$, $i = 1, \dots, n$, obtained by n independent observations $(\tilde{\mathbf{X}}_i, \tilde{Y}_i)$ acquired under IC conditions. Phase I includes also the estimation of the MFLR model unknown parameters, the PCs ψ_i^e and the matrix $\boldsymbol{\Sigma}_{\boldsymbol{\xi}^e}$ (calculated by means of the sample covariance) as well as the estimation of the control limits for both the Hotelling's T^2 and the SPE control charts. The latter can be obtained by means of the $(1 - \alpha)$ -quantiles of the empirical distribution of the two statistics, where α is chosen to control the overall type I error probability. In the monitoring phase (Phase II), the functional studentized residual of a new observation is calculated and an alarm signal is issued if at least one of the corresponding T_e^2 and SPE_e statistics violates the control limits.

3.4 Results and Discussion


In this section we show the results of the application of the FRCC to the data set described in Section 3.2. In particular, the retrospective and perspective phases, i.e. Phase I and Phase II, are described in Section 3.4 and Section 3.4, respectively. Moreover, in Section 3.4, a comparison with simpler monitoring approaches is shown.

Phase I

Phase I comprises the recovery of smooth functional data from the discrete observations for each voyage (Section 3.4), the identification of the reference data set of IC observations (Section 3.4), and, the estimation of the MFLR model as described in Section 3.3 (Section 3.4).

3. FUNCTIONAL REGRESSION CONTROL CHART FOR THE MONITORING OF SHIP CO₂ EMISSIONS

Figure 3.3. Median, across all voyages in the Phase I data set, of the generalized cross-validation (GCV) error as a function of the number of basis functions, for each functional variable.




Data smoothing

The first step of the analysis is to get smooth functional data from the discrete observations for each voyage of the ship. We use B-spline basis expansion and penalized least squares to estimate the corresponding basis coefficients. A common approach is to set a quite large number of basis functions and then select the optimal smoothing parameter by minimizing the generalized cross-validation (GCV) error (Ramsay and Silverman, 2005). However, the number of available discrete points is above 200 for each voyage and, by following this approach, the GCV criterion leads to choosing the smoothing parameter equal to zero in practice for all functional variables. This is a typical problem of overfitting, as also pointed out by Reiss and Todd Ogden (2009), which show that at finite sample sizes GCV is likely to develop multiple minima and under-smooth. Therefore, we encourage parsimony and achieve regularization by choosing a small, efficient number of basis functions, with the smoothing parameter fixed to a small positive value (i.e. 10^{-10}) to ensure identifiability. In Figure 3.3, we plot the GCV error against the number of B-spline basis functions. While increasing the number of basis functions reduces the GCV error, we select 25 basis functions for all functional variables as the elbow point of these curves.

Reference data set

Once functional data are obtained, it is necessary to identify a reference data set that can be used for model building and estimation of the control chart limits. Then, we consider a set of historical voyages as a starting point, from which observations that are not representative of the IC conditions have to be removed. Specific Phase I techniques are designed for the problem of eliminating anomalous observations from the reference data set and generally lead to the estimation of different limits from the ones calculated in the Phase II control charts. In this work we use the FRCC also in Phase I and integrate domain knowledge to establish which voyages are considered anomalous and have to be excluded. More details on the model building choices can be found in the next subsection. Figure 3.4 shows the FRCC applied to the initial data set. The x -axis label VN is a progressive counting label denotes subsequent voyages in the data set. Moreover, Figure 3.5 shows some OC studentized residual profiles

Figure 3.4. FRCC used for Phase I monitoring on the initial data set of 190 observations, to remove outliers and define the reference data set.



correctly signaled as anomalous. After an iterative process of identifying outliers, detecting anomalies, removing the observations from the data set and re-fitting the model, we end up with a reference data set of 169 voyages.

Model building

The FRCC relies on the choice of L and M in Equation (3.4), as well as K in Equation (3.11). Figure 3.6 shows the cumulative fraction of variance explained by the functional principal components in the multivariate functional covariates, the functional response, and the functional studentized residuals, respectively. In Centofanti et al. (2020b), L , M and K are chosen such that the retained functional principal components explain at least 95% of the variability in the data. Based on the results in Figure 3.6, we instead opt for a more parsimonious choice and set the thresholds for the explained variability in the data as 80%, 95%, and 95%, respectively, i.e. we select $L = 7$, $M = 1$, and $K = 8$. The corresponding actual fractions of variance explained are 81%, 96%, and 96%.

We further investigate on the interpretation of the selected functional PCs. Figure 3.7 shows the eigenfunctions of the covariance operator of the standardized multivariate functional covariates. Since they all have unit norm, we multiply them by the square root of the corresponding eigenvalues so that profiles with larger norm are PCs that explain a larger fraction of the total variance in the data. The first PC depends almost entirely on the two propeller pitch variables, the speed over ground, and the navigation time. The latter is negatively correlated with the other variables, and for all these variables their weight is almost constant over the entire functional domain. The second PC strongly depends on the

3. FUNCTIONAL REGRESSION CONTROL CHART FOR THE MONITORING OF SHIP CO₂ EMISSIONS

Figure 3.5. Some functional studentized residuals identified as OC in Phase I, plotted in red against all the other ones in the original Phase I data set, plotted in gray.




Figure 3.6. Eigenvalues of the covariance operator of the functional covariates (a), the functional response (b) and the functional studentized residuals (c), estimated on the reference data set. Vertical dashed lines indicate the selected number of components.

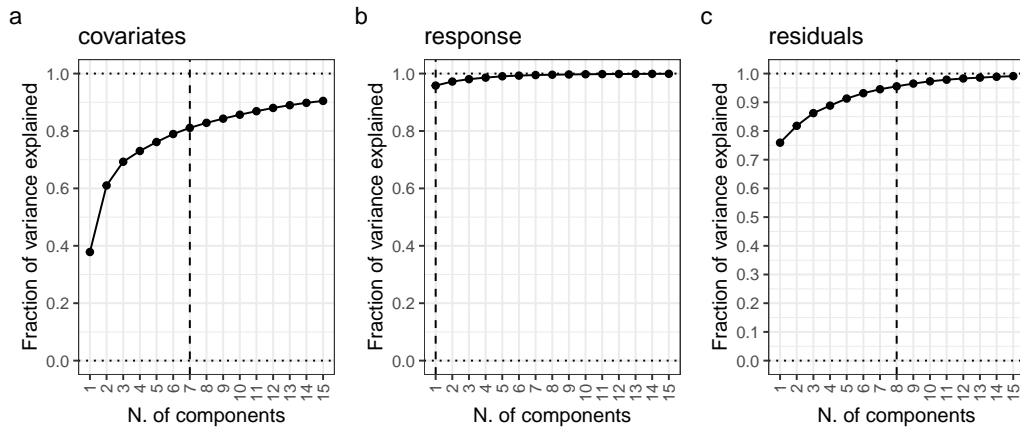
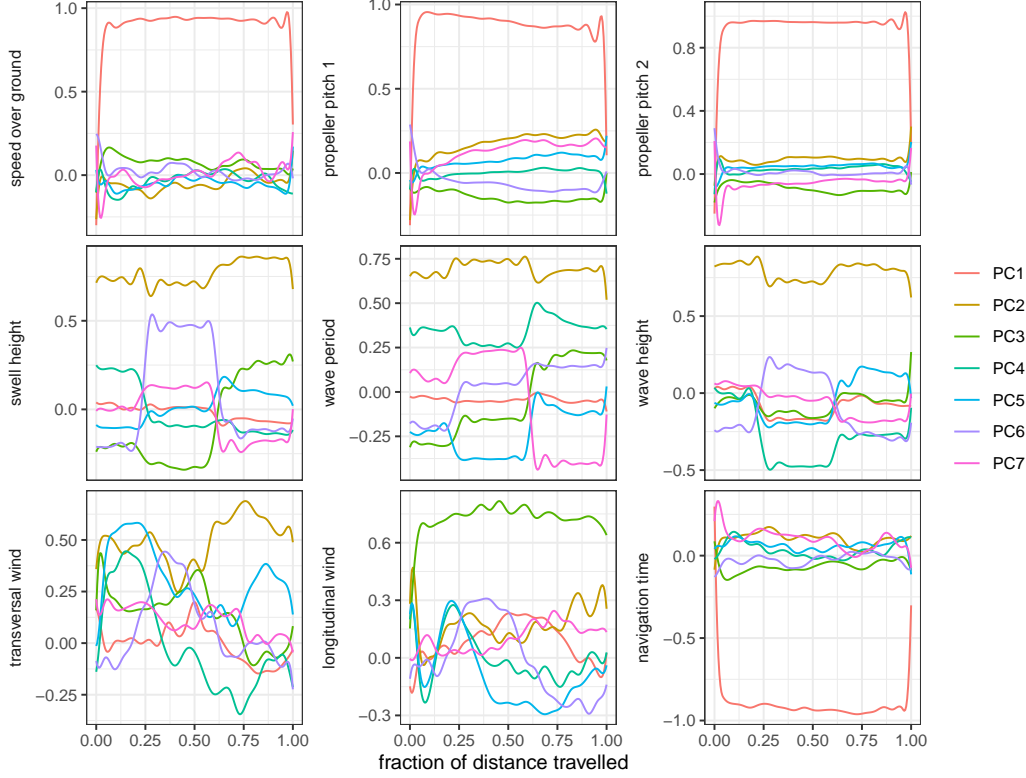


Figure 3.7. Eigenfunctions of the covariance operator of the functional covariates, estimated on the reference data set.

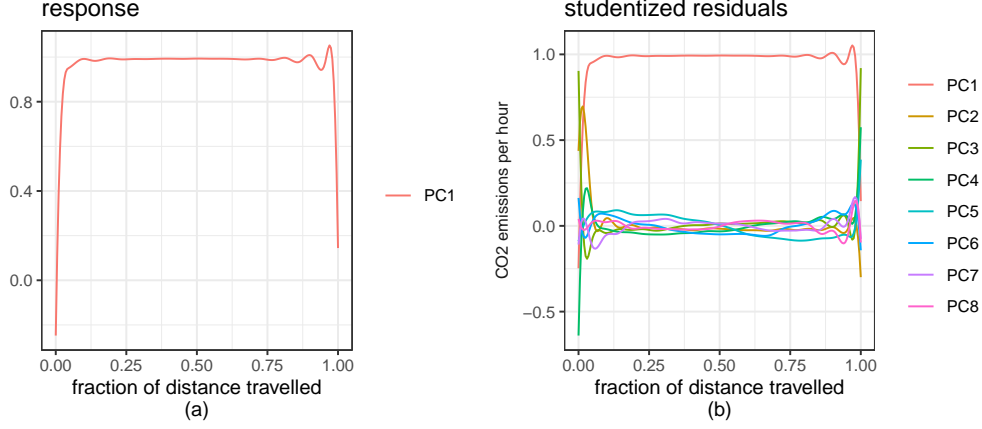


sea variables (swell height, wave height and wave period) and the transversal component of the wind. These functional variables have all positive weight, with some parts of the domain showing slightly larger weight than others. The third PC seems to mainly depend on the longitudinal component of the wind alone. To summarize, the first PC describes how fast the ship is moving, while the second and third PCs capture two distinct aspects of environmental conditions. Figure 3.8a shows the eigenfunction of the covariance operator of the standardized functional response. The first PC alone explains most of the variability in the data, indicating that, after standardization, the fuel consumption per hour is mostly constant functions over the entire domain, apart from the beginning and the end of the voyage. Figure 3.8b shows the eigenfunctions of the covariance operator of the studentized residuals. The first PC depends on the average value over the entire voyage. The second PC looks at the difference between the first and the second half of the voyage, some of the other PCs seem to assign a larger weight to the boundaries of the functional domain, however, the interpretation becomes more difficult with the following PCs.

Figure 3.9 shows the estimated functional coefficients obtained as in Equation (3.6). Since the functional response is approximated with a single functional principal component that in practice is constant over the entire domain, the functional coefficient shows only vertical bands along the direction of t . The most important predictors are the ones associated with

3. FUNCTIONAL REGRESSION CONTROL CHART FOR THE MONITORING OF SHIP CO₂ EMISSIONS

Figure 3.8. Eigenfunctions of the covariance operator of the functional response (a) and the studentized residuals (b).




the first and third PCs in the functional covariates, i.e. the two propeller pitch variables, the speed over ground (all with positive coefficients), and the navigation time and longitudinal wind component (with negative coefficients). The other environmental variables, which are strongly associated to the second PC in the functional covariates, seem to have a relatively lower weight in the regression. To summarize, we can state that voyages with larger value of the propeller pitch and speed over ground variables have a large expected value of the functional fuel consumption, while voyages with larger values of the navigation time and longitudinal wind component have a lower expected fuel consumption.

A final consideration on the obtained MFLR model regards the choice to consider the studentized functional residuals instead of the raw residuals. Figure 3.10 shows the effect of this choice. Some of the more extreme raw functional residuals, shown in Figure 3.10a, are attenuated by the studentization (Figure 3.10b). This happens because those residuals correspond to more extreme functional covariates observations as described in Section 3.3.

Phase II

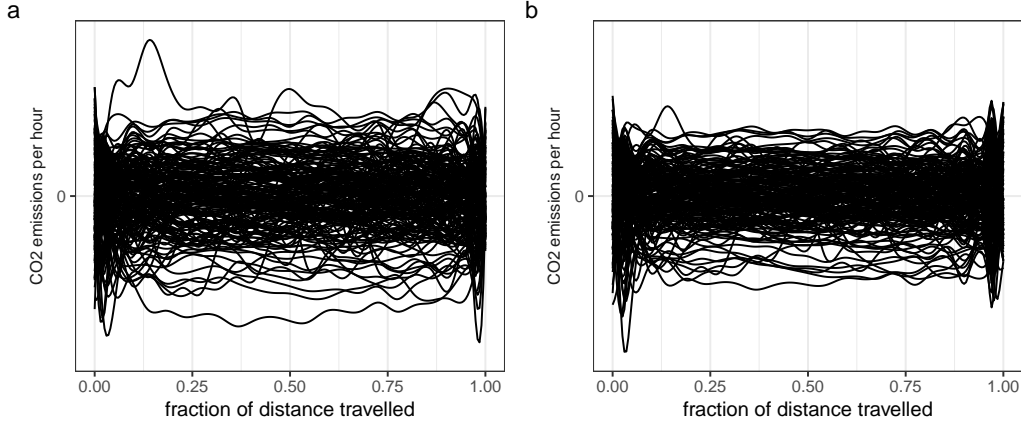
Figure 3.11 shows the FRCC used in Phase II, that is the actual monitoring phase. Points correspond to 22 subsequent voyages, each of them is denoted by a voyage number. For simplicity, we count voyages starting again from 1 to 22, even though these voyages must not be confused with the first 22 Phase I voyages in Figure 3.4. Some Phase II voyages are OC, in particular voyages 1, 3, 15 and 16 are OC in both Hotelling T^2 and SPE control charts, while voyages 20 and 21 are OC in the SPE control chart only. OC observations are generally characterized by some unexpected behavior in the CO₂ emissions that has not been predicted appropriately by the functional regression model in some specific part of the domain, or because the prediction error was moderately large for a considerable part of the voyage. Functional studentized residuals for these OC voyages are plotted in Figure 3.12 against the studentized residuals in the reference data set. Voyages 1, 3 and 15 are far above the upper control limits in both the Hotelling's T^2 and SPE control charts. In particular, the functional studentized residual of voyage 1 is signaled as OC because it is

Figure 3.9. Estimate of the regression coefficient obtained by MFLR as in Equation (3.6), based on the reference data set.



3. FUNCTIONAL REGRESSION CONTROL CHART FOR THE MONITORING OF SHIP CO₂ EMISSIONS

Figure 3.10. (a) Raw vs (b) studentized functional residuals in the reference data set.

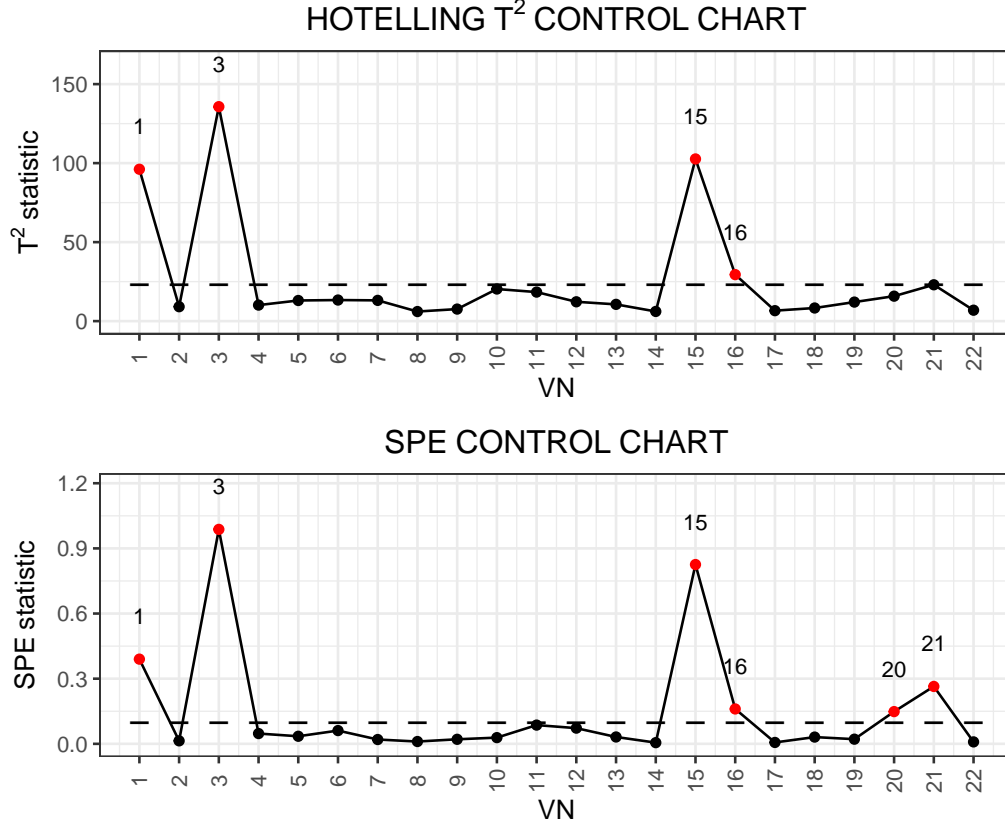


larger than usual on average and because of the large error at the end of the voyage. Voyage 3 shows a large departure from the reference profiles of the studentized residuals, in fact in the first part of the voyage the studentized residual profile is positive, while in the second part it becomes negative, with a large error at the very end of the voyage. Voyage 15 shows the same large error at the end of the voyage as in voyage 3, while for the previous part of the voyage the profile remains always negative. Voyages 16, 20 and 21 are signaled as OC, however they are closer to the upper control limits with respect to the previous voyages, in fact they show less dramatic behavior. In fact, studentized residuals for these voyages show less prominent peaks/valleys in the profiles. In particular, voyage 16 is signaled as anomalous because it shows residual profile with large positive values for almost the entire voyage. Voyage 20 looks more regular and near to zero in the middle of the voyage, with a mild peak and valley at the extremes of the voyage. Voyage 21, if considered pointwise, is almost entirely in the range of the Phase I studentized residuals. However it is signaled as OC because of a sudden profile jump in the middle of the voyage.

Comparison with Other Methods

In this section we try to show if it is actually convenient to use the FRCC rather than simpler approaches. Centofanti et al. (2020b) showed that the FRCC is more powerful than the index based (INBA) control chart, which monitors the area under the response variable, and the RESP control chart, which monitors the coefficients coming from the functional principal component decomposition of the response via Hotelling's T^2 and SPE control charts. We compare the FRCC with the these two control charts and discuss if there are different results in terms of detection of OC observations in this specific application. Figure 3.13 shows the INBA control chart, which is not able to detect any of the voyages signaled by the FRCC and seems not appropriate for this type of application. Moreover, this control chart only detects voyage 22 as OC, which is signaled also by the RESP control chart in Figure 3.14, but, on the other hand, is IC in the FRCC. Apparently, voyage 22 could be an anomalous voyage that the FRCC is not able to correctly identify. By further investigating the functional

Figure 3.11. FRCC in Phase II. Each point corresponds to a voyage and the values of the Hotelling T^2 and SPE statistics are reported. Horizontal dashes indicate upper control limits. Red points denote OC statistics.



response profile and the corresponding studentized residual (Figure 3.15), it turns out that this voyage is signaled as OC in the INBA and RESP control charts because, marginally, the CO_2 emissions were particularly low during the entire voyage. However, these low values of the response variable are predicted well by the MFLR model. Therefore, conditionally on the functional covariates, the response variable is not anomalous, in fact the functional studentized residual profile of voyage 22 is IC in the FRCC. This highlights the convenience in using the FRCC with respect to simpler approaches, when the interest is in monitoring a functional quality of interest conditionally on functional covariates having influence on it. The RESP control chart seems to correctly detect some of the voyages identified by the FRCC, i.e. voyages 1, 3, 15 and 16, but it misses voyages (i.e. 20 and 21), while it signals voyage 11 that is IC in the FRCC. Note that voyage 11 is close the upper control limits in both the RESP control chart and the FRCC.

3. FUNCTIONAL REGRESSION CONTROL CHART FOR THE MONITORING OF SHIP CO₂ EMISSIONS

Figure 3.12. Profiles of the Phase II studentized residuals signaled as OC by the FRCC, plotted as red lines against the studentized residuals in the reference data set, plotted in gray in each panel.

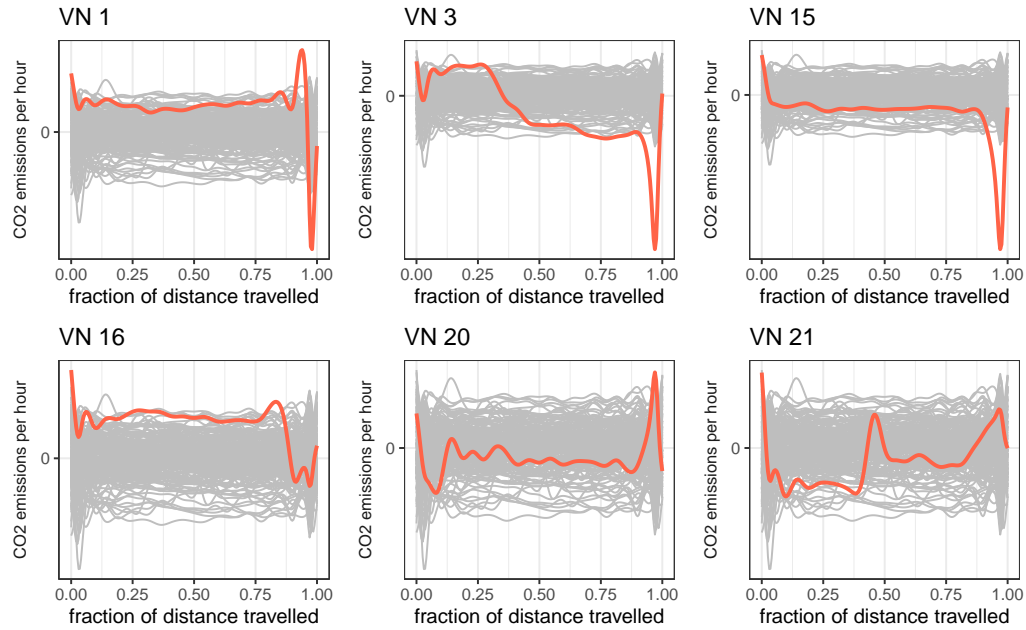


Figure 3.13. INBA control chart





Figure 3.14. RESP control chart




3.5 Conclusions

A particular instance of the functional regression control chart (FRCC) framework proposed in Centofanti et al. (2020b) is applied in this chapter to monitor the ship CO_2 emissions profiles, in order to identify special causes given the recorded values of the functional covariates that may have an influence on them. The quality characteristic is adjusted for the effects of these covariates by means of a multivariate functional linear regression (MFLR) model, based on multivariate functional principal component analysis. That is, the residuals from the MFLR model are monitored by using jointly the Hotelling's T^2 and the *SPE* control charts built on their functional principal component decomposition. The specific implementation of the FRCC relies on the use of the studentized functional residual to take into account the different residual variance at different covariate values. The proposed FRCC demonstrated to be effective in the identification of anomalous observations over the real-case study presented, which is concerned on the data collected during 2020 on a Ro-Pax ship operating in the Mediterranean Sea. Moreover, it was compared against alternative approaches available in the literature, which, however, only look at the marginal distribution

3. FUNCTIONAL REGRESSION CONTROL CHART FOR THE MONITORING OF SHIP CO₂ EMISSIONS

Figure 3.15. Profile of voyage 22 of the Phase II data set. Both the original functional response and the studentized residual are reported. Profiles of voyage 22 are plotted as red lines against the profiles in the reference data set, plotted in gray.



of the functional response, or to some specific features. The competing methods showed a lack of ability to signal some important out of control observations and, in other situations, provide with false alarms.

Finally, one important output achieved in our research is the technological transfer of the FRCC framework to the shipping company Grimaldi Group. The practical applicability of these statistical tools is further investigated by providing the energy saving department of the company with R code able to automatically import new data from the company server, to envelop mathematical and numerical details provided in the chapter, and to routinely produce automatic voyage reports for some Ro-Pax ships of interest from their fleet.

Chapter 4

Function-on-Function Regression for Assessing Production Quality in Industrial Manufacturing

Abstract

Key responses of manufacturing processes are often represented by spatially- or time-ordered data known as *functional data*. In practice, these are usually treated by extracting one or few representative scalar features from them to be used in the following analysis, but this implies the risk of discarding important information, thus leading to drawing only partial conclusions. To take into account all the information available in the measured profiles, new and more sophisticated methods, such as functional data analysis (FDA), must be used. In this chapter, that represents the first contribution in the direction of integrating FDA methods, and in particular functional regression methods, into the manufacturing field, the use function-on-function linear regression modelling is proposed. This approach is based on a finite dimensional approximation of the regression coefficient in terms of two sets of basis functions, as well as on two roughness penalties that control the degree of smoothness of the final estimator. The potential of the proposed method is demonstrated by applying it to a real-case study in powder bed fusion additive manufacturing for metals to predict the mechanical properties of an additively manufactured artifact given the particle size distribution of the powder used for its production.

4.1 Introduction

The fast technological developments of in-line sensing systems and non-contact acquisition architectures allow gathering huge amounts of data ordered by space or time from production processes. Those data are often referred to as *functional data* or *profiles*. The most common practice to handle those data, however, mainly consist in extracting and analysing only one or few representative scalar features from them, possibly based on engineering conjectures. This approach may obviously risk to discard non-trivial information available in the measured profiles and to veil knowledge on the characteristics under study. To avoid scalar feature selection from profiles, functional data analysis (FDA) developed in the statistical field (Ramsay and Silverman, 2005; Horváth and Kokoszka, 2012; Ferraty and Vieu, 2006; Kokoszka and Reimherr, 2017) can be suitably used. In particular, *functional*

regression has become one of the most suited analytic tools to be applied in real applications where it is of interest to know the relationship between a quality characteristic (referred to as *response*) and one or more independent variables (referred to as *covariates*), in which at least one of them is apt to be modelled as a function over a given domain. Excellent overviews are provided by Cuevas (2014), Morris (2015), Horváth and Kokoszka (2012) and Ramsay and Silverman (2005).

In this chapter, we focus on the setting where both the response and covariate are functional and the relationship between them is linear, that is usually referred to as *function-on-function (FoF) linear regression model*. The estimation of FoF linear regression models was originally introduced by Ramsay and Dalzell (1991) through piecewise Fourier basis. The first steps in the estimation development of FoF models are mainly due to Besse and Cardot (1996), that used spline-based approaches, and Ramsay and Silverman (2005), that proposed a general estimation method based on regularization. Then, Yao et al. (2005b) presented a method based on the functional principal components decomposition Yao et al. (2005a) of both the response and covariate and Ivanescu et al. (2015) extended penalized functional regression (PFR) of Goldsmith et al. (2011) to the FoF setting. More recent articles on the FoF model are Luo and Qi (2017) and Luo and Qi (2019).

In order to show its practical applicability and to highlight the benefits of a functional data approach, in this chapter the FoF linear regression model is eventually applied to study a real-case in the additive manufacturing (AM) field. The impact of many characteristics of raw materials on the final properties of the produced parts is in fact still an open research issue and hampers the industrialization of additive technologies (Khajavi et al., 2018; Mani et al., 2017; Slotwinski and Garboczi, 2015; Hague et al., 2004), which are, on the other hand, also very costly. In particular, the considered AM process refers to laser powder bed fusion (L-PBF) technique for metals, which enables the layer-wise production of complex shaped components through spreading and selective laser melting of subsequent metal powder layers onto a substrate plate (Gardan, 2016; Petrovic et al., 2011; Gu et al., 2012).

As is known, not all the powder is melted to form the final component. Then, the final properties of the produced parts are affected by the reuse of powder, which may have been altered, from previous production runs (Barclift et al., 2016; Cordova et al., 2019; Jacob et al., 2017). In particular, the reuse may modify the size and shape of powder particles and affect flowability and packing properties, usable layer thickness, and thus powder-processing behaviour. Specifically, the particle size is affordably measurable and very critical for the mechanical properties of metal parts produced through laser melting processes (Herzog et al., 2016; Khairallah et al., 2016; Sutton et al., 2017). The static mechanical properties are usually summarized into stress-strain curves that are very popular in materials science and engineering field (Ramberg, Walter and Osgood, William R, 1943; Davis, 2003).

In this setting, practitioners would classically settle the analyses by extracting scalar features from both the sample particle size distribution (e.g., 10th, 50th and 90th percentiles) and the observed stress-strain curve (e.g., Yield strength, ultimate tensile strength and elongation at break). In this chapter, the application of FoF linear regression methods allows instead the estimation and prediction of the complete stress-strain curve of produced part (i.e., the functional response) given the powder particle size distribution (i.e., the functional covariate) that, in turn, changes with powder reuse from previous runs.

More in general, both suppliers and users can benefit from the knowledge of particle size distribution for product specifications and manufacturing control, as well as for research and development. Indeed, the laser melting process parameters are usually tuned with respect to a well-defined distribution of the powder particle size.

In addition, the ability to reuse metal powder as many time as possible, obviously without undermining process stability and final part quality, opens up the possibility for these manufacturing processes to become increasingly popular as green technologies.

The model estimation is based on a finite dimensional approximation of the FoF model in terms of two sets of basis functions (one for the response and one for the covariate), as well as on two roughness penalties that control the degree of smoothness of the final estimator.

To quantify the model estimation uncertainty, we propose the use of a simple novel bootstrap methodology. It is broadly based on resampling with replacement of both response and covariate functional observations and is a straightforward extension of the *resampling cases* bootstrap in the classical regression setting (Efron et al., 1979; Davison and Hinkley, 1997). To date, little work has been done on bootstrap for functional data. In this regard it is worth mentioning De Castro et al. (2005) for a bootstrap method to evaluate the range of the forecasts of sulfur dioxide levels near a power plant, and Cuevas et al. (2006), for a comparison of bootstrap confidence bands (obtained with different resampling methods) of several functional estimators. The literature on bootstrap for functional regression is even poorer, with the exception of González-Manteiga and Martínez-Calvo (2011) who obtained pointwise confidence intervals by means of a bootstrap procedure for functional linear model with scalar response and functional covariate.

In Section 4.2 the proposed FoF methodology is described. In particular, the estimation method and model selection issues are presented in Section 4.2 and Section 4.2, respectively. The bootstrap method is illustrated in Section 4.2. The real-case study in the AM field is presented in Section 4.3 and conclusions are reported in Section 4.4. Eventually, Appendix contains additional details on the estimation method. All computations and plots are created by using R R Core Team (2020b).

4.2 Methodology

Let the set of observations (X_i, Y_i) , for $i = 1, \dots, n$, of the functional covariate X and the functional response Y , that are assumed to belong to the Hilbert space of square integrable functions $L^2(\mathcal{S})$ and $L^2(\mathcal{T})$, be defined on the compact intervals \mathcal{S} and \mathcal{T} , respectively. Without loss of generality, we assume that X_i s and Y_i s have zero-mean. In practice, this is obtained by subtracting the covariate and response sample means from the X_i s and Y_i s, respectively. Then, the function-on-function (FoF) linear regression model is defined as:

$$Y_i(t) = \int_{\mathcal{S}} X_i(s) \beta(s, t) ds + \varepsilon_i(t) \quad t \in \mathcal{T} \quad i = 1, \dots, n, \quad (4.1)$$

where the regression coefficient β is in $L^2(\mathcal{S} \times \mathcal{T})$, the Hilbert space of bivariate square integrable functions defined on the interval $\mathcal{S} \times \mathcal{T}$, and ε_i is a zero mean error process with covariance function $K(t_1, t_2)$, t_1 and $t_2 \in \mathcal{T}$ and are independent of X_i .

Model Estimation

The coefficient function β estimator is based on the following penalized least squares estimator:

$$\hat{\beta} = \underset{\beta \in L^2(\mathcal{S} \times \mathcal{T})}{\operatorname{argmin}} \left\{ \sum_{i=1}^n \left\| Y_i - \int_{\mathcal{S}} X_i(s) \beta(s, \cdot) ds \right\|^2 + \lambda_s \|\mathcal{L}_s^{m_s} \beta\|^2 + \lambda_t \|\mathcal{L}_t^{m_t} \beta\|^2 \right\}, \quad (4.2)$$

where $\mathcal{L}_s^{m_s}$ and $\mathcal{L}_t^{m_t}$ are the m_s -th and m_t -th order differential operators with respect to variables s and t , respectively, and $\|\cdot\|$ denotes the L^2 -norm corresponding to the inner

4. FUNCTION-ON-FUNCTION REGRESSION FOR ASSESSING PRODUCTION QUALITY IN INDUSTRIAL MANUFACTURING

product $\langle f, g \rangle = \int fg$; λ_s and λ_t are smoothing parameters that control the trade-off between the roughness of the estimator $\hat{\beta}$ and the goodness of fit. As an example, when $\lambda_s = 0$ and $\lambda_t = 0$, $\hat{\beta}$ is the usual least squares estimator, whereas for $\lambda_s \rightarrow \infty$ and $\lambda_t \rightarrow \infty$, $\hat{\beta}$ converges to a bivariate polynomial function with degree $\max(m_s, m_t) - 1$. Clearly, for any finite n , the optimization problem in Equation (4.2) has infinite solutions; and thus, restrictions on β must be placed in order to obtain a unique estimator. In this regard, by following the approach proposed by Ramsay and Silverman (2005), we consider the approximation of β as double expansion in terms of L basis functions $\{\psi_i^s\}$, defined on \mathcal{S} , and M basis functions $\{\psi_j^t\}$, defined on \mathcal{T} , that is:

$$\beta(s, t) \approx \sum_{i=1}^L \sum_{j=1}^M b_{ij} \psi_i^s(s) \psi_j^t(t) = \boldsymbol{\psi}^s(s)^T \mathbf{B} \boldsymbol{\psi}^t(t) \quad s \in \mathcal{S}, t \in \mathcal{T}, \quad (4.3)$$

where $\mathbf{B} = \{b_{ij}\} \in \mathbb{R}^{L \times M}$ is the coefficient matrix, $\boldsymbol{\psi}^s = (\psi_1^s, \dots, \psi_L^s)^T$ and $\boldsymbol{\psi}^t = (\psi_1^t, \dots, \psi_M^t)^T$. By plugging Equation (4.3) in the model of Equation (4.1), the FoF model becomes:

$$Y_i(t) = \mathbf{X}_i^T \mathbf{B} \boldsymbol{\psi}^t(t) + \varepsilon_i(t) \quad t \in \mathcal{T} \quad i = 1, \dots, n, \quad (4.4)$$

with $\mathbf{X}_i = \int_{\mathcal{S}} X_i(s) \boldsymbol{\psi}^s(s) ds$.

Thus, the problem of estimating β is reduced to the estimation of the unknown coefficient matrix \mathbf{B} . Indeed, the optimization problem in Equation (4.2) becomes:

$$\hat{\mathbf{B}} = \underset{\mathbf{B} \in \mathbb{R}^{L \times M}}{\operatorname{argmin}} \left\{ \sum_{i=1}^n \|Y_i - \mathbf{X}_i^T \mathbf{B} \boldsymbol{\psi}^t\|^2 + \lambda_s \|\mathcal{L}_s^{m_s}(\boldsymbol{\psi}^{sT} \mathbf{B} \boldsymbol{\psi}^t)\|^2 + \lambda_t \|\mathcal{L}_t^{m_t}(\boldsymbol{\psi}^{sT} \mathbf{B} \boldsymbol{\psi}^t)\|^2 \right\}, \quad (4.5)$$

where $\hat{\mathbf{B}}$ is the estimator of \mathbf{B} .

As shown in the Appendix, the estimator of \mathbf{B} is:

$$\operatorname{vec}(\hat{\mathbf{B}}) = \left[\mathbf{W}_Y \otimes \mathbf{X}^T \mathbf{X} + \lambda_s \mathbf{W}_Y \otimes \mathbf{R}_X + \lambda_t \mathbf{R}_Y \otimes \mathbf{W}_X \right]^{-1} (\mathbf{I} \otimes \mathbf{X}^T) \operatorname{vec}(\mathbf{Y}), \quad (4.6)$$

where $\mathbf{W}_X = \int_{\mathcal{S}} \boldsymbol{\psi}^s(s) \boldsymbol{\psi}^s(s)^T ds$, $\mathbf{W}_Y = \int_{\mathcal{T}} \boldsymbol{\psi}^t(t) \boldsymbol{\psi}^t(t)^T dt$, $\mathbf{R}_X = \int_{\mathcal{S}} \mathcal{L}_s^{m_s}[\boldsymbol{\psi}^s(s)] \mathcal{L}_s^{m_s}[\boldsymbol{\psi}^s(s)]^T ds$, $\mathbf{R}_Y = \int_{\mathcal{T}} \mathcal{L}_t^{m_t}[\boldsymbol{\psi}^t(t)] \mathcal{L}_t^{m_t}[\boldsymbol{\psi}^t(t)]^T dt$, $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)^T$ and $\mathbf{Y} = (\mathbf{Y}_1, \dots, \mathbf{Y}_n)^T$, with $\mathbf{Y}_i = \int_{\mathcal{T}} Y_i(t) \boldsymbol{\psi}^t(t) dt$.

Given the matrices $\mathbf{A} \in \mathbb{R}^{j \times k}$ and $\mathbf{B} \in \mathbb{R}^{l \times m}$, $\operatorname{vec}(\mathbf{A})$ indicates the vector of length jk obtained by writing the matrix \mathbf{A} as a column-wise vector, whereas $\mathbf{A} \otimes \mathbf{B}$ represents the Kronecker product matrix of dimensions $jl \times km$. Then, an estimator of β is obtained as follows:

$$\hat{\beta}(s, t) = \boldsymbol{\psi}^s(s)^T \hat{\mathbf{B}} \boldsymbol{\psi}^t(t) \quad s \in \mathcal{S}, t \in \mathcal{T}. \quad (4.7)$$

Particularly, the prediction \hat{Y}_* at given new realization X_* of the covariate X is defined as:

$$\hat{Y}_* = \mathbf{X}_*^T \hat{\mathbf{B}} \boldsymbol{\psi}^t, \quad (4.8)$$

with $\mathbf{X}_* = \int_{\mathcal{S}} X_*(s) \boldsymbol{\psi}^s(s) ds$, and can be regarded as the estimated expectation of the response conditional to the covariate $X = X_*$.

Model Selection

The choice of the number of basis functions L and M is not crucial for penalized methods, because the roughness of the estimator is controlled by the smoothing parameters (Cardot et al., 2003) λ_s and λ_t , which however must be chosen carefully. To this aim, the K-fold cross-validation (Hastie et al., 2009a) is the most popular method because of its simplicity and applicability. Broadly speaking, it consists in dividing the set of observations (X_i, Y_i) , for $i = 1, \dots, n$, into K roughly equal-sized parts. For each part $k = 1, \dots, K$, the model is fit to the remaining $K - 1$ parts and the prediction error is calculated on the observation that belongs to the part k . Then, from the K estimates of the prediction error, the K-fold cross validation estimated prediction error, denoted by $CV(\lambda_s, \lambda_t)$, is obtained as a function of the smoothing parameters λ_s and λ_t that are usually made to vary in a pre-specified grid of values. The values of λ_s and λ_t are chosen such that $CV(\lambda_s, \lambda_t)$ is minimum. However, we suggest whenever possible to thoroughly inspect the $CV(\lambda_s, \lambda_t)$ function and customarily choose the values of λ_s and λ_t to achieve the better prediction performance (i.e. the corresponding value of K-fold cross validation estimated prediction error is near the minimum of $CV(\lambda_s, \lambda_t)$) and interpretability (i.e. λ_s and λ_t are large enough so that $\hat{\beta}$ results sufficiently smooth). The most common choice for K is 10 (Hastie et al., 2009a). Standard choices for the basis functions $\{\psi_i^s\}$ and $\{\psi_j^t\}$ are B-spline, Fourier and Wavelet basis (Ramsay and Silverman, 2005). In particular, we suggest, at least for the real-case study presented in Section 4.3, to use the B-splines of order 4 (i.e. *cubic B-spline*) with evenly spaced knots, due to their good properties and wide applicability (De Boor et al., 1978). Lastly, the standard choice for the order of the differential operators $\mathcal{L}_s^{m_s}$ and $\mathcal{L}_t^{m_t}$ are $m_s = 2$ and $m_t = 2$. That is the curvature of the coefficient function β is penalized in both s and t directions (Ramsay and Silverman, 2005) when cubic B-splines are used.

Bootstrap Analysis

To assess the uncertainty of the estimator $\hat{\beta}$ of the coefficient function β , we propose a simple bootstrap procedure. Bootstrap is a re-sampling technique that allows assigning a measure of accuracy to a sample estimate (Efron and Tibshirani, 1986; Davison and Hinkley, 1997). It is extremely useful when the uncertainty of an estimate is either impossible or too difficult to be gauged by analytical calculations. The proposed bootstrap procedure is the following:

1. Sample with replacement from the set of observations $(X_1, Y_1), \dots, (X_n, Y_n)$ to obtain B bootstrap samples of size n $(X_{1b}^*, Y_{1b}^*), \dots, (X_{nb}^*, Y_{nb}^*)$, for $b = 1, \dots, B$.
2. Fit the model in Equation (4.1) B times to obtain the estimates $\hat{\beta}_1^*, \dots, \hat{\beta}_B^*$ of the coefficient function β .
3. Obtain pointwise confidence intervals with confidence level $1 - \alpha$ by using the $\frac{\alpha}{2}$ and $1 - \frac{\alpha}{2}$ quantiles of the empirical distribution of $\hat{\beta}(s, t)$, estimated by using $\hat{\beta}_1^*(s, t), \dots, \hat{\beta}_B^*(s, t)$ for each $s \in \mathcal{S}$ and $t \in \mathcal{T}$.

Along with the uncertainty quantification of $\hat{\beta}$, the reliability of the prediction \hat{Y}_* in Equation (4.8), given a new realization X_* of the covariate, should be evaluated as well. Let e_1, \dots, e_n be the residuals, obtained as $e_i = Y_i - \mathbf{X}_i^T \hat{\mathbf{B}} \psi^t$. Then, we propose the following procedure:

1. Sample with replacement from the set of observations $(X_1, Y_1), \dots, (X_n, Y_n)$ to obtain B bootstrap samples of size n $(X_{1b}^*, Y_{1b}^*), \dots, (X_{nb}^*, Y_{nb}^*)$, for $b = 1, \dots, B$.

4. FUNCTION-ON-FUNCTION REGRESSION FOR ASSESSING PRODUCTION QUALITY IN INDUSTRIAL MANUFACTURING

2. Fit the model in Equation (4.1) B times to obtain the estimates $\hat{\beta}_1^*, \dots, \hat{\beta}_B^*$ of the coefficient function β and thus the estimates $\hat{B}_1^*, \dots, \hat{B}_B^*$ of the coefficient matrix B .
3. Obtain e_1^*, \dots, e_B^* by sampling with replacement from e_1, \dots, e_n .
4. Compute $\delta_1^*, \dots, \delta_B^*$ realizations of the estimated prediction error δ^* as $\delta_i^* = \mathbf{X}_*^T \hat{B}_i^* \psi^t - (\mathbf{X}_*^T \hat{B} \psi^t + e_i^*)$, where \hat{B} is as in Equation (4.5).
5. Obtain the prediction limits at level $1 - \alpha$ as $\left[\mathbf{X}_*^T \hat{B} \psi^t(t) - \Delta_{1-\frac{\alpha}{2}}(t), \mathbf{X}_*^T \hat{B} \psi^t(t) - \Delta_{\frac{\alpha}{2}}(t) \right]$ for $t \in \mathcal{T}$, where $\Delta_{\frac{\alpha}{2}}(t)$ and $\Delta_{1-\frac{\alpha}{2}}(t)$ are the $\frac{\alpha}{2}$ and $1 - \frac{\alpha}{2}$ quantiles of the empirical distribution of $\delta^*(t)$, estimated by using $\delta_1^*(t), \dots, \delta_B^*(t)$ for each $t \in \mathcal{T}$.

In this case, the quantity to be predicted is $Y_* = \mathbf{X}_*^T B \psi^t + \varepsilon_*$, where the random error ε_* is assumed independent of $\varepsilon_1, \dots, \varepsilon_n$, and the point estimator is $\hat{Y}_* = \mathbf{X}_*^T \hat{B} \psi^t$, as shown in Equation (4.8). Thus, to assess the accuracy of \hat{Y}_* , we estimate the pointwise distribution of the prediction error $\delta = \hat{Y}_* - Y_* = \mathbf{X}_*^T \hat{B} \psi^t - (\mathbf{X}_*^T B \psi^t + \varepsilon_*)$ with the pointwise empirical distribution of $\delta_i^* = \mathbf{X}_*^T \hat{B}_i^* \psi^t - (\mathbf{X}_*^T \hat{B} \psi^t + e_i^*)$. The proposed prediction limits are inspired by the bootstrap prediction limits in the classical regression setting (Davison and Hinkley, 1997).

4.3 A Real-case Study in Laser Powder Bed Fusion Additive Manufacturing

As stated before in the introduction, a real dataset in L-PBF for metals manufacturing is analysed (Del Re et al., 2018). In the latter work, the effect of the number of powder reuses on the tensile properties of the produced AlSi10Mg parts is estimated through scalar one-way analysis of variance (ANOVA). Nine consecutive runs were carried out, starting from as-received powders only, in the first, and from a mix of as-received and reused powder (coming from the previous run), in the following eight runs. The work flow used to test the effect of powder reuse is shown in Fig. 4.1, in which the variables considered also in the present chapter are framed in red.

A total of fifty-two dog-bone shaped tensile specimens (six in each AM run, except in the seventh and the last ones for which only five specimens were available) were produced and tested to measure the mechanical properties achieved under the different processing conditions due to the different number of reuses. All the specimens were produced by a EOS EOSINT M280 3D printer with optimal parameters provided by the producer through the *EOS Part Property Profile AlSi10Mg Speed 30 μ m*. An illustration of the nine printing job produced is given in Fig. 4.2. At each run, a metal powder sample was collected and examined. In particular, a Malvern MS2000 equipment for laser diffraction was used to measure the particle size distribution of each sample according to ASTM B822-17 (ASTM B822-17, 2017) standard. Tensile tests were conducted by an Instron 1185 TSTM at room temperature and with displacement rate of cross head equal to 0.0075 mm/s, according to ISO 6892:2016 (ISO 6892:2016, 2016) and ASTM E8/E8M-16a (ASTM E8/E8M-16a, 2016) standards.

Detailed size information on the powder particle size is usually provided through a histogram, even if, as stated before, more common approach refers only to the 10th, 50th,

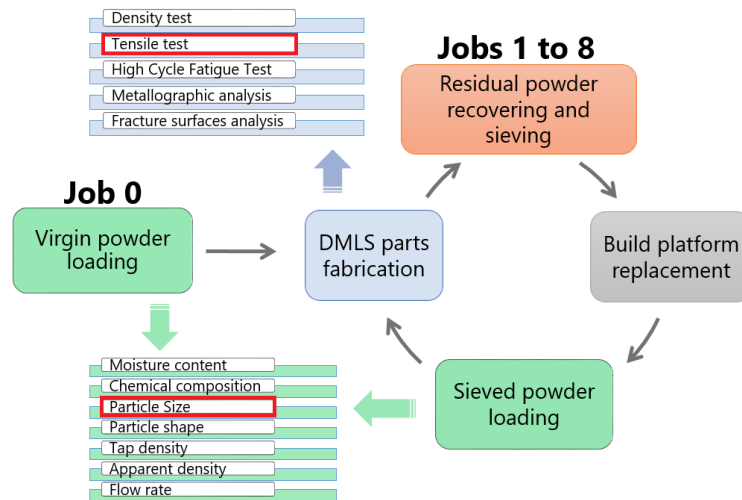


Figure 4.1. Work flow to test the effect of powder reuse in a L-PBF process(Del Re et al., 2018).




Figure 4.2. Example of the printing jobs for studying the effects of powder reuse times (Del Re et al., 2018).

and 90th sample percentiles, denoted by $D10$, $D50$, and $D90$, respectively. A dissertation on tensile strength testing and stress-strain curves is far beyond the objectives of this chapter so just brief hints about them will be given below. The presented description is mainly referred to metals. Other materials, such as plastics or ceramics, can exhibit different behaviours and kind of curves. Detailed information can be found in Davis (2003). As it is known, stress-strain curves are widely used in materials science and engineering to describe, for a given material, the relationship between the applied stress, denoted by σ , and the resulting strain (or elongation) exhibited by the material itself, denoted by ϵ . The former is calculated as the ratio of the applied load (orthogonal to the cross section) and the cross-section of the specimen and is measured in *Pascal*. The latter is the length of elongation exhibited by the material under the applied load divided by the original length of the material. Since both these lengths are usually measured in the same unit, strain can be considered unitless and is generally expressed as a percentage. A $\sigma - \epsilon$ curve provides information on the degree of ductility, or brittleness, of a material and on the maximum loads that it can withstand before it breaks. Some scalar parameters are usually derived from $\sigma - \epsilon$ curves to define the mechanical properties of the specific material (Del Re et al., 2018). The most commonly used are: the *Young's modulus* (or *elastic modulus* or *modulus of elasticity*), E , the *yield strength*, Y_S , the *ultimate tensile strength*, UTS and the *elongation at break*, $\%A$.

Implementation Details and Results

The dataset available for the analysis, consisting of 52 $\sigma - \epsilon$ and 9 PSD curves, can be rearranged by coupling the estimated probability density functions (pdf) of the particle size, denoted by X_i , and the corresponding stress-strain curves, denoted by Y_i , for $i = 1, \dots, n$. The X_i s are obtained through the kernel density estimation method Silverman (2018) proposed by Charpentier and Flachaire (2015), that ensures positiveness of domain points $s \in \mathcal{S} = [0, 120]$ of the estimated pdf's. Instead, the stress-strain curves Y_i are not necessarily defined on the same domain, and thus they need to be opportunely rescaled on a common domain $\mathcal{T} = [0, 1]$ with respect to the elongation. Covariate and rescaled response observations are displayed in Figure 4.3 for illustrative purposes.

The density functions are strictly linked to the number of metal powder reuses. We have more than one observation with the same covariate function. Strictly speaking, we have in total $n = 52$ observations instead of 54 due to two missing specimens.

The basis functions $\{\psi_i^s\}$ and $\{\psi_j^t\}$ used to approximate β in Equation (4.3) are cubic B-splines with evenly spaced knots defined on the domains \mathcal{S} and \mathcal{T} , respectively. The smoothing parameters λ_s and λ_t are set equal to 10^3 by means of 10-fold cross-validation, to trade off predictive performance and interpretability of the resulting estimator, as explained in Section 4.2. Accordingly, the truncation parameters, L and M , in Equation (4.3) are both set equal to 30. The estimator $\hat{\beta}$ is plotted in Figure 4.4 as a function of $s \in \mathcal{S}$ at different $t = 0.1, 0.2, \dots, 0.9$ along with level $1-\alpha$ pointwise bootstrap confidence intervals obtained as described in Section 4.2, with $\alpha = 0.05$ and $B = 100$.

From Figure 4.4, it appears that $\hat{\beta}$ is significantly different from zero in the shadowed regions for all $t \in \{0.1, 0.2, \dots, 0.8, 0.9\}$. Few (resp., many) small particles produce higher (resp., lower) stress values in the stress-strain curve. This behaviour is related to the effect of small particles (and thus big ones) particles on the powder packing factor, that affects how particles bind together during laser melting and therefore the resulting mechanical properties of the produced parts Sames et al. (2016).

As t increases, that is moving along the stress-strain curve, the effect of the particle size

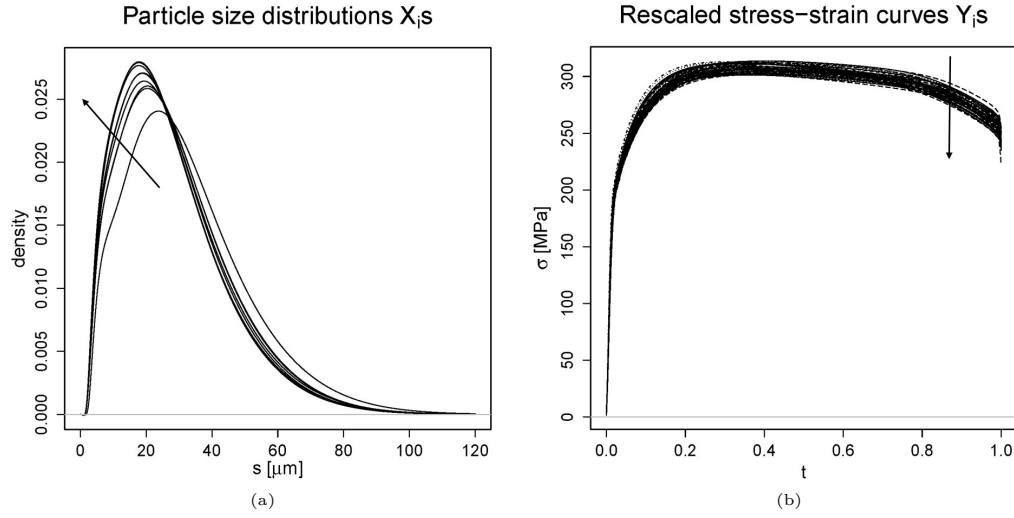


Figure 4.3. (a) Estimated pdf's of powder particle size and (b) rescaled stress-strain curves. The arrows give a broad indication of the curves evolution over number of powder reuses.




Figure 4.4. Estimation $\hat{\beta}$ of the coefficient function β evaluated at $t \in \{0.1, 0.2, \dots, 0.8, 0.9\}$, with upper and lower limits (dotted lines) of the level $1-\alpha$ pointwise bootstrap confidence intervals, for $\alpha = 0.05$. The shadowed regions are where the $\hat{\beta}$ is significantly different from zero.




Figure 4.5. On the left centered predictions (PRE) of one observation, randomly chosen, the true centered realizations of the response (TRUE) and level $1-\alpha$ pointwise bootstrap prediction upper (UL) and lower (LL) limits, for $\alpha = 0.05$. The corresponding non-centered $\sigma - \epsilon$ curve is shown on the right for σ values between 250 MPa and 320 MPa.

distribution on the stress-strain changes after the ultimate tensile strength is reached and the necking occurs, corresponding to $t \simeq 0.3 - 0.4$. From this point, the width of the shadowed region on the left (resp., right) progressively increases (resp., reduces).

For instance, to appreciate the prediction performance of the FoF linear model, the centered and the non-centered predictions of one observation, randomly selected, of the $\sigma - \epsilon$ curve are shown in Figure 4.5, along with the true centered and non-centered realizations of the response. Moreover, level $1-\alpha$ pointwise bootstrap prediction limits, calculated as described in Section 4.2 with $\alpha = 0.05$ and $B = 100$, are displayed as well.

It is clear that the FoF linear regression model estimated by the procedure illustrated in Section 4.2 provides good predictions, and thus could be used to predict the mechanical properties of an artifact produced by the powder bed fusion technique, given the particle size distribution of the powder used for its production. It is worth noting that for small values of t the observed curve is outside the pointwise bootstrap prediction limits, as expected for the 5% of the curve at the given pointwise confidence level $1-\alpha = 0.95$.

4.4 Conclusions

Quality control of additive manufacturing products represents a great challenge due to the complexity of this kind of processes, that are affected by a wide number of aspects related to raw materials, processing parameters, ambient conditions, post-processing operations, etc. To this aim, we proposed a novel approach based on function-on-function regression modelling aimed at assessing the mechanical properties of industrial parts. The potential advantage of the proposed approach was demonstrated by applying it to a real-case study in the laser powder bed fusion additive manufacturing for metals with respect to the models and methods commonly used in the manufacturing field, that need to collapse functional data into scalar features (e.g., sample mean, variance, percentiles). To the best of the authors' knowledge, the present study represents the first contribution in the direction of integrating functional regression methods into the manufacturing field. The developed regression method

is shown to well predict the $\sigma - \epsilon$ curve (i.e. the response function) of a part produced by laser power bed fusion additive manufacturing processes given the particle size distribution of the raw powdered material, that can be affordably obtained before starting a new AM production.

Irrespective of the powder reuse times, the mechanical properties of the produced parts exhibited small variability, with a potential acceptable decay up to the last printing run. The possibility of reusing metal powder, at least for a certain number of times, allows companies to minimize the material waste, thus saving money and reducing the overall environmental impact, that result mandatory in an industrial context.

A step forward could refer to the possibility of adding other covariates, not necessarily functional, which more fully describe the powder properties (e.g. chemical composition, moisture content or tap and apparent density) to increase the predictive power of the method. In addition, further analysis may address the extensibility of the proposed method to other types of profiles (for different materials or kind of inspection tests), since its appeal is likely to increase with increasing complexity of the considered profiles. Indeed, other materials, not limited to metals, or other tests can return more complex curves, for which the functional, rather than the scalar, approach may result much more powerful.

4.5 Appendix

Derivation of the Coefficient Function Estimator

The penalty $\|\mathcal{L}_s^{m_s}(\psi^{sT} \mathbf{B} \psi^t)\|^2$ in the right-hand side of Equation (4.5) may be written as:

$$\begin{aligned}
 \|\mathcal{L}_s^{m_s}(\psi^{sT} \mathbf{B} \psi^t)\|^2 &= \int_{\mathcal{S}} \int_{\mathcal{T}} \mathcal{L}_s^{m_s}[\psi^s(s)^T \mathbf{B} \psi^t(t)] \mathcal{L}_s^{m_s}[\psi^s(s)^T \mathbf{B} \psi^t(t)] ds dt \\
 &= \int_{\mathcal{S}} \int_{\mathcal{T}} \mathcal{L}_s^{m_s}[\psi^s(s)^T] \mathbf{B} \psi^t(t) \psi^t(t)^T \mathbf{B}^T \mathcal{L}_s^{m_s}[\psi^s(s)] ds dt \\
 &= \text{Tr} \left[\left(\int_{\mathcal{S}} \mathcal{L}_s^{m_s}[\psi^s(s)] \mathcal{L}_s^{m_s}[\psi^s(s)^T] ds \right) \mathbf{B} \left(\int_{\mathcal{T}} \psi^t(t) \psi^t(t)^T dt \right) \mathbf{B}^T \right] \\
 &= \text{Tr} [\mathbf{B}^T \mathbf{R}_X \mathbf{B} \mathbf{W}_Y], \tag{4.9}
 \end{aligned}$$

where \mathbf{R}_X and \mathbf{W}_Y are as in Equation (4.6) and $\text{Tr}(\mathbf{A})$ is the trace of a generic matrix \mathbf{A} . Analogously to Equation (4.9), the penalty: $\|\mathcal{L}_t^{m_t}(\psi^{sT} \mathbf{B} \psi^t)\|^2$ in the right-hand side of Equation (4.5) may be written as

$$\|\mathcal{L}_t^{m_t} \beta\|^2 = \text{Tr} [\mathbf{B}^T \mathbf{W}_X \mathbf{B} \mathbf{R}_Y], \tag{4.10}$$

4. FUNCTION-ON-FUNCTION REGRESSION FOR ASSESSING PRODUCTION QUALITY IN INDUSTRIAL MANUFACTURING

where \mathbf{W}_X and \mathbf{R}_Y are as in Equation (4.6). Moreover:

$$\begin{aligned}
\|Y_i - \mathbf{X}_i^T \mathbf{B} \psi^t\|^2 &= \int_{\mathcal{T}} \left[Y_i(t) - \mathbf{X}_i^T \mathbf{B} \psi^t(t) \right]^2 dt \\
&= \int_{\mathcal{T}} \left[Y_i(t)^2 - 2Y_i(t) \psi^t(t)^T \mathbf{B}^T \mathbf{X}_i + \mathbf{X}_i^T \mathbf{B} \psi^t(t) \psi^t(t)^T \mathbf{B}^T \mathbf{X}_i \right] dt \\
&= \int_{\mathcal{T}} Y_i(t)^2 dt - 2 \left(\int_{\mathcal{T}} Y_i(t) \psi^t(t)^T dt \right) \mathbf{B}^T \mathbf{X}_i \\
&\quad + \mathbf{X}_i^T \mathbf{B} \left(\int_{\mathcal{T}} \psi^t(t) \psi^t(t)^T dt \right) \mathbf{B}^T \mathbf{X}_i \\
&= \int_{\mathcal{T}} Y_i(t)^2 dt - 2\mathbf{Y}_i^T \mathbf{B}^T \mathbf{X}_i + \mathbf{X}_i^T \mathbf{B} \mathbf{W}_Y \mathbf{B}^T \mathbf{X}_i,
\end{aligned}$$

and, thus:

$$\sum_{i=1}^n \|Y_i - \mathbf{X}_i^T \mathbf{B} \psi^t\|^2 = \sum_{i=1}^n \int_{\mathcal{T}} Y_i(t)^2 dt - 2 \operatorname{Tr} [\mathbf{X} \mathbf{B} \mathbf{Y}^T] + \operatorname{Tr} [\mathbf{X}^T \mathbf{X} \mathbf{B} \mathbf{W}_Y \mathbf{B}^T], \quad (4.11)$$

where \mathbf{X} and \mathbf{Y} are as in Equation (4.6).

Then, the optimization problem in Equation (4.5) becomes:

$$\hat{\mathbf{B}} = \operatorname{argmin}_{\mathbf{B} \in \mathbb{R}^{L \times M}} \left\{ \sum_{i=1}^n \int_{\mathcal{T}} Y_i(t)^2 dt - 2 \operatorname{Tr} [\mathbf{X} \mathbf{B} \mathbf{Y}^T] + \operatorname{Tr} [\mathbf{X}^T \mathbf{X} \mathbf{B} \mathbf{W}_Y \mathbf{B}^T] \right. \quad (4.12)$$

$$\left. + \lambda_s \operatorname{Tr} [\mathbf{R}_X \mathbf{B} \mathbf{W}_Y \mathbf{B}^T] + \lambda_t \operatorname{Tr} [\mathbf{W}_X \mathbf{B} \mathbf{R}_Y \mathbf{B}^T] \right\}. \quad (4.13)$$

By taking the derivative of the right-hand side of Equation (4.12) with respect to \mathbf{B} and setting it equal to zero, we obtain:

$$\mathbf{X}^T \mathbf{X} \mathbf{B} \mathbf{W}_Y + \lambda_s \mathbf{R}_X \mathbf{B} \mathbf{W}_Y + \lambda_t \mathbf{W}_X \mathbf{B} \mathbf{R}_Y = \mathbf{X}^T \mathbf{Y}, \quad (4.14)$$

and, thus, by applying the vec operator on both sides of Equation (4.14), we obtain:

$$\operatorname{vec} (\mathbf{X}^T \mathbf{X} \mathbf{B} \mathbf{W}_Y) + \lambda_s \operatorname{vec} (\mathbf{R}_X \mathbf{B} \mathbf{W}_Y) + \lambda_t \operatorname{vec} (\mathbf{W}_X \mathbf{B} \mathbf{R}_Y) = \operatorname{vec} (\mathbf{X}^T \mathbf{Y}). \quad (4.15)$$

By using the fact that, for generic matrices \mathbf{A} , \mathbf{B} and \mathbf{C} of appropriate dimensions, $\operatorname{vec} (\mathbf{A} \mathbf{B} \mathbf{C}) = (\mathbf{C}^T \otimes \mathbf{A}) \operatorname{vec} (\mathbf{B})$ and $\operatorname{vec} (\mathbf{A} \mathbf{B}) = (\mathbf{I} \otimes \mathbf{A}) \operatorname{vec} (\mathbf{B})$, where \mathbf{I} is an identity matrix of appropriate dimensions, the estimator of \mathbf{B} is:

$$\operatorname{vec} (\hat{\mathbf{B}}) = \left[\mathbf{W}_Y \otimes \mathbf{X}^T \mathbf{X} + \lambda_s \mathbf{W}_Y \otimes \mathbf{R}_X + \lambda_t \mathbf{R}_Y \otimes \mathbf{W}_X \right]^{-1} (\mathbf{I} \otimes \mathbf{X}^T) \operatorname{vec} (\mathbf{Y}). \quad (4.16)$$

Chapter 5

Functional Clustering Methods for Resistance Spot Welding Process Data in the Automotive Industry

Abstract

Quality assessment of resistance spot welding (RSW) joints of metal sheets in the automotive industry is typically based on costly and lengthy off-line tests that are unfeasible on the full production, especially on large scale. However, the massive industrial digitalization triggered by the Industry 4.0 framework makes available, for every produced joint, on-line RSW process parameters, such as, in particular, the so-called dynamic resistance curve (DRC), which is recognized as the full technological signature of the spot welds. Motivated by this context, the present chapter means to show the potentiality and the practical applicability to clustering methods of the functional data approach that avoids the need for arbitrary and often controversial feature extraction to find out homogeneous groups of DRCs, which likely pertain to spot welds sharing common mechanical and metallurgical properties. We intend to provide an essential hands-on overview of the most promising functional clustering methods, and to apply the latter to the DRCs collected from the RSW process at hand, even if they could go far beyond the specific application hereby investigated. The methods analyzed are demonstrated to possibly support practitioners along the identification of the mapping relationship between process parameters and the final quality of RSW joints as well as, more specifically, along the priority assignment for off-line testing of welded spots and the welding tool wear analysis. The analysis code, that has been developed through the software environment R, and the DRC data set are made openly available online at <https://github.com/unina-sfere/funclustRSW/>.

5.1 Introduction

Resistance Spot Welding (RSW) is the most common technique employed in joining metal sheets during body-in-white assembly of automobiles (Ighodaro et al., 2016; Zhou and Cai, 2014), mainly because of its adaptability for mass production (Martín et al., 2014). Typical car body contains about 5000 spot welds joining metal sheets of different materials and thicknesses (Zhao et al., 2006). The quality of many critical spots (El-Banna et al., 2008) is routinely controlled in order to guarantee the structural integrity and solidity of welded

assemblies per vehicle (Martín et al., 2014). Quality assessment is typically based on tests performed at the end of the RSW process (off-line) on finished sub-assemblies through direct or indirect evaluation of weld-joint key characteristics (Raoelison et al., 2012). Off-line testing is, however, costly and lengthy and thus unfeasible on the full production, especially on large scale.

In the modern automotive Industry 4.0 framework, automatic acquisition systems allow to routinely control welders during running operations (on-line) through the continuous record of a large volume of process parameters. In particular, the so-called *dynamic resistance curve* (DRC) is the most important process parameter acquired on-line (Zhou and Cai, 2013) and is popularly recognized as the full technological signature of the metallurgical development of a spot weld (Dickinson et al., 1980).

In this scenario, a paramount issue constantly faced by practitioners is the identification of homogeneous groups (clusters) of spot welds based on DRC observations, in terms of mechanical and metallurgical properties. The identification of clusters with a convenient interpretation is useful for exploring the mapping relationship between process parameters and the final quality of the RSW joints produced, and, in general, for supporting the experience-based learning of any technological process. In this regard, the most common practice in industry is to analyze one or few scalar features extracted from the acquired DRC, even though feature extraction is known to be often difficult, arbitrary and risky of collapsing useful information.

On the contrary, in this chapter, each DRC observation is suitably modelled as a function defined on the time domain, i.e., as *functional datum*. Functional data analysis (FDA) (Ramsay and Silverman, 2005; Horváth and Kokoszka, 2012; Ferraty and Vieu, 2006; Kokoszka and Reimherr, 2017) is the set of methods that consider functional data as its founding elements. Clustering functional data is usually a difficult task, because of the intrinsic infinite dimensionality of the problem. A thorough overview of functional clustering methods can be found in Ramsay and Silverman (2005) and Ferraty and Vieu (2006). Then, it is worth mentioning Cuesta-Albertos and Fraiman (2007) who proposed a pure functional version of the k-means algorithm, which is very popular in the multivariate setting (Everitt et al., 2011), as an alternative to the method of Abraham et al. (2003), who instead applied the k-means algorithm to the coefficients obtained by projecting the original profiles onto a lower-dimensional subspace spanned by B-spline basis functions. Another version of k-means algorithm is that of Chiou and Li (2007), which relies on a particular distance between truncations at a given order of the functional principal components expansion (Ramsay and Silverman, 2005; Hall and Hosseini-Nasab, 2006). This version can be broadly regarded as an instance of the method proposed by Bouveyron and Jacques (2011), who modelled the functional principal components through Gaussian mixture. Some parsimony constraints on the variance parameters are also considered to define a family of parsimonious sub-models. A similar methods, which is based on a functional principal components expansion of the functional observations, was proposed by Jacques and Preda (2013). The work of James and Sugar (2003) is recognized as the first example of model-based procedure for functional data clustering, as well as the method proposed by Giacomini et al. (2013), which, in particular, relies on the wavelet decomposition of the functional observations, and is particularly appropriate for peak-like data, as opposed to methods based on splines. More recently, Delaigle et al. (2019) proposed a functional k-means algorithm able to cluster observations asymptotically perfectly. A sparse functional clustering procedure, that is clustering functional data while jointly selecting the most relevant features, was developed by Floriello and Vitelli (2017) and, in particular, by Vitelli (2019) who accounted for possible curve misalignments. For the

sake of completeness, Bayesian approaches have appeared as well (Ray and Mallick, 2006; Rodríguez et al., 2009; Rigon, 2019) in the literature, even if they are beyond the scope of this chapter.

After providing Section 5.2 with the technological background and the description of the functional DRC data set collected from the RSW process that has motivated this research, we give in Section 5.3 a deeper hands-on illustration of the most promising functional clustering methods to be applied to the DRC data set at hand. In Section 5.4, we discuss and interpret from technological viewpoint the main results obtained, even if the proposed approach could go far beyond the specific application hereby investigated. We conclude by Section 5.5 with a discussion of issues highlighted by this data set and a broader perspective of the potentiality of the proposed methods. Technical details for each of the clustering methods implemented in this chapter are presented in the Appendix.

The DRC data set and the R R Core Team (2020a) code are made openly available online Capezza et al. (2020a) to allow the reader to possibly investigate other approaches with this data set and to encourage the fruitful spread of functional data clustering methods among practitioners in industry.

5.2 Technological Background and Data Structure

The considered RSW process (Zhang and Senkara, 2011) refers to an autogenous welding process in which two overlapping steel galvanized sheets are joint together, without the use of any filler material, at discrete spots. Joints are performed by applying pressure to the weld area from two opposite sides by means of two copper electrodes. Voltage applied to the electrodes generates a current flowing between them through the material. The electrical current flows because the resistance offered by metals causes a large heat generation (Joule effect) that increases the metal temperature at the faying surfaces of the work pieces up to the melting point. Finally, due to the mechanical pressure of the electrodes, the molten metal of the metal sheets jointed cools and solidifies forming the so-called weld *nugget* (Raoelison et al., 2012; Manufacturers' Alliance, 2003).

The typical shape of a DRC acquired during this process is displayed in Figure 5.1 for illustrative purposes. In the light of Dickinson et al. (1980), it mainly depends on physical changes induced in the material by the ongoing welding process and can be roughly outlined into five stages, as well depicted by Adams et al. (2017). For the sake of conciseness, these stages can be summarized as influenced by two main concurrent effects due to (a) the metal electrical resistivity and (b) the contact area among the metal sheets to joint. These effects develop during the RSW process by means of the heat produced by the current flow and the clamping pressure generated by copper electrodes. In particular, DRC values are proportional to (a), which increases with material temperature.

On the contrary, DRC values are decreasing with b . That, in turn, is increasing with two main factors: (b.1) the deformation, due to the clamping force, of the surface asperities, that are softened by the high temperatures; and (b.2) the melting of the metal, that guarantees the sheet continuity by occupying the interstices between the work pieces to weld. So stated, the typical DRC behaviour (Figure 5.1) can be interpreted by the turnover of the effects due to (a) and (b). Specifically, DRC decreases at first because of the effect due to (b.1), which dominates effect due to (a) until the local minimum; then, conversely, DRC increases because the effect due to (a) dominates effect due to (b) until the local maximum, which represents the beginning of the nugget formation. Finally, the DRC decreases slowly to the end of the RSW process, because the effect due to factor (b.2) dominates that due to (a) to




Figure 5.1. Typical DRC behaviour.

a lesser extent. In a nutshell, DRC behaviour can be roughly outlined by one local minimum point, one local maximum point, and the resistance value at the end of the welding process.

The data set for the problem at hand consists of 538 DRCs that are plotted in Figure 5.2 and pertains to spot welds of the same type collected during RSW lab tests at Centro Ricerche Fiat (CRF). The latter have been carried out on coupons of two sheets having thickness equal to 0.7 mm and 1.3 mm and made of *FE220BH* and *FE600DP* galvanised steels, respectively. The energy was supplied in a single pulse of current. The weld time period is 237 ms. Strictly speaking, the values of electrical resistance used to obtain each DRC observation are not direct measurements, but obtained, according to the first Ohm's law (Ohm, 1827), as the ratio between the voltage at electrode tips and the current intensity measurements. For each DRC observation, these have been collected at a regular grid of 238 points equally spaced by 1 ms. In particular, the electrode tip voltage has been measured using dressed copper wires attached to the electrodes. Whereas, the current intensity has been measured by means of an air-core toroid in the primary of the welder transformer. Copper wires are checked up to ensure their integrity at the beginning of every welding cycle. Electrical resistance of the metal sheets is assumed much larger than that of the copper electrodes. That is, the copper electrode resistance does not practically affect the measurement of metal sheet resistance.

Note that the raw data plots for the 538 DRCs at hand yield shape and features coherent with those discussed with reference to Figure 5.1, but show non-negligible variability that motivates the goal of the present chapter in supporting practitioners to build homogeneous groups of DRCs. The intent is to identify through the latter spot welds having similar mechanical and metallurgical properties, and groups themselves that stand apart from one another. In particular, clustering methods, and even more their functional version, will be of great value in this regard with the ultimate goal of guiding practitioners along the priority assignment for off-line testing of welded spots and the welding tool wear analysis.




Figure 5.2. Raw data plot of spot welding DRCs.

5.3 Functional Data Clustering Approaches for Dynamic Resistance Curves

Usually, functional data consist of independent realizations X_1, \dots, X_n of a functional random variable X with values in an infinite dimensional space, which is typically taken to be $L^2(\mathcal{T})$, the separable Hilbert space of square integrable functions defined on the compact domain \mathcal{T} . In most applications, $\mathcal{T} \subset \mathbb{R}$ and represents time, however, multidimensional domains could be considered as well. Typically, X_1, \dots, X_n are not entirely available but are observed through a finite set of observation points. This means, only discrete observations $\{X_{ij}\}$ of functional observations $\{X_i\}_{i=1, \dots, n}$ at time points $\{t_{ij}, j = 1 \dots, m_i\}$ are available, being m_i the number of discrete points available for the i -th observation. The aim of the clustering analysis is to define M partitions, i.e., clusters, of the data X_1, \dots, X_n such that observations in different clusters are as dissimilar as possible and that observations within the same cluster are as similar as possible. In the rest of this section, we describe the most promising approaches for functional clustering, which can be grouped in raw-data clustering, filtering methods, adaptive methods, and distance-based methods.

Raw-data Clustering

The raw-data clustering approach consists in the clustering of discretized version $\{X_{ij}\}$ of the functional observations $\{X_i\}$ by means of classical multivariate methods. This simple approach does not need the reconstruction of the functional data and relies on well-established multivariate algorithms.

One of the most popular clustering algorithm is k-means. In the functional setting, k-means aims to partition the observations into M clusters C_1^*, \dots, C_M^* such that the within-cluster sum of squares is minimized, that is

$$\{C_1^*, \dots, C_M^*\} = \underset{C_1, \dots, C_M}{\operatorname{argmin}} \sum_{m=1}^M \sum_{\mathbf{X}_i \in C_m} (\mathbf{X}_i - \boldsymbol{\mu}_m)^T (\mathbf{X}_i - \boldsymbol{\mu}_m), \quad (5.1)$$

where C_1, \dots, C_M are all the possible observation partitions in M groups, $\mathbf{X}_i = (X_{i1}, \dots, X_{im_i})^T$, and $\boldsymbol{\mu}_m$ is the mean vector of the observations in C_m . Hierarchical

clustering (Everitt et al., 2011) produces a representation in which clusters at each level of the hierarchy is formed by all and only the clusters of the lower levels. Strategies for hierarchical clustering are mainly divided into two approaches: agglomerative (bottom-up) and divisive (top-down). The former starts at the bottom (i.e., each observation in one cluster) and at each level recursively merges a selected pair of clusters into a single cluster. The latter starts at the top (i.e., all observations in one cluster) and at each level recursively splits one of the existing clusters at that level into two new clusters. Different versions of agglomerative methods arise from the choice of the intergroup dissimilarity metric, e.g., single linkage, complete linkage, average linkage. (Ward Jr, 1963) considered hierarchical clustering procedures based on minimizing the loss of information from joining two groups. Finally, model-based clustering assumes that the data in each cluster is generated from a given probabilistic distribution and the combined data stems from a convex combination of these distributions. In all the aforementioned methods, the number of clusters M has to be determined. For k-means and hierarchical clustering, this can be done based on many indices (Charrad et al., 2012), e.g., the silhouette width (Rousseeuw, 1987), the gap statistic (Tibshirani et al., 2001), the Dunn index (Dunn, 1973); whereas, for model-based clustering, information criteria, such as the Akaike information criterion (AIC), Bayesian information criterion (BIC) as well as integrated completed likelihood (ICL) could be used.

Analysis of raw data through classical multivariate techniques has several problems, because of the high number of evaluation points and the strong correlation. This may especially affect the model-based approach, that assumes non-singular covariance matrix. Moreover, raw-data clustering approach has the drawback of not taking into account the functional nature of the data and is not suited for curves observed at different evaluation points. While this suggests to use approaches specifically designed for functional data, for comparison purposes we still propose the raw data approach to the clustering. Details on multivariate clustering methods can be found in Everitt et al. (2011), Hastie et al. (2009b) and Johnson et al. (2002).

Filtering Methods

Filtering methods rely on the reconstruction of the functional observations $\{X_i\}$ from the discrete points $\{X_{ij}\}$. The most common approach (Ramsay and Silverman, 2005) is to assume that the functional observations are embedded in a finite dimensional functional space spanned by a finite set of basis functions. In particular, each X_i can be written as

$$X_i(t) = \sum_{k=1}^K c_{ik} \phi_k(t) = \mathbf{c}_i^T \boldsymbol{\phi}(t) \quad t \in \mathcal{T} \quad i = 1, \dots, n, \quad (5.2)$$

where $\boldsymbol{\phi} = (\phi_1, \dots, \phi_K)^T$ is the vector of basis functions that span the K -dimensional subset of $L^2(\mathcal{T})$, and $\mathbf{c}_i = (c_{i1}, \dots, c_{iK})^T$ is the K -dimensional coefficient vector. The basis functions $\{\phi_j\}$ can be either pre-specified, e.g. B-spline (De Boor et al., 1978), Fourier (Ramsay and Silverman, 2005), and wavelet (Walnut, 2013), or data-adaptive, e.g. obtained using functional principal component analysis (FPCA) (Hall and Hosseini-Nasab, 2006).

In case of pre-specified basis functions, if the $\{X_{ij}\}$ are observed with measurement error, then the coefficient vector \mathbf{c}_i is usually estimated as $\hat{\mathbf{c}}_i$ via penalized least-squares, even though standard least-squares could be used as well (Ramsay and Silverman, 2005), that is

$$\hat{\mathbf{c}}_i = \underset{\mathbf{c}_i \in \mathbb{R}^K}{\operatorname{argmin}} \sum_{j=1}^{m_i} \left(X_{ij} - \mathbf{c}_i^T \boldsymbol{\phi}(t_{ij}) \right)^2 + \lambda \int_{\mathcal{T}} [D^2 \mathbf{c}_i^T \boldsymbol{\phi}(t)]^2 dt, \quad (5.3)$$

where D^2 is the second order differential operator and $\lambda > 0$ is a smoothing parameter. It measures the trade-off between fit to the data, as determined by the residual sum of squares in the first term, and smoothness of X_i , as quantified by the second term. Then, the reconstructed functional observation is

$$\hat{X}_i^{PS}(t) = \hat{\mathbf{c}}_i^T \boldsymbol{\phi}(t) \quad t \in \mathcal{T} \quad i = 1, \dots, n. \quad (5.4)$$

The choice of the smoothing parameter λ is based on the well-known trade-off between variance and bias. In particular, it is usually performed by picking the λ corresponding to the minimum value assumed by the generalized cross-validation criterion. This criterion takes into account the degrees of freedom of the estimated curve that vary according to λ (Ramsay and Silverman, 2005). Moreover, the choice of K in Equation (5.2) is not crucial (Cardot et al., 2003), until it is sufficiently large to capture local behaviours of functional data.

The FPCA provides a data-adaptive basis to obtain the functional data as in Equation (5.2). In particular, the functional observations are reconstructed, for $i = 1, \dots, n$, as

$$\hat{X}_i^{DA}(t) = \sum_{l=1}^L \xi_{il} \psi_l(t) = \boldsymbol{\xi}_i^T \boldsymbol{\psi}(t) \quad t \in \mathcal{T} \quad i = 1, \dots, n, \quad (5.5)$$

where $\boldsymbol{\xi}_i = (\xi_{i1}, \dots, \xi_{iL})^T$ is the vector of principal component scores or simply scores defined as $\xi_{il} = \int_{\mathcal{T}} \psi_l(t) X_i(t) dt$, and $\boldsymbol{\psi} = (\psi_1, \dots, \psi_L)^T$ is the vector whose elements are weight functions referred to as principal components. Principal components are defined by an iterative algorithm which at each step finds the weight function that maximizes the mean square of the scores, or their sample variance, that is

$$\psi_l = \underset{\psi}{\operatorname{argmax}} \sum_{i=1}^n \xi_{il}^2 = \sum_{i=1}^n \left(\int_{\mathcal{T}} \psi(t) X_i(t) dt \right)^2 \quad l = 1, \dots, L, \quad (5.6)$$

under the constraints: $\int_{\mathcal{T}} \psi_l(t)^2 dt = 1$ and $\int_{\mathcal{T}} \psi_i(t) \psi_j(t) dt = 0$, for $i \neq j$. The choice of the number L in Equation (5.5) of retained components depends on several necessities. Generally, the retained principal components are chosen such that they explain at least a given percentage of the total variability. However, more sophisticated methods could be used as well (Jolliffe, 2011).

In practice, reconstruction of functional observations allows one to reduce the dimensionality of the data by summarizing each curve through a finite set of parameters, that is $\{\hat{\mathbf{c}}_i\}$ or $\{\boldsymbol{\xi}_i\}$ depending on whether basis functions used are pre-specified or data-adaptive. Then, the finite set of parameters are clusterized by means of standard multivariate clustering techniques, such as k-means, hierarchical clustering or model-based clustering. As for the raw-data clustering methods, several indices could be used (Charrad et al., 2012) to choose the number M of clusters.

Adaptive Methods

The present set of methods relies on a finite dimensional representation of the functional data through basis functions (similarly to the filtering approaches) where the basis expansion coefficients are treated as random variables with cluster-specific probability distributions. This differs from the filtering methods, where the basis expansion coefficients are considered

5. FUNCTIONAL CLUSTERING METHODS FOR RESISTANCE SPOT WELDING PROCESS DATA IN THE AUTOMOTIVE INDUSTRY

as parameters to be estimated in the smoothing phase. One of the first example of adaptive method was in James and Sugar (2003), referred to as *fclust*. Similarly to the filtering approaches, if the functional observation X_i belongs to the m -th cluster among the M clusters, it is modeled through basis functions as

$$X_i(t) = \boldsymbol{\eta}_{im}^T \boldsymbol{\phi}(t) \quad t \in \mathcal{T} \quad i = 1, \dots, n, \quad (5.7)$$

where $\boldsymbol{\phi} = (\phi_1, \dots, \phi_K)^T$ are natural cubic splines, and $\boldsymbol{\eta}_{ik}$ is a vector of spline normal random coefficients defined as

$$\boldsymbol{\eta}_{im} = \boldsymbol{\mu}_m + \boldsymbol{\gamma}_i, \quad (5.8)$$

with $\boldsymbol{\mu}_m$ the coefficient vector of the m -th cluster mean, and $\boldsymbol{\gamma}_i \sim N(0, \boldsymbol{\Gamma})$ the subject-specific random effects for the i -th curve. Then, the vector of discretized values $\mathbf{X}_i = (X_{i1}, \dots, X_{im_i})^T$ is modelled as

$$\mathbf{X}_i = \mathbf{S}_i(\boldsymbol{\mu}_m + \boldsymbol{\gamma}_i) + \boldsymbol{\varepsilon}_i, \quad (5.9)$$

where $\mathbf{S}_i = (\boldsymbol{\phi}(t_{i1}), \dots, \boldsymbol{\phi}(t_{im_i}))^T$ is the realization matrix of the vector $\boldsymbol{\phi}$, and $\boldsymbol{\varepsilon}_i \sim N(0, \mathbf{R})$ is the measurement error random vector. The covariance matrix \mathbf{R} is usually set equal $\sigma^2 \mathbf{I}_{m_i}$, where \mathbf{I}_{m_i} is the size m_i identity matrix. The unknown parameters $\boldsymbol{\mu}_m$, $m = 1, \dots, M$, $\boldsymbol{\Gamma}$ and σ are estimated by maximizing the mixture likelihood in Equation (5.10), where the cluster membership vector is modeled as a multinomial random variable with parameters (π_1, \dots, π_M) , with π_m the probability of an observation to belong to the m -th cluster. Thus, the mixture likelihood is defined as

$$L(\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_M, \boldsymbol{\Gamma}, \sigma, \pi_1, \dots, \pi_M) = \prod_{i=1}^N \sum_{m=1}^M \pi_m f_m(\mathbf{X}_i), \quad (5.10)$$

where $f_m(\mathbf{X}_i)$ is the conditional density function of \mathbf{X}_i belonging to the m -th cluster, that is $\mathbf{X}_i|m \sim N(\mathbf{S}_i \boldsymbol{\mu}_m, \boldsymbol{\Sigma}_i)$, with $\boldsymbol{\Sigma}_i = \sigma^2 \mathbf{I}_{m_i} + \mathbf{S}_i \boldsymbol{\Gamma} \mathbf{S}_i^T$. The maximization is often carried out by means of the expected maximization (EM) algorithm. Once the unknown parameters have been estimated, each curve X_i is assigned to the cluster whose estimated posterior probability of cluster membership $\pi_{m|i} = \hat{f}_m(\mathbf{X}_i) \hat{\pi}_m / \sum_{j=1}^M \hat{f}_j(\mathbf{X}_i) \hat{\pi}_j$ is maximum. Moreover, the cluster mean coefficients $\boldsymbol{\mu}_m$ could be further optimally parameterized to produce useful low-dimensional representations of the curves (James and Sugar, 2003). Information criteria, such as AIC and BIC, are used to select the number M of clusters and the basis dimension K (James and Sugar, 2003).

The use of spline basis has two main drawbacks: (i) they are inappropriate when dealing with functions that show peaks and irregularities, (ii) they require heavy computational efforts and so are not suitable to represent high dimensional data. For these reasons, Giacomini et al. (2013) proposed an adaptive method based on the wavelet decomposition of the curves, referred to as *waveclust*. Similarly to James and Sugar (2003), the functional observation X_i belonging to the m -th cluster is modeled as

$$X_i(t) = \mu_m(t) + U_i(t) \quad t \in \mathcal{T} \quad i = 1, \dots, n, \quad (5.11)$$

where μ_m is the principal functional fixed effect that characterizes the m -th cluster mean and U_i is a subject-specific random deviation from μ_m . By applying discrete wavelet transform

to model in Equation (5.11), contaminated with an additional measurement error function $E_i(t)$, $t \in \mathcal{T}$, the model reduces to a linear mixed-effect one. That is,

$$\begin{pmatrix} \mathbf{c}_i^T \\ \mathbf{d}_i^T \end{pmatrix}^T = \begin{pmatrix} \boldsymbol{\alpha}_m^T \\ \boldsymbol{\beta}_m^T \end{pmatrix}^T + \begin{pmatrix} \boldsymbol{\nu}_i^T \\ \boldsymbol{\theta}_i^T \end{pmatrix}^T + \begin{pmatrix} \boldsymbol{\varepsilon}_{\mathbf{c}_i}^T \\ \boldsymbol{\varepsilon}_{\mathbf{d}_i}^T \end{pmatrix}^T, \quad (5.12)$$

where $(\boldsymbol{\alpha}_m^T, \boldsymbol{\beta}_m^T)^T$, $(\boldsymbol{\nu}_i^T, \boldsymbol{\theta}_i^T)^T$, $(\boldsymbol{\varepsilon}_{\mathbf{c}_i}^T, \boldsymbol{\varepsilon}_{\mathbf{d}_i}^T)^T$, and $(\mathbf{c}_i^T, \mathbf{d}_i^T)^T$ are the vectors of scaling and wavelet coefficients of μ_m , U_i , E_i and $X_i + E_i$, respectively; $\boldsymbol{\alpha}_m$ and $\boldsymbol{\beta}_m$ are non-random parameters, whereas $(\boldsymbol{\nu}_i^T, \boldsymbol{\theta}_i^T)^T$ and $(\boldsymbol{\varepsilon}_{\mathbf{c}_i}^T, \boldsymbol{\varepsilon}_{\mathbf{d}_i}^T)^T$ are normal random vectors with zero mean and covariance matrices \mathbf{G} and $\sigma_\varepsilon \mathbf{I}_{m_i}$, respectively. Once projected in the wavelet domain, the clustering model (5.12) resumes to a standard one with additional random effects whose variance is of particular form. Thus, parameters are estimated by maximizing the likelihood function typically using the EM algorithm. Final assignment of each curve to a cluster is performed by maximizing the posterior probability of clustering membership. The number of clusters are chosen through BIC or ICL (Giacofci et al., 2013).

The last presented adaptive method was proposed by Bouveyron and Jacques (2011), and referred to as *funHDDC*. They consider, as James and Sugar (2003), that if X_i belongs to a given cluster m , it admits the following basis expansion

$$X_i(t) = \boldsymbol{\gamma}_{im}^T \boldsymbol{\Psi}(t) \quad t \in \mathcal{T} \quad i = 1, \dots, n, \quad (5.13)$$

where $\boldsymbol{\Psi} = (\Psi_1, \dots, \Psi_K)^T$ is a given vector of basis functions, and $\boldsymbol{\gamma}_{im}$ is a k -dimensional random vector. All the functions X_i in a given cluster m are assumed to be adequately described in a low-dimensional functional latent subspace with dimension $d_m < K$ spanned by a group-specific basis function $\{\varphi_{mj}\}$. Then, for a given X_i in the cluster m , the random latent expansion coefficients $\boldsymbol{\lambda}_i = (\lambda_{i1}, \dots, \lambda_{id_m})^T$ in the group-specific basis function $\{\varphi_{mj}\}$ are linked to $\boldsymbol{\gamma}_{im}$ as

$$\boldsymbol{\gamma}_{im} = \mathbf{U}_m \boldsymbol{\lambda}_i + \boldsymbol{\varepsilon}_i, \quad (5.14)$$

where \mathbf{U}_m is the $K \times d_m$ matrix composed by the first d_m columns of the orthogonal $K \times K$ matrix \mathbf{Q}_m , whose entries are the coefficients that the linearly link $\{\Psi_k\}$ and $\{\varphi_{mj}\}$, and $\boldsymbol{\varepsilon}_i \in \mathbb{R}^K$ is an independent random noise term. By assuming that $\boldsymbol{\lambda}_i \sim N(\boldsymbol{\mu}_m, \mathbf{S}_m)$, with $\mathbf{S}_m = \text{diag}(a_{m1}, \dots, a_{md_m})$, and that $\boldsymbol{\varepsilon}_i \sim N(0, \boldsymbol{\Xi}_m)$, then

$$\boldsymbol{\gamma}_{im} \sim N(\mathbf{U}_m \boldsymbol{\mu}_m, \mathbf{Q}_m \boldsymbol{\Delta}_m \mathbf{Q}_m^T), \quad (5.15)$$

where the $K \times K$ matrix $\boldsymbol{\Delta}_m = \mathbf{Q}_m (\mathbf{U}_m \boldsymbol{\Delta}_m \mathbf{U}_m^T + \boldsymbol{\Xi}_m) \mathbf{Q}_m^T$ and the noise covariance matrix $\boldsymbol{\Xi}_m$ is chosen such that $\boldsymbol{\Delta}_m = \text{diag}(a_{m1}, \dots, a_{md_m}, b_m, \dots, b_m)$. Let us assume the cluster membership vector is modeled as a multinomial random variable with parameters (π_1, \dots, π_M) , with π_m the probability of an observation to belong to the m -th cluster. Then, the mixture likelihood is defined as

$$L(\mathbf{U}_1, \dots, \mathbf{U}_M, \boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_M, \pi_1, \dots, \pi_M, \mathbf{Q}_1, \dots, \mathbf{Q}_M, \boldsymbol{\Delta}_1, \dots, \boldsymbol{\Delta}_M) = \prod_{i=1}^N \sum_{m=1}^M \pi_m f_m(\boldsymbol{\gamma}_i), \quad (5.16)$$

where $f_m(\boldsymbol{\gamma}_i)$ is the conditional density function of \mathbf{X}_i to belong to the m -th cluster, that is $\boldsymbol{\gamma}_i|m \sim N(\mathbf{U}_m \boldsymbol{\mu}_m, \mathbf{Q}_m \boldsymbol{\Delta}_m \mathbf{Q}_m^T)$. The maximization is conveniently carried out by means of the EM algorithm. Moreover, it is possible to obtain parsimonious submodels of Equation

5. FUNCTIONAL CLUSTERING METHODS FOR RESISTANCE SPOT WELDING PROCESS DATA IN THE AUTOMOTIVE INDUSTRY

Table 5.1. Number of clusters obtained and computation time for each approach. Programs were run using a machine with an Intel Xeon 2.10 GHz processor.

Method	Number of clusters	Computation time (min)
Raw data hierarchical	2	3
Raw data k-means	3	
Raw data model-based	8	
Filtering B-spline hierarchical	3	2
Filtering B-spline k-means	3	
Filtering B-spline model-based	4	
Filtering FPCA hierarchical	3	1
Filtering FPCA k-means	3	
Filtering FPCA model-based	5	
Adaptive fclust	4	621
Adaptive curvclust	2	359
Adaptive funHDDC	2	559
Distance-based	2	2

(5.15) by constraining model parameters within or between groups. The latent subspace dimension d_m and the number of clusters M are chosen through a scree-test and BIC, respectively (Bouveyron and Jacques, 2011).

Distance-based Methods

These methods are the functional extension of classical geometric clustering algorithm to functional data, such as k-means (Cuesta-Albertos and Fraiman, 2007) and hierarchical (Ferraty and Vieu, 2006) clustering, that basically rely on the definition of proximity or dissimilarity among observations. Therefore, the extension to functional data consists in the introduction of an appropriate functional measure of proximity or dissimilarity. In this respect, several authors (Tarpey and Kinader, 2003; Ferraty and Vieu, 2006; Cuesta-Albertos and Fraiman, 2007) agree upon the use of the following measure of proximity between the curves X_i and X_j

$$d_l(X_i, X_j) = \left(\int_{\mathcal{T}} \left(X_i^{(l)} - X_j^{(l)} \right)^2 dt \right)^{1/2}, \quad (5.17)$$

where $X^{(l)}$ denotes the l -th derivative of X . In this case the number of clusters could be suitably chosen through the silhouette index (Rousseeuw, 1987).

5.4 Results and Discussion

In this section, we discuss on the results obtained by implementing the functional clustering methods presented in Section 5.3 to the DRC data set illustrated in Section 5.2. For the sake of readability, implementation details are deferred to the Appendix. The optimal number of clusters selected by each approach mentioned in Section 5.3 is reported in Table 5.1. Note that most of the methods provide similar results and identify two or three clusters. The only exceptions are some *model-based* methods, viz. *adaptive fclust*, *filtering B-spline*, *filtering FPCA*, and *raw data*, which select from four to eight clusters. In general, the larger the

number of clusters, the harder the technological interpretation, i.e., the less straightforward the discrimination of spot welds belonging to different groups. Inflation in the number of clusters is usually due to overfitting problems especially for model-based approaches applied to high-dimensional correlated data and complex variance structures. In this case, the number of parameters to be estimated can be very large and may lead to instability, no matter if the BIC criterion, that penalizes the model complexity, is used to select the optimal number of clusters. This issue may be exacerbated by the use of model-based methods on raw data (see third row of Table 5.1), which do not rely on an optimal basis representation and typically contain additional noise. In Table 5.1 it is also reported the computation time required for each approach using a machine with an Intel Xeon 2.10 GHz processor. The adaptive approaches result as the most computationally intensive, while all the others require less than 3 minutes to complete the analysis. Even if strictly dependent on the data set at hand, this information may be crucial when dealing with complex data structures in order to pick the most appropriate method to be used when computational resources are limited.

Figure 5.3 shows the DRCs coloured according to the cluster assignment provided by each method. Whereas, Figure 5.4 depicts the centroids (i.e., mean functions) for each cluster. Note that, in both figures, first, second and third rows of panels refer to clustering methods that select two, three and more than three clusters, respectively.

With reference to those figures and coherently with the features highlighted in Section 5.2, we note that DRC centroids have local minimum points with approximately the same abscissa, but different resistance values; local maximum points with approximately the same value, but different abscissa; different resistance values at the end of the functional domain. It will be in fact convenient to facilitate the following technological interpretation and insights into the industrial problem at hand to focus attention on (I) the *amplitude difference* between minimum and maximum resistance values, (II) the *phase difference* between minimum and maximum abscissas, and (III) the *final resistance value*.

The first row of panels of Figure 5.4 displays centroids associated to cluster 1 having amplitude, phase difference and final resistance value smaller than those respectively associated to cluster 2. The separation is clear as cluster 1 centroids show a local minimum value, in the first part of the functional domain, that is distinctly larger than that corresponding to cluster 2 centroids, and decrease more rapidly to lower values in the last part of the domain.

Whereas, in the second row of panels, the amplitude, phase difference and final value of centroids increase together from cluster 1 to cluster 3, except for panels *filtering FPCA k-means* and *raw k-means* that have centroids of clusters 2 and 3 with approximately the same final value. That is, cluster 1 centroids show the larger local minimum value and the more rapid decrease at the end of the functional domain, whereas the other centroids tend to be more similar.

Finally, with reference to each panel of the third row of Figure 5.4, centroids are shown to have cluster number sorted in ascending order by amplitude and phase difference, only. That is, the final resistance values do not preserve the order set by the phase difference, as in the first two rows of panels. In fact, with reference to the third row of panels of Figure 5.4, final resistance values of centroids obtained by *adaptive fclust* and *filtering B-spline model-based*, respectively depicted in the first and second panel, are sorted in ascending order with the cluster number as 1,2,4,3; whereas, those of the third panel, referring to *filtering FPCA model-based* method, are ordered as 1,2,3,5,4; and those of the fourth panel (*raw model-based* method) as 1,3,2,5,7,8,4,6.

As it typically happens (Everitt et al., 2011), also for this data set no one clustering method can be judged to be best in all circumstances. However, dealing with a small number

5. FUNCTIONAL CLUSTERING METHODS FOR RESISTANCE SPOT WELDING PROCESS DATA IN THE AUTOMOTIVE INDUSTRY

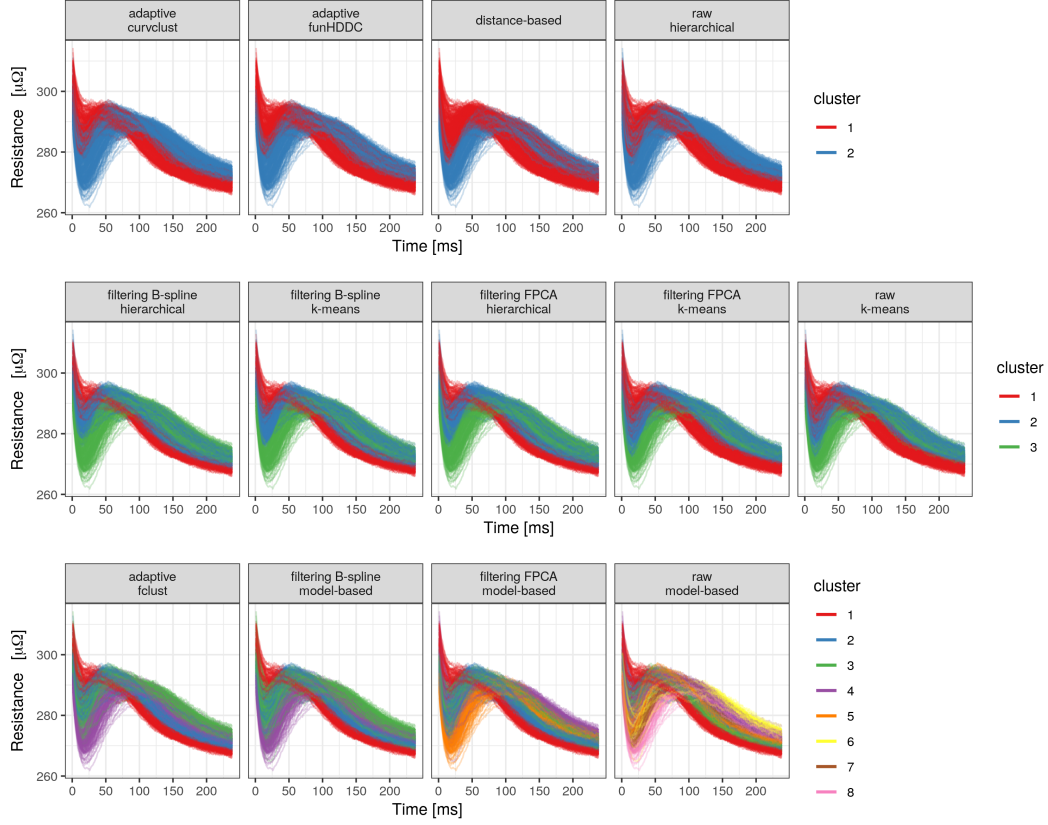


Figure 5.3. Plot of the functional data. Each panel correspond to one of the proposed clustering methods, curves are coloured accordingly to the corresponding cluster assignment. Plots are arranged such that first, second and third rows of panels refer to clustering methods that divide DRCs into two, three, and more than three clusters, respectively.

of clusters, say two or three in this case, shall provide with a clearer interpretation of groups of functions that are well distinct and more likely to lead to informative classifications. Therefore, in what follows we assume selecting three clusters as the better compromise to trade off straightforward interpretation of DRCs belonging to the same clusters and distinct characterization of each cluster.

Consistently with the technological literature (Dickinson et al., 1980; Adams et al., 2017), being the minimum point abscissa practically constant and the maximum point the landmark for the start of nugget formation (see also Section 5.2), we can state that the smaller the phase difference, and thus the larger the time interval between the local maximum and the end of the welding process, the more the heat energy supplied for nugget growth. Note that the inflection point typically located between the local minimum and maximum of the DRC (see, e.g. Figure 1) ideally represents the welding melting point. Centroids having larger phase difference shall thus characterize welding spot clusters with smaller nugget size. Unfortunately, in the real case at hand, this ideal statement may not hold because of the natural wear process of the electrodes, which induces, as conjectured by RSW process experts, a non-negligible increase of the weld area, and thus different welding conditions for

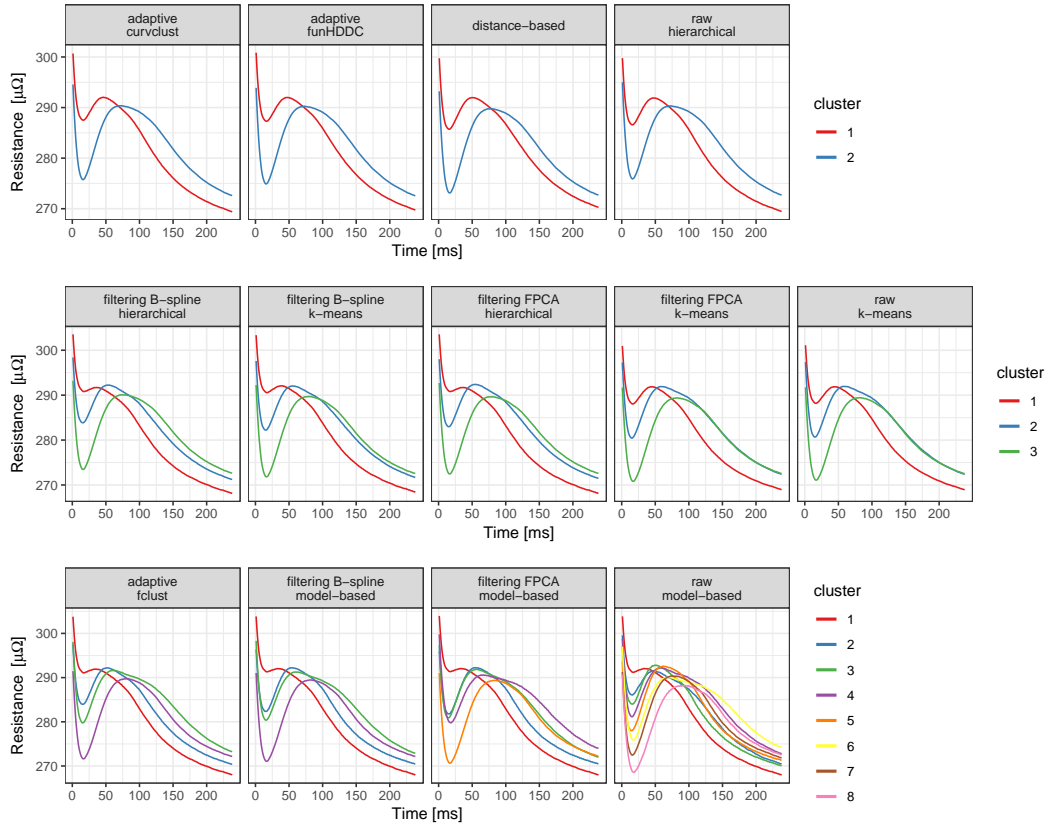


Figure 5.4. Plot of the cluster centroids. Each panel corresponds to one of the proposed clustering methods, curves are centroids of each cluster obtained with the corresponding method. Plots are arranged such that first, second and third rows of panels refer to clustering methods that divide DRCs into two, three, and more than three clusters, respectively.




Figure 5.5. Seven DRC observations, included in the 538 original DRCs, corresponding to spot welds for which qualitative information about the electrode wear status is available: just renewed (thin solid line), intermediate wear (dashed line), severe wear (dotted line).

each spot weld.

In order to explain the cluster in terms of electrode wear, in Figure 5.5 we report seven DRCs for which it has been possible to retrieve qualitative information on the wear status of the electrodes. For this purpose, without loss of generality, we refer to the *filtering B-spline hierarchical* method among those selecting three clusters, and compare the corresponding panel in the second row of Figure 5.4 with Figure 5.5. With reference to the latter figure, we may want to analyze the two extreme wear cases (thin solid and dotted lines) and conjecture that centroid of cluster 1 of Figure 5.4 shall correspond to the smaller electrode area, i.e., electrode just renewed, whereas centroid of cluster 3 to the larger one, i.e., electrode with severe wear. Even though the technological cause is different, experts' opinion is that DRCs associated to cluster 1 and those associated to cluster 3 correspond to spot welds with improper nugget diameter. In particular, for spot welds that belong to DRCs in cluster 1, the root cause is attributed to the excessive clamping pressure in the welding zone. The large clamping pressure is also confirmed by the small amplitude difference in the DRC centroid of cluster 1 (Dickinson et al., 1980).

Conversely, spot welds pertaining to cluster 3 correspond to larger electrode area, so that the clamping force generates the lower pressure in the welding zone and cannot guarantee the proper value of current density. Indeed, despite the larger amplitude difference in the DRC centroid, which proves the smaller pressure, the nugget diameter may result undersized because of the delay in the nugget formation. Therefore, it turns out that the better spot welds in terms of nugget formation achieved by DRCs pertaining to intermediate cluster(s).

This conjecture becomes more clear in the light of Figure 5.6 in which observations with qualitative electrode wear status, already displayed in Figure 5.5, are colored by cluster number (assigned through *filtering B-spline hierarchical* method) and are superimposed to the corresponding centroids, already displayed in the first panel of the second row of Figure 5.4. By Figure 5.6, the wear status of the electrode clearly appears as a determinant factor in the clustering of DRCs. This result represents an important industrial finding that confirms an expert conjecture supported by functional data made available under the Industry 4.0 paradigm. The next steps to exploit this result is to explicitly identify conditions on the wear status to signal, when the corresponding DRC is associated to a cluster that does not guarantee proper mechanical and metallurgical properties, the need e.g., for electrode break-in, renewal or substitution. To put into action this strategy, or even more complicated maintenance programs, further technological investigation and off-line quality testing should be carried out for every DRC cluster. The ultimate goal is to avoid the random sampling of the sub-assemblies to be tested off-line and to support more specific priority assignment. One could in fact imagine to assign higher priority to future spot welds having DRC observation with larger (resp., smaller) distance from cluster centroids that have revealed to refer to adequate (resp., inadequate) quality.

5.5 Conclusions

In this chapter, we tackled the issue of finding homogeneous groups of dynamic resistance curves (DRCs) coming from a resistance spot welding (RSW) process, which, in the modern automotive Industry 4.0, is of crucial relevance to better understand the effects of the process parameters on the final weld quality. To avoid loss of information caused by arbitrary scalar feature extraction, DRCs have been modelled as functional data defined on the time domain, and, accordingly, clustering methods specifically designed for functional data have been presented in a practical hands-on overview with the aim of facilitating their




Figure 5.6. The corresponding centroids (solid line), already displayed in the first panel of the second row of Figure 5.4, are superimposed to the seven DRCs with qualitative electrode wear status: just renewed (thin solid line), intermediate wear (dashed line), severe wear (dotted line), already displayed in Figure 5.5, are colored by cluster number assigned through *filtering B-spline hierarchical* method.

practical implementation. To the best of the authors' knowledge, this is the first study where functional clustering methods are applied to the whole DRC functional observations to gain technological insights on RSW processes, even if the framework used could go far beyond the specific application hereby investigated.

The effectiveness of the presented functional clustering methods is demonstrated by applying them to 538 DRCs acquired during RSW lab tests at Centro Ricerche Fiat (CRF). It turned out that the identified clusters of DRCs are strictly linked with the wear status of the electrodes, that, in turn, affects the electrode contact area, clamping pressure in the welding zone and current density, and impacts on the final quality of spot welds in terms of mechanical and metallurgical properties. Indeed, in accordance with the experts, we agree the better spot welds shall correspond to DRCs belonging intermediate clusters having proper amplitude difference and small phase difference.

A broader perspective of the results is given in supporting practitioners in the priority assignment for off-line testing of welded spots and in the electrode wear analysis. Functional clustering analysis could be in fact imagined to be embedded in a wider on-line statistical quality control framework for RSW processes, which is able to properly exploit the properties of the clusters identified. Finally, the relationship between the electrode wear and the final quality of spot welds, which has been discovered by the proposed functional clustering analysis, could be now further investigated through the specific definition of opportune quantitative variables in the direction of routinely tracing wear status.

5.6 Appendix

Implementation Details

In the following paragraphs, we provide with further details for each of the approaches implemented in this chapter on the DRC data set at hand, which should help practitioners to unbox the R code provided.

Raw data To mitigate problems due to the high dimensionality and the strong correlation of the data, for every DRC observation consisting of 238 equally spaced points, we chose to keep only 19 points, one each 13, and applied clustering methods on these sliced DRCs. When using model-based methods, the selection of the optimal number of clusters was based on the optimization of the BIC criterion. In the other cases, viz. k-means and hierarchical clustering methods, we relied on the R package `NbClust` Charrad et al. (2012), which allows the calculation of several indices. Then, the optimal number of clusters was chosen according to the majority rule, i.e., as that optimizing the largest number of criteria.

Filtering Let us consider the case when using B-spline basis, first. If we choose too many basis functions to represent profiles, we still have the same high dimensionality and correlation problems as in the raw data case. This can also make computation very slow. Since this data set is characterized by functions that are relatively smooth, we chose to regularize using a lower number of basis functions and to avoid the penalization of the integrated squared second derivative. We selected the number of basis functions on the basis of the generalized cross-validation criterion. In particular, in order to keep the number of basis functions low, after plotting the generalized cross-validation against the number of bases, we selected 12 basis functions as the elbow of the curve.

When using FPCA in the filtering approach, we applied clustering on the functional principal component scores. We first obtained smooth functions using B-spline basis expansion, with 100 basis. Then, we regularized as described in Section 5.3 by means of a penalty on the integrated squared second derivative, with smoothing parameters chosen by minimizing the generalized cross-validation criterion. Thus, we performed functional principal component analysis on the obtained functional data set and retained only the components that explain the 99% of the total variability in the data. This allowed in practice to reconstruct original functions with a strong dimension reduction. In fact, since profiles in this data set are smooth, only 6 principal components were required.

For both B-spline and FPCA basis, when using the model-based approach, the selection of the optimal number of clusters was based on optimization of the BIC criterion. In the other cases, viz. k-means and hierarchical clustering methods, we relied on the R package `NbClust` Charrad et al. (2012), which allows the calculation of several indices. The optimal number of clusters was chosen also in this case according to the majority rule.

Adaptive The *adaptive fclust* method was performed by means of the R package *fclust*, which requires the choice of the number of basis functions K , as mentioned in Section 5.3, that was set equal to 5, 10. Parameters and the number of clusters were set based on the BIC criterion.

For *waveclust*, we relied on the R package `curvclust` Giacomini et al. (2012), dedicated to model-based curve clustering. In particular, the considered models include Functional Clustering Mixed Models (FCMM, i.e., functional clustering with the presence of functional random effects), but also traditional functional clustering model (FCM, without functional random effects). Among FCMMs, several structures of the variance of the random effect can be chosen. In particular, the following alternatives are available, as mentioned in Giacomini et al. (2013): constant, group, scale-position, and group-scale-position dependent. It is also possible to decide whether to retain all coefficients or to perform individual denoising to keep coefficients which contain individual-specific information, by applying nonlinear wavelet hard thresholding before clustering. We considered all parameter combinations of the variance structures and chose the model and number of clusters that optimize the

BIC criterion. Model fitting was performed by maximum likelihood method using the EM algorithm and the stochastic EM as initialization method.

When using *funHDDC*, several parameters need to be chosen. We considered 0.2, 0.5, and 0.9 as possible values for the threshold of the Cattell' scree-test used for selecting the group-specific intrinsic dimensions d_m . Moreover, the following alternative models are available, as described in Bouveyron and Jacques (2011): $a_{kj}b_kQ_kd_k$, $a_{kj}bQ_kd_k$, $a_kb_kQ_kd_k$, $ab_kQ_kd_k$, $a_kbQ_kd_k$, abQ_kd_k . In this chapter, we considered all parameter combinations and chose the model and numbers of clusters that optimize the BIC criterion. Moreover, to avoid local minima, for each parameter combination we repeated the model fitting 20 times and kept the model with the largest log-likelihood.

Distance-based In the distance-based approach, we applied the clustering algorithm for each number of clusters by considering the distance in Equation (5.17) with $l = 0$, that is the usual L^2 distance, then, we kept the model with the best value of the silhouette index.

Part II

Chapter 6

Smooth LASSO Estimator for the Function-on-Function Linear Regression Model

Abstract

A new estimator, named S-LASSO, is proposed for the coefficient function of the function-on-function linear regression model. The S-LASSO estimator is shown to be able to increase the interpretability of the model, by better locating regions where the coefficient function is zero, and to smoothly estimate non-zero values of the coefficient function. The sparsity of the estimator is ensured by a *functional LASSO penalty*, which pointwise shrinks toward zero the coefficient function, while the smoothness is provided by two roughness penalties that penalize the curvature of the final estimator. The resulting estimator is proved to be estimation and pointwise sign consistent. Via an extensive Monte Carlo simulation study, the estimation and predictive performance of the S-LASSO estimator are shown to be better than (or at worst comparable with) competing estimators already presented in the literature before. Practical advantages of the S-LASSO estimator are illustrated through the analysis of the well known *Canadian weather* and *Swedish mortality data*.

6.1 Introduction

Functional linear regression (FLR) is the generalization of the classical multivariate regression to the context of the functional data analysis (FDA) (e.g. Ramsay and Silverman (2005); Horváth and Kokoszka (2012); Hsing and Eubank (2015); Kokoszka and Reimherr (2017)), where either the predictor or the response or both have a functional form. In particular, we study the function-on-function (FoF) linear regression model, where both the predictor and the response variable are functions and each value of the latter, for any domain point, depends on the full trajectory of the former. The model is as follows

$$Y_i(t) = \int_{\mathcal{S}} X_i(s) \beta(s, t) ds + \varepsilon_i(t) \quad t \in \mathcal{T}, \quad (6.1)$$

for $i = 1, \dots, n$. The pairs (X_i, Y_i) are independent realizations of the predictor X and the response Y , which are assumed to be smooth random process with realizations in $L^2(\mathcal{S})$ and $L^2(\mathcal{T})$, i.e., the Hilbert spaces of square integrable functions defined on the compact

sets \mathcal{S} and \mathcal{T} , respectively. Without loss of generality, the latter are also assumed with functional mean equal to zero. The functions ε_i are zero-mean random errors, independent of X_i , and have autocovariance function $K(t_1, t_2)$, t_1 and $t_2 \in \mathcal{T}$. The function β is smooth in $L^2(\mathcal{S} \times \mathcal{T})$, i.e., the Hilbert space of bivariate square integrable functions defined on the compact set $\mathcal{S} \times \mathcal{T}$, and is hereinafter referred to as *coefficient function*.

FLR analysis is a hot topic in the FDA literature. A comprehensive review of the main results is provided by Morris (2015) as well as by Ramsay and Silverman (2005); Horváth and Kokoszka (2012) and Cuevas (2014) who give worthwhile modern perspectives. Although the research efforts have been focused mainly on the case where either the predictor or the response have functional form (Cardot et al., 2003; Li et al., 2007; Hall et al., 2007), the interest in the FoF linear regression has increased in the very last years. In particular, Besse and Cardot (1996) developed a spline based approach to estimate the coefficient function β , while Ramsay and Silverman (2005) proposed an estimator assumed to be in a finite dimension tensor space spanned by two basis sets and where regularization is achieved by either truncation or roughness penalties. Yao et al. (2005b) built up an estimation method based on the principal component decomposition of the autocovariance function of both the predictor and the response based on the *principal analysis by conditional expectation* (PACE) method (Yao et al., 2005a). This estimator was extended by Chiou et al. (2014) to the case of multivariate functional responses and predictors. A general framework for the estimation of the coefficient function was proposed by Ivanescu et al. (2015) by means of the mixed model representation of the penalized regression. An extension of the ridge regression (Hastie et al., 2009a) to the FoF linear regression with an application to the Italian gas market was presented in Canale and Vantini (2016). To take into account the case when the errors ε_i are correlated, in Scheipl et al. (2015) the authors developed a general framework for additive mixed models by extending the work of Ivanescu et al. (2015).

Analogously to the classical multivariate setting, in Equation (6.1) the functional predictor X contributes linearly to the response Y through the coefficient function β , which works as a continuous weight function. Trivially note that, in the domain regions over which β is equal to zero (if any), changes in the functional predictor X do not affect the conditional value of Y . Because of the infinite dimensional nature of the FLR problem, coefficient functions that are sparse, i.e., zero valued over large parts of domain, arise very often in real applications. When the aim of the analysis is descriptive, that is the interest relies on understanding the relationship between X and Y , rather than predictive only, methods that are able to capture the sparse nature of the coefficient function may be of great practical interest. These methods are referred to as *sparse* or *interpretable* because they allow better interpreting the effects of the predictor on the response and reveal the sparse nature of β . On the contrary, the interpretation of the relationship between X and Y is often cumbersome for methods that do not focus on the sparsity of the coefficient function. In particular, here the interpretability of the model in Equation (6.1) is ultimately related to the knowledge of the parts of the domain $\mathcal{S} \times \mathcal{T}$ where β is equal or different to zero, which are hereinafter referred to as *null* and *non-null regions*, respectively.

Few works address the issue of the interpretability in FLR. In the scalar-on-function setting, James et al. (2009) proposed the FLiRTI (*Functional Linear Regression That's Interpretable*) estimator that is able to recover the sparseness of the coefficient function, by imposing L_1 -penalties on the coefficient function itself and its first two derivatives. Zhou et al. (2013) introduced an estimator obtained in two stages where the initial estimate is obtained by means of a Dantzig selector (Candes et al., 2007) refined via the group *Smoothly Clipped Absolute Deviation* (SCAD) penalty (Fan and Li, 2001). The most recent work that

addresses the issue of interpretability is that of Lin et al. (2017), who proposed a *Smooth and Locally Sparse* (SLoS) estimator of the coefficient function based on the combination of the smoothing spline method with the functional SCAD penalty.

However, to the best of the author knowledge, no effort has been made in the literature to obtain an interpretable estimator for the FoF linear regression model. In this work, we try to fill this gap by developing an interpretable estimator of the coefficient function β , named S-LASSO (Smooth plus LASSO) that is *locally sparse* (i.e., is zero on the null region) and, at the same time, *smooth* on the non-null region. The property of sparseness of the S-LASSO estimator is provided by a *functional LASSO penalty* (FLP), which is the functional generalization of the classical Least Absolute Shrinkage and Selection Operator (LASSO) (Tibshirani, 1996). Whereas, two roughness penalties, introduced in the objective function, ensure smoothness of the estimator on the non-null region. From a computational point of view, the S-LASSO estimator is obtained as the solution of a single optimization problem by means of a new version of the *orthant-wise limited-memory quasi-Newton* (OWL-QN) algorithm (Andrew and Gao, 2007), which is specifically designed to solve optimization problems involving L_1 penalties.

To give an idea of the properties of the proposed estimator, in Figure 6.1 the S-LASSO estimator is applied to four different scenarios, whose data generation is detailed in Section 6.4. In particular, in each plot the S-LASSO estimate, the true coefficient function, and the smoothing spline estimate proposed by Ramsay and Silverman (2005), referred to as SMOOTH, are shown for $t = 0.5$. For Scenario I, the true coefficient function is zero all over the domain, which means that the predictor X is independent of the response. It is clear from Figure 6.1(a) that the S-LASSO estimate successfully recovers the sparseness of β , the same cannot be said for the SMOOTH estimate, which is different from zero for all values of s . Figure 6.1(b) and Figure 6.1(c) show the coefficient function estimates for Scenario II and Scenario III, where β is zero on the edge and in the central part of the domain, respectively. Also in this case the proposed method provides an estimate which is sparse on the null region and smooth on the non-null one. Indeed, for Scenario II, the S-LASSO estimate is zero for $s \in [0, 0.2]$ and for $s \in [0.8, 1]$, whereas, it well resembles the true coefficient function in the central part of the domain. In Figure 6.1(c), the sparsity property of the S-LASSO estimator can be further appreciated, which, in fact, over the central domain region, i.e., for $s \in [0.1, 0.9]$, successfully estimates β . The SMOOTH estimate in both scenarios is not able instead to capture the sparse nature of the coefficient function. The last scenario in Figure 6.1(d) is not favourable to the proposed estimator because the true coefficient function is not sparse. However, also in this case the S-LASSO method provides satisfactory results. These four examples are provided with the preliminary purpose of giving an idea about the ability of the S-LASSO estimator to recover sparsity in the coefficient function while simultaneously estimating the relationship between X and Y over the non-null region. The performance of the S-LASSO estimator will be deeply analysed in Section 6.4.

The chapter is organized as follows. In Section 6.2, the S-LASSO estimator is presented. In Section 6.3, asymptotic properties of the S-LASSO estimator are discussed in terms of consistency and pointwise sign consistency. In Section 6.4, by means of a Monte Carlo simulation study, the S-LASSO estimator is compared, in terms of estimation error and prediction accuracy, with competing estimators already proposed in the literature. In Section 6.5, the potential of the S-LASSO estimator are demonstrated with respect of two benchmark datasets: the *Canadian weather* and *Swedish mortality data*. Proofs, data generation procedures in the simulation study, and, algorithm description are given in the Supplementary Materials.

6. SMOOTH LASSO ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL

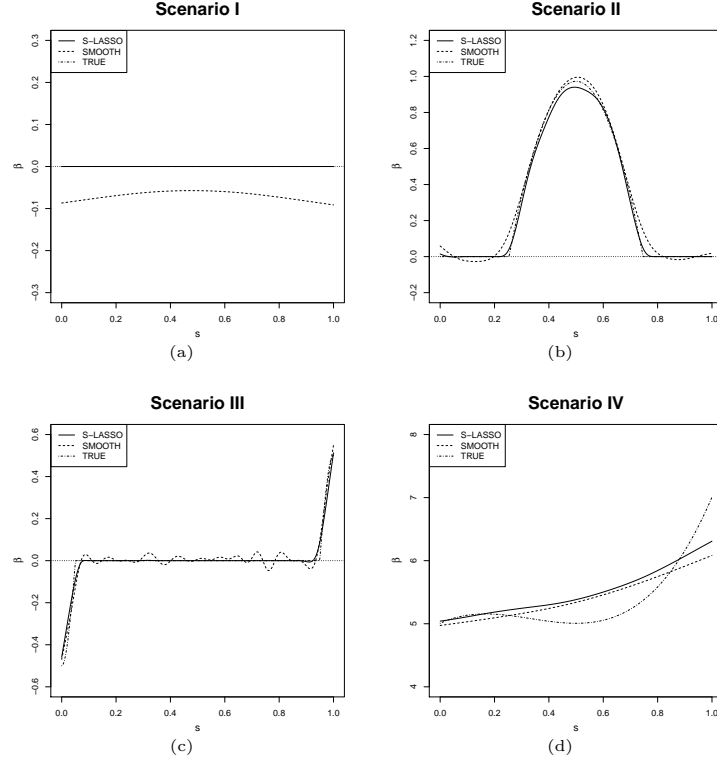


Figure 6.1. The S-LASSO (—) and SMOOTH (· · · · ·) estimates along with the true coefficient function, referred to as TRUE (---) at $t = 0.5$ for Scenario I (a), Scenario II (b), Scenario III (c) and Scenario IV (d) in the simulation study detailed in Section 6.4.

6.2 Methodology

In Section 6.2, we briefly describe the smoothing spline estimator. Readers who are already familiar with this approach may skip to the next subsection. In Section 6.2, the S-LASSO estimator definition is given along with details on both computational issues and model selection.

The Smoothing Spline Estimator

The smoothing spline estimator of the FoF linear regression model (Ramsay and Silverman, 2005) is the first key component of the S-LASSO estimator. It is based on the assumption that the coefficient function β may be well approximated by an element in the tensor product space generated by two spline function spaces, where a spline is a function defined piecewise by polynomials. Well-known basis functions for the spline space are the B-splines. A B-spline basis is a set of spline functions uniquely defined by an order k and a non-decreasing sequence of $M + 2$ knots, that we hereby assume to be equally spaced in a general domain \mathcal{D} . Cubic B-splines are B-splines of order $k = 4$. Each B-spline function is a positive polynomial of degree $k - 1$ over each subinterval defined by the knot sequence and is non-zero over no more than k of these subintervals (i.e., the compact support property). In our setting, besides the

computational advantage (Hastie et al., 2009a), the compact support property is fundamental because it allows one to link the values of β over a given region to the B-splines with support in the same region and to discard all the B-splines that are outside that region. Thorough descriptions of splines and B-splines are in De Boor et al. (1978) and Schumaker (2007).

The smoothing spline estimator (Ramsay and Silverman, 2005) is defined as

$$\hat{\beta}_S = \underset{\alpha \in \mathbb{S}_{k_1, k_2, M_1, M_2}}{\operatorname{argmin}} \left\{ \sum_{i=1}^n \left\| Y_i - \int_{\mathcal{S}} X_i(s) \alpha(s, \cdot) ds \right\|^2 + \lambda_s \|\mathcal{L}_s^{m_s} \alpha\|^2 + \lambda_t \|\mathcal{L}_t^{m_t} \alpha\|^2 \right\}, \quad (6.2)$$

where $\mathbb{S}_{k_1, k_2, M_1, M_2}$ is the tensor product space generated by the sets of B-splines of orders k_1 and k_2 associated with the non-decreasing sequences of $M_1 + 2$ and $M_2 + 2$ knots defined on \mathcal{S} and \mathcal{T} , respectively. $\mathcal{L}_s^{m_s}$ and $\mathcal{L}_t^{m_t}$, with $m_s \leq k_1 - 1$ and $m_t \leq k_2 - 1$, are the m_s -th and m_t -th order linear differential operators applied to α with respect to the variables s and t , respectively. The symbol $\|\cdot\|$ denotes the L^2 -norm corresponding to the inner product $\langle f, g \rangle = \int fg$. The parameters $\lambda_s, \lambda_t \geq 0$ are generally referred to as *roughness parameters*. The aim of the second and third terms on the right-hand side of Equation (6.2) is that of penalizing features along s and t directions. A common practice, when dealing with cubic splines, is to choose $m_s = 2$ and $m_t = 2$, which results in the penalization of the curvature of the final estimator. When $\lambda_s = \lambda_t = 0$, the wiggleness of the estimator is not penalized and the resulting estimator is the one that minimizes the sum of squared errors. On the contrary, as $\lambda_s \rightarrow \infty$ and $\lambda_t \rightarrow \infty$, $\hat{\beta}_S$ converges to a bivariate polynomial with degree equal to $|\max(m_s, m_t) - 1|$. However, there is no guarantee that $\hat{\beta}_S$ is a sparse estimator, i.e., it is exactly equal to zero in some part of the domain $\mathcal{S} \times \mathcal{T}$.

The S-LASSO Estimator

Based on the smoothing spline estimator of Equation (6.2), the S-LASSO estimator is defined as follows

$$\hat{\beta}_{SL} = \underset{\alpha \in \mathbb{S}_{k_1, k_2, M_1, M_2}}{\operatorname{argmin}} \left\{ \sum_{i=1}^n \left\| Y_i - \int_{\mathcal{S}} X_i(s) \alpha(s, \cdot) ds \right\|^2 + \lambda_s \|\mathcal{L}_s^{m_s} \alpha\|^2 + \lambda_t \|\mathcal{L}_t^{m_t} \alpha\|^2 + \lambda_L \int_{\mathcal{S}} \int_{\mathcal{T}} |\alpha(s, t)| ds dt \right\}. \quad (6.3)$$

The last term in the right-hand side of Equation (6.3) is the extension of the LASSO penalty (Tibshirani, 1996) to the FoF linear regression setting, which is referred to as *functional LASSO penalty* (FLP). In particular, by starting from the multivariate LASSO penalty, the FLP is built by substituting summation with integration and by taking into account the continuity of the domain. The FLP is able to pointwise shrink the value of the coefficient function for each $s, t \in \mathcal{S} \times \mathcal{T}$. Due to the property of the absolute value function of being singular at zero, some of these values are shrunken exactly to zero. Thus, the FLP allows $\hat{\beta}_{SL}$ to be exactly zero over the null region. The constant $\lambda_L \geq 0$ is usually called *regularization parameter* and controls the degree of shrinkage towards zero of the FLP. The larger this value, the larger the shrinkage effect and the domain portion where the resulting estimator is zero. Moreover, the FLP is expected to be able to improve the prediction accuracy (in terms of expected mean square error) by introducing a bias in the final estimator.

The second and third terms on the right-hand side of Equation (6.3) represent the two roughness penalties introduced in Equation (6.2) to control the smoothness of the coefficient function estimator. It is worth noting that, in general, the estimator smoothness may be also controlled by opportunely choosing the dimension of the space $\mathbb{S}_{k_1, k_2, M_1, M_2}$, that is,

6. SMOOTH LASSO ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL

by fixing k_1 and k_2 , and choosing M_1 and M_2 (Ramsay and Silverman, 2005). However, this strategy is not suitable in this case. To obtain a sparse estimator, the dimension of the space $\mathbb{S}_{k_1, k_2, M_1, M_2}$ must be in fact as large as possible. In this way, the value of β in a given region is strictly related to the coefficients of the B-spline functions defined on the same part of the domain and, thus, they tend to be zero in the null region. On the contrary, when the dimension of $\mathbb{S}_{k_1, k_2, M_1, M_2}$ is small, there is a larger probability that some B-spline functions have support both in the null and non-null regions and, thus the corresponding B-spline coefficients result different from zero. Therefore, we find suitable the use of the two roughness penalty terms in Equation (6.3).

To compute the S-LASSO estimator, let us consider the space $\mathbb{S}_{k_1, k_2, M_1, M_2}$ generated by the two sets of B-splines $\boldsymbol{\psi}^s = (\psi_1^s, \dots, \psi_{M_1+k_1}^s)^T$ and $\boldsymbol{\psi}^t = (\psi_1^t, \dots, \psi_{M_2+k_2}^t)^T$, of order k_1 and k_2 and non-decreasing knots sequences $\Delta^s = \{s_0, s_1, \dots, s_{M_1}, s_{M_1+1}\}$ and $\Delta^t = \{t_0, t_1, \dots, t_{M_2}, t_{M_2+1}\}$, defined on $\mathcal{S} = [s_0, s_{M_1+1}]$ and $\mathcal{T} = [t_0, t_{M_2+1}]$, respectively. Similarly to the standard smoothing spline estimator, by performing the minimization in Equation (6.3) over $\alpha \in \mathbb{S}_{k_1, k_2, M_1, M_2}$, we implicitly assume that β can be suitably approximated by an element in $\mathbb{S}_{k_1, k_2, M_1, M_2}$, that is

$$\beta(s, t) \approx \tilde{\beta}(s, t) \doteq \sum_{i=1}^{M_1+k_1} \sum_{j=1}^{M_2+k_2} b_{ij} \psi_i^s(s) \psi_j^t(t) = \boldsymbol{\psi}^s(s)^T \mathbf{B} \boldsymbol{\psi}^t(t) \quad s \in \mathcal{S}, t \in \mathcal{T}, \quad (6.4)$$

where $\mathbf{B} = \{b_{ij}\} \in \mathbb{R}^{M_1+k_1 \times M_2+k_2}$ and b_{ij} are scalar coefficients. So stated, the problem of estimating β has been reduced to the estimation of the unknown coefficient matrix \mathbf{B} . Let $\alpha(s, t) = \boldsymbol{\psi}^s(s)^T \mathbf{B}_\alpha \boldsymbol{\psi}^t(t)$, $s \in \mathcal{S}, t \in \mathcal{T}$, in $\mathbb{S}_{k_1, k_2, M_1, M_2}$, where $\mathbf{B}_\alpha = \{b_{\alpha, ij}\} \in \mathbb{R}^{M_1+k_1 \times M_2+k_2}$. Then, the first term of the right-hand side of Equation (6.3) may be rewritten as

$$\sum_{i=1}^n \|Y_i - \int_{\mathcal{S}} X_i(s) \alpha(s, \cdot) ds\|^2 = \sum_{i=1}^n \int_{\mathcal{T}} Y_i(t)^2 dt - 2 \operatorname{Tr} [\mathbf{X} \mathbf{B}_\alpha \mathbf{Y}^T] + \operatorname{Tr} [\mathbf{X}^T \mathbf{X} \mathbf{B}_\alpha \mathbf{W}_t \mathbf{B}_\alpha^T], \quad (6.5)$$

whereas, the roughness penalties on the left side of Equation (6.3) become

$$\lambda_s \|\mathcal{L}_s^{m_s} \alpha\|^2 = \lambda_s \operatorname{Tr} [\mathbf{B}_\alpha^T \mathbf{R}_s \mathbf{B}_\alpha \mathbf{W}_t] \quad \lambda_t \|\mathcal{L}_t^{m_t} \alpha\|^2 = \lambda_t \operatorname{Tr} [\mathbf{B}_\alpha^T \mathbf{W}_s \mathbf{B}_\alpha \mathbf{R}_t], \quad (6.6)$$

where $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)^T$, with $\mathbf{X}_i = \int_{\mathcal{S}} X_i(s) \boldsymbol{\psi}^s(s) ds$, $\mathbf{Y} = (\mathbf{Y}_1, \dots, \mathbf{Y}_n)^T$ with $\mathbf{Y}_i = \int_{\mathcal{T}} Y_i(t) \boldsymbol{\psi}^t(t) dt$, $\mathbf{W}_s = \int_{\mathcal{S}} \boldsymbol{\psi}^s(s) \boldsymbol{\psi}^s(s)^T ds$, $\mathbf{W}_t = \int_{\mathcal{T}} \boldsymbol{\psi}^t(t) \boldsymbol{\psi}^t(t)^T dt$, $\mathbf{R}_s = \int_{\mathcal{S}} \mathcal{L}_s^{m_s} [\boldsymbol{\psi}^s(s)] \mathcal{L}_s^{m_s} [\boldsymbol{\psi}^s(s)]^T ds$ and $\mathbf{R}_t = \int_{\mathcal{T}} \mathcal{L}_t^{m_t} [\boldsymbol{\psi}^t(t)] \mathcal{L}_t^{m_t} [\boldsymbol{\psi}^t(t)]^T dt$. The term $\operatorname{Tr} [\mathbf{A}]$ denotes the trace of a square matrix \mathbf{A} .

Standard optimization algorithms for L_1 -regularized objective functions are designed for the case where the absolute value enters the problem as a linear function of the parameters. However, in the optimization problem in Equation (6.3), the FLP particularizes as follows

$$\lambda_L \int_{\mathcal{S}} \int_{\mathcal{T}} |\alpha(s, t)| ds dt = \lambda_L \int_{\mathcal{S}} \int_{\mathcal{T}} |\boldsymbol{\psi}^s(s)^T \mathbf{B}_\alpha \boldsymbol{\psi}^t(t)| ds dt, \quad (6.7)$$

which is not a linear function of the absolute value of the coefficient matrix $|\mathbf{B}_\alpha|$, because the absolute value is instead applied to a linear combination of the parameters. Therefore, in order to be able to use optimization algorithms for L_1 -regularized objective functions by the following theorem, we provide a practical way to approximate the FLP as a linear function of $|\mathbf{B}_\alpha|$, and thus extremely simplify the computation.

Theorem 1. Let $\mathbb{S}_{k_1, k_2, \Delta_{1,e}, \Delta_{2,e}} = \text{span}\{B_{i_1} B_{i_2}\}_{i_1=1, i_2=1}^{M_1+k_1, M_2+k_2}$, with $\{B_{i_j}\}$ the set of B-splines of orders k_j and non-decreasing evenly spaced knots sequences $\Delta_j = \{x_{j,0}, x_{j,1}, \dots, x_{j,M_j}, x_{j,M_j+1}\}$ defined on the compact set $\mathcal{D}_j = [x_{j,0}, x_{j,M_j+1}]$ and $\Delta_{j,e}$ the extended partitions corresponding to Δ_j defined as $\Delta_{j,e} = \{y_{j,l}\}_{l=1}^{M_j+2k_j}$ where $y_{j,1}, \dots, y_{j,k_j} = x_{j,0}$, $y_{j,1+k_j}, \dots, y_{j,M_j+k_j} = x_{j,1}, \dots, x_{j,M_j}$ and $y_{j,M_j+1+k_j}, \dots, y_{j,M_j+2k_j} = x_{j,M_j+1}$, for $j = 1, 2$. Then, for $f(z_1, z_2) = \sum_{i_1=1}^{M_1+k_1} \sum_{i_2=1}^{M_2+k_2} c_{i_1 i_2} B_{i_1}(z_1) B_{i_2}(z_2) \in \mathbb{S}_{k_1, k_2, \Delta_{1,e}, \Delta_{2,e}}$, with $z_1 \in \mathcal{D}_1$ and $z_2 \in \mathcal{D}_2$,

$$0 \leq \|f\|_{\ell^1, \Delta_{1,e}, \Delta_{2,e}} - \|f\|_{L^1} = O\left(\frac{1}{M_1}\right) + O\left(\frac{1}{M_2}\right), \quad (6.8)$$

where

$$\|f\|_{\ell^1, \Delta_{2,e}, \Delta_{1,e}} = \sum_{i_1=1}^{M_1+k_1} \sum_{i_2=1}^{M_2+k_2} |c_{i_1 i_2}| \frac{(y_{1,i_1+k_1} - y_{1,i_1})(y_{2,i_2+k_2} - y_{2,i_2})}{k_1 k_2}, \quad (6.9)$$

and

$$\|f\|_{L^1} = \int_{\mathcal{D}_1} \int_{\mathcal{D}_2} |f(z_1, z_2)| dz_1 dz_2. \quad (6.10)$$

The interpretation of the above theorem is quite simple. It basically says that for large values of M_1 and M_2 , $\|f\|_{L^1}$ is well approximated from the top by $\|f\|_{\ell^1, \Delta_{2,e}, \Delta_{1,e}}$ and the approximation error tends to zero as $M_1, M_2 \rightarrow \infty$. By using this result, the FLP can be approximated as follows

$$\lambda_L \int_{\mathcal{S}} \int_{\mathcal{T}} |\alpha(s, t)| ds dt \approx \lambda_L \sum_{i=1}^{M_1+k_1} \sum_{j=1}^{M_2+k_2} |b_{\alpha, ij}| \frac{(s_{i+k_1}^e - s_i^e)(t_{j+k_2}^e - t_j^e)}{k_1 k_2} = \lambda_L \mathbf{w}_s^T |\mathbf{B}_\alpha| \mathbf{w}_t, \quad (6.11)$$

where $\{s_i^e\}$ and $\{t_i^e\}$ are the extended partitions associated with Δ^s and Δ^t , respectively,

$$\mathbf{w}_s = \left[\frac{(s_{1+k_1}^e - s_1^e)}{k_1}, \dots, \frac{(s_{M_1+2k_1}^e - s_{M_1+k_1}^e)}{k_1} \right]^T \text{ and } \mathbf{w}_t = \left[\frac{(t_{1+k_2}^e - t_1^e)}{k_2}, \dots, \frac{(t_{M_2+2k_2}^e - t_{M_2+k_2}^e)}{k_2} \right]^T.$$

Therefore, upon using the approximation in Equation (6.11), Equation (6.5) and Equation (6.6), the optimization problem in Equation (6.3) becomes

$$\begin{aligned} \hat{\mathbf{B}}_{SL} \approx \underset{\mathbf{B}_\alpha \in \mathbb{R}^{(M_1+k_1) \times (M_2+k_2)}}{\text{argmin}} \left\{ \sum_{i=1}^n \int_{\mathcal{T}} Y_i(t)^2 dt - 2 \text{Tr} [\mathbf{X} \mathbf{B}_\alpha \mathbf{Y}^T] + \text{Tr} [\mathbf{X}^T \mathbf{X} \mathbf{B}_\alpha \mathbf{W}_t \mathbf{B}_\alpha^T] \right. \\ \left. + \lambda_s \text{Tr} [\mathbf{B}_\alpha^T \mathbf{R}_s \mathbf{B}_\alpha \mathbf{W}_t] + \lambda_t \text{Tr} [\mathbf{B}_\alpha^T \mathbf{W}_s \mathbf{B}_\alpha \mathbf{R}_t] + \lambda_L \mathbf{w}_s^T |\mathbf{B}_\alpha| \mathbf{w}_t \right\}. \end{aligned} \quad (6.12)$$

Then, the coefficient β is estimated by $\hat{\beta}_{SL}(s, t) = \boldsymbol{\psi}^s(s)^T \hat{\mathbf{B}}_{SL} \boldsymbol{\psi}^t(t)$ for $s \in \mathcal{S}$ and $t \in \mathcal{T}$. Note that, in Equation (6.12), the FLP is approximated through a weighted linear combination of the absolute values of the coefficients which strictly resembles the multivariate LASSO penalty applied to the basis expansion coefficients, i.e. $\lambda_L \sum_{i=1}^{M_1+k_1} \sum_{j=1}^{M_2+k_2} |b_{\alpha, ij}|$. However, the presence of \mathbf{w}_s and \mathbf{w}_t in the FLP approximation in Equation (6.12) is crucial because it allows the penalty to differently shrink coefficients among B-splines. That is, it avoids that the absolute values of coefficients corresponding to B-splines strictly localized are weighted as

the absolute value of coefficients of spreader basis in the computation of the penalty. This is the direct consequence of the fact that the proposed approximation is a better approximation of the FLP than the multivariate LASSO penalty applied to the coefficients.

The optimization problem with L_1 -regularized loss in Equation (6.12) is (i) convex, being sum or integral of convex function; and (ii) has a unique solution given some general conditions on the matrix $\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X}$ (with \otimes the Kronecker product). See Section 6.3 for further details. Unfortunately, the objective function is not differentiable in zero, and thus it has not a closed-form solution. In view of this, general purpose gradient-based optimization algorithms—as for instance the *L-BFGS* quasi-Newton method (Nocedal and Wright, 2006)—and classical optimization algorithms for solving LASSO problems—such as coordinate descent (Friedman et al., 2010) and least-angle regression (LARS) (Efron et al., 2004)—are not suitable. In contrast, we found very promising a modified version of the *orthant-wise limited-memory quasi-Newton* (OWL-QN) algorithm proposed by Andrew and Gao (2007). The OWL-QN algorithm is based on the fact that the L_1 norm is differentiable for the set of points named *orthant* in which each coordinate never changes sign, being a linear function of its argument. In each orthant, the second-order behaviour of an objective function of the form $f(\mathbf{x}) = l(\mathbf{x}) + C\|\mathbf{x}\|_1$, to be minimized, is determined by l alone. The function $l : \mathbb{R}^r \rightarrow \mathbb{R}$ is convex, bounded below, continuously differentiable with continuously differentiable gradient ∇l , $\mathbf{x} = (x_1, \dots, x_r)^T$, C is a given positive constant, and $\|\cdot\|_1$ is the usual L_1 norm. Therefore, Andrew and Gao (2007) propose to derive a quadratic approximation of the function l that is valid for some orthant containing the current point and then to search for the minimum of the approximation, by constraining the solution in the orthant where the approximation is valid. There may be several orthants containing or adjacent to a given point. The choice of the orthant to explore is based on the *pseudo-gradient* $\diamond f(\mathbf{x})$ of f at \mathbf{x} , whose components are defined as

$$\diamond_i f(\mathbf{x}) = \begin{cases} \frac{\partial l(\mathbf{x})}{\partial x_i} + C \text{sign}(x_i) & \text{if } |x_i| > 0 \\ \frac{\partial l(\mathbf{x})}{\partial x_i} + C & \text{if } x_i = 0, \frac{\partial l(\mathbf{x})}{\partial x_i} < -C \\ \frac{\partial l(\mathbf{x})}{\partial x_i} - C & \text{if } x_i = 0, \frac{\partial l(\mathbf{x})}{\partial x_i} > C \\ 0 & \text{otherwise,} \end{cases} \quad (6.13)$$

where $\text{sign}(\cdot)$ denotes the usual sign function. However, the objective function of the optimization problem in Equation (6.12) is in the form $f^*(\mathbf{x}) = l(\mathbf{x}) + C\|\mathbf{D}\mathbf{x}\|_1$, with $\mathbf{D} = \{d_i\} \in \mathbb{R}^{r \times r}$ a diagonal matrix of positive weights. To take into account these weights, the OWL-QN algorithm must be implemented with a different *pseudo-gradient* $\diamond f^*(\mathbf{x})$ whose components are defined as

$$\diamond_i f^*(\mathbf{x}) = \begin{cases} \frac{\partial l(\mathbf{x})}{\partial x_i} + d_i C \text{sign}(x_i) & \text{if } |x_i| > 0 \\ \frac{\partial l(\mathbf{x})}{\partial x_i} + d_i C & \text{if } x_i = 0, \frac{\partial l(\mathbf{x})}{\partial x_i} < -C \\ \frac{\partial l(\mathbf{x})}{\partial x_i} - d_i C & \text{if } x_i = 0, \frac{\partial l(\mathbf{x})}{\partial x_i} > C \\ 0 & \text{otherwise.} \end{cases} \quad (6.14)$$

A more detailed description of the OWL-QN algorithm for objective functions in the form $l(\mathbf{x}) + C\|\mathbf{D}\mathbf{x}\|_1$ is given in the Supplementary Materials. Note that, the optimization

problem in Equation (6.12) can be rewritten by vectorization as

$$\hat{\mathbf{b}}_{SL} \approx \hat{\mathbf{b}}_{app} = \underset{\mathbf{b}_\alpha \in \mathbb{R}^{(M_1+k_1)(M_2+k_2)}}{\operatorname{argmin}} \left\{ -2 \operatorname{vec} \left(\mathbf{X}^T \mathbf{Y} \right)^T \mathbf{b}_\alpha + \mathbf{b}_\alpha^T \left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) \mathbf{b}_\alpha + \lambda_s \mathbf{b}_\alpha^T \mathbf{L}_{wr} \mathbf{b}_\alpha + \lambda_t \mathbf{b}_\alpha^T \mathbf{L}_{rw} \mathbf{b}_\alpha + \lambda_L \|\mathbf{W}_{st} \mathbf{b}_\alpha\|_1 \right\}, \quad (6.15)$$

where $\hat{\mathbf{b}}_{SL} = \operatorname{vec}(\hat{\mathbf{B}}_{SL})$, $\mathbf{L}_{rw} \doteq (\mathbf{R}_t \otimes \mathbf{W}_s)$ and $\mathbf{L}_{wr} \doteq (\mathbf{W}_t \otimes \mathbf{R}_s)$, and \mathbf{W}_{st} is the diagonal matrix whose diagonal elements are $\mathbf{w}_s^T \otimes \mathbf{w}_t^T$. Moreover, for generic a matrix $\mathbf{A} \in \mathbb{R}^{j \times k}$, $\operatorname{vec}(\mathbf{A})$ indicates the vector of length jk obtained by writing the matrix \mathbf{A} as a vector column-wise. Therefore, the OWL-QN with pseudo-gradient as in Equation (6.14) can be straightforwardly applied.

In the following, we summarize all the parameters that need to be set to obtain the S-LASSO estimator. The orders k_1 and k_2 should be chosen with respect to the degree of smoothness we want to achieve, and the computational efforts. The larger the values of k_1 and k_2 , the smoother the resulting estimator will be. A standard choice are cubic B-splines, with equally spaced knot sequences, i.e., $k_1 = k_2 = 4$. As stated before, M_1 and M_2 should be as large as possible to ensure that the null region is correctly captured and the approximation in Equation (6.11) is valid, with respect to the maximum computational efforts. Finally, at given k_1 , k_2 , M_1 , and M_2 , the optimal values of λ_s , λ_t and λ_L can be selected as those that minimize the estimated prediction error function $CV(\lambda_s, \lambda_t, \lambda_L)$, i.e., $CV(\lambda_s, \lambda_t, \lambda_L)$, over a grid of candidate values (Hastie et al., 2009a). However, although this choice could be optimal for the prediction performance, it may affect the interpretability of the model. Much more interpretable models, with a slight decrease in predictive performance, may in fact exist. The *k-standard error* rule, which is a generalization of the *one-standard error* rule (Hastie et al., 2009a), may be a more reasonable choice. That is, to choose the most sparse model whose error is no more than k standard errors above the error of the best model. In practice, as sparseness is controlled by the parameter λ_L , we first find the best model in terms of estimated prediction error at given λ_L and then, among the selected models, we apply the *k-standard error* rule. This rule may be particularly useful when $CV(\lambda_s, \lambda_t, \lambda_L)$ is flat with respect to λ_L . In this case, it chooses the simplest model among those achieving similar estimated prediction errors.

6.3 Theoretical Properties of the S-LASSO Estimator

In this section, we provide some theoretical results on the S-LASSO estimator, under some regularity assumptions, i.e., the estimation consistency (Theorem 2) and the pointwise sign consistency (Theorem 3) of $\hat{\beta}_{SL}$. All proofs are in the Supplementary Materials. The following regularity conditions are assumed.

C 1. $\|\mathbf{X}\|_2$ is almost surely bounded, i.e., $\|\mathbf{X}\|_2 \leq c < \infty$.

C 2. β is in the Hölder space $C^{p', \nu}(\mathcal{S} \times \mathcal{T})$ defined as the set of functions f on $\mathcal{S} \times \mathcal{T}$ having continuous partial and mixed derivatives up to order p' and such that the partial and mixed derivatives of order p' are Hölder continuous, i.e., $|f^{(p')}(\mathbf{x}_1) - f^{(p')}(\mathbf{x}_2)| \leq c \|\mathbf{x}_1 - \mathbf{x}_2\|^\nu$, for some constant c , integer p' and $\nu \in [0, 1]$, and for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{S} \times \mathcal{T}$, where $f^{(p')}$ is the partial and mixed derivatives of order p' . Moreover, let $p \doteq p' + \nu$ such that $3/2 < p \leq k_1 - 1$ and $3/2 < p \leq k_2 - 1$.

6. SMOOTH LASSO ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL

C 3. $M_1 = o\left(n^{1/4}\right)$, $M_2 = o\left(n^{1/4}\right)$, $M_1 = \omega\left(n^{\frac{1}{2p+1}}\right)$, $M_2 = \omega\left(n^{\frac{1}{2p+1}}\right)$, where $a_n = \omega(b_n)$ means $\frac{a_n}{b_n} \rightarrow \infty$ for $n \rightarrow \infty$,

C 4. There exist two positive constants b and B such that

$$b \leq \Lambda_{\min}\left(\mathbf{W}_t \otimes n^{-1} \mathbf{X}^T \mathbf{X}\right) \leq \Lambda_{\max}\left(\mathbf{W}_t \otimes n^{-1} \mathbf{X}^T \mathbf{X}\right) \leq B, \quad (6.16)$$

where $\Lambda_{\min}(\mathbf{M})$ and $\Lambda_{\max}(\mathbf{M})$ denote the minimum and maximum eigenvalues of the matrix \mathbf{M} .

C 5. $\lambda_s = o\left(M_1^{-2m_s+1}\right)$, $\lambda_t = o\left(M_2^{-2m_t+1}\right)$.

C.1 and C.2 are the anoulogus of (H1) and (H2) in Cardot et al. (2003) for a bivariate regression function. C.2 ensures that β is sufficiently smooth. C.3 provides information on the growth rate of the number of knots M_1 and M_2 , which are strictly related to the sample size n . C.4 is the anologus of condition (F) of Fan et al. (2004) and assumes that the matrix $(\mathbf{W}_t \otimes n^{-1} \mathbf{X}^T \mathbf{X})$ has reasonably good behaviour, whereas, C.5 provides guidance on the choice of the parameters λ_s and λ_t .

Theorem 2 shows that with probability tending to one there exists a solution of the optimization problem in Equation (6.3) that converges to $\tilde{\beta}$, chosen such that $\|\beta - \tilde{\beta}\|_{\infty} = O(M_1^{-p}) + O(M_2^{-p})$. To prove Theorem 2, in addition to C.1-C.5, the following condition is considered.

C 6. $\lambda_L = o\left(M_1^{-1} M_2^{-1}\right)$.

The first result is about the convergence rate of $\hat{\beta}_{SL}$ to β in terms of L_{∞} -norm.

Theorem 2. Under assumptions C.1-C.6, there exists a unique solution $\hat{\beta}_{SL}$ of the optimization problem in Equation (6.3), such that

$$\|\hat{\beta}_{SL} - \beta\|_{\infty} = O_p\left(M_1^{1/2} M_2^{1/2} n^{-1/2}\right). \quad (6.17)$$

According to the above theorem, there exists an estimator $\hat{\beta}_{SL}$ of β that is root- $n/M_1 M_2$ consistent.

Before stating Theorem 3, let us define with $\mathbf{b}_{(1)}$ the vector whose entries are the q non-zero elements of \mathbf{b} that are and with $\mathbf{b}_{(2)}$ the vector whose entries are the $(M_1 + k_1)(M_2 + k_2) - q$ elements of \mathbf{b} that are equal to zero. In what follows, we assume, without loss of generality, that $\mathbf{b} = \begin{bmatrix} \mathbf{b}_{(1)}^T & \mathbf{b}_{(2)}^T \end{bmatrix}^T$ and that a matrix $\mathbf{A}_l \in \mathbb{R}^{(M_1+k_1)(M_2+k_2) \times (M_1+k_1)(M_2+k_2)}$ can be expressed in block-wise form as

$$\mathbf{A}_l = \begin{bmatrix} \mathbf{A}_{l,11} \in \mathbb{R}^{q \times q} & \mathbf{A}_{l,12} \in \mathbb{R}^{q \times (M_1+k_1)(M_2+k_2)-q} \\ \mathbf{A}_{l,21} \in \mathbb{R}^{(M_1+k_1)(M_2+k_2)-q \times q} & \mathbf{A}_{l,22} \in \mathbb{R}^{(M_1+k_1)(M_2+k_2)-q \times (M_1+k_1)(M_2+k_2)-q} \end{bmatrix}.$$

To prove Theorem 3, in addition to C.1-C.5, the following conditions are considered.

C 7. (*S-LASSO irrepresentable condition (SL-IC)*) There exists λ_s , λ_t , λ_L , and a constant $\eta > 0$ such that, element-wise,

$$\begin{aligned} & \left| \mathbf{W}_{st,21}^{-1} \left\{ \left[\left(\mathbf{W}_t \otimes n^{-1} \mathbf{X}^T \mathbf{X} \right)_{21} + n^{-1} \lambda_s \mathbf{L}_{wr,21} + n^{-1} \lambda_t \mathbf{L}_{rw,21} \right] \right. \right. \\ & \quad \left. \left[\left(\mathbf{W}_t \otimes n^{-1} \mathbf{X}^T \mathbf{X} \right)_{11} + n^{-1} \lambda_s \mathbf{L}_{wr,11} + n^{-1} \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \right. \\ & \quad \left. \left[\mathbf{W}_{st,11} \text{sign}(\mathbf{b}_{\alpha(1)}) + 2\lambda_L^{-1} \lambda_s \mathbf{L}_{wr,11} \mathbf{b}_{(1)} + 2\lambda_L^{-1} \lambda_t \mathbf{L}_{rw,11} \mathbf{b}_{(1)} \right] \right. \\ & \quad \left. \left. - 2\lambda_L^{-1} \lambda_s \mathbf{L}_{wr,21} \mathbf{b}_{(1)} - 2\lambda_L^{-1} \lambda_t \mathbf{L}_{rw,21} \mathbf{b}_{(1)} \right\} \right| \leq 1 - \eta. \end{aligned}$$

C 8. The functions $\varepsilon_i(t)$ in Equation (6.1) are zero mean Gaussian random processes with autocovariance function $K(t_1, t_2)$, t_1 and $t_2 \in \mathcal{T}$, independent of X_i .

C 9. Given $\rho \doteq \min \left| \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} \mathbf{b}_{(1)} \right] \right|$ and $C_{\min} \doteq \Lambda_{\min} \left[\left(\mathbf{W}_t \otimes n^{-1} \mathbf{X}^T \mathbf{X} \right)_{11} \right]$, $\Lambda_{\min}(\mathbf{W}_t) M_2 \rightarrow c_w$ as $n \rightarrow \infty$, with $0 < c_w < \infty$, and the parameters λ_s , λ_t and λ_L are chosen such that

1.

$$\frac{M_1^2 M_2^2 \log[(M_1 + k_1)(M_2 + k_2) - q]}{\lambda_L^2} \left[nc^2 + \frac{\lambda_s^2 \Lambda_{\max}^2(\mathbf{L}_{wr})}{nC_{\min}} + \frac{\lambda_t^2 \Lambda_{\max}^2(\mathbf{L}_{rw})}{nC_{\min}} \right] = o(1),$$

2.

$$\begin{aligned} & \frac{1}{\rho} \left\{ \sqrt{\frac{M_1 M_2 \log(q)}{nC_{\min}}} \right. \\ & \quad \left. + \frac{\lambda_L}{n M_1 M_2} \Lambda_{\min}^{-1} \left[\left(\mathbf{W}_t \otimes n^{-1} \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s n^{-1} \mathbf{L}_{wr,11} + \lambda_t n^{-1} \mathbf{L}_{rw,11} \right] \|\text{sign}(\mathbf{b}_{(1)})\|_2 \right\} = o(1). \end{aligned}$$

The SL-IC in C.7 is the straightforward generalization to the problem in Equation (6.3) of the elastic irrepresentable condition described in Jia and Yu (2010). It is a consequence of the standard Karush–Kuhn–Tucker (KKT) conditions applied to the optimization problem in Equation (6.12). C.8 gives some conditions on the relationship of λ_s , λ_t , and λ_L with M_1 , M_2 and n . In the classical setting, an estimator is sign selection consistent if it has the same sign of the true parameter with probability tending to one. Analogously, we say that an estimator of the coefficient function β is pointwise sign consistent if, in each point of the domain, it has the same sign of β with probability tending to one. The following theorem states that, under opportune assumptions, the S-LASSO estimator is pointwise sign consistent.

Theorem 3. Under assumptions C.1-C.5 and C.7-C.9, $\hat{\beta}_{SL}$ is pointwise sign consistent, that is, for all $s \in \mathcal{S}$ and $t \in \mathcal{T}$,

$$\Pr \left\{ \text{sign} \left[\hat{\beta}_{SL}(s, t) \right] = \text{sign} [\beta(s, t)] \right\} \rightarrow 1, \quad (6.18)$$

as $n \rightarrow \infty$.

This theorem is the functional extension of the sign consistency result for the multivariate LASSO estimator (Zou and Zhang, 2009).




Figure 6.2. True coefficient function β for Scenario II a, Scenario III b and Scenario IV c in the simulation study.

6.4 Simulation Study

In this section, we conduct a Monte Carlo simulation study to explore the performance of the S-LASSO estimator. We consider four scenarios whose corresponding coefficient functions are depicted in Figure 6.2. Note that the coefficient function for Scenario I is not shown because it is zero all over the domain. In Scenario II and III, β is sparse, indeed, it is zero on the edge and in the central part of the domain, respectively. Scenario IV corresponds to a non-sparse setting, which is not expected to be favourable to the S-LASSO estimator. Further details on the data generation are given in the Supplementary Materials.

For each scenario, we generate 100 datasets composed of a training set with sample size n and a test set T with size N equal to 4000 that are used to estimate the coefficient function and to test its predictive performance. This is repeated for three different sample sizes $n = 150, 500, 1000$. As in Lin et al. (2017), we consider the integrated squared error (ISE) to assess the quality of the estimator $\hat{\beta}$ of the coefficient function β . In particular, the ISE over the null region (ISE_0) and the non-null region (ISE_1) are defined as

$$\text{ISE}_0 = \frac{1}{A_0} \int \int_{N(\beta)} \left(\hat{\beta}(s, t) - \beta(s, t) \right)^2 ds dt, \quad \text{ISE}_1 = \frac{1}{A_1} \int \int_{NN(\beta)} \left(\hat{\beta}(s, t) - \beta(s, t) \right)^2 ds dt, \quad (6.19)$$

where A_0 and A_1 are the measures of the null ($N(\beta)$) and non-null ($NN(\beta)$) regions, respectively. The ISE_0 and the ISE_1 are indicators of the estimation error of $\hat{\beta}$ over both the null and the non-null regions. Moreover, predictive performance is measured through the prediction mean squared error (PMSE), defined as

$$\text{PMSE} = \frac{1}{N} \sum_{(X, Y) \in T} \int_0^1 \left(Y(t) - \int_0^1 X(s) \hat{\beta}(s, t) ds \right)^2 dt, \quad (6.20)$$

where $\hat{\beta}$ is obtained through the observations in the training set. The observations in the test set are centred by means of the sample mean functions estimated through the training set observations.

The S-LASSO estimator is compared with four different estimators of β which are already present in the literature of the FoF linear regression model estimation. The first two are




Figure 6.3. The integrated squared error on the null region (ISE_0) along with $\pm 0.5(\text{standard error})$ for the TRU(), SMOOTH(.....), PCA(---), RIDGE(- - -), and S-LASSO (—) estimators.

those proposed by Ramsay and Silverman (2005), where the coefficient function estimator is assumed to be in a finite dimension tensor space with regularization achieved either by choosing the dimension of the tensor space or by introducing roughness penalties. They will be referred to as TRU and SMOOTH estimators, respectively. The third and fourth ones are those proposed by Yao et al. (2005b), based on the functional principal components analysis (referred to as PCA), and by Canale and Vantini (2016), based on a ridge-type penalization (referred to as RIDGE). The TRU, SMOOTH and S-LASSO are computed by using cubic B-splines with evenly space knot sequences. The dimensions of the B-spline sets that generate the tensor product space for the SMOOTH and S-LASSO estimator are both set equal to 60. All the tuning parameters of the five considered estimators are chosen by means of 10-fold cross-validation, viz., the dimension of the tensor basis space for the TRU, the roughness penalties for the SMOOTH, the numbers of retained principal components for the PCA, the penalization parameter for the RIDGE and λ_s , λ_t and λ_L for the S-LASSO. In particular the 10-fold cross-validation for the S-LASSO method is applied with the 0.5-standard deviation rule.

The performance of the estimators in terms of ISE_0 is displayed in Figure 6.3. It is not surprising that the estimation error of β over $N(\beta)$ of the S-LASSO estimator is significantly smaller than those of the other estimators, being the capability of recovering sparseness of β its main feature. In Scenario I, the RIDGE estimator is the only one that performs comparably to the S-LASSO estimator. This is in accordance with the multivariate setting where it is well known that, when the response is independent of the covariates, the ridge estimator is able to shrink all the coefficients towards zero. The TRU, SMOOTH, and PCA estimators have difficulties to correctly identify $N(\beta)$ for all sample sizes. Nevertheless, their performance is very poor at $n = 150$. In Scenario II, the S-LASSO estimator is still the best one to estimate β over $N(\beta)$. However, in this case, the RIDGE estimator performance is unsatisfactory and is mainly caused by the lack of smoothness control that makes the estimator over-rough, especially at small n . Among the competitor estimators, the SMOOTH one has the best performance. In Scenario III, results are similar to those of Scenario II, even if the TRU estimator appears as the best alternative. Both PCA and RIDGE estimators are not able to successfully recover sparseness of β for $n = 150$. For the former, the cause is the number of observations not sufficient to capture the covariance structure of the data, whereas for the latter, it is due to the excessive roughness of the estimator.

Results in terms of ISE_1 are summarized in Figure 6.4. It is worth noting that, in this case, as expected the performance of the S-LASSO estimator is generally worse than that

6. SMOOTH LASSO ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL




Figure 6.4. The integrated squared error on the non-null region (ISE_1) along with $\pm 0.5(\text{standard error})$ for the TRU (· · · · ·), SMOOTH (· · · · ·), PCA (· · · · ·), RIDGE (· · · · ·), and S-LASSO (—) estimators.

of the SMOOTH estimator. In some cases, it is worse than that of the TRU estimator as well. However, in Scenario II performance differences between the S-LASSO estimator and TRU or SMOOTH estimators become negligible as sample size increases. The PCA and RIDGE estimators are always less efficient. The results are similar for Scenario III, where the performance of the S-LASSO estimator is comparable with that of the SMOOTH estimator. By comparing to the classical LASSO method, the behaviour of the S-LASSO estimator — in terms of ISE_1 — is not surprising. Indeed, it is well known that LASSO method does nice variable selection, even if it tends to overshrink the estimators of the non-null coefficients (Fan et al., 2004; James and Radchenko, 2009). By looking at the result for Scenario II and III, we surmise that this phenomenon arises in the FoF linear regression model as well. Finally, in Scenario IV, where β is always different from zero, the S-LASSO estimator performs comparably to the SMOOTH (i.e., the S-LASSO estimator with $\lambda_L = 0$). In this case β is not sparse and, thus, the FLP does not help.

Figure 6.5 shows PMSE averages and corresponding standard errors for all the considered estimators. Since PMSE is strictly related to the ISE_0 and the ISE_1 , results are totally consistent with those of Figure 6.3 and Figure 6.4. In particular, the S-LASSO estimator outperforms all the competitor ones in favorable scenarios (viz., Scenario I, II, and III), being the corresponding PMSE lower than that achieved by the other competing estimators. In these scenarios, although the performance of the S-LASSO estimator in terms of ISE_1 is not excellent, the clear superiority in terms of ISE_0 compensates and gives rise to smaller PMSE. Otherwise, for Scenario IV, where the coefficient function is not sparse, the performance of the S-LASSO estimator is very similar to that of the SMOOTH estimator, which is the best one in this case. This is encouraging, because, it proves that the performance of the S-LASSO estimator does not dramatically decline in less favourable scenarios.

In summary, the S-LASSO estimator outperforms the competitors both in terms of estimation error on the null region and prediction accuracy on a new dataset, as well as that it is able to estimate competitively the coefficient function on the non-null region. On the other hand, in order to achieve sparseness, the S-LASSO tends to overshrink the estimator of the coefficient function on the non-null region. This means that, as in the classical setting (James and Radchenko, 2009), there is a trade-off between the ability of recovering sparseness and the estimation accuracy on the non-null region of the final estimator. Moreover, even when the coefficient function is not sparse (Scenario IV), the proposed estimator demonstrates to have both good prediction and estimation performance. This is another key property of the proposed estimator that, encourages practitioners to use the S-LASSO estimator

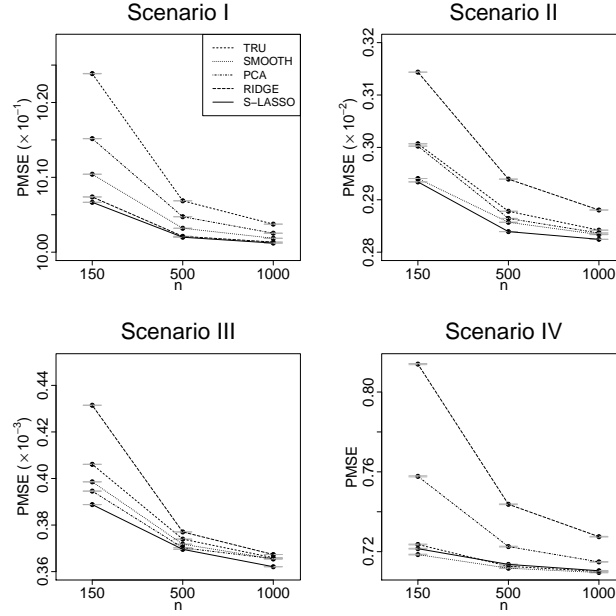


Figure 6.5. The prediction mean squared error (PMSE) along with $\pm 0.5(\text{standard error})$ for the TRU(.....), SMOOTH(.....), PCA(-.-.-), RIDGE(- - -), and S-LASSO (—) estimators.

even when there is not prior knowledge about the shape of the coefficient function. Finally, it should be noticed that, in scenarios similar to those analysed, the PCA and RIDGE estimators should not be preferred with respect to the TRU, SMOOTH and S-LASSO ones.

6.5 Real-Data Examples

In this section, we analyse two real-data examples. We aim to confirm that the S-LASSO estimator has advantages in terms of both prediction accuracy and interpretability, over the SMOOTH estimator, which has been demonstrated in Section 6.4 to be the best alternative among the competitors. The datasets used in the examples are the *Canadian weather* and *Swedish mortality*. Both are classical benchmark functional data sets thoroughly studied in the literature.

Canadian Weather Data

The Canadian weather data have been studied by Ramsay and Silverman (2005) and Sun et al. (2018). The data set contains the daily mean temperature curves, measured in Celsius degree, and the log-scale of the daily rainfall profiles, measured in millimeter, recorded at 35 cities in Canada. Both temperature and rainfall profiles are obtained by averaging over the years 1960 through 1994. Figure 6.6 shows the profiles. The aim is to predict the log-daily rainfall based on the daily temperature using the model reported in Equation (6.1). Figure 6.7 shows the S-LASSO and SMOOTH estimates of the coefficient function β . The SMOOTH estimate is obtained using a Fourier basis—to take into account the periodicity of the data—and roughness penalties were chosen by using 10-fold cross-validation over an

6. SMOOTH LASSO ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL




Figure 6.6. Daily mean temperature and log-daily rainfall profiles at 35 cities in Canada over the year.

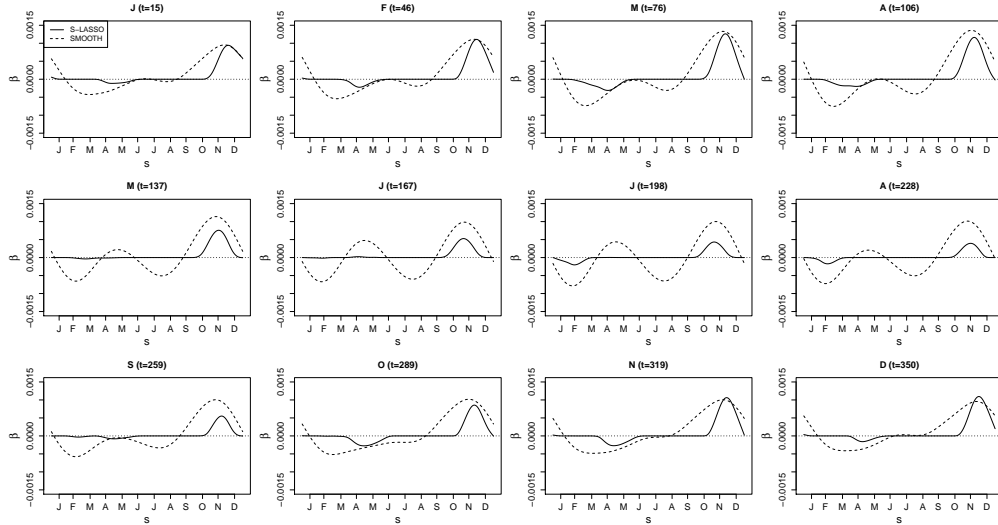


Figure 6.7. S-LASSO (—) and SMOOTH (-----) estimates of the coefficient functions β at different months for the Canadian weather data.

opportune grid of values. 10-fold cross-validation is used to set the parameters λ_s , λ_t and λ_L as well.

The S-LASSO estimates is roughly zero over large domain portions. In particular, except for values from July through August, it is always zero in summer months (i.e., late June, July, August and September) and in January and February. This suggests in those months rainfalls are not significantly influenced by daily temperature throughout the year. Otherwise, temperature in fall months (i.e., October, November and December) gives strong positive contribution on the daily rainfalls. In other words, the higher (the lower) the temperature in October, November and December, the heavier (the lighter) the precipitations throughout the year. It is interesting that the S-LASSO estimate in spring months (i.e., March, April and May) is negative for values of t from January through April, and from October through December. This suggests that the higher (the lower) the temperature in the spring the lighter (the heavier) the daily rainfalls from October through April. Finally, it is evidenced




Figure 6.8. For the Canadian weather data, a R_g^2 from permuting the response 500 times, where the black line corresponds to the observed R_g^2 and the grey line to the 95th sample percentile; b the black line is the observed R^2 and the grey line is the pointwise 95th sample percentile curve.

a small influence of the temperature in February on precipitations in July and August. It is worth noting that the S-LASSO estimate is more interpretable than the SMOOTH estimates, which does not allow for a straightforward interpretation. Moreover, the S-LASSO estimate appears to have, even if slightly, better prediction performance than the SMOOTH one. Indeed, 10-fold cross-validation mean squared errors are 22.314 and 22.365, respectively.

Finally, we perform two permutation tests to assess the statistical significance of the S-LASSO estimator. The first test is based on the *global functional coefficient of determination*

defined as $R_g^2 \doteq \int_{\mathcal{T}} \frac{\text{Var}[\mathbb{E}(Y(t)|X)]}{\text{Var}[Y(t)]} dt$ (Horváth and Kokoszka, 2012), with $\mathcal{T} = [0, 365]$. In

Figure 6.8a the solid black line indicates the observed R_g^2 that is equal to 0.55. The bold points represent 500 R_g^2 values obtained by means of random permutations of the response variable. Whereas, the grey line correspond to the 95th sample percentile. All 100 values of R^2 as well as the value of the 95th sample percentile is far below 0.55, which gives a strong evidence of a significant relationship between rainfalls and temperature, globally.

By a second test, we aim to analyse the pointwise statistical significance, i.e., for each $t \in \mathcal{T}$. It is based on the *pointwise functional coefficient of determination* defined as

$R^2(t) \doteq \frac{\text{Var}[\mathbb{E}(Y(t)|X)]}{\text{Var}[Y(t)]}$ for $t \in \mathcal{T}$ (Horváth and Kokoszka, 2012). Figure 6.8b shows the

observed R^2 (solid black line) along with the pointwise 95th sample percentile curve. The latter has been obtained by means of 500 R^2 values produced by randomly permuting the response variable. The observed R^2 is far above the 95th sample percentile curve, except for some summer months (viz., July and August). As global conclusion, we can state that the temperature has a large influence on the rainfalls in autumn, winter and spring.

Swedish Mortality Data

The Swedish mortality data, available from the Human Mortality Database (<http://mortality.org>), are regarded as a very reliable dataset on long-term longitudinal mortalities. In particular, we focus on the log-hazard rate functions of the Swedish females mortality data for year-of-birth cohorts that refer to females born in the years 1751-1894 with ages 0-80. The value of a log-hazard rate function at a given age is the natural logarithm of the

6. SMOOTH LASSO ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL




Figure 6.9. Log-hazard rates as a function of age for Swedish female cohorts born in the years 1751-1894.

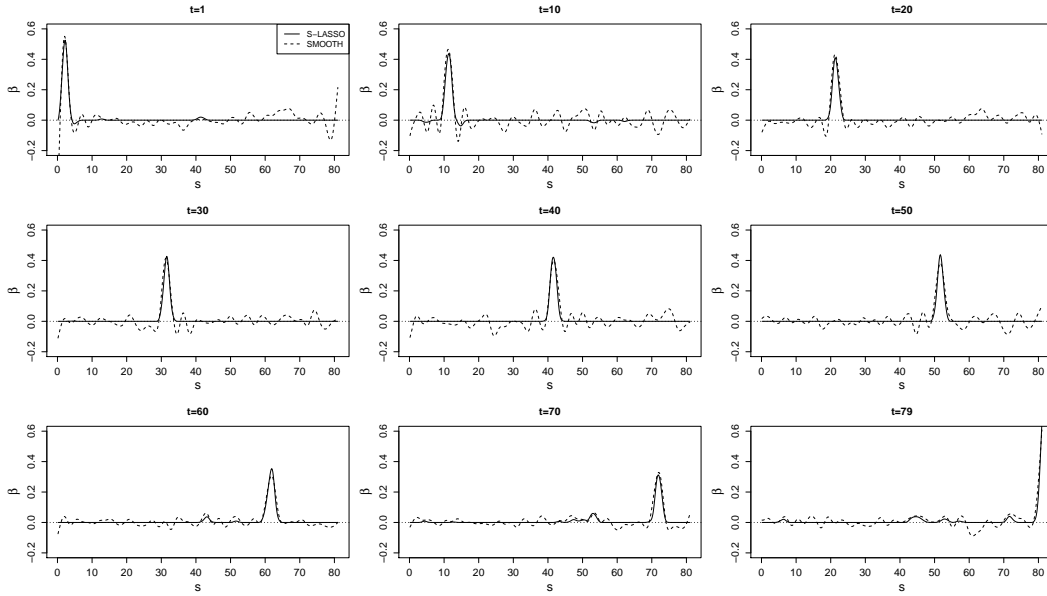


Figure 6.10. S-LASSO (—) and SMOOTH (-----) estimates of the coefficient functions β at different values of t for the Swedish mortality data.

ratio of the number of females who died at that age and the number of females alive with the same age. Note that, those data have been analysed also by Chiou and Müller (2009) and Ramsay et al. (2009). Figure 6.9 shows the 144 log-hazard functions.

The aim of the analysis is to explore the relationship of the log-hazard rate function for a given year with the log-hazard rate curve of the previous year by means of the model reported in Equation (6.1). Our interest is to identify what features of the log-hazard rate functions for a given year influence the log-hazard rate of the following year.

Figure 6.10 shows the S-LASSO and SMOOTH estimates of coefficient function β . The unknown parameters to obtain the SMOOTH and S-LASSO estimates are chosen as in the Canadian weather example, but in this case B-splines are used for both estimators. The S-LASSO estimate is zero almost over all the domain except for few regions. In particular,




Figure 6.11. For the Swedish mortality data, a R_g^2 from permuting the response 500 times, where the black line corresponds to the observed R_g^2 and the grey line to the 95th sample percentile; b the black line is the observed R^2 and the grey line is the pointwise 95th sample percentile curve.

at given t , the S-LASSO estimate is different from zero in an interval located right after that age. This can likely support the conjecture that if an event influences the mortality of the Swedish female at a given age, it impacts on the death rate below that age born in the following years. Nevertheless, this expected dependence is poorly pointed out by the SMOOTH estimator, where this behaviour is confounded by less meaningful periodic components. It is interesting to note that the S-LASSO estimate at high values of t is slightly different from zero for ages ranging from 40 to 60. This shows that if an event affecting the death rate occurs in that range, the log-hazard functions of the following cohorts will be influenced at high ages (i.e., corresponding to high values of t). On the contrary, the wiggle of the SMOOTH estimate does not allow drawing such conclusions.

Finally, we perform the two permutation test already described in the Canadian weather data example. Figure 6.11 shows the results. Both the observed R_g^2 and R^2 are far above the 95th sample percentile (Figure 6.11a) and the pointwise 95th sample percentile curve (Figure 6.11b) respectively. This significantly evidences a relation between two consecutive log-hazard rate functions for all ages.

6.6 Conclusions

The LASSO is one of the most used and popular method to estimate coefficients in classical linear regression models as it ensures both prediction accuracy and interpretability of the phenomenon under study (by simultaneously performing variable selection). In this chapter, we propose the S-LASSO estimator for the coefficient function of the function-on-function (FoF) linear regression model, which is an extension to the functional setting of the multivariate LASSO estimator. As the latter, the S-LASSO estimator is able to increase both the prediction accuracy of the estimated model, via continuous shrinking, and the interpretability, by identifying the null region of the regression coefficient, i.e., the region where the coefficient function is exactly zero.

The S-LASSO estimator is obtained by combining several elements: the *functional LASSO penalty* (FLP), which has the task of shrinking towards zero the estimator on the null region; the B-splines, which are essential to ensure sparsity of the estimator because of the compact

support property; and two roughness penalties, which are needed to ensure smoothness of the estimator on the non-null region, also when the number of B-splines escalates. We proved that the S-LASSO estimator is both estimation and point-wise sign consistent, i.e., the estimation error in terms of L^2 -norm goes to zero in probability and the S-LASSO estimator has the same sign of the true coefficient function with probability tending to one. Moreover, we showed via an extensive Monte Carlo simulation study that, with respect to other methods that have already appeared in the literature, the S-LASSO estimator is much more interpretable, on the one hand, and has still good estimation and appealing predictive performance, on the other. However, consistently with the behaviour of the classical LASSO estimator (Fan et al., 2004), the S-LASSO estimator is found sometimes to over-shrink the coefficient function over the non-null region.

To the best of the authors knowledge, this is the first work that addresses the issue of interpretability, intended as sparseness of the coefficient function, in the FoF linear regression setting. However, although the FLP produces an estimator with good properties, other penalties, e.g. the *SCAD* (Fan and Li, 2001) and *adaptive LASSO* (Zou, 2006), properly adapted to the functional setting, may guarantee even better performance both in terms of interpretability and prediction accuracy, and are, indeed, subjects of ongoing research.

6.7 Supplementary Materials

Orthant-Wise Limited-memory Quasi-Newton Algorithm

Let consider a loss function $l : \mathbb{R}^r \rightarrow \mathbb{R}$ that is convex, bounded below and continuously differentiable with continuously differentiable gradient ∇l . Then, the objective function to be minimized is $f(\mathbf{x}) = l(\mathbf{x}) + C\|\mathbf{D}\mathbf{x}\|_1$, where $\mathbf{x} = (x_1, \dots, x_r)^T$, $\mathbf{D} = \{d_i\} \in \mathbb{R}^{r \times r}$ is a diagonal matrix of positive weights, C is a given positive constant, and $\|\cdot\|_1$ is the usual L^1 norm. The OWL-QN explores points

$$\mathbf{x}^{k+1} = \pi_{\boldsymbol{\xi}^k}(\mathbf{x}^k + \alpha \mathbf{p}^k). \quad (6.21)$$

The function $\pi_{\mathbf{y}} : \mathbb{R}^r \rightarrow \mathbb{R}^r$, for $\mathbf{y} = (y_1, \dots, y_r)^T$, applied to generic vector $\mathbf{z} = (z_1, \dots, z_r)^T$, is defined as the vector whose components are

$$\pi_{\mathbf{y}i}(\mathbf{z}) = \begin{cases} z_i & \text{if } \text{sign}(z_i) = \text{sign}(y_i) \\ 0 & \text{otherwise,} \end{cases}$$

and it is the projection operator of \mathbf{z} on the orthant defined by \mathbf{y} . The vector $\boldsymbol{\xi}^k$ is in $\{-1, 0, 1\}^r$ and has components

$$\boldsymbol{\xi}_i^k = \begin{cases} \text{sign}(x_i^k) & \text{if } x_i^k \neq 0 \\ \text{sign}(-\diamond_i f^*(\mathbf{x}^k)) & \text{if } x_i^k = 0, \end{cases}$$

where $\diamond_i f^*(\mathbf{x}^k)$ are the components of the pseudo-gradient of f in \mathbf{x}^k as defined in Equation (6.14). The vector $\mathbf{p}^k \in \mathbb{R}^r$ is defined as

$$\mathbf{p}^k = \pi_{\mathbf{v}^k}(\mathbf{H}_k \mathbf{v}^k),$$

where $\mathbf{v}^k = -\diamond f^*(\mathbf{x}^k)$ and \mathbf{H}_k is the *L-BFGS* approximation to the inverse Hessian of the loss l (Nocedal and Wright, 2006). Finally, the step size α is chosen by using a variation

backtracking line search method (Nocedal and Wright, 2006). That is, for $n = 1, 2, \dots$, accept the first step size $\alpha = \alpha_0 0.5^n$, with $\alpha_0 = 1/\|\mathbf{v}^k\|$, such that

$$f(\mathbf{x}^{k+1}) \leq f(\mathbf{x}^k) - \gamma (\mathbf{x}^{k+1} - \mathbf{x}^k)^T \mathbf{v}^k \quad (6.22)$$

where γ is a parameter that controls the accuracy of the line search method. Given a starting point \mathbf{x}^0 , the OWL-QN finds a sequence of approximate solutions of the objective function thorough Equation (6.21) until a stopping condition is met. A pseudo code of the algorithm is given in Algorithm 1. All the details about the rationale of the aforementioned steps are in Andrew and Gao (2007).

Algorithm 1 OWL-QN

```

Set  $k = 0$ 
Choose the initial point  $\mathbf{x}^0$ 
repeat
  Compute  $\mathbf{v}^k = -\diamond f^*(\mathbf{x}^k)$ 
  Compute  $\mathbf{H}_k$  with L-BFGS
  Compute  $\mathbf{d}_k = \mathbf{H}_k \mathbf{v}^k$ 
  Compute  $\mathbf{p}^k = \pi_{\mathbf{v}^k}(\mathbf{d}_k)$ 
  Find  $\mathbf{x}^{k+1}$  by using the line search method in Equation (6.22)
  Set  $k = k + 1$ 
until the stopping condition is met

```

Details on the Data Generation in the Simulation Study

Four different scenarios are considered

- **Scenario I** The coefficient function is zero all over the domain, i.e., $\beta(s, t) = 0$, if $(s, t) \in [0, 1] \times [0, 1]$.

- **Scenario II** β is different from zero in the central part of the domain, i.e.,

$$\beta(s, t) = \begin{cases} -\left(\frac{s-0.5}{0.25}\right)^2 - \left(\frac{t-0.5}{0.25}\right)^2 + 1 & \text{if } 0.5 - 0.25\sqrt{1 - (t-0.5)^2} \leq s \leq 0.5 + 0.25\sqrt{1 - (t-0.5)^2} \\ 0 & \text{otherwise.} \end{cases} \quad (6.23)$$

- **Scenario III** β is different from zero on the edge of the domain, i.e.,

$$\beta(s, t) = \begin{cases} 0.5(1-t) \sin \left[10\pi \left(t - 1.05 + \sqrt{1 - (s-0.5)^2} \right) \right] & \text{if } t \leq 1.05 - \sqrt{1 - (s-0.5)^2} \\ 0.5 \sin \left(10\pi \left(s + 1.05 + \sqrt{1 - (t-0.5)^2} \right) \right) & \text{if } s \leq -0.05 - \sqrt{1 - (t-0.5)^2} \\ 0 & \text{otherwise.} \end{cases} \quad (6.24)$$

- **Scenario IV** β is non-null everywhere.

$$\beta(s, t) = \left(\frac{t-0.5}{0.5}\right)^3 + \left(\frac{s-0.5}{0.5}\right)^3 + \left(\frac{t-0.5}{0.5}\right)^2 + \left(\frac{s-0.5}{0.5}\right)^2 + 5. \quad (6.25)$$

This scenario is not expected to be favourable to the S-LASSO estimator.

6. SMOOTH LASSO ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL

The independent observations of the covariates X_i and errors ε_i are generated as $X_i = \sum_{j=1}^{32} x_{ij} \psi_i^x$ and $\varepsilon_i = k \sum_{j=1}^{20} e_{ij} \psi_i^\varepsilon$, where the coefficients x_{ij} and e_{ij} are independent realizations of standard normal random variable, and $\psi_i^x(s)$ and $\psi_i^\varepsilon(s)$ are cubic B-splines with evenly spaced knot sequence (the numbers of basis has been randomly chosen between 10 and 50). In Scenario I, the constant k is chosen equal to 1; whereas, in Scenario II, Scenario III and Scenario IV, it is chosen such that the modified signal-to-noise ratio function $MSN \doteq \text{Var}[E(Y_i|X_i)] + \max \text{Var}[E(Y_i|X_i)] / \text{Var}(\varepsilon_i)$ is equal to 4.

Proof of Theorems

Preliminaries Let $\mathbb{S}_{k_1, k_2, \Delta_{1,e} \Delta_{2,e}} = \text{span}\{\psi_{i_1} \psi_{i_2}\}_{i_1=1, i_2=1}^{M_1+k_1, M_2+k_2}$, with $\{\psi_{i_j}\}$ the set of B-splines of orders k_j and non-decreasing knots sequences $\Delta_j = \{x_{j,0}, x_{j,1}, \dots, x_{j,M_j}, x_{j,M_j+1}\}$ defined on the compact set $\mathcal{D}_j = [x_{j,0}, x_{j,M_j+1}]$ and $\Delta_{j,e}$ the extended partitions corresponding to Δ_j defined as $\Delta_{j,e} = \{y_{j,l}\}_{l=1}^{M_j+2k_j}$ where $y_{j,1}, \dots, y_{j,k_j} = x_{j,0}$, $y_{j,1+k_j}, \dots, y_{j,M_j+k_j} = x_{j,1}, \dots, x_{j,M_j}$ and $y_{j,M_j+1+k_j}, \dots, y_{j,M_j+2k_j} = x_{j,M_j+1}$, for $j = 1, 2$. Let $\delta_j \doteq M_j^{-1}$, for $j = 1, 2$. Moreover, the capital letter C with different subscripts indicates a positive constant. For convenience, we assume in what follows that $\mathcal{D}_1 = \mathcal{D}_2 = [0, 1]$, but the same arguments can be easily extended to general domains. Let $f \in L^p(\Omega)$, with Ω a bounded open set in \mathbb{R}^d , the usual d -dimensional Euclidian space, $\mathbf{a} = (a_1, \dots, a_n)^T$, with $a_i \in \mathbb{R}$, and $\mathbf{h} = (h_1, \dots, h_d)^T$, with $h_i \in \mathbb{R}$, then,

$$\omega_{E_{\mathbf{a}}}(f, \mathbf{h})_p \doteq \sum_{\gamma \in E_{\mathbf{a}}} \omega_{\gamma}(f, \mathbf{h})_p$$

where $E_{\mathbf{a}} = \{a_i \mathbf{e}_i\}_{i=1}^n$, with \mathbf{e}_i the unit vector in the i -th direction, and $\omega_{\gamma}(f, \mathbf{h})_p$ is the γ -modulus of smoothness in the p -norm of f defined as

$$\omega_{\gamma}(f, \mathbf{h})_p \doteq \sup_{0 \leq t \leq \mathbf{h}} \|\Delta_{\mathbf{t}}^{\gamma}(f)\|_{L^p(\Omega)}$$

is the γ -modulus of smoothness in the p -norm of f , where $\gamma = \{\gamma_1, \dots, \gamma_d\}$ is a multi-index, $\|\cdot\|_{L^p(\Omega)}$ is the usual L^p -norm, and

$$\Delta_{\mathbf{t}}^{\gamma}(f) \doteq \Delta_{t_1}^{\gamma_1} \dots \Delta_{t_d}^{\gamma_d}(f),$$

with $\Delta_{t_i}^{\gamma_i}$ the usual γ_i -forward difference of step length t_i applied to the i -th variable of f (Schumaker, 2007). Moreover, $D^{\alpha} \doteq D_1^{\alpha_1} \dots D_d^{\alpha_d}$, where $D_i^{\alpha_i}$ stands for the α_i derivative in the i -th variable and $\alpha = \{\alpha_1, \dots, \alpha_d\}$ is a multi-index.

Proof of Theorem 1 The proof that $\|f\|_{\ell^1, \Delta_{2,e}, \Delta_{1,e}} - \|f\|_{L^1} \geq 0$ is easy and thus is skipped. To prove that $\|f\|_{\ell^1, \Delta_{1,e}, \Delta_{2,e}} - \|f\|_{L^1} = O\left(\frac{1}{M_1}\right) + O\left(\frac{1}{M_2}\right)$, let define two new knot sequences $\{\tau_{i,j}\}_{j=1}^{M_i+k_i+2}$ with $\tau_{i,1} = 0$, $\tau_{i,j} = \sum_{l=1}^{j-1} (y_{i,k_i+l} - y_{i,l}) / k_i$, $\tau_{i,M_i+k_i+2} = 1$,

$i = 1, 2$, then,

$$\begin{aligned}
& \|f\|_{\ell^1, \Delta_{2,e}, \Delta_{1,e}} - \|f\|_{L^1} \\
&= \sum_{i_1=1}^{M_1+k_1} \sum_{i_2=1}^{M_2+k_2} |c_{i_1 i_2}| \frac{(y_{1,i_1+k_1} - y_{1,i_1})(y_{2,i_2+k_2} - y_{2,i_2})}{k_1 k_2} - \int_{[0,1]} \int_{[0,1]} |f(z_1, z_2)| dz_1 dz_2 \\
&= \int_{[0,1]} \int_{[0,1]} \left[\sum_{i_1=1}^{M_1+k_1} \sum_{i_2=1}^{M_2+k_2} |c_{i_1 i_2}| I_{[(\tau_{1,i_1}, \tau_{1,i_1+1}) \times (\tau_{2,i_2}, \tau_{2,i_2+1})]}(z_1, z_2) - |f(z_1, z_2)| \right] dz_1 dz_2 \\
&\leq \sup_{(z_1, z_2) \in [0,1] \times [0,1]} \left| \sum_{i_1=1}^{M_1+k_1} \sum_{i_2=1}^{M_2+k_2} |c_{i_1 i_2}| I_{[(\tau_{1,i_1}, \tau_{1,i_1+1}) \times (\tau_{2,i_2}, \tau_{2,i_2+1})]}(z_1, z_2) - |f(z_1, z_2)| \right| \\
&= \max_{i_1 i_2} \sup_{(z_1, z_2) \in [(\tau_{1,i_1}, \tau_{1,i_1+1}) \times (\tau_{2,i_2}, \tau_{2,i_2+1})]} \|c_{i_1 i_2} - f(z_1, z_2)\| \\
&\leq \max_{i_1 i_2} \sup_{(z_1, z_2) \in [(\tau_{1,i_1}, \tau_{1,i_1+1}) \times (\tau_{2,i_2}, \tau_{2,i_2+1})]} |c_{i_1 i_2} - f(z_1, z_2)| \\
&\leq \max_{i_1 i_2} \sup_{(z_1, z_2) \in [(y_{1,i_1}, y_{1,i_1+k_1}) \times (y_{2,i_2}, y_{2,i_2+k_2})]} |c_{i_1 i_2} - f(z_1, z_2)|.
\end{aligned}$$

where $I_{[a \times b]}(z_1, z_2) = 1$ for $(z_1, z_2) \in [a \times b]$ and zero elsewhere. The last inequality follows because it is always true that $[(\tau_{1,i_1}, \tau_{1,i_1+1}) \times (\tau_{2,i_2}, \tau_{2,i_2+1})] \subseteq [(y_{1,i_1}, y_{1,i_1+k_1}) \times (y_{2,i_2}, y_{2,i_2+k_2})]$. From Schumaker (2007) (page 489, Equation (12.25)) and following De Boor et al. (1978) (page 132) we have

$$|c_{i_1 i_2} - f(z_1, z_2)| \leq C_1 C_2 \sup_{\substack{y_{1,i_1} \leq u_1, u_2 \leq y_{1,i_1+k_1} \\ y_{2,i_2} \leq v_1, v_2 \leq y_{2,i_2+k_2}}} |f(u_1, v_1) - f(u_2, v_2)|, \quad (6.26)$$

for $(z_1, z_2) \in [y_{1,i_1}, y_{1,i_1+k_1}] \times [y_{2,i_2}, y_{2,i_2+k_2}]$. Thus,

$$\|f\|_{\ell^1, \Delta_{2,e}, \Delta_{1,e}} - \|f\|_{L^1} \leq C_1 C_2 \omega_{(1,1)} \left[f, (k_1 \delta_1, k_2 \delta_2)^T \right] \leq k_1 k_2 C_1 C_2 \omega_{(1,1)} \left[f, (\delta_1, \delta_2)^T \right].$$

Note that

$$\begin{aligned}
0 &\leq \omega_{(1,1)} \left[f, (\delta_1, \delta_2)^T \right] \leq C_3 \delta_1 \omega_{(0,1)} \left[D^{(1,0)} f, (\delta_1, 0)^T \right] \\
0 &\leq \omega_{(1,1)} \left[f, (\delta_1, \delta_2)^T \right] \leq C_4 \delta_2 \omega_{(1,0)} \left[D^{(0,1)} f, (0, \delta_2)^T \right],
\end{aligned}$$

thus,

$$\begin{aligned}
\omega_{(1,1)} \left[f, (\delta_1, \delta_2)^T \right] &\leq 2^{-1} C_3 \delta_1 \omega_{(0,1)} \left[D^{(1,0)} f, (\delta_1, 0)^T \right] + 2^{-1} C_4 \delta_2 \omega_{(1,0)} \left[D^{(0,1)} f, (0, \delta_2)^T \right] \\
&\leq C_3 \delta_1 \|D^{(1,0)} f\|_\infty + C_4 \delta_2 \|D^{(0,1)} f\|_\infty \\
&\leq \max \left(C_3 \|D^{(1,0)} f\|_\infty, C_4 \|D^{(0,1)} f\|_\infty \right) (\delta_1 + \delta_2). \quad (6.27)
\end{aligned}$$

Therefore, by combining Equation (6.26) and Equation (6.27), Theorem 1 is proved.

6. SMOOTH LASSO ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL

Proof of Theorem 2 Let $\hat{\beta}_{SL} = \hat{\mathbf{b}}_{SL}^T \Psi$ and $\tilde{\beta} = \mathbf{b}^T \Psi$, where $\Psi = \left(\psi_1^s \psi_1^t, \psi_1^s \psi_2^t, \dots, \psi_1^s \psi_{(M_2+k_2)}^t, \psi_2^s \psi_1^t, \dots, \psi_{(M_1+k_1)}^s \psi_{(M_2+k_2)}^t \right)^T$, $\hat{\mathbf{b}}_{app}$ the solution of the optimization problem in Equation (6.12), and $\mathbf{b} = \text{vec}(\mathbf{B})$. By condition C.4, it can be easily proven that given Ψ , both $\hat{\mathbf{b}}_{SL}$ and $\hat{\mathbf{b}}_{app}$ exist and are unique. Note that, by triangular inequality, we have

$$\begin{aligned} \|\hat{\beta}_{SL} - \beta\|_\infty &\leq \|\hat{\beta}_{SL} - \tilde{\beta}\|_\infty + \|\tilde{\beta} - \beta\|_\infty \\ &= \left\| \left(\hat{\mathbf{b}}_{SL} - \mathbf{b} \right)^T \Psi \right\|_\infty + \|\tilde{\beta} - \beta\|_\infty \\ &\leq \|\hat{\mathbf{b}}_{SL} - \mathbf{b}\|_\infty \sup_{(s,t) \in [0,1] \times [0,1]} \sum_{i=1}^{(M_1+k_1)} \sum_{j=1}^{(M_2+k_2)} |\psi_i^s(s) \psi_j^t(t)| + \|\tilde{\beta} - \beta\|_\infty \\ &= \|\hat{\mathbf{b}}_{SL} - \mathbf{b}\|_\infty + \|\tilde{\beta} - \beta\|_\infty \\ &\leq \|\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app}\|_\infty + \|\hat{\mathbf{b}}_{app} - \mathbf{b}\|_\infty + \|\tilde{\beta} - \beta\|_\infty \\ &\leq \|\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app}\|_2 + \|\hat{\mathbf{b}}_{app} - \mathbf{b}\|_2 + \|\tilde{\beta} - \beta\|_\infty, \end{aligned} \quad (6.28)$$

because, $\sum_{i=1}^{(M_1+k_1)} \sum_{j=1}^{(M_2+k_2)} |\psi_i^s(s) \psi_j^t(t)| = 1$ for all $(s, t) \in [0, 1] \times [0, 1]$. To prove Theorem 2, we focus on the three terms on the last line of Equation (6.28), separately.

Let consider $\|\tilde{\beta} - \beta\|_\infty$. Let $\boldsymbol{\delta} = (\delta_1, \delta_2)^T$ and $\mathbf{d} = (k_1, k_2)^T$, then, from Schumaker (2007) (page 492 Equation (12.39)) $\tilde{\beta}$ can be chosen such that

$$\begin{aligned} \|\beta - \tilde{\beta}\|_\infty &\leq C_1 \omega_{Ed}(\beta, \boldsymbol{\delta})_\infty \\ &= C_1 \left(\omega_{(k_1, 0)}(\beta, \boldsymbol{\delta})_\infty + \omega_{(0, k_2)}(\beta, \boldsymbol{\delta})_\infty \right) \\ &\leq C_1 \left(\delta_1^{p'} \omega_{(k_1-p', 0)} \left(D^{(p', 0)} \beta, \boldsymbol{\delta} \right)_\infty + \delta_2^{p'} \omega_{(0, k_2-p')} \left(D^{(0, p')} \beta, \boldsymbol{\delta} \right)_\infty \right) \\ &\leq C_1 \left(\delta_1^{p'} 2^{k_1-p'-1} \omega_{(1, 0)} \left(D_1^{p'} \beta, \boldsymbol{\delta} \right)_\infty + \delta_2^{p'} 2^{k_2-p'-1} \omega_{(0, 1)} \left(D_2^{p'} \beta, \boldsymbol{\delta} \right)_\infty \right) \\ &\leq C_1 \left(\delta_1^{p'} 2^{k_1-p'-1} C_2 \delta_1^\nu + \delta_2^{p'} 2^{k_2-p'-1} C_3 \delta_2^\nu \right) \\ &\leq C_1 \max \left(2^{k_1-p'-1} C_2, 2^{k_2-p'-1} C_3 \right) \left(\delta_1^{p'+\nu} + \delta_2^{p'+\nu} \right) \\ &= C_4 \left(\delta_1^p + \delta_2^p \right), \end{aligned}$$

where properties of the modulus of smoothness have been used and condition C.2. Thus,

$$\|\beta - \tilde{\beta}\|_\infty = O(\delta_1^p) + O(\delta_2^p) = O(M_1^{-p}) + O(M_2^{-p}). \quad (6.29)$$

Let focus on $\|\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app}\|_2$. Let $L(\mathbf{b}_\alpha) \doteq \ell(\mathbf{b}_\alpha) + \lambda_L \|\alpha\|_{L^1}$ and $L_{app}(\mathbf{b}_\alpha) \doteq \ell(\mathbf{b}_\alpha) + \lambda_L \|\alpha\|_{\ell^1, \Delta_{1,e}, \Delta_{2,e}}$ with

$$\ell(\mathbf{b}_\alpha) \doteq \sum_{i=1}^n \int_{\mathcal{T}} Y_i(t)^2 dt - 2 \text{vec} \left(\mathbf{X}^T \mathbf{Y} \right)^T \mathbf{b}_\alpha + \mathbf{b}_\alpha^T \left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) \mathbf{b}_\alpha + \lambda_s \mathbf{b}_\alpha^T \mathbf{L}_{wr} \mathbf{b}_\alpha + \lambda_t \mathbf{b}_\alpha^T \mathbf{L}_{rw} \mathbf{b}_\alpha.$$

Moreover, the Hessian functions with respect to \mathbf{b}_α are

$$\nabla^2 L(\mathbf{b}_\alpha) = \nabla^2 L_{app}(\mathbf{b}_\alpha) = 2 \left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) + 2\lambda_s \mathbf{L}_{wr} + 2\lambda_t \mathbf{L}_{rw}.$$

From Theorem 1, we have that for each \mathbf{b}_α

$$0 \leq L_{app}(\mathbf{b}_\alpha) - L(\mathbf{b}_\alpha) = \|\alpha\|_{\ell^1, \Delta_{1,e}, \Delta_{2,e}} - \|\alpha\|_{L^1} \leq C_5 \left(M_1^{-1} + M_2^{-1} \right),$$

Thus,

$$L_{app}(\hat{\mathbf{b}}_{SL}) - L(\hat{\mathbf{b}}_{SL}) \leq C_5 \left(M_1^{-1} + M_2^{-1} \right).$$

Let $\hat{\mathbf{b}}_{app} = \operatorname{argmin}_{\mathbf{b}_\alpha} L_{app}(\mathbf{b}_\alpha)$, then by Taylor expansion

$$\begin{aligned} L_{app}(\hat{\mathbf{b}}_{SL}) - L(\hat{\mathbf{b}}_{SL}) &\geq L_{app}(\hat{\mathbf{b}}_{app}) + \frac{1}{2} (\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app})^T \nabla^2 L_{app}(\hat{\mathbf{b}}_{app}) (\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app}) - L(\hat{\mathbf{b}}_{SL}) \\ &= L_{app}(\hat{\mathbf{b}}_{app}) - L(\hat{\mathbf{b}}_{SL}) + \frac{1}{2} (\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app})^T \nabla^2 L_{app}(\hat{\mathbf{b}}_{app}) (\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app}) \\ &\geq \frac{1}{2} (\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app})^T \nabla^2 L_{app}(\hat{\mathbf{b}}_{app}) (\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app}) \end{aligned}$$

Thus, by the property of the quadratic form,

$$(\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app})^T \nabla^2 L_{app}(\hat{\mathbf{b}}_{app}) (\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app}) \geq \Lambda_{min} \left[\nabla^2 L_{app}(\hat{\mathbf{b}}_{app}) \right] \|\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app}\|_2^2,$$

where $\Lambda_{min}(\mathbf{A})$ denotes the minimum eigenvalues of the matrix \mathbf{A} , that is

$$\begin{aligned} \Lambda_{min} \left[\nabla^2 L_{app}(\hat{\mathbf{b}}_{app}) \right] &= 2\Lambda_{min}(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X}) + 2\lambda_s \Lambda_{min}(\mathbf{L}_{wr}) + 2\lambda_t \Lambda_{min}(\mathbf{L}_{rw}) \\ &= 2n \left[\frac{\Lambda_{min}(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X})}{n} + \frac{\lambda_s}{n} \Lambda_{min}(\mathbf{L}_{wr}) + \frac{\lambda_t}{n} \Lambda_{min}(\mathbf{L}_{rw}) \right]. \end{aligned}$$

Therefore,

$$\Lambda_{min} \left[\nabla^2 L_{app}(\hat{\mathbf{b}}_{app}) \right] \|\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app}\|_2^2 \leq C_6 \left(M_1^{-1} + M_2^{-1} \right),$$

and, by combining this with conditions C.4 and C.5, we have

$$\|\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app}\|_2 \leq \frac{C_7}{\sqrt{bn}} \left(M_1^{-1} + M_2^{-1} \right)^{1/2} \leq \frac{C_7}{\sqrt{bn}} \left(M_1^{-1/2} + M_2^{-1/2} \right),$$

and

$$\|\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app}\|_2 = O(M_1^{-1/2} n^{-1/2}) + O(M_2^{-1/2} n^{-1/2}). \quad (6.30)$$

Turn the attention to $\|\hat{\mathbf{b}}_{app} - \mathbf{b}\|_2$. The following arguments are based on Zou and Zhang (2009). Let define $\hat{\mathbf{b}}_{app,0}$ as the solution of the optimization problem in Equation (6.12) when $\lambda_L = 0$, and

$$Q(\mathbf{b}_\alpha) \doteq \sum_{i=1}^n \int_{\mathcal{T}} Y_i(t)^2 dt - 2 \operatorname{vec}(\mathbf{X}^T \mathbf{Y})^T \mathbf{b}_\alpha + \mathbf{b}_\alpha^T (\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X}) \mathbf{b}_\alpha.$$

Then, by the definition of $\hat{\mathbf{b}}_{app}$, we have

$$\begin{aligned} &Q(\hat{\mathbf{b}}_{app,0}) + \lambda_s \hat{\mathbf{b}}_{app,0}^T \mathbf{L}_{wr} \hat{\mathbf{b}}_{app,0} + \lambda_t \hat{\mathbf{b}}_{app,0}^T \mathbf{L}_{rw} \hat{\mathbf{b}}_{app,0} + \lambda_L \|\mathbf{W}_{st} \hat{\mathbf{b}}_{app,0}\|_1 \\ &\geq Q(\hat{\mathbf{b}}_{app}) + \lambda_s \hat{\mathbf{b}}_{app}^T \mathbf{L}_{wr} \hat{\mathbf{b}}_{app} + \lambda_t \hat{\mathbf{b}}_{app}^T \mathbf{L}_{rw} \hat{\mathbf{b}}_{app} + \lambda_L \|\mathbf{W}_{st} \hat{\mathbf{b}}_{app}\|_1. \end{aligned}$$

By rearranging the above inequality, we get

$$\begin{aligned} \lambda_L \left[\|\mathbf{W}_{st} \hat{\mathbf{b}}_{app,0}\|_1 - \|\mathbf{W}_{st} \hat{\mathbf{b}}_{app}\|_1 \right] &\geq Q(\hat{\mathbf{b}}_{app}) + \lambda_s \hat{\mathbf{b}}_{app}^T \mathbf{L}_{wr} \hat{\mathbf{b}}_{app} + \lambda_t \hat{\mathbf{b}}_{app}^T \mathbf{L}_{rw} \hat{\mathbf{b}}_{app} \\ &\quad - \left\{ Q(\hat{\mathbf{b}}_{app,0}) + \lambda_s \hat{\mathbf{b}}_{app,0}^T \mathbf{L}_{wr} \hat{\mathbf{b}}_{app,0} + \lambda_t \hat{\mathbf{b}}_{app,0}^T \mathbf{L}_{rw} \hat{\mathbf{b}}_{app,0} \right\}. \end{aligned} \quad (6.31)$$

Note that

$$\begin{aligned} \|\mathbf{W}_{st} \hat{\mathbf{b}}_{app,0}\|_1 - \|\mathbf{W}_{st} \hat{\mathbf{b}}_{app}\|_1 &\leq \|\mathbf{W}_{st} \hat{\mathbf{b}}_{app,0} - \mathbf{W}_{st} \hat{\mathbf{b}}_{app}\|_1 \\ &\leq \|\mathbf{W}_{st}\|_1 \|\hat{\mathbf{b}}_{app,0} - \hat{\mathbf{b}}_{app}\|_1 \\ &\leq (M_1 + k_1)^{1/2} (M_2 + k_2)^{1/2} \|\hat{\mathbf{b}}_{app,0} - \hat{\mathbf{b}}_{app}\|_2, \end{aligned} \quad (6.32)$$

where $\|\mathbf{W}_{st}\|_1$ is the induced matrix norm corresponding to $\|\cdot\|_1$, defined as the maximum of the absolute values of the column sums of \mathbf{W}_{st} , which is less or equal to one. Note that

$$\text{vec}(\mathbf{X}^T \mathbf{Y}) = (\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X}) \hat{\mathbf{b}}_{app,0} + \lambda_s \mathbf{L}_{wr} \hat{\mathbf{b}}_{app,0} + \lambda_t \mathbf{L}_{rw} \hat{\mathbf{b}}_{app,0},$$

then, we have

$$\begin{aligned} &Q(\hat{\mathbf{b}}_{app,0}) + \lambda_s \hat{\mathbf{b}}_{app,0}^T \mathbf{L}_{wr} \hat{\mathbf{b}}_{app,0} + \lambda_t \hat{\mathbf{b}}_{app,0}^T \mathbf{L}_{rw} \hat{\mathbf{b}}_{app,0} - \left\{ Q(\hat{\mathbf{b}}_{app}) + \lambda_s \hat{\mathbf{b}}_{app}^T \mathbf{L}_{wr} \hat{\mathbf{b}}_{app} + \lambda_t \hat{\mathbf{b}}_{app}^T \mathbf{L}_{rw} \hat{\mathbf{b}}_{app} \right\} \\ &= (\hat{\mathbf{b}}_{app} - \hat{\mathbf{b}}_{app,0})^T \left[(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X}) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right] (\hat{\mathbf{b}}_{app} - \hat{\mathbf{b}}_{app,0}). \end{aligned}$$

Thus, by using Equation (6.31) and Equation (6.32), we have

$$\begin{aligned} &\Lambda_{min} \left[(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X}) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right] \|\hat{\mathbf{b}}_{app} - \hat{\mathbf{b}}_{app,0}\|_2^2 \\ &\leq (\hat{\mathbf{b}}_{app} - \hat{\mathbf{b}}_{app,0})^T \left[(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X}) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right] (\hat{\mathbf{b}}_{app} - \hat{\mathbf{b}}_{app,0}) \\ &\leq \lambda_L (M_1 + k_1)^{1/2} (M_2 + k_2)^{1/2} \|\hat{\mathbf{b}}_{app,0} - \hat{\mathbf{b}}_{app}\|_2. \end{aligned}$$

Then,

$$\|\hat{\mathbf{b}}_{app,0} - \hat{\mathbf{b}}_{app}\|_2 \leq \frac{\lambda_L (M_1 + k_1)^{1/2} (M_2 + k_2)^{1/2}}{\Lambda_{min} \left[(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X}) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right]}. \quad (6.33)$$

Moreover, let $\boldsymbol{\epsilon} = \text{vec} \left(\int_{\mathcal{T}} \boldsymbol{\varepsilon}(t) \boldsymbol{\psi}^t(t)^T dt \right) = \int_{\mathcal{T}} (\boldsymbol{\psi}^t(t) \otimes \mathbf{I}) \boldsymbol{\varepsilon}(t) dt$, with $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_n)^T$, then

$$\begin{aligned} \hat{\mathbf{b}}_{app,0} - \mathbf{b} &= -(\lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw}) \left[(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X}) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right]^{-1} \mathbf{b} \\ &\quad + \left[(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X}) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right]^{-1} (\mathbf{I} \otimes \mathbf{X}^T) \boldsymbol{\epsilon}, \end{aligned}$$

which implies

$$\begin{aligned}
\mathbb{E} \left(\|\hat{\mathbf{b}}_{app,0} - \mathbf{b}\|_2^2 \right) &\leq 2\lambda_s^2 \|\mathbf{L}_{wr}\| \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right]^{-1} \mathbf{b} \|_2^2 \\
&\quad + 2\lambda_t^2 \|\mathbf{L}_{rw}\| \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right]^{-1} \mathbf{b} \|_2^2 \\
&\quad + 2 \mathbb{E} \left(\left\| \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right]^{-1} \left(\mathbf{I} \otimes \mathbf{X}^T \right) \boldsymbol{\epsilon} \right\|_2^2 \right) \\
&\leq 2\lambda_s^2 \Lambda_{max}^2(\mathbf{L}_{wr}) \left\| \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right]^{-1} \mathbf{b} \right\|_2^2 \\
&\quad + 2\lambda_t^2 \Lambda_{max}^2(\mathbf{L}_{rw}) \left\| \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right]^{-1} \mathbf{b} \right\|_2^2 \\
&\quad + 2 \mathbb{E} \left(\left\| \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right]^{-1} \left(\mathbf{I} \otimes \mathbf{X}^T \right) \boldsymbol{\epsilon} \right\|_2^2 \right) \\
&\leq 2\lambda_s^2 \Lambda_{max}^2(\mathbf{L}_{wr}) \Lambda_{min}^{-2} \left(\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right) \|\mathbf{b}\|_2^2 \\
&\quad + 2\lambda_t^2 \Lambda_{max}^2(\mathbf{L}_{rw}) \Lambda_{min}^{-2} \left(\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right) \|\mathbf{b}\|_2^2 \\
&\quad + 2\Lambda_{min}^{-2} \left(\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right) \mathbb{E} \left(\boldsymbol{\epsilon}^T \left(\mathbf{I} \otimes \mathbf{X} \mathbf{X}^T \right) \boldsymbol{\epsilon} \right) \\
&\leq 2\Lambda_{min}^{-2} \left(\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right) \left(\lambda_s^2 \Lambda_{max}^2(\mathbf{L}_{wr}) \|\mathbf{b}\|_2^2 \right. \\
&\quad \left. + \lambda_t^2 \Lambda_{max}^2(\mathbf{L}_{rw}) \|\mathbf{b}\|_2^2 + \text{Tr} \left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) \|K\|_\infty \right) \\
&\leq 2\Lambda_{min}^{-2} \left(\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right) \left(\lambda_s^2 \Lambda_{max}^2(\mathbf{L}_{wr}) \|\mathbf{b}\|_2^2 \right. \\
&\quad \left. + \lambda_t^2 \Lambda_{max}^2(\mathbf{L}_{rw}) \|\mathbf{b}\|_2^2 + (M_1 + k_1)(M_2 + k_2) \Lambda_{max} \left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) \|K\|_\infty \right),
\end{aligned} \tag{6.34}$$

because

$$\begin{aligned}
 & \mathbb{E} \left[\boldsymbol{\epsilon}^T \left(\mathbf{I} \otimes \mathbf{X} \mathbf{X}^T \right) \boldsymbol{\epsilon} \right] = \mathbb{E} \left[\text{Tr} \left(\boldsymbol{\epsilon} \boldsymbol{\epsilon}^T \left(\mathbf{I} \otimes \mathbf{X} \mathbf{X}^T \right) \right) \right] = \mathbb{E} \left[\text{Tr} \left(\boldsymbol{\epsilon} \boldsymbol{\epsilon}^T \left(\mathbf{I} \otimes \mathbf{X} \mathbf{X}^T \right) \right) \right] \\
 &= \mathbb{E} \left[\text{Tr} \left(\int_{[0,1]} \int_{[0,1]} \left(\boldsymbol{\psi}^t(t_1) \otimes \mathbf{I} \right) \boldsymbol{\varepsilon}(t_1) \boldsymbol{\varepsilon}(t_2)^T \left(\boldsymbol{\psi}^t(t_2)^T \otimes \mathbf{I} \right) dt_1 dt_2 \left(\mathbf{I} \otimes \mathbf{X} \mathbf{X}^T \right) \right) \right] \\
 &= \text{Tr} \left[\int_{[0,1]} \int_{[0,1]} K(t_1, t_2) \left(\boldsymbol{\psi}^t(t_1) \otimes \mathbf{I} \right) \left(\boldsymbol{\psi}^t(t_2)^T \otimes \mathbf{I} \right) dt_1 dt_2 \left(\mathbf{I} \otimes \mathbf{X} \mathbf{X}^T \right) \right] \\
 &= \int_{[0,1]} \int_{[0,1]} K(t_1, t_2) \text{Tr} \left(\boldsymbol{\psi}^t(t_1) \boldsymbol{\psi}^t(t_2)^T \otimes \mathbf{X} \mathbf{X}^T \right) dt_1 dt_2 \\
 &\leq \|K\|_\infty \int_{[0,1]} \int_{[0,1]} \text{Tr} \left(\boldsymbol{\psi}^t(t_1) \boldsymbol{\psi}^t(t_2)^T \otimes \mathbf{X} \mathbf{X}^T \right) dt_1 dt_2 \\
 &\leq \|K\|_\infty \text{Tr} \left(\int_{[0,1]} \boldsymbol{\psi}^t(t_1) dt_1 \int_{[0,1]} \boldsymbol{\psi}^t(t_2)^T dt_2 \right) \text{Tr} \left(\mathbf{X} \mathbf{X}^T \right) \\
 &\leq \|K\|_\infty \text{Tr} \left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right), \tag{6.35}
 \end{aligned}$$

given that for the Holder inequality $\left(\int_{[0,1]} \psi_i^t(t_1) dt_1 \right)^2 \leq \int_{[0,1]} (\psi_i^t(t_1))^2 dt_1$ for $i = 1, \dots, M_2 + k_2$.

By combining the two inequalities in (6.33) and (6.34), we have

$$\begin{aligned}
 & \mathbb{E} \left(\|\hat{\mathbf{b}}_{app} - \mathbf{b}\|_2^2 \right) \\
 &\leq 2 \mathbb{E} \left(\|\hat{\mathbf{b}}_{app,0} - \mathbf{b}\|_2^2 \right) + 2 \mathbb{E} \left(\|\hat{\mathbf{b}}_{app} - \hat{\mathbf{b}}_{app,0}\|_2^2 \right) \\
 &\leq \frac{4\lambda_s^2 \Lambda_{max}^2(\mathbf{L}_{wr}) \|\mathbf{b}\|_2^2 + 4\lambda_t^2 \Lambda_{max}^2(\mathbf{L}_{rw}) \|\mathbf{b}\|_2^2}{\Lambda_{min}^2 \left[(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X}) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right]} \\
 &+ \frac{4(M_1 + k_1)(M_2 + k_2) \Lambda_{max}(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X}) \|K\|_\infty + 2\lambda_L^2 (M_1 + k_1)(M_2 + k_2)}{\Lambda_{min}^2 \left[(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X}) + \lambda_s \mathbf{L}_{wr} + \lambda_t \mathbf{L}_{rw} \right]} \\
 &\leq \frac{4\lambda_s^2 \Lambda_{max}^2(\mathbf{L}_{wr}) \|\mathbf{b}\|_2^2 + 4\lambda_t^2 \Lambda_{max}^2(\mathbf{L}_{rw}) \|\mathbf{b}\|_2^2}{n^2 \Lambda_{min}^2 \left[(\mathbf{W}_t \otimes n^{-1} \mathbf{X}^T \mathbf{X}) + n^{-1} \lambda_s \mathbf{L}_{wr} + n^{-1} \lambda_t \mathbf{L}_{rw} \right]} \\
 &+ \frac{4n(M_1 + k_1)(M_2 + k_2) B \|K\|_\infty + 2\lambda_L^2 (M_1 + k_1)(M_2 + k_2)}{n^2 \Lambda_{min}^2 \left[(\mathbf{W}_t \otimes n^{-1} \mathbf{X}^T \mathbf{X}) + n^{-1} \lambda_s \mathbf{L}_{wr} + n^{-1} \lambda_t \mathbf{L}_{rw} \right]}.
 \end{aligned}$$

Under C.5 and C.6, we have

$$\|\hat{\mathbf{b}}_{app} - \mathbf{b}\|_2 = O_p \left(M_1^{1/2} M_2^{1/2} n^{-1/2} \right). \tag{6.36}$$

Thus, from Equation (6.28), (6.30), and (6.36), we have that by choosing $\tilde{\beta}$ such that

Equation (6.29) is satisfied, there exist a unique $\hat{\beta}_{SL}$ with

$$\begin{aligned} & \|\hat{\beta}_{SL} - \beta\|_\infty \\ & \leq \|\hat{\mathbf{b}}_{SL} - \hat{\mathbf{b}}_{app}\|_2 + \|\hat{\mathbf{b}}_{app} - \mathbf{b}\|_2 + \|\tilde{\beta} - \beta\|_\infty \\ & = O(M_1^{-1/2}n^{-1/2}) + O(M_2^{-1/2}n^{-1/2}) + O_p\left(M_1^{1/2}M_2^{1/2}n^{-1/2}\right) + O(M_1^{-p}) + O(M_2^{-p}), \end{aligned}$$

and by C.3, we have

$$\|\hat{\beta}_{SL} - \beta\|_\infty = O_p\left(M_1^{1/2}M_2^{1/2}n^{-1/2}\right).$$

Proof of Theorem 3

Note that, for $s \in \mathcal{S}$ and $t \in \mathcal{T}$

$$\begin{aligned} & \text{sign}(\beta(s, t)) \\ & = \text{sign}\left(\beta(s, t) - \tilde{\beta}(s, t) + \tilde{\beta}(s, t)\right) \\ & = \text{sign}\left(\text{sign}\left(\beta(s, t) - \tilde{\beta}(s, t)\right)|\beta(s, t) - \tilde{\beta}(s, t)| + \tilde{\beta}(s, t)\right) \\ & = \text{sign}\left(\text{sign}\left(\beta(s, t) - \tilde{\beta}(s, t)\right)|\beta(s, t) - \tilde{\beta}(s, t)| + \text{sign}\left(\tilde{\beta}(s, t) - b_{i_1 i_2}\right)|\tilde{\beta}(s, t) - b_{i_1 i_2}| + b_{i_1 i_2}\right) \\ & \leq \text{sign}\left(|\beta(s, t) - \tilde{\beta}(s, t)| + |\tilde{\beta}(s, t) - b_{i_1 i_2}| + b_{i_1 i_2}\right) \end{aligned}$$

where $i_1 \in (1, \dots, M_1 + k_1)$ and $i_2 \in (1, \dots, M_2 + k_2)$ are chosen such that $(s, t) \in \left[(y_{1, i_1}, y_{1, i_1 + k_1}) \times (y_{2, i_2}, y_{2, i_2 + k_2})\right]$, where $b_{i_1 i_2}$ are the elements of $\mathbf{b} = (b_{11}, \dots, b_{(M_1 + k_1)1}, b_{12}, \dots, b_{(M_1 + k_1)(M_2 + k_2 - 1)}, b_{1(M_2 + k_2)}, \dots, b_{(M_1 + k_1)(M_2 + k_2)})^T$. By using the results in both in Theorem 1 (Equation (6.26) and Equation (6.27)) and Theorem 2 (Equation (6.29)), we have that

$$\text{sign}(\beta(s, t)) = \text{sign}\left(b_{i_1 i_2} + O\left(\frac{1}{M_1}\right) + O\left(\frac{1}{M_2}\right) + O(M_1^{-p}) + O(M_2^{-p})\right),$$

and, thus, by C.3

$$\text{sign}(\beta(s, t)) = \text{sign}(b_{i_1 i_2}),$$

for i_1, i_2 such that $(s, t) = (y_{1, i_1}, y_{2, i_2})$. By the same arguments, by using results in Theorem 2 (Equation (6.30)), under conditions C.3, C.4 and C.5, we have

$$\text{sign}(\hat{\beta}(s, t)) = \text{sign}(\hat{b}_{app, i_1 i_2}),$$

for i_1, i_2 such that $(s, t) = (y_{1, i_1}, y_{2, i_2})$, where $\hat{b}_{app, i_1 i_2}$ are the entries of $\hat{\mathbf{b}}_{app} = (\hat{b}_{app, 11}, \dots, \hat{b}_{app, (M_1 + k_1)1}, \hat{b}_{app, 12}, \dots, \hat{b}_{app, (M_1 + k_1)(M_2 + k_2 - 1)}, \hat{b}_{app, 1(M_2 + k_2)}, \dots, \hat{b}_{app, (M_1 + k_1)(M_2 + k_2)})^T$. Therefore, to prove Theorem 2, we must show that

$$\text{sign}(b_{i_1 i_2}) = \text{sign}(\hat{b}_{app, i_1 i_2}),$$

for each i_1, i_2 . Or equivalently

$$\text{sign}(\mathbf{b}) = \text{sign}(\hat{\mathbf{b}}_{app}), \quad (6.37)$$

where the sign function is applied element-wise.

Before proving Equation (6.37), let define with $\mathbf{b}_{(1)}$ the vector whose entries are the q elements of \mathbf{b} that are non-zero and with $\mathbf{b}_{(2)}$ the vector whose entries are the $(M_1 + k_1)(M_2 + k_2) - q$ elements of \mathbf{b} that are equal to zero. In what follows, we assume, without loss of generality, that $\mathbf{b} = \begin{bmatrix} \mathbf{b}_{(1)}^T & \mathbf{b}_{(2)}^T \end{bmatrix}^T$. The subscripts (1) and (2) applied to a generic vector correspond to the first q and the last $(M_1 + k_1)(M_2 + k_2) - q$ elements, whereas applied to a generic matrix correspond to the matrix composed by the first q and the last $(M_1 + k_1)(M_2 + k_2) - q$ columns. Moreover, a generic matrix $\mathbf{A}_l \in \mathbb{R}^{(M_1 + k_1)(M_2 + k_2) \times (M_1 + k_1)(M_2 + k_2)}$ can be expressed in a block-wise form as

$$\mathbf{A}_l = \begin{bmatrix} \mathbf{A}_{l,11} \in \mathbb{R}^{q \times q} & \mathbf{A}_{l,12} \in \mathbb{R}^{q \times (M_1 + k_1)(M_2 + k_2) - q} \\ \mathbf{A}_{l,21} \in \mathbb{R}^{(M_1 + k_1)(M_2 + k_2) - q \times q} & \mathbf{A}_{l,22} \in \mathbb{R}^{(M_1 + k_1)(M_2 + k_2) - q \times (M_1 + k_1)(M_2 + k_2) - q} \end{bmatrix}.$$

The following arguments follow closely those of Zou and Zhang (2009). To prove (6.37), we need the following lemma.

Lemma 1. *For any positive λ_s, λ_t and λ_L , there is a $\hat{\mathbf{b}}_{app}$ in Equation (6.12) that satisfies Equation (6.37), if and only if, element-wise,*

$$\begin{aligned} & |\mathbf{W}_{st,21}^{-1} \left\{ 2 \left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{21} + 2\lambda_s \mathbf{L}_{wr,21} + 2\lambda_t \mathbf{L}_{rw,21} \right\} \\ & \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \\ & \left[-\lambda_s \mathbf{L}_{wr,11} \mathbf{b}_{(1)} - \lambda_t \mathbf{L}_{rw,11} \mathbf{b}_{(1)} + (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \boldsymbol{\epsilon} - 2^{-1} \lambda_L \mathbf{W}_{st,11} \text{sign}(\mathbf{b}_{(1)}) \right] \\ & - 2 (\mathbf{I} \otimes \mathbf{X})_{(2)}^T \boldsymbol{\epsilon} + 2\lambda_s \mathbf{L}_{wr,21} \mathbf{b}_{(1)} + 2\lambda_t \mathbf{L}_{rw,21} \mathbf{b}_{(1)} \Big\} \leq \lambda_L \end{aligned}, \quad (6.38)$$

$$\begin{aligned} & \left| \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \right. \\ & \left. \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} \mathbf{b}_{(1)} + (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \boldsymbol{\epsilon} - 2^{-1} \lambda_L \mathbf{W}_{st,11} \text{sign}(\mathbf{b}_{(1)}) \right] \right| > 0. \end{aligned} \quad (6.39)$$

Proof By standard Karush–Kuhn–Tucker (KKT) conditions for optimality in convex program $\hat{\mathbf{b}}_{app}$ is optimal if and only if

$$-2 \text{vec}(\mathbf{X}^T \mathbf{Y}) + 2 \left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) \hat{\mathbf{b}}_{app} + 2\lambda_s \mathbf{L}_{wr} \hat{\mathbf{b}}_{app} + 2\lambda_t \mathbf{L}_{rw} \hat{\mathbf{b}}_{app} + \lambda_L \mathbf{W}_{st} \hat{\mathbf{z}} = 0, \quad (6.40)$$

with

$$\hat{\mathbf{z}} = \begin{cases} \text{sign}(\hat{b}_{app,i}) & \hat{b}_{app,i} \neq 0, \\ r \in [-1, 1] & \text{otherwise.} \end{cases}$$

On the other hand, $\hat{\mathbf{b}}_{app}$ is sign consistent if and only if $\hat{\mathbf{b}}_{app(2)} = 0$, $\hat{\mathbf{b}}_{app(1)} \neq 0$ and $\hat{\mathbf{z}}_{(1)} = \text{sign}(\mathbf{b}_{(1)})$, $|\hat{\mathbf{z}}_{(2)}| \leq 1$, and, thus, if and only if

$$\begin{aligned} & \left[2 \left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{21} + 2\lambda_s \mathbf{L}_{wr,21} + 2\lambda_t \mathbf{L}_{rw,21} \right] \left(\hat{\mathbf{b}}_{app(1)} - \mathbf{b}_{(1)} \right) - 2 \left(\mathbf{I} \otimes \mathbf{X} \right)_{(2)}^T \boldsymbol{\epsilon} \\ & + 2\lambda_s \mathbf{L}_{wr,21} \mathbf{b}_{(1)} + 2\lambda_t \mathbf{L}_{rw,21} \mathbf{b}_{(1)} = -\lambda_L \mathbf{W}_{st,21} \hat{\mathbf{z}}_{(2)}, \\ & 2 \left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} \left(\hat{\mathbf{b}}_{app(1)} - \mathbf{b}_{(1)} \right) - 2 \left(\mathbf{I} \otimes \mathbf{X} \right)_{(1)}^T \boldsymbol{\epsilon} \\ & + 2\lambda_s \mathbf{L}_{wr,11} \hat{\mathbf{b}}_{app(1)} + 2\lambda_t \mathbf{L}_{rw,11} \hat{\mathbf{b}}_{app(1)} = -\lambda \mathbf{W}_{st,11} \text{sign}(\mathbf{b}_{\alpha_L(1)}). \end{aligned}$$

By solving for $\hat{\mathbf{b}}_{app(1)}$ and $\hat{\mathbf{z}}_{(2)}$, we get

$$\begin{aligned} -\lambda_L \hat{\mathbf{z}}_{(2)} &= \mathbf{W}_{st,21}^{-1} \left[2 \left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{21} + 2\lambda_s \mathbf{L}_{wr,21} + 2\lambda_t \mathbf{L}_{rw,21} \right] \\ & \quad \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \\ & \quad \left[-\lambda_s \mathbf{L}_{wr,11} \mathbf{b}_{(1)} - \lambda_t \mathbf{L}_{rw,11} \mathbf{b}_{(1)} + \left(\mathbf{I} \otimes \mathbf{X} \right)_{(1)}^T \boldsymbol{\epsilon} - 2^{-1} \lambda_L \mathbf{W}_{st,11} \text{sign}(\mathbf{b}_{(1)}) \right] \\ & \quad - 2 \left(\mathbf{I} \otimes \mathbf{X} \right)_{(2)}^T \boldsymbol{\epsilon} + 2\lambda_s \mathbf{L}_{wr,21} \mathbf{b}_{(1)} + 2\lambda_t \mathbf{L}_{rw,21} \mathbf{b}_{(1)}, \\ \hat{\mathbf{b}}_{app(1)} &= \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \\ & \quad \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} \mathbf{b}_{(1)} + \left(\mathbf{I} \otimes \mathbf{X} \right)_{(1)}^T \boldsymbol{\epsilon} - 2^{-1} \lambda_L \mathbf{W}_{st,11} \text{sign}(\mathbf{b}_{(1)}) \right]. \end{aligned}$$

Lemma 1 follows by applying the conditions $\hat{\mathbf{b}}_{app(1)} \neq 0$ and $|\hat{\mathbf{z}}_{(2)}| \leq 1$. \square

To prove Theorem 3, we have to verify (6.38) and (6.39) in Lemma 1. Denote by \mathbf{e}_i the vector with 1 in the i -th position and zeroes elsewhere. For each index $i \in S = \{1, \dots, q\}$ and $j \in S^c = \{q+1, \dots, (M_1 + k_1)(M_2 + k_2)\}$, define the following random variables:

$$\begin{aligned} U_i &\doteq \mathbf{e}_i^T \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \left[\left(\mathbf{I} \otimes \mathbf{X} \right)_{(1)}^T \boldsymbol{\epsilon} - 2^{-1} \lambda_L \mathbf{W}_{st,11} \text{sign}(\mathbf{b}_{(1)}) \right], \\ V_j &\doteq \mathbf{W}_{st,jj}^{-1} \left\{ \left[2 \left(\mathbf{I} \otimes \mathbf{X} \right)_j^T \left(\mathbf{W}_t \otimes \mathbf{I} \right) \left(\mathbf{I} \otimes \mathbf{X} \right)_{(1)} + 2\lambda_s \mathbf{L}_{wr,j} + 2\lambda_t \mathbf{L}_{rw,j} \right] \right. \\ & \quad \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \\ & \quad \left[2^{-1} \lambda_L \mathbf{W}_{st,11} \text{sign}(\mathbf{b}_{(1)}) + \lambda_s \mathbf{L}_{wr,11} \mathbf{b}_{(1)} + \lambda_t \mathbf{L}_{rw,11} \mathbf{b}_{(1)} - \left(\mathbf{I} \otimes \mathbf{X} \right)_{(1)}^T \boldsymbol{\epsilon} \right] \\ & \quad \left. - 2\lambda_s \mathbf{L}_{wr,j} \mathbf{b}_{(1)} - 2\lambda_t \mathbf{L}_{rw,j} \mathbf{b}_{(1)} + 2 \left(\mathbf{I} \otimes \mathbf{X} \right)_j^T \boldsymbol{\epsilon} \right\}, \end{aligned}$$

where $\left(\mathbf{I} \otimes \mathbf{X} \right)_j$ is the j -th column of $\left(\mathbf{I} \otimes \mathbf{X} \right)$, $\mathbf{L}_{wr,j}$ and $\mathbf{L}_{rw,j}$ are the j -th rows of $\mathbf{L}_{wr,21}$ and $\mathbf{L}_{rw,21}$ respectively, and $\mathbf{W}_{st,jj}$ is the j -th diagonal element of $\mathbf{W}_{st,12}$. Then

6. SMOOTH LASSO ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL

conditions (6.38) and (6.39) hold if and only if the events $M(U) \doteq \{\max_{i \in S} |U_i| < \rho\}$ and $M(V) \doteq \{\max_{j \in S^c} |V_j| \leq \lambda_L\}$ hold simultaneously, with $\rho \doteq \min \left| \left[(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X})_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \left[(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X})_{11} \mathbf{b}_{(1)} \right] \right|$. Now, we analyse the events $M(U)$ and $M(V)$ separately. To analyse $M(V)$, we define

$$\begin{aligned} \mu_j &\doteq W_{st,jj}^{-1} \left\{ \left[(\mathbf{I} \otimes \mathbf{X})_j^T (\mathbf{W}_t \otimes \mathbf{I}) (\mathbf{I} \otimes \mathbf{X})_{(1)} + \lambda_s \mathbf{L}_{wr,j} + \lambda_t \mathbf{L}_{rw,j} \right] \right. \\ &\quad \left[(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X})_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \\ &\quad \left(\lambda_L \mathbf{W}_{st,11} \text{sign}(\mathbf{b}_{(1)}) + 2\lambda_s \mathbf{L}_{wr,11} \mathbf{b}_{(1)} + 2\lambda_t \mathbf{L}_{rw,11} \mathbf{b}_{(1)} \right) \\ &\quad \left. - 2\lambda_s \mathbf{L}_{wr,j} \mathbf{b}_{(1)} - 2\lambda_t \mathbf{L}_{rw,j} \mathbf{b}_{(1)} \right\}, \\ \tilde{V}_j &\doteq W_{st,jj}^{-1} \left\{ \left(2 (\mathbf{I} \otimes \mathbf{X})_j^T \left[\mathbf{I} - (\mathbf{W}_t \otimes \mathbf{I}) (\mathbf{I} \otimes \mathbf{X})_{(1)} \right] \right. \right. \\ &\quad \left. \left[(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X})_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right] \epsilon \\ &\quad - (2\lambda_s \mathbf{L}_{wr,j} + 2\lambda_t \mathbf{L}_{rw,j}) \\ &\quad \left. \left[(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X})_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right\}, \end{aligned}$$

then $V_j = \mu_j + \tilde{V}_j$. From the SL-IC (C.7), we have that, element-wise

$$\begin{aligned} &\left| \mathbf{W}_{st,21}^{-1} \left\{ \left[(\mathbf{W}_t \otimes n^{-1} \mathbf{X}^T \mathbf{X})_{21} + n^{-1} \lambda_s \mathbf{L}_{wr,21} + n^{-1} \lambda_t \mathbf{L}_{rw,21} \right] \right. \right. \\ &\quad \left. \left[(\mathbf{W}_t \otimes n^{-1} \mathbf{X}^T \mathbf{X})_{11} + n^{-1} \lambda_s \mathbf{L}_{wr,11} + n^{-1} \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \right. \\ &\quad \left. \left[\mathbf{W}_{st,11} \text{sign}(\mathbf{b}_{(1)}) + 2\lambda_L^{-1} \lambda_s \mathbf{L}_{wr,11} \mathbf{b}_{(1)} + 2\lambda_L^{-1} \lambda_t \mathbf{L}_{rw,11} \mathbf{b}_{(1)} \right] \right. \\ &\quad \left. \left. - 2\lambda_L^{-1} \lambda_s \mathbf{L}_{wr,21} \mathbf{b}_{(1)} - 2\lambda_L^{-1} \lambda_t \mathbf{L}_{rw,21} \mathbf{b}_{(1)} \right\} \right| \leq 1 - \eta, \end{aligned}$$

then, $|E(V_j)| = |\mu_j| \leq (1 - \eta) \lambda_L$ for $j = q + 1, \dots, (M_1 + k_1)(M_2 + k_2)$. $M(V)$ holds if and only if $\frac{\max_{j \in S^c} V_j}{\lambda_L} \leq 1$ and $\frac{\min_{j \in S^c} V_j}{\lambda_L} \geq -1$. Since

$$\begin{aligned} \frac{\max_{j \in S^c} V_j}{\lambda_L} &= \frac{\max_{j \in S^c} \mu_j + \tilde{V}_j}{\lambda_L} \leq (1 - \eta) + \frac{\max_{j \in S^c} \tilde{V}_j}{\lambda_L}, \\ \frac{\min_{j \in S^c} V_j}{\lambda_L} &= \frac{\min_{j \in S^c} \mu_j + \tilde{V}_j}{\lambda_L} \geq -(1 - \eta) + \frac{\min_{j \in S^c} \tilde{V}_j}{\lambda_L}, \end{aligned}$$

$M(V)$ holds if and only if

$$\Pr \left\{ \frac{\max_{j \in S^c} \tilde{V}_j}{\lambda_L} > \eta, \text{ or } \frac{\min_{j \in S^c} \tilde{V}_j}{\lambda_L} < -\eta \right\} \rightarrow 0.$$

Under assumption C.8, $\tilde{V} = (\tilde{V}_{q+1}, \dots, \tilde{V}_{(M_1+k_1)(M_2+k_2)})$ is a mean zero Gaussian random vector and by applying Markov's inequality and Gaussian comparison results (Ledoux and Talagrand, 2013), we have

$$\Pr \left\{ \frac{\max_{j \in S^c} \tilde{V}_j}{\lambda_L} > \eta \right\} \leq \frac{\mathbb{E} \left(\max_{j \in S^c} |\tilde{V}_j| \right)}{\lambda_L \eta} \leq \frac{3\sqrt{\log [(M_1 + k_1)(M_2 + k_2) - q]}}{\lambda_L \eta} \max_{j \in S^c} \sqrt{E \left(\tilde{V}_j^2 \right)}. \quad (6.41)$$

Moreover,

$$\begin{aligned} & E \left(\tilde{V}_j^2 \right) / 4 \\ & \leq 3W_{st,jj}^{-2} \text{Tr} \left\{ E \left(\epsilon \epsilon^T \right) \right. \\ & \quad \left[\mathbf{I} - (\mathbf{W}_t \otimes \mathbf{I}) (\mathbf{I} \otimes \mathbf{X})_{(1)} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right]^T (\mathbf{I} \otimes \mathbf{X})_j \\ & \quad (\mathbf{I} \otimes \mathbf{X})_j^T \left[\mathbf{I} - (\mathbf{W}_t \otimes \mathbf{I}) (\mathbf{I} \otimes \mathbf{X})_{(1)} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)} \left. \right\} \\ & \quad + 3W_{st,jj}^{-2} \lambda_s^2 \text{Tr} \left\{ E \left(\epsilon \epsilon^T \right) \right. \\ & \quad \left[\left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right]^T \mathbf{L}_{wr,j}^T \\ & \quad \mathbf{L}_{wr,j} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \left. \right\} \\ & \quad + 3W_{st,jj}^{-2} \lambda_t^2 \text{Tr} \left\{ E \left(\epsilon \epsilon^T \right) \right. \\ & \quad \left[\left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right]^T \mathbf{L}_{rw,j}^T \\ & \quad \mathbf{L}_{rw,j} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \left. \right\} \end{aligned}$$

Note that, from Theorem 2 (Equation (6.35)),

$$\begin{aligned} \text{Tr} \left\{ E \left(\epsilon \epsilon^T \right) \mathbf{A} \right\} & \leq \|K\|_\infty \text{Tr} [(\mathbf{W}_t \otimes \mathbf{I}) (\mathbf{A})] \\ & \leq \|K\|_\infty \text{Tr} (\mathbf{W}_t) \text{Tr} (\mathbf{I}) \text{Tr} (\mathbf{A}) \\ & \leq \|K\|_\infty (M_1 + k_1) (M_2 + k_2) \text{Tr} (\mathbf{A}) \end{aligned} \quad (6.42)$$

where \mathbf{A} is a generic positive semidefinite matrix of suitable dimensions. We used the property that for two positive definite matrices \mathbf{A} and \mathbf{B} of suitable dimensions, $\text{Tr} (\mathbf{AB}) \leq$

$\text{Tr}(\mathbf{A}) \text{Tr}(\mathbf{B})$. Thus,

$$\begin{aligned}
 & E\left(\tilde{V}_j^2\right)/4 \\
 & \leq 3W_{st,jj}^{-2} \|K\|_\infty (M_1 + k_1)(M_2 + k_2) (\mathbf{I} \otimes \mathbf{X})_j^T \left[\mathbf{I} - (\mathbf{W}_t \otimes \mathbf{I})(\mathbf{I} \otimes \mathbf{X})_{(1)} \right. \\
 & \quad \left. \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right]^2 (\mathbf{I} \otimes \mathbf{X})_j \\
 & + 3W_{st,jj}^{-2} \|K\|_\infty (M_1 + k_1)(M_2 + k_2) \lambda_s^2 \\
 & \quad \mathbf{L}_{wr,j} \left[\left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right]^2 \mathbf{L}_{wr,j}^T \\
 & + 3W_{st,jj}^{-2} \|K\|_\infty (M_1 + k_1)(M_2 + k_2) \lambda_t^2 \\
 & \quad \mathbf{L}_{rw,j} \left[\left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right]^2 \mathbf{L}_{rw,j}^T, \quad (6.43)
 \end{aligned}$$

where for a generic matrix \mathbf{A} , $\mathbf{A}^2 = \mathbf{A}\mathbf{A}^T$. Note that

$$\begin{aligned}
 & (\mathbf{I} \otimes \mathbf{X})_j^T \left[\mathbf{I} - (\mathbf{W}_t \otimes \mathbf{I})(\mathbf{I} \otimes \mathbf{X})_{(1)} \right. \\
 & \quad \left. \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right]^2 (\mathbf{I} \otimes \mathbf{X})_j \\
 & \leq (\mathbf{I} \otimes \mathbf{X})_j^T \left[\mathbf{I} - 2(\mathbf{W}_t \otimes \mathbf{I})(\mathbf{I} \otimes \mathbf{X})_{(1)} \right. \\
 & \quad \left. \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right] (\mathbf{I} \otimes \mathbf{X})_j \\
 & + (\mathbf{I} \otimes \mathbf{X})_j^T \left[(\mathbf{I} \otimes \mathbf{X})_{(1)} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{W}_t \otimes \mathbf{I}) \right. \\
 & \quad \left. \left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right] (\mathbf{I} \otimes \mathbf{X})_j + \\
 & + (\mathbf{I} \otimes \mathbf{X})_j^T \left[(\mathbf{I} \otimes \mathbf{X})_{(1)} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{W}_t \otimes \mathbf{I}) \right. \\
 & \quad \left. \lambda_s \mathbf{L}_{wr,11} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right] (\mathbf{I} \otimes \mathbf{X})_j \\
 & + (\mathbf{I} \otimes \mathbf{X})_j^T \left[(\mathbf{I} \otimes \mathbf{X})_{(1)} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{W}_t \otimes \mathbf{I}) \right. \\
 & \quad \left. \lambda_t \mathbf{L}_{rw,11} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right] (\mathbf{I} \otimes \mathbf{X})_j \\
 & = (\mathbf{I} \otimes \mathbf{X})_j^T \left[\mathbf{I} - (\mathbf{W}_t \otimes \mathbf{I})(\mathbf{I} \otimes \mathbf{X})_{(1)} \right. \\
 & \quad \left. \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right] (\mathbf{I} \otimes \mathbf{X})_j \\
 & \leq (\mathbf{I} \otimes \mathbf{X})_j^T (\mathbf{I} \otimes \mathbf{X})_j \leq n \max(\mathbf{X}_{(j)})^2 \leq nc^2 \quad (6.44)
 \end{aligned}$$

where the last inequality follows from condition C.1, and the max of a vector corresponds to its maximum entry.

Moreover,

$$\begin{aligned}
& \lambda_s^2 \mathbf{L}_{wr,j} \left[\left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right]^2 \mathbf{L}_{wr,j}^T \\
&= \lambda_s^2 \mathbf{L}_{wr,j} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{W}_t \otimes \mathbf{I})^{-1} \mathbf{L}_{wr,j}^T \\
&- \lambda_s^2 \mathbf{L}_{wr,j} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \lambda_s (\mathbf{W}_t \otimes \mathbf{I})^{-1} \mathbf{L}_{wr,11} \\
&\left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \mathbf{L}_{wr,j}^T \\
&- \lambda_s^2 \mathbf{L}_{wr,j} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \lambda_t (\mathbf{W}_t \otimes \mathbf{I})^{-1} \mathbf{L}_{rw,11} \\
&\left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \mathbf{L}_{wr,j}^T \\
&\leq \lambda_s^2 \mathbf{L}_{wr,j} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{W}_t \otimes \mathbf{I})^{-1} \mathbf{L}_{wr,j}^T \\
&\leq \frac{\lambda_s^2}{nC_{min} \Lambda_{min}(\mathbf{W}_t) M_2 (M_1 + k_1) M_2^{-1}} \mathbf{L}_{wr,j} \mathbf{L}_{wr,j}^T \leq C_1 \frac{\lambda_s^2 \Lambda_{max}^2(\mathbf{L}_{wr})}{nC_{min} c_w}, \tag{6.45}
\end{aligned}$$

with $C_{min} \doteq \Lambda_{min} \left[\left(\mathbf{W}_t \otimes n^{-1} \mathbf{X}^T \mathbf{X} \right)_{11} \right]$, where the last inequality follows because $\mathbf{L}_{wr,j} \mathbf{L}_{wr,j}^T = \mathbf{e}_i^T \mathbf{L}_{wr}^T \mathbf{L}_{wr} \mathbf{e}_i$ and because $\Lambda_{max}(\mathbf{A}^T \mathbf{A}) = \Lambda_{max}^2(\mathbf{A})$ for a generic matrix \mathbf{A} . Moreover, we used conditions C.3, C.5 and C.9. For the same arguments,

$$\begin{aligned}
& \lambda_t^2 \mathbf{L}_{rw,j} \left[\left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \right]^2 \mathbf{L}_{rw,j}^T \\
&\leq C_1 \frac{\lambda_t^2}{nC_{min} c_w} \mathbf{L}_{rw,j} \mathbf{L}_{rw,j}^T \leq C_1 \frac{\lambda_s^2 \Lambda_{max}^2(\mathbf{L}_{rw})}{nC_{min} c_w}. \tag{6.46}
\end{aligned}$$

Thus, by combining (6.41), (6.43), (6.44), (6.45) and (6.46), we have

$$\begin{aligned}
& \Pr \left\{ \frac{\max_{j \in S^c} \tilde{V}_j}{\lambda} > \eta \right\} \\
&\leq \frac{6\sqrt{3}C_2 M_1 M_2 \|K\|_x \sqrt{\log[(M_1 + k_1)(M_2 + k_2) - q]}}{\lambda_L \eta} \sqrt{nc^2 + C_1 \frac{\lambda_s^2 \Lambda_{max}^2(\mathbf{L}_{wr})}{nC_{min} c_w} + C_1 \frac{\lambda_s^2 \Lambda_{max}^2(\mathbf{L}_{rw})}{nC_{min} c_w}},
\end{aligned}$$

because $W_{st,jj}^{-2} \leq M_1^2 M_2^2$. Thus, under condition C.9 (a), $\Pr \left\{ \frac{\max_{j \in S^c} \tilde{V}_j}{\lambda} > \eta \right\} \rightarrow 0$ and hence condition (6.38) holds.

To analyse the event $M(U)$, let

$$\mathbf{Z}_i = \mathbf{e}_i^T \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \boldsymbol{\epsilon},$$

then,

$$\begin{aligned}
& \max_i |U_i| \\
&= \max_i |Z_i - \frac{1}{2} \mathbf{e}_i^T \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \lambda_L \mathbf{W}_{st,11} \text{sign}(\mathbf{b}_{(1)})| \\
&\leq \max_i |Z_i| + \frac{1}{2} \lambda_L \left\| \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \mathbf{W}_{st,11} \text{sign}(\mathbf{b}_{(1)}) \right\|_2 \\
&\leq \max_i |Z_i| + \frac{1}{2} \lambda_L \sqrt{\Lambda_{max} \left(\left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} \mathbf{W}_{st,11} \right)^2} \|\text{sign}(\mathbf{b}_{(1)})\|_2 \\
&\leq \max_i |Z_i| + \frac{1}{2} \lambda_L \Lambda_{min}^{-1} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right] \Lambda_{max}(\mathbf{W}_{st,11}) \|\text{sign}(\mathbf{b}_{(1)})\|_2 \\
&\leq \max_i |Z_i| + \frac{1}{2} \frac{\lambda_L}{M_1 M_2} \Lambda_{min}^{-1} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right] \|\text{sign}(\mathbf{b}_{(1)})\|_2
\end{aligned}$$

Under condition C.5, we have

$$\Lambda_{min} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right] \geq n C_{min},$$

and, from (6.42)

$$\begin{aligned}
& \text{Var}(Z_i) \\
&\leq \|K\|_\infty (M_1 + k_1) (M_2 + k_2) \mathbf{e}_i^T \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{I} \otimes \mathbf{X})_{(1)}^T \\
&(\mathbf{I} \otimes \mathbf{X})_{(1)} \left(\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right)^{-1} \mathbf{e}_i \\
&\leq \|K\|_\infty (M_1 + k_1) (M_2 + k_2) \mathbf{e}_i^T \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right]^{-1} (\mathbf{W}_t \otimes \mathbf{I})^{-1} \mathbf{e}_i \\
&\leq C_1 \frac{\|K\|_\infty (M_1 + k_1) (M_2 + k_2)}{n C_{min} c_w}.
\end{aligned}$$

By condition C.8, $Z = (Z_1, \dots, Z_q)^T$ is a Gaussian random vector, by using the Gaussian comparison result we have,

$$E \left(\max_i |Z_i| \right) \leq 3 \sqrt{C_1 \frac{\|K\|_\infty (M_1 + k_1) (M_2 + k_2) \log(q)}{n C_{min} c_w}} \leq 3 \sqrt{\frac{C_1 C_2 \|K\|_\infty M_1 M_2 \log(q)}{n C_{min} c_w}}.$$

Then,

$$\begin{aligned}
& \Pr\{\max_i |U_i| \geq \rho\} \\
& \leq \Pr\left\{\frac{1}{\rho} \left[\max_i |Z_i| \right. \right. \\
& \quad \left. \left. + \frac{1}{2} \frac{\lambda_L}{M_1 M_2} \Lambda_{min}^{-1} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right] \|\text{sign}(\mathbf{b}_{(1)})\|_2 \right] \geq 1 \right\} \\
& \leq \frac{1}{\rho} \left\{ E \left(\max_i |Z_i| \right) \right. \\
& \quad \left. + \frac{1}{2} \frac{\lambda_L}{M_1 M_2} \Lambda_{min}^{-1} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right] \|\text{sign}(\mathbf{b}_{(1)})\|_2 \right\} \\
& \leq \frac{1}{\rho} \left\{ 3 \sqrt{\frac{C_1 C_2 \|K\|_\infty M_1 M_2 \log(q)}{n C_{min} c_w}} \right. \\
& \quad \left. + \frac{1}{2} \frac{\lambda_L}{M_1 M_2} \Lambda_{min}^{-1} \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right)_{11} + \lambda_s \mathbf{L}_{wr,11} + \lambda_t \mathbf{L}_{rw,11} \right] \|\text{sign}(\mathbf{b}_{(1)})\|_2 \right\}.
\end{aligned}$$

Therefore, condition C.9 (b) guarantees that (6.39) holds.

Chapter 7

Adaptive Smoothing Spline Estimator for the Function-on-Function Linear Regression Model

Abstract

In this chapter, we propose an adaptive smoothing spline (AdaSS) estimator for the function-on-function linear regression model where each value of the response, at any domain point, depends on the full trajectory of the predictor. The AdaSS estimator is obtained by the optimization of an objective function with two spatially adaptive penalties, based on initial estimates of the partial derivatives of the regression coefficient function. This allows the proposed estimator to adapt more easily to the true coefficient function over regions of large curvature and not to be undersmoothed over the remaining part of the domain. A novel evolutionary algorithm is developed ad hoc to obtain the optimization tuning parameters. Extensive Monte Carlo simulations have been carried out to compare the AdaSS estimator with competitors that have already appeared in the literature before. The results show that our proposal mostly outperforms the competitor in terms of estimation and prediction accuracy. Lastly, those advantages are illustrated also on two real-data benchmark examples.

7.1 Introduction

Complex datasets are increasingly available due to advancements in technology and computational power and have stimulated significant methodological developments. In this regard, functional data analysis (FDA) addresses the issue of dealing with data that can be modeled as functions defined on a compact domain. FDA is a thriving area of statistics and, for a comprehensive overview, the reader could refer to Ramsay and Silverman (2005); Hsing and Eubank (2015); Horváth and Kokoszka (2012); Kokoszka and Reimherr (2017); Ferraty and Vieu (2006). In particular, the generalization of the classical multivariate regression analysis to the case where the predictor and/or the response have a functional form is referred to as functional regression and is illustrated e.g., in Morris (2015) and Ramsay and Silverman (2005). Most of the functional regression methods have been developed for models with scalar response and functional predictors (scalar-on-function regression) or functional response and scalar predictors (function-on-scalar regression). Some results may be found in Cardot et al. (2003); James (2002); Yao and Müller (2010); Müller and Stadtmüller (2005); Scheipl

7. ADAPTIVE SMOOTHING SPLINE ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL

et al. (2015); Ivanescu et al. (2015); Hullait et al. (2020). Models where both the response and the predictor are functions, namely function-on-function (FoF) regression, have been far less studied until now. In this work, we study FoF linear regression models, where the response variable function, at any domain point, depends linearly on the full trajectory of the predictor. That is,

$$Y_i(t) = \int_{\mathcal{S}} X_i(s) \beta(s, t) ds + \varepsilon_i(t) \quad t \in \mathcal{T}, \quad (7.1)$$

for $i = 1, \dots, n$. The pairs (X_i, Y_i) are independent realizations of the predictor X and the response Y , which are assumed to be smooth random process with realizations in $L^2(\mathcal{S})$ and $L^2(\mathcal{T})$, i.e., the Hilbert spaces of square integrable functions defined on the compact sets \mathcal{S} and \mathcal{T} , respectively. Without loss of generality, the latter are also assumed with functional mean equal to zero. The functions ε_i are zero-mean random errors, independent of X_i . The function β is smooth in $L^2(\mathcal{S} \times \mathcal{T})$, i.e., the Hilbert space of bivariate square integrable functions defined on the closed intervals $\mathcal{S} \times \mathcal{T}$, and is hereinafter referred to as *coefficient function*. For each $t \in \mathcal{T}$, the contribution of $X_i(\cdot)$ to the conditional value of $Y_i(t)$ is generated by $\beta(\cdot, t)$, which works as continuous set of weights of the predictor evaluations. Different methods to estimate β in (7.1) have been proposed in the literature. Ramsay and Silverman (2005) assume the estimator of β to be in a finite dimension tensor space spanned by two basis sets and where regularization is achieved by either truncation or roughness penalties. (The latter is the foundation of the method proposed in this chapter as we will see below.) Yao et al. (2005b) assume the estimator of β to be in a tensor product space generated by the eigenfunctions of the covariance functions of the predictor X and the response Y , estimated by using the principal analysis by conditional expectation (PACE) method (Yao et al., 2005a). More recently, Luo and Qi (2017) propose an estimation method of the FoF linear model with multiple functional predictors based on a finite-dimensional approximation of the mean response obtained by solving a penalized generalized functional eigenvalue problem. Qi and Luo (2018) generalize the method in Luo and Qi (2017) to the high dimensional case, where the number of covariates is much larger than the sample size (i.e., $p \gg n$). In order to improve model flexibility and prediction accuracy, Luo and Qi (2019) consider a FoF regression model with interaction and quadratic effects. A nonlinear FoF additive regression model with multiple functional predictors is proposed by Qi and Luo (2019).

One of the most used estimation method is the *smoothing spline estimator* $\hat{\beta}_{SS}$ introduced by Ramsay and Silverman (2005). It is obtained as the solution of the following optimization problem

$$\begin{aligned} \hat{\beta}_{SS} = \operatorname{argmin}_{\alpha \in \mathbb{S}_{k_1, k_2, M_1, M_2}} \left\{ \sum_{i=1}^n \int_{\mathcal{T}} \left[Y_i(t) - \int_{\mathcal{S}} X_i(s) \alpha(s, t) ds \right]^2 dt \right. \\ \left. + \lambda_s \int_{\mathcal{S}} \int_{\mathcal{T}} (\mathcal{L}_s^{m_s} \alpha(s, t))^2 ds dt + \lambda_t \int_{\mathcal{S}} \int_{\mathcal{T}} (\mathcal{L}_t^{m_t} \alpha(s, t))^2 ds dt \right\}, \end{aligned} \quad (7.2)$$

where $\mathbb{S}_{k_1, k_2, M_1, M_2}$ is the tensor product space generated by the sets of B-splines of orders k_1 and k_2 associated with the non-decreasing sequences of $M_1 + 2$ and $M_2 + 2$ knots defined on \mathcal{S} and \mathcal{T} , respectively. The operators $\mathcal{L}_s^{m_s}$ and $\mathcal{L}_t^{m_t}$, with $m_s \leq k_1 - 1$ and $m_t \leq k_2 - 1$, are the m_s th and m_t th order linear differential operators applied to α with respect to the

variables s and t , respectively. The two penalty terms on the right-hand side of (7.2) measure the roughness of the function α . The positive constants λ_s and λ_t are generally referred to as *roughness parameters* and trade off smoothness and goodness of fit of the estimator. The higher their values, the smoother the estimator of the coefficient function.

Note that the two penalty terms on the right-side hand of (7.2) do not depend on s and t . Therefore, the estimator $\hat{\beta}_{SS}$ may suffer from over and under smoothing when, for instance, the true coefficient function β is wiggly or peaked only in some parts of the domain. To solve this problem, we consider two adaptive roughness parameters that are allowed to vary on the domain $\mathcal{S} \times \mathcal{T}$. In this way, more flexible estimators can be obtained to improve the estimation of the coefficient function.

Methods that use adaptive roughness parameters are very popular and well established in the field of nonparametric regression, and are referred to as *adaptive* methods. In particular, the smoothing spline estimator for nonparametric regression (Wahba, 1990; Green and Silverman, 1993; Eubank, 1999; Gu, 2013) has been extended by different authors to take into account the non-uniform smoothness along the domain of the function to be estimated (Ruppert and Carroll, 2000; Pintore et al., 2006; Storlie et al., 2010; Wang et al., 2013; Yang and Hong, 2017).

In this chapter, a *spatially adaptive* estimator is proposed as the solutions of the following minimization problem

$$\begin{aligned} \operatorname{argmin}_{\alpha \in \mathbb{S}_{k_1, k_2, M_1, M_2}} \left\{ \sum_{i=1}^n \int_{\mathcal{T}} \left[Y_i(t) - \int_{\mathcal{S}} X_i(s) \alpha(s, t) ds \right]^2 dt \right. \\ \left. + \int_{\mathcal{S}} \int_{\mathcal{T}} \lambda_s(s, t) (\mathcal{L}_s^{m_s} \alpha(s, t))^2 ds dt + \int_{\mathcal{S}} \int_{\mathcal{T}} \lambda_t(s, t) (\mathcal{L}_t^{m_t} \alpha(s, t))^2 ds dt \right\}, \end{aligned} \quad (7.3)$$

where the two roughness parameters $\lambda_s(s, t)$ and $\lambda_t(s, t)$ are functions that produce different amount of penalty, and, thus, allow the estimator to spatially adapt, i.e., to accommodate varying degrees of roughness over the domain $\mathcal{S} \times \mathcal{T}$. Therefore, the model may accommodate the local behavior of β by imposing a heavier penalty in regions of lower smoothness. Because $\lambda_s(s, t)$ and $\lambda_t(s, t)$ are intrinsically infinite dimensional, their specification could be rather complicated without further assumptions.

The proposed estimator is applied to FoF linear regression model reported in (7.1), and is referred to as adaptive smoothing spline (AdaSS) estimator. It is obtained as the solution of the optimization problem in (7.3), with $\lambda_s(s, t)$ and $\lambda_t(s, t)$ chosen based on an initial estimate of the partial derivatives $\mathcal{L}_s^{m_s} \alpha(s, t)$ and $\mathcal{L}_t^{m_t} \alpha(s, t)$. The rationale behind this choice is to allow the contribution of $\lambda_s(s, t)$ and $\lambda_t(s, t)$, to the penalties in (7.3), to be small over regions where the initial estimate has large m_s th and m_t th curvatures (i.e., partial derivatives), respectively. This can be regarded as an extension to the FoF linear regression model of the idea of Storlie et al. (2010) and Abramovich and Steinberg (1996). Moreover, to overcome some limitations of the most famous grid-search method (Bergstra et al., 2011), a new evolutionary algorithm is proposed for the choice of the unknown parameters, needed to compute the AdaSS estimator.

The rest of the chapter is organized as follows. In Section 7.2, the proposed estimator is presented. Computational issues involved in the AdaSS estimator calculation are discussed in Section 7.2 and Section 7.2. In Section 7.3, by means of a Monte Carlo simulation study, the performance of the proposed estimator are compared with those achieved by competing estimators already appeared in the literature. Lastly, two real-data examples are presented

in Section 7.4 to illustrate the practical applicability of the proposed estimator. Conclusions are in Section 7.5.

7.2 The Adaptive Smoothing Spline Estimator

The Estimator

The AdaSS estimator $\hat{\beta}_{AdaSS}$ is defined as the solution of the optimization problem in (7.3) where the two roughness parameters $\lambda_s(s, t)$ and $\lambda_t(s, t)$ are as follows

$$\lambda_s(s, t) = \lambda_s^{AdaSS} \frac{1}{\left(|\widehat{\beta_s^{m_s}}(s, t)| + \delta_s\right)^{\gamma_s}}, \quad \lambda_t(s, t) = \lambda_t^{AdaSS} \frac{1}{\left(|\widehat{\beta_t^{m_t}}(s, t)| + \delta_t\right)^{\gamma_t}}$$

that is,

$$\begin{aligned} \hat{\beta}_{AdaSS} = \underset{\alpha \in \mathbb{S}_{k_1, k_2, M_1, M_2}}{\operatorname{argmin}} \bigg\{ & \sum_{i=1}^n \int_{\mathcal{T}} \left[Y_i(t) - \int_{\mathcal{S}} X_i(s) \alpha(s, t) ds \right]^2 dt + \lambda_s^{AdaSS} \int_{\mathcal{S}} \int_{\mathcal{T}} \frac{1}{\left(|\widehat{\beta_s^{m_s}}(s, t)| + \delta_s\right)^{\gamma_s}} (\mathcal{L}_s^{m_s} \alpha(s, t))^2 ds dt \\ & + \lambda_t^{AdaSS} \int_{\mathcal{S}} \int_{\mathcal{T}} \frac{1}{\left(|\widehat{\beta_t^{m_t}}(s, t)| + \delta_t\right)^{\gamma_t}} (\mathcal{L}_t^{m_t} \alpha(s, t))^2 ds dt \bigg\}, \quad (7.4) \end{aligned}$$

for some tuning parameters $\lambda_s^{AdaSS}, \delta_s, \gamma_s, \lambda_t^{AdaSS}, \delta_t, \gamma_t \geq 0$ and $\widehat{\beta_s^{m_s}}$ and $\widehat{\beta_t^{m_t}}$ initial estimates of $\mathcal{L}_s^{m_s} \beta$ and $\mathcal{L}_t^{m_t} \beta$, respectively. Note that the two roughness parameters λ_s and λ_t assume large values over domain regions where $\widehat{\beta_s^{m_s}}$ and $\widehat{\beta_t^{m_t}}$ are small. Therefore, in the right-hand side of (7.4), $(\mathcal{L}_s^{m_s} \alpha)^2$ and $(\mathcal{L}_t^{m_t} \alpha)^2$ are weighted through the inverse of $|\widehat{\beta_s^{m_s}}|$ and $|\widehat{\beta_t^{m_t}}|$. That is, over domain regions where $\widehat{\beta_s^{m_s}}$ and $\widehat{\beta_t^{m_t}}$ are small, $(\mathcal{L}_s^{m_s} \alpha)^2$ and $(\mathcal{L}_t^{m_t} \alpha)^2$ have larger weights than over those regions where $\widehat{\beta_s^{m_s}}$ and $\widehat{\beta_t^{m_t}}$ are large. For this reasons, the final estimator is able to adapt to the coefficient function over regions of large curvature without over smoothing it over regions where the m_s th and m_t th curvatures are small.

The constants δ_s and δ_t allow $\hat{\beta}_{AdaSS}$ not to have m_s th and m_t th-order inflection points at the same location of $\widehat{\beta_s^{m_s}}$ and $\widehat{\beta_t^{m_t}}$, respectively. Indeed, when δ_s and δ_t are set to zero, where $\widehat{\beta_s^{m_s}} = 0$ and $\widehat{\beta_t^{m_t}} = 0$ (m_s th and m_t th-order inflection points), the corresponding penalties go to infinite, and, thus, $\mathcal{L}_s^{m_s} \alpha(s, t)$ and $\mathcal{L}_t^{m_t} \alpha(s, t)$ become zero in accordance with the minimization problem. Therefore, the presence of δ_s and δ_t makes $\hat{\beta}_{AdaSS}$ more robust against the choice of the initial estimate of the linear differential operators applied to β with respect to s and t . Finally, γ_s and γ_t control the amount of weight placed in $\widehat{\beta_s^{m_s}}$ and $\widehat{\beta_t^{m_t}}$, whereas λ_s^{AdaSS} and λ_t^{AdaSS} are smoothing parameters. The solution of the optimization problem in (7.4) can be obtained in closed form if the penalty terms are approximated as described in Section 7.2. There are several choices for the initial estimates $\widehat{\beta_s^{m_s}}$ and $\widehat{\beta_t^{m_t}}$. As in Abramovich and Steinberg (1996), we suggest to apply the m_s th and m_t th order linear differential operator to the smoothing spline estimator $\hat{\beta}_{SS}$ in (7.2).

The Derivation of the AdaSS Estimator

The minimization in (7.2) is carried out over $\alpha \in \mathbb{S}_{k_1, k_2, M_1, M_2}$. This implicitly means that we are approximating β as follows

$$\beta(s, t) \approx \tilde{\beta}(s, t) = \sum_{i=1}^{M_1+k_1} \sum_{j=1}^{M_2+k_2} b_{ij} \psi_i^s(s) \psi_j^t(t) = \boldsymbol{\psi}^s(s)^T \mathbf{B} \boldsymbol{\psi}^t(t) \quad s \in \mathcal{S}, t \in \mathcal{T}, \quad (7.5)$$

where $\mathbf{B} = \{b_{ij}\} \in \mathbb{R}^{M_1+k_1 \times M_2+k_2}$. The two sets $\boldsymbol{\psi}^s = (\psi_1^s, \dots, \psi_{M_1+k_1}^s)^T$ and $\boldsymbol{\psi}^t = (\psi_1^t, \dots, \psi_{M_2+k_2}^t)^T$ are B-spline functions of order k_1 and k_2 and non-decreasing knots sequences $\Delta^s = \{s_0, s_1, \dots, s_{M_1}, s_{M_1+1}\}$ and $\Delta^t = \{t_0, t_1, \dots, t_{M_2}, t_{M_2+1}\}$, defined on $\mathcal{S} = [s_0, s_{M_1+1}]$ and $\mathcal{T} = [t_0, t_{M_2+1}]$, respectively, that generate $\mathbb{S}_{k_1, k_2, M_1, M_2}$. Thus, estimating β in (7.2) means estimating \mathbf{B} . Let $\alpha(s, t) = \boldsymbol{\psi}^s(s)^T \mathbf{B}_\alpha \boldsymbol{\psi}^t(t)$, $s \in \mathcal{S}, t \in \mathcal{T}$, in $\mathbb{S}_{k_1, k_2, M_1, M_2}$, where $\mathbf{B}_\alpha = \{b_{\alpha, ij}\} \in \mathbb{R}^{M_1+k_1 \times M_2+k_2}$. Then, the first term of the right-hand side of (7.4) may be rewritten as (see Ramsay and Silverman (2005), pag 291-293, for the derivation)

$$\sum_{i=1}^n \int_{\mathcal{T}} \left[Y_i(t) - \int_{\mathcal{S}} X_i(s) \alpha(s, t) ds \right]^2 dt = \sum_{i=1}^n \int_{\mathcal{T}} Y_i(t)^2 dt - 2 \operatorname{Tr} [\mathbf{X} \mathbf{B}_\alpha \mathbf{Y}^T] + \operatorname{Tr} [\mathbf{X}^T \mathbf{X} \mathbf{B}_\alpha \mathbf{W}_t \mathbf{B}_\alpha^T], \quad (7.6)$$

where $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)^T$, with $\mathbf{X}_i = \int_{\mathcal{S}} X_i(s) \boldsymbol{\psi}^s(s) ds$, $\mathbf{Y} = (\mathbf{Y}_1, \dots, \mathbf{Y}_n)^T$ with $\mathbf{Y}_i = \int_{\mathcal{T}} Y_i(t) \boldsymbol{\psi}^t(t) dt$, and $\mathbf{W}_t = \int_{\mathcal{T}} \boldsymbol{\psi}^t(t) \boldsymbol{\psi}^t(t)^T dt$. The term $\operatorname{Tr}[\mathbf{A}]$ denotes the trace of a square matrix \mathbf{A} .

In order to simplify the integrals in the two penalty terms on the right-hand side of (7.4), and thus obtain a linear form in \mathbf{B}_α , we consider, for $s \in \mathcal{S}$ and $t \in \mathcal{T}$, the following approximations of $\widehat{\beta}_s^{m_s}$ and $\widehat{\beta}_t^{m_t}$

$$\widehat{\beta}_s^{m_s}(s, t) \approx \sum_{i=0}^{L_s} \sum_{j=0}^{L_t} \widehat{\beta}_s^{m_s}(\tau_{s,i+1}, \tau_{t,j+1}) I_{[(\tau_{s,i}, \tau_{s,i+1}) \times (\tau_{t,j}, \tau_{t,j+1})]}(s, t), \quad (7.7)$$

and

$$\widehat{\beta}_t^{m_t}(s, t) \approx \sum_{i=0}^{L_s} \sum_{j=0}^{L_t} \widehat{\beta}_t^{m_t}(\tau_{s,i+1}, \tau_{t,j+1}) I_{[(\tau_{s,i}, \tau_{s,i+1}) \times (\tau_{t,j}, \tau_{t,j+1})]}(s, t), \quad (7.8)$$

where $\Theta^s = \{\tau_{s,0}, \tau_{s,1}, \dots, \tau_{s,L_s}, \tau_{s,L_s+1}\}$ and $\Theta^t = \{\tau_{t,0}, \tau_{t,1}, \dots, \tau_{t,L_t}, \tau_{t,L_t+1}\}$ are non increasing knot sequences with $\tau_{s,0} = s_0$, $\tau_{s,L_s+1} = s_{M_1+1}$, $\tau_{t,0} = t_0$, $\tau_{t,L_t+1} = t_{M_2+1}$, and $I_{[a \times b]}(z_1, z_2) = 1$ for $(z_1, z_2) \in [a \times b]$ and zero elsewhere. In (7.7) and (7.8), we are assuming that $\widehat{\beta}_s^{m_s}$ and $\widehat{\beta}_t^{m_t}$ are well approximated by a piecewise constant function, whose values are constant on rectangles defined by the two knot sequences Θ^s and Θ^t . It can be easily proved, by following Schumaker (2007) (pag. 491, Theorem 12.7), that the approximation error in both cases goes to zero as the mesh widths $\bar{\delta}^s = \max_i (\tau_{s,i+1} - \tau_{s,i})$ and $\bar{\delta}^t = \max_j (\tau_{t,j+1} - \tau_{t,j})$ go to zero. Therefore, $\widehat{\beta}_s^{m_s}$ and $\widehat{\beta}_t^{m_t}$ can be exactly recovered by uniformly increasing the number of knots L_s and L_t . In this way, the two penalties on the

7. ADAPTIVE SMOOTHING SPLINE ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL

right-hand side of (7.4) can be rewritten as (Appendix)

$$\lambda_s^{AdaSS} \int_S \int_{\mathcal{T}} \frac{1}{\left(|\widehat{\beta_s^{m_s}}(s, t)| + \delta_s \right)^{\gamma_s}} \left(\mathcal{L}_s^{m_s} \alpha(s, t) \right)^2 ds dt \approx \lambda_s^{AdaSS} \sum_{i=1}^{L_s+1} \sum_{j=1}^{L_t+1} d_{ij}^s \text{Tr} \left[\mathbf{B}_\alpha^T \mathbf{R}_{s,i} \mathbf{B}_\alpha \mathbf{W}_{t,j} \right] \quad (7.9)$$

and

$$\lambda_t^{AdaSS} \int_S \int_{\mathcal{T}} \frac{1}{\left(|\widehat{\beta_t^{m_t}}(s, t)| + \delta_t \right)^{\gamma_t}} \left(\mathcal{L}_t^{m_t} \alpha(s, t) \right)^2 ds dt \approx \lambda_t^{AdaSS} \sum_{i=1}^{L_s+1} \sum_{j=1}^{L_t+1} d_{ij}^t \text{Tr} \left[\mathbf{B}_\alpha^T \mathbf{W}_{s,i} \mathbf{B}_\alpha \mathbf{R}_{t,j} \right], \quad (7.10)$$

where $\mathbf{W}_{s,i} = \int_{[\tau_{s,i-1}, \tau_{s,i}]} \boldsymbol{\psi}^s(s) \boldsymbol{\psi}^s(s)^T ds$, $\mathbf{W}_{t,j} = \int_{[\tau_{t,j-1}, \tau_{t,j}]} \boldsymbol{\psi}^t(t) \boldsymbol{\psi}^t(t)^T dt$,
 $\mathbf{R}_{s,i} = \int_{[\tau_{s,i-1}, \tau_{s,i}]} \mathcal{L}_s^{m_s} [\boldsymbol{\psi}^s(s)] \mathcal{L}_s^{m_s} [\boldsymbol{\psi}^s(s)]^T ds$ and $\mathbf{R}_{t,j} = \int_{[\tau_{t,j-1}, \tau_{t,j}]} \mathcal{L}_t^{m_t} [\boldsymbol{\psi}^t(t)] \mathcal{L}_t^{m_t} [\boldsymbol{\psi}^t(t)]^T dt$, and $d_{ij}^s = \left\{ \frac{1}{\left(|\widehat{\beta_s^{m_s}}(\tau_{s,i}, \tau_{t,j})| + \delta_s \right)^{\gamma_s}} \right\}$ and $d_{ij}^t = \left\{ \frac{1}{\left(|\widehat{\beta_t^{m_t}}(\tau_{s,i}, \tau_{t,j})| + \delta_t \right)^{\gamma_t}} \right\}$, for $i = 1, \dots, L_s + 1$ and $j = 1, \dots, L_t + 1$.

The optimization problem in (7.4) can be then approximated with the following

$$\begin{aligned} \hat{\mathbf{B}}_{AS} \approx \underset{\mathbf{B}_\alpha \in \mathbb{R}^{(M_1+k_1) \times (M_2+k_2)}}{\text{argmin}} \left\{ \sum_{i=1}^n \int_{\mathcal{T}} Y_i(t)^2 dt - 2 \text{Tr} \left[\mathbf{X} \mathbf{B}_\alpha \mathbf{Y}^T \right] + \text{Tr} \left[\mathbf{X}^T \mathbf{X} \mathbf{B}_\alpha \mathbf{W}_t \mathbf{B}_\alpha^T \right] \right. \\ \left. + \sum_{i=1}^{L_s+1} \sum_{j=1}^{L_t+1} \left(\lambda_s^{AdaSS} d_{ij}^s \text{Tr} \left[\mathbf{B}_\alpha^T \mathbf{R}_{s,i} \mathbf{B}_\alpha \mathbf{W}_{t,j} \right] + \lambda_t^{AdaSS} d_{ij}^t \text{Tr} \left[\mathbf{B}_\alpha^T \mathbf{W}_{s,i} \mathbf{B}_\alpha \mathbf{R}_{t,j} \right] \right) \right\}, \end{aligned} \quad (7.11)$$

or by vectorization as

$$\begin{aligned} \hat{\mathbf{b}}_{AS} \approx \underset{\mathbf{b}_\alpha \in \mathbb{R}^{(M_1+k_1)(M_2+k_2)}}{\text{argmin}} \left\{ -2 \text{vec} \left(\mathbf{X}^T \mathbf{Y} \right)^T \mathbf{b}_\alpha + \mathbf{b}_\alpha^T \left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) \mathbf{b}_\alpha \right. \\ \left. + \sum_{i=1}^{L_s+1} \sum_{j=1}^{L_t+1} \left(\lambda_s^{AdaSS} d_{ij}^s \mathbf{b}_\alpha^T \mathbf{L}_{wr,ij} \mathbf{b}_\alpha + \lambda_t^{AdaSS} d_{ij}^t \mathbf{b}_\alpha^T \mathbf{L}_{rw,ij} \mathbf{b}_\alpha \right) \right\}, \end{aligned} \quad (7.12)$$

where $\hat{\mathbf{b}}_{AS} = \text{vec}(\hat{\mathbf{B}}_{AS})$, $\mathbf{L}_{rw,ij} = (\mathbf{R}_{t,j} \otimes \mathbf{W}_{s,i})$ and $\mathbf{L}_{wr,ij} = (\mathbf{W}_{t,j} \otimes \mathbf{R}_{s,i})$, for $i = 1, \dots, L_s + 1$ and $j = 1, \dots, L_t + 1$. For a matrix $\mathbf{A} \in \mathbb{R}^{j \times k}$, $\text{vec}(\mathbf{A})$ indicates the vector of length jk obtained by writing the matrix \mathbf{A} as a vector column-wise, and \otimes is the Kronecker product. Then, the minimizer of the optimization problem in (7.12) has the following expression

$$\begin{aligned} \hat{\mathbf{b}}_{AdaSS} \approx \left[\left(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X} \right) + \sum_{i=1}^{L_s+1} \sum_{j=1}^{L_t+1} \left(\lambda_s^{AdaSS} d_{ij}^s \mathbf{L}_{wr,ij} + \lambda_t^{AdaSS} d_{ij}^t \mathbf{L}_{rw,ij} \right) \right]^{-1} \text{vec} \left(\mathbf{X}^T \mathbf{Y} \right) \\ = \mathbf{K}^{-1} \text{vec} \left(\mathbf{X}^T \mathbf{Y} \right) \end{aligned} \quad (7.13)$$

The identifiability of β , i.e., the uniqueness of $\hat{\mathbf{b}}_{AdaSS}$, comes from the fact that the inverse of \mathbf{K} exists. This is guaranteed with probability tending to one as the sample size increases, under the condition that the covariance operator of X is strictly positive, i.e., its kernel is empty (Prchal and Sarda, 2007). In Equation (7.13), this reverts into the condition that $\mathbf{X}^T \mathbf{X}$ is positive definite. Moreover, Scheipl and Greven (2016) show that identifiability still holds also in case of rank deficiency of $(\mathbf{W}_t \otimes \mathbf{X}^T \mathbf{X})$ if, and only if, the kernel of the covariate covariance operator does not overlap that of the roughness penalties.

To obtain $\hat{\mathbf{b}}_{AdaSS}$ in (7.13) the tuning parameters $\lambda_s^{AdaSS}, \delta_s, \gamma_s, \lambda_t^{AdaSS}, \delta_t, \gamma_t$ must be opportunely chosen. This issue is discussed in Section 7.2.

The Algorithm for the Parameter Selection

There are some tuning parameters in the optimization problem (7.12) that must be chosen to obtain the AdaSS estimator. Usually, the tensor product space $\mathbb{S}_{k_1, k_2, M_1, M_2}$ is chosen with $k_1 = k_2 = 4$, i.e., cubic B-splines, and equally spaced knot sequences. Although the choice of M_1 and M_2 is not crucial (Cardot et al., 2003), it should allow the final estimator to capture the local behaviour of the coefficient function β , that is, M_1 and M_2 should be sufficiently large. The smoothness of the final estimator is controlled by the two penalty terms on the right-hand side of (7.12).

The tuning parameters $\lambda_s^{AdaSS}, \delta_s, \gamma_s, \lambda_t^{AdaSS}, \delta_t, \gamma_t$ could be fixed by using the conventional K -fold cross validation (CV) (Hastie et al., 2009a), where the combination of parameters to be explored is chosen by means of the classic grid search method (Hastie et al., 2009a). That is an exhaustive searching through a manually specified subset of the tuning parameter space (Bergstra and Bengio, 2012). Although, in our setting, grid search is embarrassingly parallel (Herlihy and Shavit, 2011), it is not scalable because it suffers from the curse of dimensionality. However, even if this is beyond the scope of the present work, note that the number of combinations to explore grows exponentially with the number of tuning parameters and makes unsuitable the application of the proposed method to the FoF linear model in the case of multiple predictors. Then, to facilitate the use of the proposed method by practitioners, in what follows, we proposed a novel evolutionary algorithm for tuning parameter selection, referred to as *evolutionary algorithm for adaptive smoothing estimator* (EAASS) inspired by the *population based training* (PBT) introduced by Jaderberg et al. (2017). The PBT algorithm was introduced to address the issue of hyperparameter optimization for neural networks. It bridges and extends parallel search method (e.g., grid search and random search) with sequential optimization method (e.g., hand tuning and Bayesian optimization). The former runs many parallel optimization processes, for different combinations of hyperparameter values, and, then chooses the combination that shows the best performance. The latter performs several steps of few parallel optimizations, where, at each step, information coming from the previous step is used to identify the combinations of hyperparameter values to explore. For further details on the PBT algorithm the readers should refer to Jaderberg et al. (2017), where the authors demonstrated its effectiveness and wide applicability. The pseudo code of the EAASS algorithm is given in Algorithm 2.

The first step is the identification of an initial population \mathcal{P} of tuning parameter combinations p_i s. This can be done, for each combination and each tuning parameter, by randomly selecting a value in a pre-specified range. Then, the set \mathcal{V} of estimated prediction errors v_i s corresponding to \mathcal{P} is obtained by means of K -fold CV. We choose a subset \mathcal{Q} of \mathcal{P} , by following a given exploitation strategy and, thus, the corresponding subset \mathcal{Z} of \mathcal{V} . A typical exploitation strategy is the *truncation selection*, where the worse $r\%$, for $0 \leq r \leq 100$, of \mathcal{P} ,

7. ADAPTIVE SMOOTHING SPLINE ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL

Algorithm 2 EAASS algorithm

- 1: Choose the initial population $\mathcal{P} = \{p_i\}$ of combinations of tuning parameter values
 - 2: Obtain the set $\mathcal{V} = \{v_i\}$ of estimated prediction errors corresponding to \mathcal{P}
 - 3: **repeat**
 - 4: Identify the set $\mathcal{Q} \subseteq \mathcal{P}$ and the corresponding $\mathcal{Z} \subseteq \mathcal{V}$ ► *exploitation*
 - 5: **for** $p_i \in \mathcal{Q}$ **do** ► *exploration*
 - 6: Obtain the new combination of tuning parameter values, p'_i
 - 7: Obtain the new estimated prediction error v'_i corresponding to p'_i
 - 8: **end for**
 - 9: Define $\mathcal{Q}' = \{p'_i\}$ and $\mathcal{Z}' = \{v'_i\}$
 - 10: Set $\mathcal{P} = \mathcal{P} \setminus \mathcal{Q} \cup \mathcal{Q}'$ and $\mathcal{V} = \mathcal{V} \setminus \mathcal{Z} \cup \mathcal{Z}'$
 - 11: **until** The stopping condition is met
 - 12: Return $p_i \in \mathcal{P}$ with the highest $v_i \in \mathcal{V}$
-

in terms of estimated prediction error, is substituted by elements randomly sampled from the remaining $(100 - r)\%$ part of the current population (Jaderberg et al., 2017). Then the following step consists of an exploration strategy where the tuning parameter combinations in \mathcal{Q} are substituted by new ones. The simulation study in Section 7.3 and the real-data Examples in Section 7.4 are based on a *perturbation* where each tuning parameter value of the given combination is randomly perturbed by a factor of 1.2 or 0.8. The exploitation and exploration phases are repeated until a stopping condition is met, e.g, maximum number of iterations. Other exploration and exploitation strategies can be found in Bäck et al. (1997). At last, the selected tuning parameter combination is obtained as an element of \mathcal{P} that achieves the lowest estimated prediction error. As a remark, in our trials the AdaSS estimator works quite well with $\delta_s = \delta_s^* \max |\widehat{\beta}_s^{m_s}(s, t)|$ and $\delta_t = \delta_t^* \max |\widehat{\beta}_t^{m_t}(s, t)|$, for $0 \leq \delta_s^*, \delta_t^* \leq 0.1$.

7.3 Simulation Study

In this section, the performance of the AdaSS estimator is assessed on several simulated datasets. In particular, we compare the AdaSS estimator with cubic B-splines and $m_s = m_t = 2$ with five competing methods that represent the state of the art in the FoF liner regression model estimation. The first two are those proposed by Ramsay and Silverman (2005). The first one, hereinafter referred to as SMOOTH estimator, is the smoothing spline estimator described in (7.2), whereas, the second one, hereinafter referred to as TRU estimator, assumes that the coefficient function is in a finite dimensional tensor product space generate by two sets of B-splines with regularization achieved by choosing the space dimension. Then, we consider also the estimator proposed by Yao et al. (2005b) and Canale and Vantini (2016). The former is based on the functional principal component decomposition, and is hereinafter referred to as PCA estimator, while the latter relies on a ridge type penalization, hereinafter referred to as RIDGE estimator. Lastly, as the fifth alternative, we explore the estimator proposed by Luo and Qi (2017), hereinafter referred to as SIGCOMP. Moreover, the AdaSS estimator with cubic B-splines and $m_s = m_t = 2$ is considered. For illustrative purposes, we also consider a version of the AdaSS estimator, referred to AdaSStrue, whose roughness parameters are calculated by assuming that the true coefficient function is known. Obviously, the AdaSStrue has not a practical meaning because

the true coefficient function is never known. However, it allows one to better understand the influence of the initial estimates of the partial derivatives on the AdaSS performance. All the unknown parameters of the competing methods considered are chosen by means of 10-fold CV. The tuning parameters of the AdaSS and AdaSStrue estimators are chosen through the EAASS algorithm. The set \mathcal{P} is obtained by using 10-fold CV, the exploitation and exploration phases are as described in Section 7.2 and a maximum number of iterations equal to 15 is set as stopping condition. For each simulation, a training sample of n observations is generated along with a test set T of size $N = 4000$. They are used to estimate β and to test the predictive performance of the estimated model, respectively. Three different sample sizes are considered, viz., $n = 100, 500, 1000$. The estimation accuracy of the estimators are assessed by using the *integrated squared error* (ISE) defined as

$$\text{ISE} = \frac{1}{A} \int_{\mathcal{S}} \int_{\mathcal{T}} \left(\hat{\beta}(s, t) - \beta(s, t) \right)^2 ds dt, \quad (7.14)$$

where A is the measure of $\mathcal{S} \times \mathcal{T}$. The ISE aims to measure the estimation error of $\hat{\beta}$ with respect to β . Whereas, the predictive accuracy is measured through the *prediction mean squared error* (PMSE) defined as

$$\text{PMSE} = \frac{1}{N} \sum_{(X, Y) \in T} \int_{\mathcal{T}} \left(Y(t) - \int_{\mathcal{S}} X(s) \hat{\beta}(s, t) ds \right)^2 dt. \quad (7.15)$$

The observations in the test set are centred by subtracting to each observation the corresponding sample mean function estimated in the training set. The observations in the training and test sets are obtained as follows. The covariate X_i and the errors ε_i are generated as linear combination of cubic B-splines, Ψ_i^x and Ψ_i^ε , with evenly spaced knots, i.e., $X_i = \sum_{j=1}^{32} x_{ij} \Psi_i^x$ and $\varepsilon_i = k \sum_{j=1}^{20} e_{ij} \Psi_i^\varepsilon$. The coefficients x_{ij} and e_{ij} , for $i = 1, \dots, n$; $j = 1, \dots, 32$ and $j = 1, \dots, 20$, are independent realizations of standard normal random variable and the numbers of basis have been randomly chosen between 10 and 50. The constant k is chosen such that the signal-to-noise ratio $SN \doteq \int_{\mathcal{T}} \text{Var}_X[E(Y_i|X_i)] / \int_{\mathcal{T}} \text{Var}(\varepsilon_i)$ is equal to 4, where Var_X is the variance with respect to the random covariate X . Then, given the coefficient function β , the response Y_i is obtained.

Mexican Hat Function

The Mexican hat function is a linear function with a sharp smoothness variation in central part of the domain. In this case, the coefficient function β is defined as

$$\beta(s, t) = -1 + 1.5s + 1.5t + 0.05\phi(s, t), \quad s, t \in [0, 1] \times [0, 1]$$

where ϕ is a multivariate normal distribution with mean $\boldsymbol{\mu} = (0.6, 0.6)^T$ and diagonal covariance matrix $\boldsymbol{\Sigma} = \text{diag}(0.001, 0.001)$. Figure 7.1 displays the AdaSS and the SMOOTH estimates along with the true coefficient function for a randomly selected simulation run. The proposed estimator tends to be smoother on the flat region and is able to better capture the peak in the coefficient function (at $t \approx 0.6$) than the SMOOTH estimate. The latter, to perform reasonably well along the whole domain, selects tuning parameters that may be not sufficiently small on the peaky region, or not sufficiently large over the flat region. This is also confirmed by the graphical appeal of the AdaSS estimate with respect to the competitor ones. In Figure 7.2 and top of Table 7.1, the values of ISE and PMSE achieved by the

7. ADAPTIVE SMOOTHING SPLINE ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL

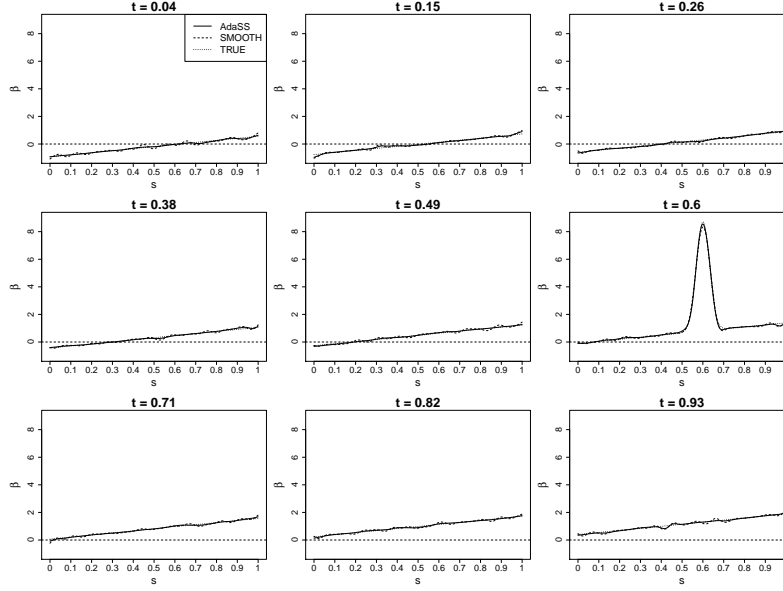


Figure 7.1. AdaSS (solid line) and SMOOTH (dashed line) estimates of the coefficient functions and the TRUE coefficient function β (dotted line) for different values of t in the case of the Mexican hat function.




Figure 7.2. a The integrated squared error (ISE) and b the prediction mean squared error (PMSE) \pm standard error for the TRU, SMOOTH, PCA, RIDGE, SIGCOMP, AdaSS and AdaSStrue estimators in the case of the Mexican hat function.

AdaSS, AdaSStrue, and competitor estimators are shown as functions of the sample size n . Without considering the AdaSStrue estimator, the AdaSS estimator yields the lowest ISE for all sample sizes, and thus has the lowest estimation error. In terms of PMSE, it is the best one for $n = 150$, whereas for $n = 500, 1000$ it performs comparably with SIGCOMP and PCA estimators. The performance of the AdaSStrue and AdaSS estimators is very similar in terms of ISE, whereas the AdaSStrue shows a lower PMSE. However, as expected, the effect of the knowledge of the true coefficient function tends to disappear as n increases, because the partial derivatives estimates become more accurate.

Table 7.1. The integrated squared error (ISE) and the prediction mean squared error (PMSE) for the TRU, SMOOTH, PCA, RIDGE, SIGCOMP, AdaSS and AdaSStrue estimators. The numbers outside the parentheses are the averages over 100 Monte Carlo replications, and the numbers inside parentheses are the corresponding standard errors. The values corresponding to the AdaSStrue estimator are emphasized to underline the fact that they rely on the knowledge of the true coefficient function, which is unlikely in real applications. In bold are marked the lowest values among the AdaSS and the competitors.

	$n = 100$		$n = 500$		$n = 1000$	
	ISE ($\times 10^{-1}$)	PMSE ($\times 10^{-2}$)	ISE ($\times 10^{-1}$)	PMSE ($\times 10^{-2}$)	ISE ($\times 10^{-1}$)	PMSE ($\times 10^{-2}$)
Mexican hat						
TRU	0.4063(0.0059)	0.3575(0.0011)	0.1384(0.0020)	0.3143(0.0007)	0.0660(0.0011)	0.3031 (0.0005)
SMOOTH	0.2191(0.0020)	0.3382(0.0007)	0.0917(0.0008)	0.3088(0.0005)	0.0564(0.0006)	0.3027 (0.0005)
PCA	0.2519(0.0068)	0.3234(0.0007)	0.0681(0.0013)	0.3030(0.0005)	0.0368(0.0008)	0.2995 (0.0005)
RIDGE	0.8813(0.0083)	0.3629(0.0008)	0.3542(0.0041)	0.3157(0.0006)	0.1847(0.0022)	0.3056 (0.0005)
SIGCOMP	0.1465(0.0026)	0.3192(0.0006)	0.0532(0.0006)	0.3026 (0.0005)	0.0358(0.0004)	0.2999 (0.0005)
AdaSS	0.0856 (0.0023)	0.3171 (0.0007)	0.0359 (0.0010)	0.3027(0.0005)	0.0217 (0.0007)	0.2994 (0.0005)
AdaSStrue	<i>0.0726(0.0176)</i>	<i>0.3080(0.0007)</i>	<i>0.0399(0.0153)</i>	<i>0.2994(0.0005)</i>	<i>0.0188(0.0048)</i>	<i>0.2977 (0.0005)</i>
Dampened harmonic						
TRU	0.2851 (0.0050)	0.5403 (0.0014)	0.0983 (0.0010)	0.5051 (0.0010)	0.0651 (0.0009)	0.4960 (0.0010)
SMOOTH	0.2288 (0.0042)	0.5391 (0.0013)	0.0836 (0.0007)	0.5032 (0.0010)	0.0555 (0.0005)	0.4936 (0.0010)
PCA	0.3710 (0.0093)	0.5259 (0.0012)	0.1100 (0.0020)	0.4994 (0.0010)	0.0594 (0.0011)	0.4915 (0.0010)
RIDGE	1.4221 (0.0135)	0.5925 (0.0016)	0.6082 (0.0076)	0.5203 (0.0011)	0.3271 (0.0038)	0.5014 (0.0010)
SIGCOMP	0.2541 (0.0045)	0.5221 (0.0012)	0.1235 (0.0013)	0.5018 (0.0010)	0.0942 (0.0009)	0.4950 (0.0010)
AdaSS	0.1749 (0.0038)	0.5241 (0.0012)	0.0695 (0.0012)	0.4997 (0.0010)	0.0461 (0.0008)	0.4918 (0.0010)
AdaSStrue	<i>0.1504 (0.0030)</i>	<i>0.5179 (0.0012)</i>	<i>0.0744 (0.0018)</i>	<i>0.4985 (0.0010)</i>	<i>0.0582 (0.0022)</i>	<i>0.4912 (0.0010)</i>
Rapid change						
TRU	1.9910(0.0278)	4.0461(0.0001)	0.9178(0.0100)	3.7583(0.0001)	0.6020(0.0074)	3.6989 (0.0001)
SMOOTH	1.2961(0.0133)	3.9427(0.0001)	0.5738(0.0046)	3.7205(0.0001)	0.3590(0.0027)	3.6787 (0.0001)
PCA	5.1052(0.0971)	4.3070(0.0001)	1.5870(0.0271)	3.7978(0.0001)	0.8383(0.0125)	3.7141 (0.0001)
RIDGE	10.4781(0.1059)	4.4295(0.0001)	4.1991(0.0537)	3.8459(0.0001)	2.2250(0.0278)	3.7356 (0.0001)
SIGCOMP	1.7129(0.0209)	4.0352(0.0001)	0.8615(0.0234)	3.7702(0.0001)	0.8552(0.0167)	3.7428 (0.0001)
AdaSS	1.0482 (0.0166)	3.8737 (0.0001)	0.4526 (0.0077)	3.6928 (0.0001)	0.2916 (0.0044)	3.6662 (0.0001)
AdaSStrue	<i>0.8181(0.0191)</i>	<i>3.8274(0.0001)</i>	<i>0.3434(0.0080)</i>	<i>3.6759(0.0001)</i>	<i>0.2114(0.0050)</i>	<i>3.6541 (0.0001)</i>

Dampened Harmonic Motion Function

This simulation scenario considers as coefficient function β the dampened harmonic motion function, also known as the *spring function* in the engineering literature. It is characterized by a sinusoidal behaviour with exponentially decreasing amplitude, that is

$$\beta(s, t) = 1 + 5 \exp[-5(s + t)] [\cos(10\pi s) + \cos(10\pi t)], \quad s, t \in [0, 1] \times [0, 1].$$

Figure 7.3 displays the AdaSS and the SMOOTH estimates along with the true coefficient function. Also in this scenario, the AdaSS estimates is smoother than the SMOOTH estimates in regions of small curvature. But, it is more flexible where the coefficient function is more wiggly. Note that intuitively, the SMOOTH estimator trades off its smoothness over the whole domain. Indeed, it over-smooths at small values of s and t and under-smooths elsewhere.

In Figure 7.4 and in the second tier of Table 7.1, values of the ISE and PMSE for the AdaSS, AdaSStrue, and competitor estimators are shown as function of the sample size n . Even in this case, the AdaSS estimator achieves the lowest ISE for all sample sizes, and thus, the lowest estimation error, without taking into account the AdaSStrue estimator. Strictly speaking, in terms of PMSE, note that the proposed estimator is not always the best choice, but it shows only a small difference with best methods, viz., PCA and SIGCOMP estimators. In this case, the AdaSS and AdaSStrue performance is very similar for $n = 500, 1000$, whereas, for $n = 150$, the AdaSStrue performs slightly better especially in terms of PMSE.

7. ADAPTIVE SMOOTHING SPLINE ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL

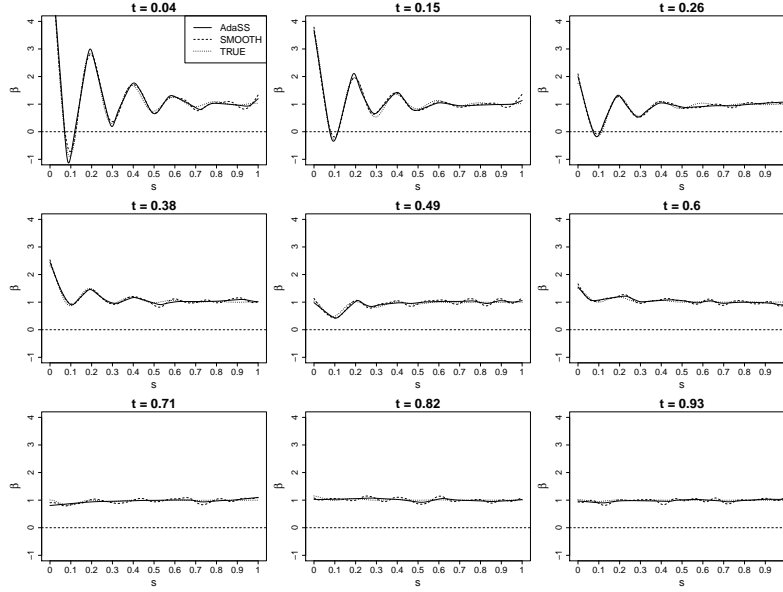


Figure 7.3. AdaSS (solid line) and SMOOTH (dashed line) estimates of the coefficient functions and the TRUE coefficient function β (dotted line) for different values of t in the case of the dampened harmonic motion function.




Figure 7.4. a The integrated squared error (ISE) and b the prediction mean squared error (PMSE) \pm standard error for the TRU, SMOOTH, PCA, SIGCOMP, AdaSS and AdaSStrue estimators in the case of the dampened harmonic motion function. The Ridge estimator is not considered due to its too different performance.




Figure 7.5. AdaSS (solid line) and SMOOTH (dashed line) estimates of the coefficient functions and the TRUE coefficient function β (dotted line) for different values of t in the case of the rapid change function.

Rapid Change Function

In this scenario, the true coefficient function β is obtained by the rapid change function, that is

$$\beta(s, t) = 1 - \frac{5}{1 + \exp[10(s + t - 0.2)]} + \frac{5}{1 + \exp[75(s + t - 0.8)]}, \quad s, t \in [0, 1] \times [0, 1].$$

Figure 7.5 shows the AdaSS and SMOOTH estimate when β is the rapid change function. The SMOOTH estimate is rougher than the AdaSS one in regions that are far from the rapid change point. On the contrary, the AdaSS estimate is able to be smoother in the flat region and to be as accurate as the SMOOTH estimate near the rapid change point.

In Figure 7.6 and the third tier of Table 7.1, values of the ISE and PMSE for the AdaSS, AdaSStrue, and competitor estimators are shown for sample sizes $n = 150, 500, 1000$. The AdaSS estimator outperforms the competitors, both in terms of ISE and PMSE. The performance of the AdaSStrue estimator is slightly better than that of the AdaSS one and this difference in performance reduces as n increases.

7.4 Real-data Examples

In this section, two real datasets, namely *Swedish mortality* and *ship CO₂ emission* datasets, are considered to assess the performance of the AdaSS estimator in real applications.

Swedish Mortality Dataset

The Swedish mortality dataset (available from the Human Mortality Database —<http://mortality.org>—) is very well known in the functional literature as a benchmark dataset.

7. ADAPTIVE SMOOTHING SPLINE ESTIMATOR FOR THE FUNCTION-ON-FUNCTION LINEAR REGRESSION MODEL

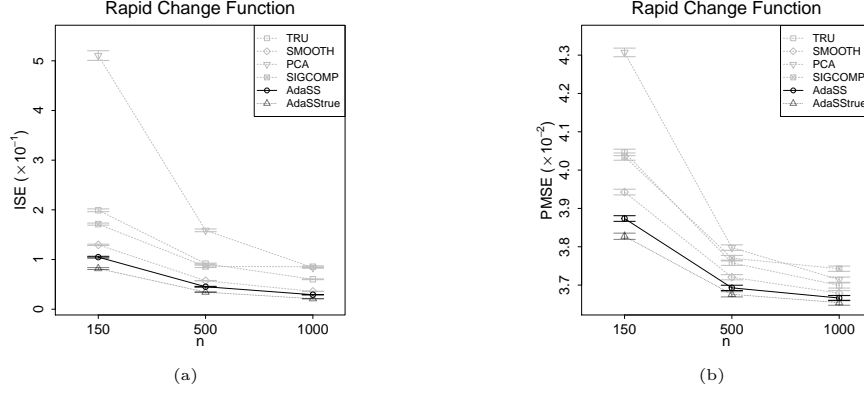


Figure 7.6. a The integrated squared error (ISE) and b the prediction mean squared error (PMSE) \pm standard error for the TRU, SMOOTH, PCA, SIGCOMP, AdaSS and AdaSStrue estimators in the case of the rapid change function. The Ridge estimator is not considered due to its too different performance.

It has been analysed by Chiou and Müller (2009) and Ramsay et al. (2009), among others. In this analysis, we consider the log-hazard rate functions of the Swedish females mortality data for year-of-birth cohorts that refer to females born in the years 1751-1935 with ages 0-80. The value of a log-hazard rate function at a given age is the natural logarithm of the ratio of females died at that age and the number of females alive with the same age. The 184 considered log-hazard rate functions (Chiou and Müller, 2009) are shown in Figure 7.7. Without loss of generality, they have been normalized to the domain $[0, 1]$.




Figure 7.7. Log-hazard rate functions for Swedish female cohorts from 1751 to 1935.

The functions from 1751 (1752) to 1934 (1935) are considered as observations X_i (Y_i) of the predictor (response) in (7.1), $i = 1, \dots, 184$. In this way, the relationship between two consecutive log-hazard rate functions becomes the focus of the analysis. To assess the predictive performance of the methods considered in the simulation study (Section 7.3), for 100 times, 166 observations out of 184 are randomly chosen, as training set, to fit the model. The 18 remaining ones are used as test set to calculate the PMSE. The averages and

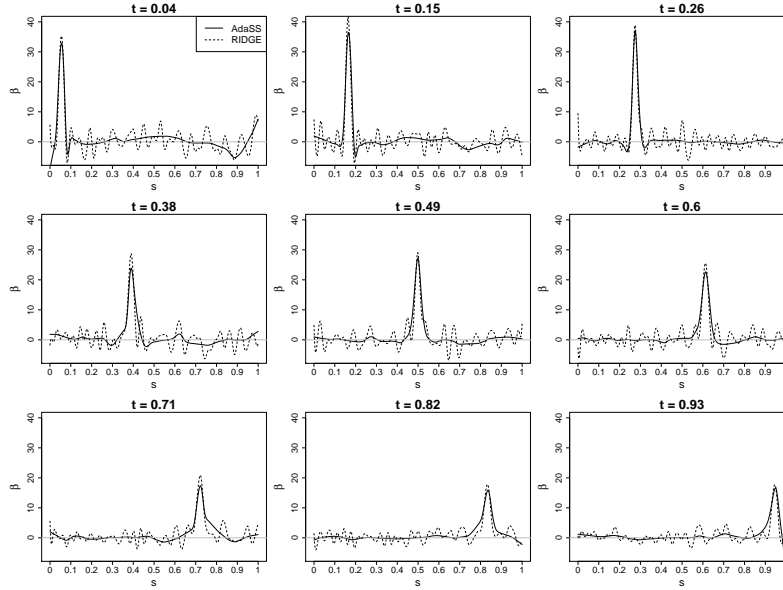


Figure 7.8. AdaSS (solid line) and RIDGE (dashed line) estimates of the coefficient functions for different values of t in the Swedish Mortality dataset.

standard deviations of PMSEs are shown in the first line of Table 7.2. The AdaSS estimator outperforms all the competitors. Only the RIDGE estimator has comparable predictive performance. Figure 7.8 shows the AdaSS estimates along with the RIDGE estimates that

Table 7.2. The prediction mean squared error (PMSE) for the TRU, SMOOTH, PCA, RIDGE, SIGCOMP, and AdaSS estimators. The numbers outside the parentheses are the averages of the PMSE over 100 replications, and the numbers inside parentheses are the corresponding standard errors.

	TRU	SMOOTH	PCA	RIDGE	SIGCOMP	AdaSS
Swedish mortality ($\times 10^{-2}$)	0.7373 (0.0000)	0.5938 (0.0000)	0.6131 (0.0000)	0.5749 (0.0000)	1.0173 (0.0000)	0.5706 (0.0000)
Ship CO ₂ emission	0.1019 (0.0008)	0.0814 (0.0007)	0.0689 (0.0008)	0.0625 (0.0007)	0.1033 (0.0013)	0.0771 (0.0007)

represents the best competitor methods in terms of PMSE. The proposed estimator has slightly better performance than the competitor, but, at the same time, it is much more interpretable. In fact, it is much smoother where the coefficient function seem to be mostly flat and successfully captures the pattern of β in the peak region. On the contrary, the RIDGE estimates is particularly rough over region of low curvature.

Ship CO₂ Emission Dataset

The ship CO₂ emission dataset has been thoroughly studied in the very last years (Lepore et al., 2018; Reis et al., 2020; Capezza et al., 2020; Centofanti et al., 2020b). It was provided by the shipping company Grimaldi Group to address aspects related to the issue of monitoring fuel consumptions or CO₂ emissions for a Ro-Pax ship that sails along a route in the Mediterranean Sea. In particular, we focus on the study of the relation between the *fuel consumption per hour* (FCPH), assumed as the response, and the *speed over ground* (SOG), assumed as predictor. The observations considered were recorded from 2015 to 2017.




Figure 7.9. SOG and FCPH observations from a Ro-Pax ship.

Figure 7.9 shows the 44 available observations of SOG and FCPH (Centofanti et al., 2020b). Similarly to the Swedish mortality dataset, the prediction performance of the methods are assessed by randomly chosen 40 out of 44 observations to fit the model and by using the 4 remaining observations to compute the PMSE. This is repeated 100 times. The averages and standard deviations of the PMSEs are listed in the second line of Table 7.2. The AdaSS estimator is, in this case, outperformed by the RIDGE estimator, which achieves the lowest PMSE. However, as shown in Figure 7.10, it is able both to well estimate the coefficient function over peaky regions, as the RIDGE estimator, and to smoothly adapt over the remaining part of the domain. Also the PCA estimator achieves smaller PMSE than that of the proposed estimator. However, the PCA estimator is even rougher than the RIDGE estimator and, thus, it is not shown in Figure 7.10.

7.5 Conclusions

In this chapter, the AdaSS estimator is proposed for the function-on-function linear regression model where each value of the response, at any domain point, depends linearly on the full trajectory of the predictor. The introduction of two adaptive smoothing penalties, based on initial estimate of its partial derivatives, allows the proposed estimator to better adapt to the coefficient function. By means of a simulation study, the proposed estimator has proven favourable performance with respect to those achieved by the five competitors already appeared in the literature before, both in terms of estimation and prediction error. The adaptive feature of the AdaSS estimator is advantageous for the interpretability of the results with respect to the competitors. Moreover, its performance has shown to be competitive also with respect to the case where the true coefficient function is known. Finally, the proposed estimator has been successfully applied to real-data examples considered, viz., the Swedish mortality and ship CO₂ emission datasets. However, some challenges are still open. Even though the proposed evolutionary algorithm has shown to perform particularly well both in the simulation study and the real-data examples, the choice of the tuning parameters still remains in fact a critical issue, because of the curse of dimensionality. This could be even more problematic in the perspective of extending the AdaSS estimator to the FoF regression model with multiple predictors.

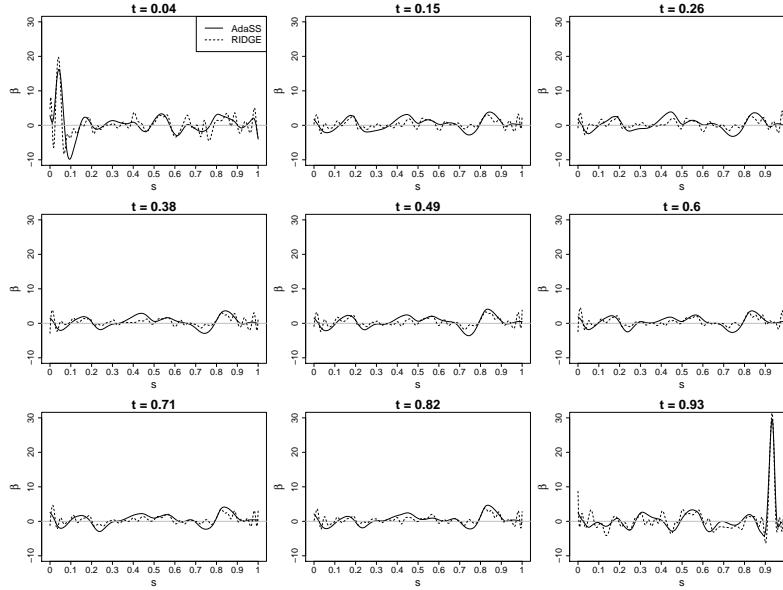


Figure 7.10. AdaSS (solid line) and RIDGE (dashed line) estimates of the coefficient functions for different values of t in the ship CO₂ emission dataset.

7.6 Appendix

Approximation of the Two Penalty Terms for the AdaSS Estimator Derivation

In this section the approximations of (7.9) and (7.10) are obtained. For the first penalty, by using (7.7), we have

$$\begin{aligned}
 & \lambda_s^{AdaSS} \int_S \int_{\mathcal{T}} \frac{1}{\left(|\widehat{\beta_s^{m_s}}(s, t)| + \delta_s \right)^{\gamma_s}} \left(\mathcal{L}_s^{m_s} \alpha(s, t) \right)^2 ds dt \\
 & \approx \lambda_s^{AdaSS} \int_S \int_{\mathcal{T}} \frac{1}{\left(\left| \sum_{i=1}^{L_s+1} \sum_{j=1}^{L_t+1} \widehat{\beta_s^{m_s}}(\tau_{s,i}, \tau_{t,j}) I_{[(\tau_{s,i-1}, \tau_{s,i}) \times (\tau_{t,j-1}, \tau_{t,j})]}(s, t) \right| + \delta_s \right)^{\gamma_s}} \left(\mathcal{L}_s^{m_s} \alpha(s, t) \right)^2 ds dt \\
 & = \lambda_s^{AdaSS} \sum_{i=1}^{L_s+1} \sum_{j=1}^{L_t+1} \int_{[\tau_{s,i-1}, \tau_{s,i}]} \int_{[\tau_{t,j-1}, \tau_{t,j}]} \frac{1}{\left(|\widehat{\beta_s^{m_s}}(\tau_{s,i}, \tau_{t,j})| + \delta_s \right)^{\gamma_s}} \left(\mathcal{L}_s^{m_s} \alpha(s, t) \right)^2 ds dt \\
 & = \lambda_s^{AdaSS} \sum_{i=1}^{L_s+1} \sum_{j=1}^{L_t+1} d_{ij}^s \int_{[\tau_{s,i-1}, \tau_{s,i}]} \int_{[\tau_{t,j-1}, \tau_{t,j}]} \left(\mathcal{L}_s^{m_s} \alpha(s, t) \right)^2 ds dt, \quad (7.16)
 \end{aligned}$$

7. ADAPTIVE SMOOTHING SPLINE ESTIMATOR FOR THE FUNCTION-ON-FUNCTION
LINEAR REGRESSION MODEL

where $d_{ij}^s = \left\{ \frac{1}{\left(|\widehat{\beta_s^{m_s}}(\tau_{s,i}, \tau_{t,j})| + \delta_s \right)^{\gamma_s}} \right\}$. Then, for (7.5), and following Ramsay and Silverman (2005), pag. 292,

$$\lambda_s^{AdaSS} \sum_{i=1}^{L_s+1} \sum_{j=1}^{L_t+1} d_{ij}^s \int_{[\tau_{s,i-1}, \tau_{s,i}]} \int_{[\tau_{t,j-1}, \tau_{t,j}]} \left(\mathcal{L}_s^{m_s} \alpha(s, t) \right)^2 ds dt = \lambda_t^{AdaSS} \sum_{i=1}^{L_s+1} \sum_{j=1}^{L_t+1} d_{ij}^t \text{Tr} \left[\mathbf{B}_\alpha^T \mathbf{W}_{s,i} \mathbf{B}_\alpha \mathbf{R}_{t,j} \right] \quad (7.17)$$

where $\mathbf{W}_{s,i} = \int_{[\tau_{s,i-1}, \tau_{s,i}]} \boldsymbol{\psi}^s(s) \boldsymbol{\psi}^s(s)^T ds$, $\mathbf{R}_{t,j} = \int_{[\tau_{t,j-1}, \tau_{t,j}]} \mathcal{L}_t^{m_t} [\boldsymbol{\psi}^t(t)] \mathcal{L}_t^{m_t} [\boldsymbol{\psi}^t(t)]^T dt$, and $d_{ij}^s = \left\{ \frac{1}{\left(|\widehat{\beta_s^{m_s}}(\tau_{s,i}, \tau_{t,j})| + \delta_s \right)^{\gamma_s}} \right\}$, for $i = 1, \dots, L_s + 1$ and $j = 1, \dots, L_t + 1$. Thus, (7.9) is demonstrated, the arguments are analogous for (7.10).

Chapter 8

Sparse and Smooth Functional Data Clustering

Abstract

A new model-based procedure is developed for sparse clustering of functional data that aims to classify a sample of curves into homogeneous groups while jointly detecting the most informative portions of domain. The proposed method is referred to as sparse and smooth functional clustering (SaS-Funclust) and relies on a general functional Gaussian mixture model whose parameters are estimated by maximizing a log-likelihood function penalized with a functional adaptive pairwise penalty and a roughness penalty. The former allows identifying the noninformative portion of domain by shrinking the means of separated clusters to some common values, whereas the latter improves the interpretability by imposing some degree of smoothing to the estimated cluster means. The model is estimated via an expectation-conditional maximization algorithm paired with a cross-validation procedure. Through a Monte Carlo simulation study, the SaS-Funclust method is shown to outperform other methods already appeared in the literature, both in terms of clustering performance and interpretability. Finally, three real-data examples are presented to demonstrate the favourable performance of the proposed method. The SaS-Funclust method is implemented in the R package `sasfunclust`, available on CRAN.

8.1 Introduction

In the last years, due to advances in technology and computational power, most of the data collected by practitioners and scientists in many fields bring information about curves or surfaces that are apt to be modelled as functional data, i.e., continuous random functions defined on a compact domain. A thorough overview of functional data analysis (FDA) techniques can be found in Ramsay and Silverman (2005); Ramsay et al. (2009); Horváth and Kokoszka (2012); Hsing and Eubank (2015) and Kokoszka and Reimherr (2017). As in the classical (non-functional) statistical literature, cluster analysis is an important topic in FDA, with many applications in various fields. The primary concern of functional clustering techniques is to classify a sample of data into homogeneous groups of curves, without having any prior knowledge about the true underlying clustering structure. The clustering of functional data is generally a difficult task because of the infinite dimensionality of the problem. For this reason, methods for functional data clustering have received a lot of

attention in recent years, and different approaches have been proposed and discussed in the last decade. To the best of authors' knowledge, the most used approach is the filtering approach (Jacques and Preda, 2014), which relies on the reduction of the infinite dimensional problem by approximating functional data in a finite dimensional space and, then, uses traditional clustering tools on the basis expansion coefficients. Along this line, Abraham et al. (2003) propose an advanced version of the k-means algorithm to the coefficients obtained by projecting the original profiles onto a lower-dimensional subspace spanned by B-spline basis functions. A similar method is proposed by Rossi et al. (2004) who apply a Self-Organizing Map (SOM) on the resulting coefficient instead of the k-means algorithm. Elaborating on this path, Serban and Wasserman (2005) present a technique for the nonparametric estimation and clustering of a large number of functional data that is still based on the k-means algorithm applied to the basis expansion coefficients obtained through smoothing techniques. A step forward is moved by Chiou and Li (2007), who introduce the k-centers functional clustering method to account, differently from the previous methods, for both the means and the mode of variation differentials between clusters by predicting cluster membership with a reclassification step.

Instead of considering the basis expansion coefficients as parameters, a different idea is that of using a model-based approach where coefficients are treated as random variables themselves with a cluster-specific probability distribution. The seminal work of James and Sugar (2003) is the first one to develop a flexible model-based procedure to cluster functional data based on a random effects model for the coefficients. This allows for borrowing strength across curves and, thus, for superior results when data contain a large number of sparsely sampled curves. More recently, Bouveyron and Jacques (2011) propose a new functional clustering method, which is referred to as funHDDC and based on a functional latent Gaussian mixture model, to fit the functional data in group-specific functional subspaces. By constraining model parameters within and between groups, they obtain a family of parsimonious models that allow for more flexibility. Analogously, Jacques and Preda (2013) assume cluster-specific Gaussian distribution on the principal components resulting from the Karhunen–Loeve expansion of the curves, and Giacomini et al. (2013) propose to use a Gaussian mixture model on the wavelet decomposition of the curves, which turns out to be particularly appropriate for peak-like data, as opposed to methods based on splines.

In the multivariate cluster analysis, some attributes could be, however, completely noninformative for uncovering the clustering structure of interest. As an example, this often happens in high-dimensional problems, i.e., where the number of variables is considerably larger than the number of observations. In this setting, the task of identifying the features in which respect true clusters differ the most is of great interest to achieve (a) a more accurate identification of the groups, as noninformative features may hide the true clustering structure, and (b) an higher interpretability of the analysis, by imputing the presence of the clustering structure to a small number of features. More in general, the methods capable of selecting informative features and eliminating noninformative ones are referred to as *sparse*. Such class of methods can be usually reconducted and regarded as variable selection methods. Sparse clustering has received increasing attention in the recent literature. Based on conventional heuristic clustering algorithms, Friedman and Meulman (2004) develop a new procedure to automatically detect subgroups of objects, which preferentially cluster on subsets of features. Witten and Tibshirani (2010) elaborate a novel clustering framework based on an adaptively chosen subset of features that are selected by means of a lasso-type penalty. In terms of model-based approaches, the method introduced by Raftery and Dean (2006) is able to sequentially compare nested models through approximate Bayes factor and

to select the informative features. Maugis et al. (2009) improve this method by considering the noninformative features as independent from some informative ones.

It is moreover worth mentioning quite promising variable selection approaches that make use of a regularization framework. The seminal work in this direction is that of Pan and Shen (2007), who introduce a penalized likelihood approach with an L_1 penalty function, which is able to automatically achieve variable selection via thresholding and delivering a sparse solution. Similarly, Wang and Zhu (2008) suggest a solution by replacing the L_1 penalty with either the L_∞ penalty or the hierarchical penalization function, which take into account the fact that cluster means, corresponding to the same feature, can be treated as grouped. Xie et al. (2008) also account for grouped parameters through the use of two planes of grouping, named vertical and horizontal grouping. In all sparse clustering methods just mentioned, a feature is selected if it is informative for at least one pair of clusters and eliminated otherwise, i.e., if it is noninformative for all clusters. However, some variables could be informative only for specific pairs of clusters. For this reason, Guo et al. (2010) propose a pairwise fusion penalty that penalizes, for each feature, the differences between all pairs of cluster means and fuses only the non separated clusters.

Only recently, the notion of sparseness has been translated into a functional data clustering framework. Specifically, sparse functional clustering methods aim to cluster the curves while jointly detecting the most informative portion of domain to the clustering in order to improve both the accuracy and the interpretability of the analysis. Based on the idea of Chen et al. (2014), Floriello and Vitelli (2017) propose a sparse functional clustering method based on the estimation of a suitable weight function that is capable of identifying the informative part of the domain. Vitelli (2019) proposes a novel framework for sparse functional clustering that also embeds an alignment step. To the best of the authors' knowledge, these are the only works that propose sparse functional clustering methods so far.

In this chapter, we present a model-based procedure for the sparse clustering of functional data, named sparse and smooth functional clustering (SaS-Funclust) method, where the basic idea is to provide accurate and interpretable cluster analysis. Specifically, the parameters of a general functional Gaussian mixture model are estimated by maximizing a penalized version of the log-likelihood function, where a functional adaptive pairwise fusion penalty, the functional extension of the penalty proposed by Guo et al. (2010), is introduced. Firstly, it penalizes the pointwise differences between all pairs of cluster functional means and locally shrinks the means of cluster pairs to some common values. Secondly, a roughness penalty on cluster functional means is considered to further improve the interpretability of the cluster analysis. Therefore, the SaS-Funclust method gains the ability to detect, for each cluster pair, the informative portion of domain to the clustering, hereinafter always intended in terms of mean differences. If a specific mean pair is fused over a portion of the domain, it is labelled as noninformative to the clustering of that pair. Otherwise, it is labelled as informative. In other words, the proposed method is able to detect portions of domain that are noninformative *pairwise*, i.e., for at least a specific cluster pair, differently from the method proposed by Floriello and Vitelli (2017) that is only able to detect portions of domain that are noninformative *overall*, i.e., for all the cluster pairs simultaneously. Moreover, the model-based fashion of the proposed method provides greater flexibility than the latter, which basically relies on k-means clustering. A specific expectation-conditional maximization (ECM) algorithm is designed to perform the maximization of the penalized log-likelihood function, which is a non-trivial problem, and a cross-validation based procedure is proposed to select the appropriate model. The method presented in this chapter is implemented in the R package `sasfuncclust`, openly available on CRAN.




Figure 8.1. True and estimated cluster means obtained through the SaS-Funclust method for three different simulated data sets with (a) two, (b) three and (c) four clusters generated as described in Section 8.3.

To give a general idea of the sparseness property of the proposed method, Figure 8.1 shows the cluster means estimated by the latter for three different simulated data sets with (a) two, (b) three and (c) four clusters. Data are generated as described in Section 8.3 and supplementary materials. In Figure 8.1(a), the estimated means are correctly fused over $t \in (0.2, 1.0]$. Hence, the proposed method is shown to be able to identify the informative portion of domain $[0.0, 0.2]$, for the unique pair of clusters and not for all. In Figure 8.1(b) and Figure 8.1(c), several cluster pairs are available, because the number of clusters is larger than 2, and, thus, a given portion of domain could be informative for a specific pair of clusters. In Figure 8.1(b), the informative portion of domain for each pair of clusters is correctly recovered. The estimated cluster means are indeed pairwise fused over approximately the same portion of domain as the true cluster means pairs. Note that, for the clusters whose true means are equal over $t \in (0.2, 1.0]$, the SaS-Funclust method identifies the informative portion of domain roughly in $[0.0, 0.2]$. In Figure 8.1(c), the sparseness property of the SaS-Funclust method is even more striking. In this case, in the face of many cluster pairs, the proposed method is still able to successfully detect the informative portion of domain. The properties of the proposed method will be deepened in Section 8.3.

The remainder of this chapter is organized as follows. Section 8.2 presents the proposed methodology. Specifically, Section 8.2 and 8.2 introduce the general functional Gaussian mixture model and the penalized maximum likelihood estimator, respectively. Whereas, the optimization algorithm and model selection considerations are discussed in Section 8.2 and Section 8.2, respectively. Properties and performance of the SaS-Funclust method are assessed through a wide Monte Carlo simulation study presented in Section 8.3. Section 8.4 illustrates the potentiality of the SaS-Funclust method by means of of three real datasets: the Berkeley Growth Study, the Canadian weather, and the ICOSAF project data.

8.2 The SaS-Funclust Method for Functional Clustering

A General Functional Clustering Model

The SaS-Funclust method is based on the general functional clustering model introduced by James and Sugar (2003). Suppose that N observations are spread among G unknown clusters, and that the probability of each observation to belong to the g th cluster is π_g , $\sum_{g=1}^G \pi_g = 1$. Moreover, let us denote with $\mathbf{Z}_i = (Z_{1i}, \dots, Z_{Gi})^T$ the unknown component-label vector

corresponding to the i th observation, where Z_{gi} equals 1 if the i th observation is in the g th cluster and 0 otherwise. Then, let us assume that for each i observation, $i = 1, \dots, N$ in the cluster $g = 1, \dots, G$, it is available a vector $\mathbf{Y}_i = (y_{i1}, \dots, y_{in_i})^T$ of size n_i , which can differ across observations, of observed values of a function g_i over the time points t_{i1}, \dots, t_{in_i} . The function g_i is assumed a Gaussian random process with mean μ_g , covariance ω_g , and values in $L^2(\mathcal{T})$, the separable Hilbert space of square integrable functions defined on the compact domain \mathcal{T} . We assume that, conditionally on that $Z_{gi} = 1$, \mathbf{Y}_i is modelled as

$$\mathbf{Y}_i = \mathbf{g}_i + \boldsymbol{\epsilon}_i, \quad i = 1, \dots, N,$$

where $\mathbf{g}_i = (g_i(t_{i1}), \dots, g_i(t_{in_i}))^T$ contains the values of the function g_i at t_{i1}, \dots, t_{in_i} and $\boldsymbol{\epsilon}_i$ is a vector of measurement errors that are mutually independent and normally distributed with mean 0 and constant variance σ_e^2 . Let us suppose also that the unknown component-label vector \mathbf{Z}_i has a multinomial distribution, which consists of one draw on g categories with probabilities π_1, \dots, π_G . Then, for every i , the unconditional density function $f(\cdot)$ of \mathbf{Y}_i is

$$f(\mathbf{Y}_i) = \sum_{g=1}^G \pi_g \psi(\mathbf{Y}_i; \boldsymbol{\mu}_{gi}, \boldsymbol{\Omega}_{gi} + \mathbf{I}\sigma_e^2), \quad (8.1)$$

where $\boldsymbol{\mu}_{gi} = (\mu_g(t_{i1}), \dots, \mu_g(t_{in_i}))^T$, $\boldsymbol{\Omega}_{gi} = \{\omega_g(t_{ki}, t_{li})\}_{k,l=1,\dots,n_i}$, \mathbf{I} is the identity matrix, and $\psi(\cdot; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ is the multivariate Gaussian density distribution with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$. The model in Equation (8.1) is the classical G -component Gaussian mixture model (McLachlan and Peel, 2004).

As discussed in James and Sugar (2003), it is necessary to impose some structure curves g_i because both the curves could be observed at different time domain points and the dimensionality of the model in Equation (8.1) could be too high in comparison to the sample size. Therefore, similarly to the filtering approach for clustering, we assume that each function g_i , for $i = 1, \dots, N$, may be represented in terms of a q -dimensional set of basis functions $\boldsymbol{\Phi} = (\phi_1, \dots, \phi_q)^T$, that is

$$g_i(t) = \boldsymbol{\eta}_i^T \boldsymbol{\Phi}(t), \quad t \in \mathcal{T}, \quad (8.2)$$

where $\boldsymbol{\eta}_i = (\eta_{i1}, \dots, \eta_{iq})^T$ are vectors of basis coefficients. Then, $\boldsymbol{\eta}_i$ are modelled as Gaussian random vectors, that is, given that $Z_{gi} = 1$,

$$\boldsymbol{\eta}_i = \boldsymbol{\mu}_g + \boldsymbol{\gamma}_{ig}, \quad (8.3)$$

where $\boldsymbol{\mu}_g = (\mu_{g1}, \dots, \mu_{gq})^T$ are q -dimensional vectors and $\boldsymbol{\gamma}_{ig}$ are Gaussian random vectors with zero mean and covariance $\boldsymbol{\Gamma}_g$. With these assumption the unconditional density function $f(\cdot)$ of \mathbf{Y}_i in Equation (8.1) becomes

$$f(\mathbf{Y}_i) = \sum_{g=1}^G \pi_g \psi(\mathbf{Y}_i; \mathbf{S}_i \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_{ig}),$$

where $\mathbf{S}_i = (\boldsymbol{\Phi}(t_{i1}), \dots, \boldsymbol{\Phi}(t_{in_i}))^T$ is the basis matrix for the i th curve and $\boldsymbol{\Sigma}_{ig} = \mathbf{S}_i \boldsymbol{\Gamma}_g \mathbf{S}_i^T + \mathbf{I}\sigma_e^2$. Therefore, the log-likelihood function corresponding to $\mathbf{Y}_1, \dots, \mathbf{Y}_N$ is given by

$$L(\boldsymbol{\Theta} | \mathbf{Y}_1, \dots, \mathbf{Y}_N) = \sum_{i=1}^N \log \sum_{g=1}^G \pi_g \psi(\mathbf{Y}_i; \mathbf{S}_i \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_{ig}), \quad (8.4)$$

where $\Theta = \{\pi_g, \mu_g, \Gamma_g, \sigma_e^2\}_{g=1, \dots, G}$ is the parameter set of interest. Based on an estimate $\hat{\Theta} = \{\hat{\pi}_g, \hat{\mu}_g, \hat{\Gamma}_g, \hat{\sigma}_e^2\}_{g=1, \dots, G}$, an observation \mathbf{Y}^* is assigned to the cluster g that achieves the largest posterior probability estimate $\hat{\pi}_g \psi(\mathbf{Y}^*; \mathbf{S}_i \hat{\mu}_g, \hat{\Sigma}_{ig})$, with $\hat{\Sigma}_{ig} = \mathbf{S}_i \hat{\Gamma}_g \mathbf{S}_i^T + \mathbf{I} \hat{\sigma}_e^2$.

The Penalized Maximum Likelihood Estimator

James and Sugar (2003) propose to estimate Θ through the maximum likelihood estimator (MLE), which is the maximizer of the log-likelihood function in Equation (8.4). In this work, we propose a different estimator $\hat{\Theta}_{PMLE}$ of Θ that is the maximizer of the following penalized log-likelihood

$$L_p(\Theta | \mathbf{Y}_1, \dots, \mathbf{Y}_N) = \sum_{i=1}^N \log \sum_{g=1}^G \pi_g \psi(\mathbf{Y}_i; \mathbf{S}_i \mu_g, \Sigma_{ig}) - \mathcal{P}(\mu_g), \quad (8.5)$$

where $\mathcal{P}(\cdot)$ is a penalty function defined as

$$\mathcal{P}(\mu_g) = \lambda_L \sum_{1 \leq g \leq g' \leq G} \int_{\mathcal{T}} \tau_{g,g'}(t) |\mu_g(t) - \mu_{g'}(t)| dt + \lambda_s \sum_{g=1}^G \int_{\mathcal{T}} \left(\mu_g^{(s)}(t) \right)^2 dt, \quad (8.6)$$

where $\lambda_L, \lambda_s \geq 0$ are tuning parameters, and $\tau_{g,g'}$ are prespecified weight functions. The symbol $f^{(s)}(\cdot)$ denotes the s th-order derivative of f if the latter a function or the entries of f if it is a vector. Note that in Equation (8.6) each function g_i may be represented as in Equation (8.2), then it follows that

$$\mathcal{P}(\mu_g) = \lambda_L \sum_{1 \leq g \leq g' \leq G} \int_{\mathcal{T}} \tau_{g,g'}(t) |\mu_g^T \Phi(t) - \mu_{g'}^T \Phi(t)| dt + \lambda_s \sum_{g=1}^G \int_{\mathcal{T}} \left(\mu_g^T \Phi^{(s)}(t) \right)^2 dt, \quad (8.7)$$

The first element of the right-hand side of Equation (8.6) is the functional extension of the penalty introduced by Guo et al. (2010) and is referred to as functional adaptive pairwise fusion penalty (FAPFP). The aim of the FAPFP is to shrink the differences between every pair of cluster means for each value of $t \in \mathcal{T}$. Due to the property of the absolute value function of being singular at zero, some of these differences are shrunk exactly to zero. In particular, the FAPFP allows pair of cluster means to be equal over specific portion of domain that are, thus, noninformative for separating the means of that pair of clusters.

The choice of the weight function $\tau_{g,g'}$ in Equation (8.6) and Equation (8.6) should be based on the idea that if a given portion of domain is informative for separating the means of the corresponding pair of clusters, then, the values of $\tau_{g,g'}$ over that portion should be small. In this way, the absolute difference $|\mu_g(\cdot) - \mu_{g'}(\cdot)|$ is penalized more over noninformative portions of domain than over informative ones. Following the standard practice for the adaptive penalties (Zou, 2006), we propose to use

$$\tau_{g,g'}(t) = |\tilde{\mu}_g(t) - \tilde{\mu}_{g'}(t)|^{-1} \quad t \in \mathcal{T}, \quad (8.8)$$

where $\tilde{\mu}_g$ are initial estimates of the cluster means.

Finally, the term $\lambda_s \sum_{g=1}^G \int_{\mathcal{T}} \left(\mu_g^{(s)}(t) \right)^2 dt$ is a smoothness penalty that penalizes the s th derivative of the cluster means. This term aims to further improve the interpretability of the results by constraining, of a magnitude quantified by λ_s , the cluster means to own a certain

degree of smoothness, measured by the derivative order s . Following the common practice in FDA (Ramsay and Silverman, 2005), we suggest to penalize the cluster mean curvature by setting $s = 2$, which implies that the basis functions chosen are differentiable at least s times. As a remark, the penalization in Equation (8.5) is applied only to the parameter vectors corresponding to the cluster means, i.e., to $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_G$. The reason is that, in this work, we consider the case where a portion domain is defined as informative only in terms of cluster mean differences. However, portions of domain could be informative also in terms of differences in cluster covariances, which together with the means uniquely identify each cluster.

The Penalty Approximation and the Optimization Algorithm

To perform the maximization of the penalized log-likelihood in Equation (8.5), the penalty $\mathcal{P}(\cdot)$, defined as in Equation (8.6), can be written as

$$\mathcal{P}(\boldsymbol{\mu}_g) = \lambda_L \sum_{1 \leq g \leq g' \leq G} \int_{\mathcal{T}} |(\tilde{\boldsymbol{\mu}}_g^T - \tilde{\boldsymbol{\mu}}_{g'}^T)^T \boldsymbol{\Phi}(t)|^{-1} |(\boldsymbol{\mu}_g^T - \boldsymbol{\mu}_{g'}^T)^T \boldsymbol{\Phi}(t)| dt + \lambda_s \sum_{g=1}^G \boldsymbol{\mu}_g^T \mathbf{W} \boldsymbol{\mu}_g, \quad (8.9)$$

where the weight functions $\tau_{g,g'}(t)$ are expressed as in Equation (8.8), and the initial estimates of the cluster means are represented through the set of basis functions $\boldsymbol{\Phi}$ as $\tilde{\boldsymbol{\mu}}_g(t) = \tilde{\boldsymbol{\mu}}_g^T \boldsymbol{\Phi}(t)$, $t \in \mathcal{T}$, with $\tilde{\boldsymbol{\mu}}_g = (\tilde{\mu}_{g1}, \dots, \tilde{\mu}_{gq})^T$. The matrix \mathbf{W} is equal to $\int_{\mathcal{T}} \boldsymbol{\Phi}^{(s)}(t) (\boldsymbol{\Phi}^{(s)}(t))^T dt$. A great simplification of the optimization problem can be achieved if the first element on the right-hand side of Equation (8.9) can be expressed as linear function of $|\boldsymbol{\mu}_g^T - \boldsymbol{\mu}_{g'}^T|$. The following theorem provides a practical way to rewrite the first term of the right-hand side of Equation (8.9) as linear function of $|\boldsymbol{\mu}_g^T - \boldsymbol{\mu}_{g'}^T|$, when $\boldsymbol{\Phi}$ is a set of B-splines (De Boor et al., 1978; Schumaker, 2007).

Theorem 4. *Let $\boldsymbol{\Phi} = (\phi_1, \dots, \phi_q)^T$ be the set of B-splines of order k and non-decreasing knots sequences $\{x_0, x_1, \dots, x_{M_j}, x_{M+1}\}$ defined on the compact set $\mathcal{T} = [x_0, x_{M+1}]$, with $q = M + k$, and $\{\tau_j\}_{j=1}^{q+1}$ being a sequence with $\tau_1 = x_0$, $\tau_j = \tau_{j-1} + (x_{\min(M+1,j)} - x_{\max(0,j-1-k)})/k$, $\tau_{q+1} = x_{M+1}$. Then, for each function $f(t) = \sum_{i=1}^q c_i \phi_i(t)$, $t \in \mathcal{T}$, where $c_i \in \mathbb{R}$, the function $\tilde{f}(t) = \sum_{i=1}^q c_i I_{[\tau_i, \tau_{i+1}]}(t)$, $t \in \mathcal{T}$, where $I_{[\tau_i, \tau_{i+1}]}(t) = 1$ for $t \in [\tau_i, \tau_{i+1}]$ and zero elsewhere, is such that*

$$\sup_{t \in \mathcal{T}} |f(t) - \tilde{f}(t)| = O(\delta), \quad (8.10)$$

where $\delta = \max_i |x_{i+1} - x_i|$, that is $f(t) - \tilde{f}(t)$ converges uniformly to the zero function.

Theorem 4, whose proof is deferred to the supplementary materials, basically states that when δ is small, f is well approximated by \tilde{f} . In other words, the approximation error $|f - \tilde{f}|$ can be made arbitrarily small by increasing the number of knots. If we further assume the knots sequence evenly spaced, δ turns out to be equal to $1/M$. These considerations allow us to approximate $|(\boldsymbol{\mu}_g^T - \boldsymbol{\mu}_{g'}^T)^T \boldsymbol{\Phi}(t)|$ and $|(\tilde{\boldsymbol{\mu}}_g^T - \tilde{\boldsymbol{\mu}}_{g'}^T)^T \boldsymbol{\Phi}(t)|$, respectively, as follows

$$|(\boldsymbol{\mu}_g - \boldsymbol{\mu}_{g'})^T \boldsymbol{\Phi}(t)| \approx |\boldsymbol{\mu}_g - \boldsymbol{\mu}_{g'}|^T \mathbf{I}(t), \quad |(\tilde{\boldsymbol{\mu}}_g - \tilde{\boldsymbol{\mu}}_{g'})^T \boldsymbol{\Phi}(t)| \approx |\tilde{\boldsymbol{\mu}}_g - \tilde{\boldsymbol{\mu}}_{g'}|^T \mathbf{I}(t), \quad \forall t \in \mathcal{T}, \quad (8.11)$$

where $\mathbf{I} = \left(I_{[\tau_1, \tau_2]}, \dots, I_{[\tau_q, \tau_{q+1}]} \right)^T$. Thus, Equation (8.9) can be rewritten as

$$\mathcal{P}(\boldsymbol{\mu}_g) = \lambda_L \sum_{1 \leq g \leq g' \leq G} \tilde{\mathbf{M}} |\boldsymbol{\mu}_g - \boldsymbol{\mu}_{g'}| + \lambda_s \sum_{g=1}^G \boldsymbol{\mu}_g^T \mathbf{W} \boldsymbol{\mu}_g, \quad (8.12)$$

where $\tilde{\mathbf{M}} = \text{diag} \left(\frac{\tau_2 - \tau_1}{|\tilde{\mu}_{g1} - \tilde{\mu}_{g'1}|}, \dots, \frac{\tau_2 - \tau_1}{|\tilde{\mu}_{gq} - \tilde{\mu}_{g'q}|} \right)$ is the diagonal matrix with diagonal entries $\frac{\tau_2 - \tau_1}{|\tilde{\mu}_{g1} - \tilde{\mu}_{g'1}|}, \dots, \frac{\tau_2 - \tau_1}{|\tilde{\mu}_{gq} - \tilde{\mu}_{g'q}|}$.

The goodness of the approximations in Equation (8.11) depends on the cardinality q of the set of B-splines Φ , which, thus, should be as large as possible. However, the number of parameters in Equation (8.1), which depends quadratically on q , becomes very large even for moderate values of q . To mitigate this issue, one may further assume equal and diagonal coefficient covariance matrices across all clusters, that is $\mathbf{\Gamma}_1 = \dots = \mathbf{\Gamma}_G = \mathbf{\Gamma} = \text{diag}(\sigma_1^2, \dots, \sigma_q^2)$. As a remark, with this assumption, we are implicitly assuming that clusters are separated only by their mean values, and, thus, informative portion of domain are identified only by cluster mean differences and not in terms of covariances.

The penalized log-likelihood function in Equation (8.5) becomes

$$L_p(\boldsymbol{\Theta} | \mathbf{Y}_1, \dots, \mathbf{Y}_N) = \sum_{i=1}^N \log \sum_{g=1}^G \pi_g \psi(\mathbf{Y}_i; \mathbf{S}_i \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_i) - \lambda_L \sum_{1 \leq g \leq g' \leq G} \tilde{\mathbf{M}} |\boldsymbol{\mu}_g - \boldsymbol{\mu}_{g'}| - \lambda_s \sum_{g=1}^G \boldsymbol{\mu}_g^T \mathbf{W} \boldsymbol{\mu}_g, \quad (8.13)$$

with $\boldsymbol{\Sigma}_i = \mathbf{S}_i \mathbf{\Gamma} \mathbf{S}_i^T + \mathbf{I} \sigma_e^2$. The maximization of this objective function is a nontrivial problem. A specifically designed algorithm is proposed, which is a modification of the expectation maximization (EM) algorithm proposed by James and Sugar (2003). By treating the component-label vectors \mathbf{Z}_i (defined at the beginning of Section 8.2) and γ_{ig} in Equation (8.3) as missing data, the complete penalized log-likelihood is given by

$$L_{cp}(\boldsymbol{\Theta} | \mathbf{Y}_1, \dots, \mathbf{Y}_N) = \sum_{i=1}^N \sum_{g=1}^G Z_{gi} \left[\log \pi_g + \log \psi(\gamma_{ig}, 0, \mathbf{\Gamma}) + \log \psi(\mathbf{Y}_i; \mathbf{S}_i (\boldsymbol{\mu}_g + \gamma_{ig}), \sigma_e^2 \mathbf{I}) \right] - \lambda_L \sum_{1 \leq g \leq g' \leq G} \tilde{\mathbf{M}} |\boldsymbol{\mu}_g - \boldsymbol{\mu}_{g'}| + \lambda_s \sum_{g=1}^G \boldsymbol{\mu}_g^T \mathbf{W} \boldsymbol{\mu}_g. \quad (8.14)$$

The EM algorithm consists in the maximization, at each iteration $t = 0, 1, 2, \dots$, of the expected value of L_{cp} , calculated with respect the joint distribution of \mathbf{Z}_i and γ_{ig} , given $\mathbf{Y}_1, \dots, \mathbf{Y}_N$ and the current parameter estimates $\hat{\boldsymbol{\Theta}}^{(t)} = \{\hat{\pi}_g^{(t)}, \hat{\boldsymbol{\mu}}_g^{(t)}, \hat{\mathbf{\Gamma}}^{(t)} = \text{diag}(\hat{\sigma}_1^{2(t)}, \dots, \hat{\sigma}_q^{2(t)})\}_{g=1, \dots, G}$. The algorithm stops when a pre specified stopping condition is met. At each t , the expected value of L_{cp} as a function the probability of membership π_1, \dots, π_G is then maximized by setting

$$\hat{\pi}_g^{(t+1)} = \frac{1}{N} \sum_{i=1}^N \hat{\pi}_{g|i}^{(t+1)},$$

with $\hat{\pi}_{g|i}^{(t+1)} = \mathbb{E} \left(Z_{ig} = 1 | \mathbf{Y}_i, \hat{\boldsymbol{\Theta}}^{(t)} \right) = \frac{\hat{\pi}^{(t)} \psi(\mathbf{Y}_i; \mathbf{S}_i \hat{\boldsymbol{\mu}}_g^{(t)}, \hat{\boldsymbol{\Sigma}}_i^{(t)})}{\sum_{g'=1}^G \hat{\pi}_{g'}^{(t)} \psi(\mathbf{Y}_i; \mathbf{S}_i \hat{\boldsymbol{\mu}}_{g'}^{(t)}, \hat{\boldsymbol{\Sigma}}_{g'}^{(t)})}$. With respect to $\sigma_1^2, \dots, \sigma_q^2$,

L_{cp} is maximized by

$$\hat{\sigma}_j^{2(t+1)} = \frac{1}{N} \sum_{i=1}^N \sum_{g=1}^G \hat{\pi}_{g|i}^{(t+1)} \mathbb{E} \left(\gamma_{ig(j)}^2 | \mathbf{Y}_i, Z_{gi} = 1, \hat{\boldsymbol{\Theta}}^{(t)} \right) \quad j = 1, \dots, q,$$

where $\gamma_{ig(j)}^2$ indicates the j th entry of $\boldsymbol{\gamma}_{ig}^2$. The value of $\mathbb{E} \left(\gamma_{ig(j)}^2 | \mathbf{Y}_i, Z_{gi} = 1, \hat{\boldsymbol{\Theta}}^{(t)} \right)$ can be calculated by using the property that the (conditional) distribution of $\boldsymbol{\gamma}_{ig}$ given $\mathbf{Y}_i, Z_{gi} = 1, \hat{\boldsymbol{\Theta}}^{(t)}$ is Gaussian with mean $\hat{\mathbf{\Gamma}}^{(t)} \mathbf{S}_i^T \left(\mathbf{S}_i \hat{\mathbf{\Gamma}}^{(t)} \mathbf{S}_i^T + \mathbf{I} \hat{\sigma}^{2(t)} \right)^{-1} \left(\mathbf{Y}_i - \mathbf{S}_i \hat{\boldsymbol{\mu}}_g^{(t)} \right)$ and covariance $\hat{\mathbf{\Gamma}}^{(t)} - \hat{\mathbf{\Gamma}}^{(t)} \mathbf{S}_i^T \left(\mathbf{S}_i \hat{\mathbf{\Gamma}}^{(t)} \mathbf{S}_i^T + \mathbf{I} \hat{\sigma}^{2(t)} \right)^{-1} \mathbf{S}_i \hat{\mathbf{\Gamma}}^{(t)}$. Then, σ_e^2 is updated as

$$\begin{aligned} \hat{\sigma}_e^{2(t+1)} = \frac{1}{\sum_{i=1}^N n_i} \sum_{i=1}^N \sum_{g=1}^G \left[\hat{\pi}_{g|i}^{(t+1)} \left(\mathbf{Y}_i - \mathbf{S}_i \hat{\boldsymbol{\mu}}_g^{(t)} - \mathbf{S}_i \hat{\boldsymbol{\gamma}}_{ig}^{(t)} \right) \right]^T \left(\mathbf{Y}_i - \mathbf{S}_i \hat{\boldsymbol{\mu}}_g^{(t)} - \mathbf{S}_i \hat{\boldsymbol{\gamma}}_{ig}^{(t)} \right) \\ - \mathbf{S}_i \text{Cov} \left(\boldsymbol{\gamma}_{ig} | \mathbf{Y}_i, Z_{gi} = 1, \hat{\boldsymbol{\Theta}}^{(t)} \right) \mathbf{S}_i^T \right], \end{aligned}$$

where $\hat{\boldsymbol{\gamma}}_{ig}^{(t)} = \mathbb{E} \left(\boldsymbol{\gamma}_{ig} | \mathbf{Y}_i, Z_{gi} = 1, \hat{\boldsymbol{\Theta}}^{(t)} \right)$.

The mean vectors $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_G$ that maximize the conditional expectation of L_{cp} are the solution of the following optimization problem

$$\begin{aligned} \hat{\boldsymbol{\mu}}_1^{(t+1)}, \dots, \hat{\boldsymbol{\mu}}_G^{(t+1)} = \underset{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_G}{\text{argmin}} \frac{1}{2} \sum_{i=1}^N \sum_{g=1}^G \hat{\pi}_{g|i}^{(t+1)} \frac{1}{\hat{\sigma}_e^{(t)}} \left(\mathbf{Y}_i - \mathbf{S}_i \left(\boldsymbol{\mu}_g + \hat{\boldsymbol{\gamma}}_{ig}^{(t)} \right) \right)^T \left(\mathbf{Y}_i - \mathbf{S}_i \left(\boldsymbol{\mu}_g + \hat{\boldsymbol{\gamma}}_{ig}^{(t)} \right) \right) \\ + \lambda_L \sum_{1 \leq g \leq g' \leq G} \tilde{\mathbf{M}} |\boldsymbol{\mu}_g - \boldsymbol{\mu}_{g'}| + \lambda_s \sum_{g=1}^G \boldsymbol{\mu}_g^T \mathbf{W} \boldsymbol{\mu}_g. \quad (8.15) \end{aligned}$$

The optimization problem in Equation (8.15) is a difficult task of the non differentiability of the absolute value function in zero, and, it has not a closed form solution. However, following the idea of Fan and Li (2001), it can be solved by means of the standard local quadratic approximation method, i.e., by iteratively solving the following quadratic optimization problem for $s = 0, 1, 2, \dots$

$$\begin{aligned} \hat{\boldsymbol{\mu}}_1^{(t+1,s+1)}, \dots, \hat{\boldsymbol{\mu}}_G^{(t+1,s+1)} = \underset{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_G}{\text{argmin}} \frac{1}{2} \sum_{i=1}^N \sum_{g=1}^G \hat{\pi}_{g|i}^{(t+1)} \frac{1}{\hat{\sigma}_e^{(t)}} \left(\mathbf{Y}_i - \mathbf{S}_i \left(\boldsymbol{\mu}_g + \hat{\boldsymbol{\gamma}}_{ig}^{(t)} \right) \right)^T \left(\mathbf{Y}_i - \mathbf{S}_i \left(\boldsymbol{\mu}_g + \hat{\boldsymbol{\gamma}}_{ig}^{(t)} \right) \right) \\ + \lambda_L \sum_{1 \leq g \leq g' \leq G} |\boldsymbol{\mu}_g - \boldsymbol{\mu}_{g'}|^T \tilde{\mathbf{M}} \mathbf{D}^{(s)} |\boldsymbol{\mu}_g - \boldsymbol{\mu}_{g'}| + \lambda_s \sum_{g=1}^G \boldsymbol{\mu}_g^T \mathbf{W} \boldsymbol{\mu}_g, \quad (8.16) \end{aligned}$$

where $\mathbf{D}^{(s)} = \text{diag} \left(\frac{1}{2|\hat{\mu}_{g1}^{(t+1,s)} - \hat{\mu}_{g'1}^{(t+1,s)}|}, \dots, \frac{1}{2|\hat{\mu}_{gq}^{(t+1,s)} - \hat{\mu}_{g'q}^{(t+1,s)}|} \right)$, and $\hat{\boldsymbol{\mu}}_1^{(t+1,0)} = \hat{\boldsymbol{\mu}}_1^{(t)}, \dots, \hat{\boldsymbol{\mu}}_G^{(t+1,0)} = \hat{\boldsymbol{\mu}}_G^{(t)}$. Equation (8.16) is based on the following approximation (Fan and Li, 2001)

$$|\mu_{gi} - \mu_{g'i}| \approx \frac{|\mu_{gi} - \mu_{g'i}|}{2|\hat{\mu}_{gq}^{(t+1,s)} - \hat{\mu}_{g'q}^{(t+1,s)}|} + \frac{1}{2} |\hat{\mu}_{gq}^{(t+1,s)} - \hat{\mu}_{g'q}^{(t+1,s)}|. \quad (8.17)$$

The solution to the original problem in Equation (8.15) can be satisfactorily approximated by the solution at iteration s^* of the optimization problem in Equation (8.17) when a pre specified stopping condition is met, i.e., $\hat{\boldsymbol{\mu}}_1^{(t+1)} = \hat{\boldsymbol{\mu}}_1^{(t+1,s^*)}, \dots, \hat{\boldsymbol{\mu}}_G^{(t+1)} = \hat{\boldsymbol{\mu}}_G^{(t+1,s^*)}$. For numerical stability, a reasonable suggestion is to set a lower bound on $|\hat{\mu}_{gi}^{(t+1,s)} - \hat{\mu}_{g'i}^{(t+1,s)}|$, and then to shrink to zero all the estimates below the lower bound. It is worth noting that the proposed modification to the algorithm of James and Sugar (2003) falls within the class of the expectation conditional maximization (ECM) algorithms (Meng and Rubin, 1993). Based on the convergence property of the ECM algorithms, which also holds for the local quadratic approximation in variable selection problems (Fan and Li, 2001; Hunter and Li, 2005), the proposed algorithm can be proved to converge to a stationary point of the penalized log-likelihood in Equation (8.13).

Model Selection

The proposed SaS-Funclust method requires the choice several hyper-parameters viz., the number of clusters G , tuning parameters λ_s, λ_L , dimension q and the order k of the set of B-spline functions Φ as well as the knot locations. A standard choice for Φ is the cubic B-splines (i.e., $k = 4$) with equally spaced knot sequence, which enjoy the optimal property of interpolation (De Boor et al., 1978). Moreover, the dimension q should be set as large as possible to reduce, to the greatest possible extent, the approximation error in Equation (8.11). This facilitates the estimated cluster means to successfully capture the local feature of the true cluster means. Unfortunately, the larger the value of q , the higher the complexity of the model in Equation (8.1), i.e., the number of parameters to estimate. The presence of the smoothness penalty on $\boldsymbol{\mu}_g$, as well as the constraint imposed on $\boldsymbol{\Gamma}_g$, allows one to control the complexity of the model and, thus, to prevent over-fitting issues. The choice of G, λ_s , and λ_L may be based on a K -fold cross-validation procedure. Based on observations divided into K equal-sized disjoint subsets $f_1, \dots, f_k, \dots, f_K$, G, λ_s , and λ_L are chosen as the maximizers of the following function

$$CV(G, \lambda_s, \lambda_L) = \frac{1}{K} \sum_{k=1}^K \sum_{i \in f_k} \log \sum_{g=1}^G \hat{\pi}_g^{-f_k} \psi \left(\mathbf{Y}_i; \mathbf{S}_i \hat{\boldsymbol{\mu}}_g^{-f_k}, \hat{\boldsymbol{\Sigma}}_i^{-f_k} \right), \quad (8.18)$$

where $\hat{\pi}_g^{-f_k}, \hat{\boldsymbol{\mu}}_g^{-f_k}$ and $\hat{\boldsymbol{\Sigma}}_i^{-f_k}$ denote the SaS-Funclust estimates of $\pi_g, \boldsymbol{\mu}_g$ and $\boldsymbol{\Sigma}_i$ obtained by leaving out the observations in the k -th subset f_k . Usually, the CV function is numerically calculated over a finite grid of values. As in the multivariate regression setting, the uncertainty of the CV function in estimating the log-likelihood observed for an out-of-sample observation is taken into account by means of the so called m -standard deviation rule. This heuristic rule suggests to pick up the most parsimonious model among those achieving values of the CV function that are no more than m standard errors below the maximum of Equation (8.18). Note that, in this problem, parsimony is reflected into large λ_s, λ_L and small G . By elaborating on the m -standard deviation rule, we propose to firstly choose G for each value of λ_s, λ_L , with $m = m_1$; secondly, at fixed G , choose λ_s for each λ_L , with $m = m_2$; thirdly, to choose λ_L at fixed λ_s and G , by using $m = m_3$. In this way, the estimated model is not unnecessarily complex and achieves predictive performance that is comparable to that of the best model (i.e., the one that maximizes the CV function in Equation (8.18)). As a remark, although the component-wise procedure proposed to choose λ_s, λ_L and G proves itself to be very effective in the simulation study of Section 8.3, we recommend whenever possible

to directly plot and inspect the CV curve as a function of G , λ_s , and λ_L and to use any information available from the specific application.

8.3 Simulation Study

In this section, the performance of the SaS-Funclust method is assessed by means of an extensive Monte Carlo simulation study. The SaS-Funclust method, implemented through the R package `sasfunclust`, is compared with the following methods that have already appeared in the literature before. In particular, we refer to the method proposed by Giacomini et al. (2013) as `curvclust`, and to that proposed by Bouveyron and Jacques (2011) as `funHDDC`. These methods are implemented through the homonymous R packages `curvclust` (Giacomini et al., 2012) and `funHDDC` (Schmutz and Bouveyron, 2019). In addition, we consider also filtering approaches based on two main steps. The first step consists in the estimation of the functions g_i by means of either smoothing B-splines or functional principal component analysis (Ramsay and Silverman, 2005); whereas the second step aims to apply standard clustering algorithms, viz. hierarchical, k-means and finite mixture model clustering methods (Everitt et al., 2011), on either the resulting B-spline coefficients or the functional principal components scores. Filtering approaches based on the smoothing B-splines and the hierarchical, k-means and finite mixture model clustering methods will be hereinafter referred to as B-HC, B-KM and B-FMM, respectively, whereas methods based on the functional principal component analysis and the hierarchical, k-means and finite mixture model clustering methods are referred to as FPCA-HC, FPCA-KM and FPCA-FMM. Finally, we evaluate also the method presented by Ieva et al. (2013), which is referred to as DIS-KM and it basically consists in the application of the k-means clustering to the L^2 distances among the observed curves. The number of clusters is selected through the Bayesian information criterion (BIC) for the `curvclust` and `funHDDC` methods, as suggested by Giacomini et al. (2013) and Bouveyron and Jacques (2011), respectively; whereas the silhouette index (Rousseeuw, 1987) is used for the DIS-KM method. The majority rule applied to several validity indices (Charrad et al., 2012) is used to determine the number of clusters for all the filtering approaches. The number of clusters and the tuning parameters needed to implement the SaS-Funclust method are determined through the CV based procedure described in Section 8.2 with $q = 30$, $K = 5$, $m_1 = m_3 = 0.5$, and $m_2 = 0$. The values of m_1 and m_3 ensure parsimony in the choice of λ_L and G , whereas for picking λ_s the m -standard deviation is not applied. The initial values of the parameters for the ECM algorithm are chosen by applying the k-means algorithm on the coefficients estimated through smoothing B-spline.

The performance of the clustering procedures in selecting the proper number of clusters and identifying the clustering structure, when the true number of cluster is known, is assessed separately. In particular, the former is measured through the mean number of selected clusters, whereas the latter is compared through the adjusted Rand index (Hubert and Arabie, 1985) denoted by *aRand*. This index accounts for the agreement between true data partitions and clustering results corrected by chance, based on the number of paired objects that are either in the same group or in different groups in both partitions. The *aRand* yields values between 0 and 1. The larger its value, the higher the similarity between the two partitions.

Three different scenarios are analysed where data are generated from $G_t = 2, 3, 4$ clusters and referred to as Scenario I, II and III, respectively. For each scenario, the considered methods are evaluated by assessing the performance over 100 independently simulated datasets where measurement errors are generated with five different values of standard

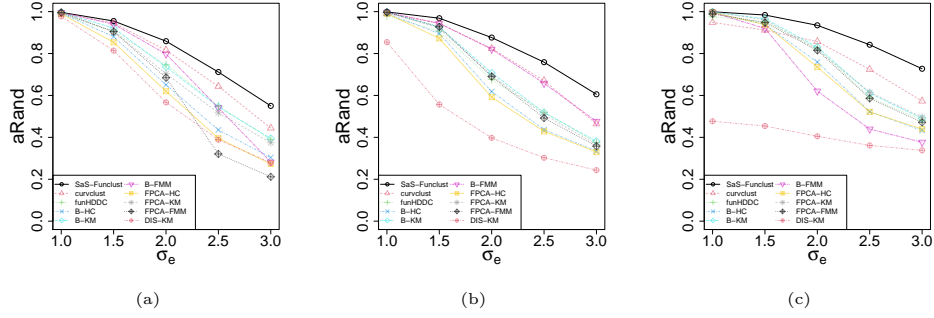


Figure 8.2. Average $aRand$ index for (a) Scenario I, (b) Scenario II, and (c) Scenario III as a function of σ_e when the true number of clusters is known.

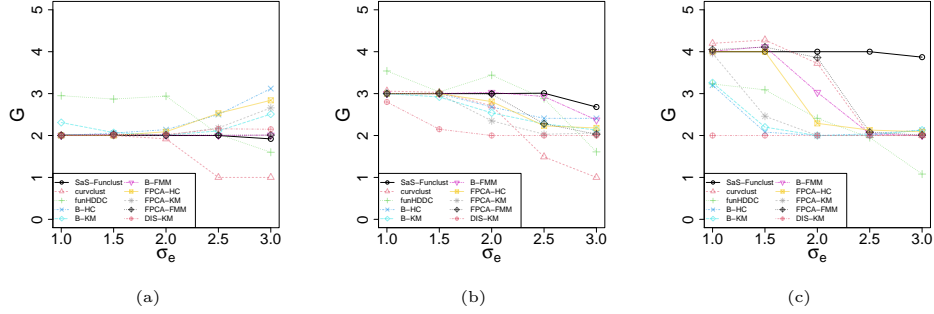


Figure 8.3. Average selected number of clusters G for (a) Scenario I, (b) Scenario II, and (c) Scenario III as a function of σ_e .

error $\sigma_e = 1, 1.5, 2, 2.5, 3$. From Scenario I to Scenario III, the portion of domain that is noninformative for *all* cluster pairs decreases, whereas, the portions of domain that are informative for *specific* cluster pairs increases. Further details about the data generation process are provided in the supplementary materials.

Figure 8.2 shows the average $aRand$ index values for Scenario I, through III as a function of the standard error σ_e . In Scenario I, at small values of σ_e , all methods perform comparably and provide clustering partitions with $aRand$ values very close to 1, which corresponds to the perfect cluster identification. However, as σ_e increases, the SaS-Funclust method turns out to be the best method, closely followed by the curvclust method. Also the B-FMM performs very well, except when $\sigma_e = 3.0$. In Scenario II and III, the SaS-Funclust method is still the best, followed by the curvclust and B-FMM case in Scenario II and only by the curvclust method in Scenario III. Note that in these scenarios, the DIS-KM underperforms also in the most favourable cases as a consequence of the lesser capacity of the L^2 distance to recover the true clustering structure.

Figure 8.3 shows the mean number of selected clusters in all scenarios. It is clear that the SaS-Funclust method is able to identify the true number of clusters much better than the competitors in all the considered scenarios. In particular, Scenario II highlights that, especially for large measurement error σ_e , the competing methods reduce their complexity and select, on average, a number of clusters smaller than the true number of clusters $G_t = 3$.

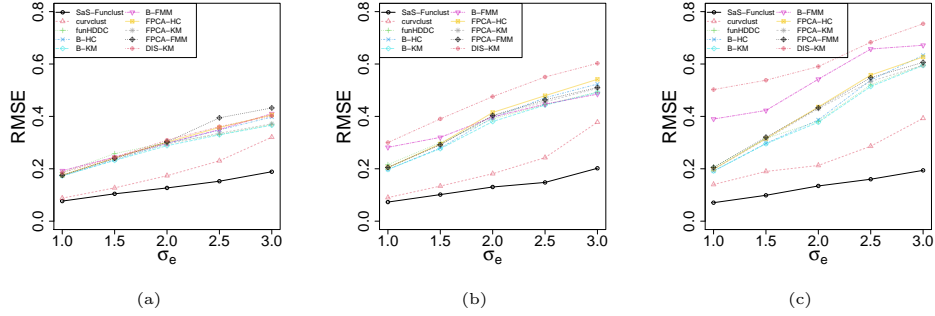


Figure 8.4. Average selected number of clusters G for Scenario I (a), Scenario II (b), and, Scenario III (c) as a function of σ_e .

Table 8.1. Average fractions of correctly identified noninformative portions of domain by the SaS-Funclust method for each scenario.

	Scenario I	Scenario II	Scenario III	
σ_e	1.0	0.9956	0.9901	0.9782
	1.5	0.9921	0.9844	0.9627
	2.0	0.9846	0.9589	0.9389
	2.5	0.9565	0.9373	0.8942
	3.0	0.8821	0.8760	0.8024

This is evident in Scenario III, where the competing methods select, on average, a number of clusters $G = 2$ for $\sigma_e = 2.5, 3.0$, which is smaller than $G_t = 4$.

Figure 8.4 and Table 8.1 highlight the ability of the SaS-Funclust method in recovering the true cluster means and detecting the informative portions of domain. The average root mean squared error calculated as $RMSE = \left[\frac{1}{G} \sum_{g=1}^G \int_{\mathcal{T}} \left(\mu_g(t) - \hat{\mu}_g(t) \right)^2 dt \right]^{1/2}$, with $\hat{\mu}_g(t) = \hat{\boldsymbol{\mu}}_g^T \boldsymbol{\Phi}(t)$, $t \in \mathcal{T}$, is plotted in Figure 8.4 for each method as a function of σ_e in all three scenarios. By this figure, the SaS-Funclust method outperforms the competitors in each scenario, especially for large measurement errors, even though the curvclust method shows comparable performance. Table 8.1 reports, for each σ_e and scenario, the average fractions of correctly identified noninformative portions of domain by the SaS-Funclust method, which can be regarded as a measure of the interpretability (i.e., sparseness) of the proposed solution. In more detail, each entry of the table is obtained as the mean of the average fraction of correctly identified noninformative portions of domain, over the 100 generated datasets, for each pair of clusters, weighted by the size of the corresponding true noninformative portions of domain. In Scenario I, it trivially coincides with the average, because the true number of clusters is $G_t = 2$. The proposed method is clearly able to provide an interpretable clustering. The fraction of correctly identified noninformative portions of domain is almost larger than or equal to 0.90 for $\sigma_e \leq 2.5$ and decreases to 0.80 for $\sigma_e = 3.5$. It is worth noting that when $\sigma_e = 1.0$, the pairs of clusters in each scenario are correctly fused over almost all the noninformative portion of domain in terms of mean differences. This confirms what is shown in Figure 8.1 of Section 8.1.

Table 8.2. The values of the *aRand* index for all the clustering methods with respect to gender difference based grouping and the SaS-Funclust partition for the Berkeley growth study dataset

	SaS-Funclust	curvclust	funHDDC	B-HC	B-KM	B-FMM	FPCA-KM	FPCA-HC	FPCA-FMM	DIS-KM
Gender difference based grouping	0.58	0.51	0.61	0.20	0.58	0.58	0.58	0.58	0.58	0.58
SaS-Funclust	-	0.83	0.96	0.37	1.00	1.00	1.00	1.00	1.00	1.00

8.4 Real-data Examples

Berkeley Growth Study Data

In this section, the SaS-Funclust method is applied to the growth dataset from the Berkeley growth study (Tuddenham, 1954), which is available in the R package *fda* (Ramsay et al., 2020). In this study, the heights of 54 girls and 39 boys were measured 31 times at age 1 through 18. The aim of the analysis is to cluster the growth curves and compare the results with the partition based on the gender difference. This problem has been already addressed by Chiou and Li (2007); Jacques and Preda (2013); Floriello and Vitelli (2017). In particular, we focus on the growth velocities from age 2 to 17, whose discrete values are estimated through the central differences method applied to the growth curves. Figure 8.5(a) shows the interpolating growth velocity curves for all the individuals.

In view of the analysis objective, all clustering methods described in Section 8.3 are applied by setting $G = 2$. As shown in the first row of Table 8.2, all clustering methods, excluded the B-HC, perform similarly in terms of the *aRand* index with respect to the gender difference partition. Moreover, by looking at the second row of Table 8.2, which shows the *aRand* index with respect to the SaS-Funclust partition, the competing methods provide partitions very similar to the SaS-Funclust one.

As expected, the SaS-Funclust method allows for a more interpretable analysis. Figure 8.5 shows (b) the estimated cluster means and (c) the clustered growth curves for the SaS-Funclust method. The estimated cluster means are fused over the first portion of the domain, whereas they are separated over the remaining portions. This implies that the two identified clusters are not different on average over the first portion of domain which can be, thus, regarded as noninformative. Separation between the two group arises over the remaining informative portion of domain, where two sharp peaks of growth velocity arise, instead. The latter peaks are known in the medical literature as pubertal spurts, in which respect the attained results indicate two main timing/duration groups. In particular, male pubertal spurt happens later and lasts longer than female one. Nevertheless, some individuals show unusual growth patterns that are not captured by the cluster analysis. Additionally, the estimated cluster means from the competing methods, not shown here, do not allow for a similar straightforward interpretation.

Canadian Weather Data

The Canadian weather dataset contains the daily mean temperature curves, measured in Celsius degree, recorded at 35 cities in Canada. The temperature profiles are obtained by averaging over the years 1960 through 1994. This is a benchmark dataset available in the R package *fda* (Ramsay et al., 2020) that has been already studied by Ramsay and Silverman (2005); Centofanti et al. (2020). Figure 8.6(a) displays the interpolating profiles, where, for computational reasons, temperature curves are sampled each five days.

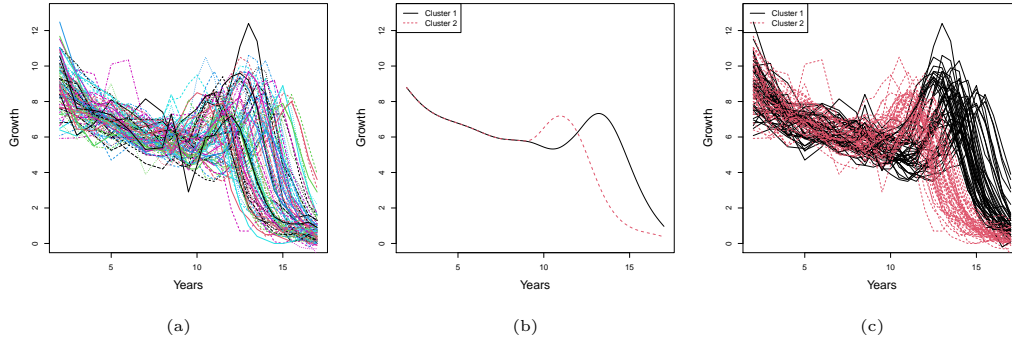


Figure 8.5. (a) Growth velocities of 54 girls and 39 boys in the Berkeley growth study dataset; (b) estimated cluster curve means and (c) curve clusters for the SaS-Funclust method in the Berkeley growth study dataset.

Table 8.3. The values of the $aRand$ index for all the clustering methods with respect to climate zones grouping and the SaS-Funclust partition for the Canadian weather dataset

	SaS-Funclust	curvclust	funHDDC	B-HC	B-KM	B-FMM	FPCA-KM	FPCA-HC	FPCA-FMM	DIS-KM
Climate zones grouping	0.37	0.24	0.21	0.38	0.21	0.33	0.22	0.30	0.17	0.27
SaS-Funclust	-	0.50	0.35	0.86	0.35	0.93	0.59	0.72	0.43	0.40

The ultimate goal of the cluster analysis applied to these curves is the geographical interpretation of the results. In particular, all methods analysed in Section 8.3 are applied by setting $G = 4$ in order to try to recover the grouping of 4 climate zones, viz., Atlantic, Pacific, Continental, Arctic (Jacques and Preda, 2013). The first row of Table 8.3 shows the $aRand$ index values of the resulting clusters calculated with respect to the 4-climate-zone grouping. Although the SaS-Funclust and the B-HC methods achieve the largest $aRand$ in this case, $aRand$ values are in all cases inadequately low, which indicates the clustering structure disagrees with such grouping. That is, different method performance cannot properly evaluated by using the 4-climate-zone grouping. The second row of Table 8.3 reports the $aRand$ index for all the competing methods calculated with respect to the SaS-Funclust method. As expected, the proposed clustering agrees with filtering methods based on B-splines, while mostly disagrees with the others.

In terms of interpretability, Figure 8.6 shows (b) the estimated cluster means and (c) geographical distribution of the curves in the clusters obtained by the SaS-Funclust method. From Figure 8.6(b), the estimated means for clusters 1, 2 and 4 are shown to fuse approximately from day 100 through 250. This is a strong evidence that the mean temperature in this period of the year is not significantly different among zones in cluster 1, 2 and 4. Hence, this portion of domain turns out to be noninformative for the separation of these clusters, whereas the mean temperature is different for the rest of the year. A different pattern is followed by the curves in cluster 3, which shows significantly smaller mean temperature all over the year. The geographical displacement of the temperature profiles coloured by the clusters identified through the SaS-Funclust method is reported in Figure 8.6(c). Observations in cluster 1, 2 and 3 correspond to Pacific, Atlantic and southern continental stations and show similar mean temperature patterns only over the middle days of the year. Observations in cluster 3, which correspond to northern stations, show lower mean temperature. This nice and plausible interpretation of this well-known

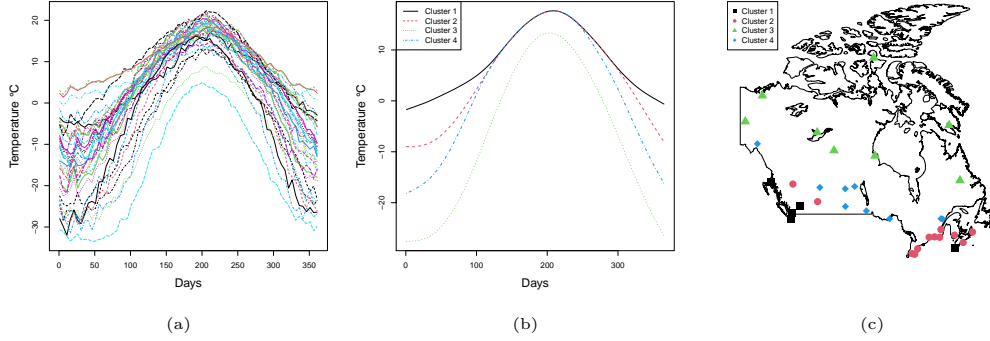


Figure 8.6. (a) Daily mean temperature profiles at 35 cities in Canada over the year in the Canadian weather dataset; (b) estimated cluster curve mean and (c) geographical displacement of the curves pertaining to clusters obtained through SaS-Funclust method.

Table 8.4. $aRand$ index calculated on the ICOSAF project dataset for all competing method partitions with respect to the SaS-Funclust one.

	curvclust	funHDDC	B-HC	B-KM	B-FMM	FPCA-HC	FPCA-KM	FPCA-FMM	DIS-KM
SaS-Funclust	0.00	0.46	0.41	0.35	0.46	0.44	0.27	0.55	0.56

real-data example is not possible by means of any competing method.

ICOSAF Project Data

The ICOSAF project dataset contains 538 dynamic resistance curves (DRCs), collected during resistance spot welding lab tests at Centro Ricerche Fiat in 2019. The DRCs are collected over a regular grid of 238 points equally spaced by 1 ms. Further details on this dataset can be found in Capezza et al. (2020b) and the data are publicly available online at <https://github.com/unina-sfere/funclustRSW/>. In this example, we focus on the first derivative of the DRCs, estimated by means of the central differences method applied to the DRC values sampled each 2 ms. Figure 8.7(a) shows the first derivative of the DRCs defined, without loss of generality, on the domain $[0, 1]$. In this setting, the aim of the analysis is to cluster DRCs to identify homogenous groups of spot welds that share common mechanical and metallurgical properties. Differently from the previous datasets, no information are available about a reasonable partition of the DRCs. Therefore, based on the considerations provided by Capezza et al. (2020b) and on cluster number selection methods described for the SaS-Funclust and competing methods in Section 8.2 and 8.3, respectively, we set $G = 3$. Table 8.4 shows $aRand$ values obtained for all method pairs with respect to the SaS-Funclust partition. In this case, the SaS-Funclust method provides partitions that are more similar to those obtained through the FPCA-based methods than those obtained with the B-splines filtering approaches. However, the clusters identified by the SaS-Funclust method do not resemble those of the other methods. It is worth noting that, for this dataset, even if results are not reported here, the partition obtained by curvclust differs dramatically from the others and does not provide meaningful clusters.

Also in this case, the SaS-Funclust method allows for an insightful interpretation of

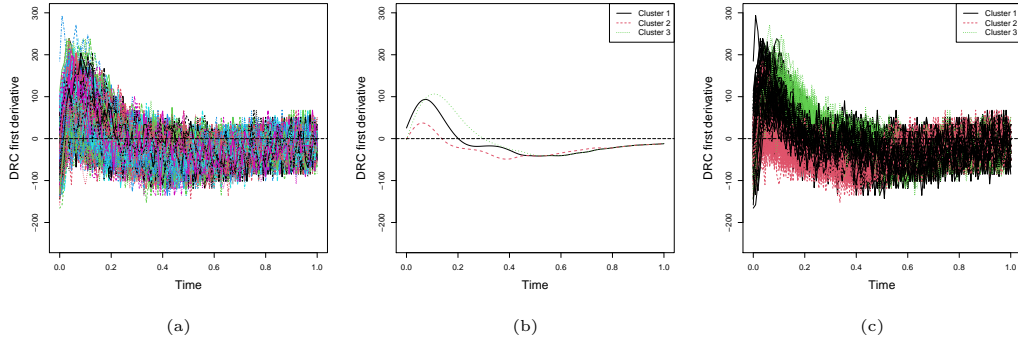


Figure 8.7. (a) First derivatives of the 538 DRCs in the ICOSAF project dataset; (b) estimated cluster curve means and (c) curve clusters for the SaS-Funclust method in the ICOSAF project dataset.

the results. The estimated cluster means and the corresponding clustered curves obtained through the SaS-Funclust method, displayed in Figure 8.7 (b) and (c), confirm the ability of the proposed method to fuse cluster means, as it is clear over the second part of the domain. In particular, the mean of cluster 1 and 3 are fused from 0.5 to 1, which accounts for the comparable decreasing rate of the DRCs over these clusters. Differently, the mean of cluster 2 is fused with other cluster means between 0.8 and 1, only. This indicates that between 0.5 and 0.8 DRCs of cluster 3 decrease with a rate that is different from that of DRCs included in other clusters. Differences between cluster 2 and clusters 1 and 3 are plainly visible also in the first part of the domain, where DRCs of cluster 2 show lower average velocity. Note also that DRCs of cluster 2 reach their peaks (i.e., zeros of the first derivative) earlier than those of clusters 3 and 1.

8.5 Conclusions and Discussions

This chapter presented the SaS-Funclust method, a new approach to the sparse clustering of functional data. Differently from methods that have already appeared in the literature before, it was shown to be capable of successfully detecting where cluster pairs are separated. In many applications, this involves limited portions of domain, which are referred to as informative, and thus, the proposed method allows for a more accurate and interpretable cluster analysis. The SaS-Funclust method can be considered as belonging to the model-based clustering procedures with parameters of a general functional Gaussian mixture model estimated by maximizing a penalized version of the log-likelihood function. The key element is the functional adaptive pairwise fusion penalty that, by locally shrinking mean differences, allows pairs of cluster means to be exactly equal over portions of domain where cluster pairs are not well separated, referred to as noninformative. In addition, a smoothness penalty is introduced to further improve cluster interpretability. The penalized log-likelihood function was maximized by means of a specifically designed expectation-conditional expectation algorithm, and model selection was addressed through a cross-validation technique. An extensive Monte Carlo simulation study showed the favourable performance of the proposed method over several competing methods both in terms of clustering accuracy and interpretability. Lastly, real-data examples further demonstrated the practical advantages of the proposed method, which provided, thanks to its sparseness property, new insightful and interpretable solutions to

cluster analysis. In the Berkeley growth study example, the SaS-Funclust method highlighted that growth velocity curves of boys and girls show different pubertal spurt, which happens later and last longer for male than female. Whereas, in the Canadian weather example, the mean temperatures over the Pacific, Atlantic and southern continental regions were found to be equal over the middle days of the year and different otherwise. Moreover, the proposed method was applied to the ICOSAF project dataset, where, differently from the previous datasets, no information are available about a reasonable partition. The SaS-Funclust method also in this case identified homogenous groups of spot welds that showed, only during the first part of the process, differences in the rate of change of dynamic resistance curves, which are likely to be responsible of distinct mechanical and metallurgical properties of the spot welds.

As closing remarks, we can envisage several important extensions to refine the proposed method. Regarding the structure of the functional clustering model, the assumption of a common diagonal coefficient covariance matrix across all clusters may be too restrictive and may result in a poor fit. Unfortunately, more flexible covariance structures dramatically increase the number of parameters to be estimated, already enlarged to achieve sparseness, in the SaS-Funclust method. For this reason, regularization framework shall necessarily be addressed to avoid overfitting, possibly either by constraining the covariance structure, as done in this chapter, or by means of shrinkage estimators. However, the choice of the best approach still remains not straightforward. Furthermore, the covariance structure of the measurement errors could be modified to include more complex relationships, and the model can be extended also by including covariates (James and Sugar, 2003).

Another natural extension of the SaS-Funclust method that is worth considering is the integration of a proper pairwise penalty applied to the covariance functions, useful in those settings where portions of domain are informative for the clustering also in terms of covariance functions. Unfortunately, the choice of such penalty and the resulting computational issues are non-trivial and need for additional careful investigation.

8.6 Supplementary Materials

Proof of Theorem 1

We have that

$$\sup_{t \in \mathcal{T}} |f(t) - \tilde{f}(t)| = \sup_{t \in \mathcal{T}} |f(t) - \sum_{i=1}^q c_i I_{[\tau_i, \tau_{i+1}]}(t)| = \max_i \sup_{t \in [\tau_i, \tau_{i+1}]} |f(t) - c_i|. \quad (8.19)$$

Following De Boor et al. (1978), we have that for $t \in [\tau_i, \tau_{i+1}]$ and $i = 1, \dots, q$,

$$|f(t) - c_i| \leq D \sup_{s \in [\tau_i, \tau_{i+1}]} |f(t) - f(s)| \leq D \sup_{s, z \in [\tau_i, \tau_{i+1}]} |f(z) - f(s)| \leq D \sup_{|s-t| \leq \delta_\tau} |f(t) - f(s)|, \quad (8.20)$$

where $\delta_\tau = \max_i |\tau_{i+1} - \tau_i|$. The term $\sup_{|s-t| \leq \delta_\tau} |f(t) - f(s)|$ is the modulus of continuity of f on \mathcal{T} , and from Schumaker (2007), the following inequality holds

$$\sup_{|s-t| \leq \delta_\tau} |f(t) - f(s)| \leq \delta_\tau \sup_{t \in \mathcal{T}} |f^{(1)}(t)| \leq \delta \sup_{t \in \mathcal{T}} |f^{(1)}(t)|, \quad (8.21)$$

Table 8.5. Coefficient mean vectors for each scenario and cluster.

	$\boldsymbol{\eta}_i$	Cluster 1	Cluster 2	Cluster 3	Cluster 4
Scenario I	$\eta_{i1}, \dots, \eta_{i5}$	1.5	-1.5	-	-
	$\eta_{i6}, \dots, \eta_{i30}$	0	0	-	-
Scenario II	$\eta_{i1}, \dots, \eta_{i5}$	3	0	0	-
	$\eta_{i6}, \dots, \eta_{i10}$	1.5	1.5	-1.5	-
	$\eta_{i11}, \dots, \eta_{i30}$	0	0	0	-
Scenario III	$\eta_{i1}, \dots, \eta_{i5}$	1.5	1.5	-1.5	-1.5
	$\eta_{i6}, \dots, \eta_{i10}$	3	0	0	-3
	$\eta_{i11}, \dots, \eta_{i15}$	1.5	1.5	-1.5	-1.5
	$\eta_{i16}, \dots, \eta_{i30}$	0	0	0	0

where the last inequality follows because $0 \leq \delta_\tau \leq \delta$. Therefore, upon using Equation (8.19), (8.20), and (8.21), it readily follows that

$$\sup_{t \in \mathcal{T}} |f(t) - \tilde{f}(t)| \leq D\delta \sup_{t \in \mathcal{T}} |f^{(1)}(t)|, \quad (8.22)$$

which proves the theorem.

Details on Data Generation of the Monte Carlo Simulation Study

Data are generated from $G_t = 2, 3, 4$ clusters and referred to as Scenario I, II and III, respectively. For each scenario, the considered methods are evaluated by assessing the performance over 100 independently generated datasets. From each cluster, 200 observations are generated over the domain $\mathcal{T} = [0, 1]$. The true functions are obtained as $g_i = \boldsymbol{\eta}_i^T \boldsymbol{\Theta}_B$, with $\boldsymbol{\Theta}_B$ representing a set of 30 evenly spaced knot cubic B-splines. The coefficients $\boldsymbol{\eta}_i = (\eta_{i1}, \dots, \eta_{i30})^T$ are Gaussian random coefficients with mean vectors (depending on both the scenario and the cluster considered) reported in Table 8.5, and covariance matrix $\boldsymbol{\Gamma} = \sigma_c^2 \mathbf{I}$, where $\sigma_c = 0.5$ and \mathbf{I} denotes the identity matrix. Then, to obtain the contaminated-with-error observations \mathbf{Y}_i , the true functions g_i are evaluated over a grid of 50 points, with the addition of measurement errors independently generated as Gaussian random variables with zero mean and five different values of standard error $\sigma_e = 1, 1.5, 2, 2.5, 3$.

Part III

Chapter 9

Robust Functional ANOVA with Application to Additive Manufacturing

Abstract

With the advance of modern technology, more and more data are being recorded continuously and, thus, are apt to be modeled as functional data. In some applications, the interest relies on identifying the presence of significant differences, in terms of functional mean, among groups of a functional data, defined by different conditions. This problem is the functional extension of the classical analysis of variance, that is, the functional analysis of variance (FANOVA). When dealing with real data, it is common that the functional sample under study is contaminated by some outliers, which can strongly bias the analysis. In order to take into account the possible presence of anomalous functional observations, a new robust method to address the multiway FANOVA problem is proposed which reduces the weights of outliers on the results of the analysis. It is a permutation test whose test statistics relies on the functional equivariant M -estimator, the functional extension of the classical robust M -estimator. By means of an extensive Monte Carlo simulation study, the proposed test is compared with some alternatives already present in the literature in both one-way and two-way designs. The performance of the proposed approach is demonstrated in the framework of a motivating real case-study in additive manufacturing, also known as 3D-printing. The real case study deals with the analysis of by-products of the process, i.e., spatter ejections, measured through high-speed video imaging and translated into a functional form. The aim of the robust FANOVA consists of identifying significant effects of relevant process factors and their interactions on such by-products, since spatter-related dynamics have the potential to be used as "signatures" of the process stability and quality.

9.1 Introduction

The development of data acquisition methods allow the analysis of complex systems in several operating conditions as never before. Several examples may be found in the current Industry 4.0 framework, that is being reshaping the variety of signals and measurements that can be gathered during manufacturing processes. Experimental data are more and more characterized by complex and novel formats, like images, videos, dense point clouds, etc. This evolution has been enabled also by new industrial production processes like additive manufacturing, or 3D printing. Indeed, thanks to the ability of manufacturing a product on

a layer by layer basis, it is possible to look and capture phenomena and events occurring in every layer, through measurement tools that may range from in-line thermography to machine vision, spectrometry, etc. (Colosimo et al., 2018; Colosimo and Grasso, 2020; Everton et al., 2016). As a consequence, the focus of many applications in industrial statistics is being moving from product quality characteristics to in-line process measurements, thanks to enhanced sensing and monitoring capabilities. Moreover, novel production paradigms are characterized by several controllable factors and complex process dynamics that impose the need for effective and efficient experimental approaches to determine optimal process conditions but also to gather a deeper comprehension of underlying physical phenomena.

One example, which also motivates the present study, regards the analysis of by-products in a metal additive manufacturing process known as laser powder bed fusion (L-PBF). L-PBF is an additive process suitable to produce metal parts exploiting a laser beam to selectively melt a thin layer of metal powder. The process is repeated layer by layer, with the material solidified in one layer being welded to the material in underneath layers, enabling the fabrication of products with complex geometries and innovative properties (Gibson et al., 2014). In addition to the local melting and consequent material solidification, the laser-material interaction yields also process by-products. They consist of spatters ejected from the melted area (also called "melt pool") or its surroundings, and a partial material vaporisation, also known as "plume" (Young et al., 2020). Various authors showed that the information enclosed by such by-products represents a possible "signature" of the process state in terms of its layer-by-layer stability and the final quality of produced parts (Yang et al., 2020; Andani et al., 2017; Repossini et al., 2017; Ly et al., 2017; Bidare et al., 2018). Because of this, the analysis of process by-products has gathered an increasing interest in the last years. The way in which spatters originate during the process and the way in which they spread in space above the melted area can be captured via in-line high-speed machine vision. However, in order to investigate the effect of various controllable process factors on the spatter behaviour, raw video image data shall be translated into a data format more convenient for the analysis of variance. One possible way to address this task is to refer to a functional data format, by representing the amount of ejected spatters as a function defined on a compact domain. In this example, the compact domain consists of the bi-dimensional space corresponding to the field of view of the machine vision equipment. An example of such data format transformation is depicted in Fig. 9.1, where the left panel shows one raw video frame acquired during the L-PBF process and the right panel shows a function of how spatters were spread in space in that video frame. This function will be referred to as "spatter intensity" function in this study, as it maps the amount of spatters observed in any region of the bi-dimension video frame space, (s, t) . The term "intensity" here refers to the occurrence of spatters in a given location. Therefore, a high spatter intensity at given spatial coordinates (s, t) means that a large amount of spatters was captured in the video image stream in that specific location.

The example shown in Fig. 9.1 is just one of many real applications where a functional data representation may be suitable to deal with complex patterns and data types. A functional representation similar to the one in Fig. 9.1 may be used in many other processes where spatters and hot ejections are generated, like welding or laser cutting. However, the representation of observation units in terms of functions in a 1D, 2D or higher dimensional domain has a much wider and general validity, not limited to manufacturing applications. The family of statistical methods suitable to tackle this problem is known as functional data analysis (FDA). For a comprehensive overview of FDA methods and applications we refer the reader to Ramsay and Silverman (2005); Horváth and Kokoszka (2012); Kokoszka and




Figure 9.1. Example of a video frame acquired during an L-PBF process showing ejected spatters as bright spots (left panel) and corresponding spatter intensity function (right panel).

Reimherr (2017) and, for further theoretical insights, to Hsing and Eubank (2015); Bosq (2012).

A classical statistical problem consists of identifying the differences among group functional means in a sample when some experimental conditions vary. In the literature, this problem is known as functional analysis of variance (FANOVA) that is the FDA extension of the classical (non-functional) ANOVA problem. In (Ramsay and Silverman, 2005), the authors proposed a functional ANOVA test, based on a pointwise F -test statistic, that relies on the normality assumption of the error function. If the observed statistics is larger than the critical value, calculated as a percentile of the Fisher distribution, for each domain value, then we safely reject the hypothesis of no differences among the groups. Cuevas et al. (2004) proposed a FANOVA test based on the integrated squared difference among group functional means, for both the homoscedastic and heteroscedastic cases. The L^2 -norm-based test proposed by Faraway (1997); Zhang et al. (2007) uses a statistic based on the integrated squared differences between the group means and the global mean, whose distribution is approximately proportional to a chi-squared random variable. Shen and Faraway (2004); Zhang (2011) proposed an F -type test based on the fraction of the sum of the integrated squared differences between the group means and the global mean, and, the sum of the integrated squared differences between the functional observations and the group means. Under certain conditions, this statistic has a Fisher distribution. Bootstrap versions of both L^2 -norm-based and F -type tests were proposed in Zhang (2013). Finally, Zhang and Liang (2014) introduced a globalized version of the pointwise F -test. Note that, all the aforementioned works deal with the one-way FANOVA design. The multi-way functional ANOVA design has been much less studied than the one-way counterpart. In particular, Brumback and Rice (1998); Guo (2002); Gu (2013) proposed tests that are able to deal with more complicated designs that rely on the use of smoothing splines (SS-ANOVA). A simple technique was proposed by Cuesta-Albertos and Febrero-Bande (2010) that transform functional data into univariate data by means random projections. Pini et al. (2017) proposed a non-parametric domain-selective multi-way functional ANOVA able to identify the specific




Figure 9.2. Example of a spatter intensity function (left panel), one cross-section of the spatter intensity function at $t = 0.75$ (central panel) and a superimposition of cross-sections corresponding to different experimental realizations of the spatter intensity function, where an outlying pattern is highlighted with a thick black line (right panel)

subdomains where group functional means differ.

In this study, we address the functional analysis of variance in the presence of nuisance effects associated to outlying patterns in the experimental dataset. The proposed real case study in additive manufacturing highlights the need for novel and effective methods in this framework. Indeed, additive manufacturing, like many other manufacturing processes and thermal treatments, is characterized by complex dynamics and many transient and local phenomena that not only affect the natural variability of the measured quantities, but could also lead to outlying patterns. In the motivating case study considered in this chapter, an outlying spatter ejection behaviour may be observed as a consequence of a variety of possible root causes, ranging from transient fluctuations of the gas flow, local inhomogeneities of the powder bed, laser beam attenuation caused by evaporated material, etc. Fig. 9.2 shows an example of an outlying pattern in the spatter intensity function. For sake of graphical clarity, functions corresponding to different realizations under the same experimental treatment are compared by looking at their cross-sections at a fixed coordinate t . The cross-section shown with a solid thick line in Fig. 9.2 represents an outlying spatter behaviour, consisting of a lower amount of spatters spread in space, possibly caused by a transient laser beam attenuation that occurred at a given point in time. Additional details about the real case study can be found in Section 4.

From a design of experiments perspective, outlying patterns like the one in Fig. 9.2 represent a nuisance for the analysis or results, as they may inflate the variability and mask effects of potential interest. From a statistical process monitoring perspective, instead, outliers are commonly drivers of relevant information being potential indicators of anomalies and flaws. In this study we refer to the former perspective, aiming at proposing an effective approach for the analysis of variance in the presence of outliers that contaminate the experimental functional data. Due to the many different dynamics involved in the process, determining whether an experimental point is an outlier and identifying its root cause can be a difficult task, but similar challenges can be faced in many different manufacturing applications, due to the complex nature of the response variables and the complex underlying physical phenomena.

All the one-way and multi-way FANOVA design cited above differently combine in a quadratic fashion the functional mean to obtain the statistics. However, as in the case of finite dimensional data, it has been shown that the functional mean, as well as quadratic forms, are highly sensitive to the presence of outliers and other anomalies.

Hubert et al. (2015) set up a taxonomy of functional outliers. To deal with outliers, the *diagnostic* and the *robust* approaches are the two common alternatives. The diagnostic approach (that is based on standard estimates after the removal of sample units identified as outliers) is often criticized as lacking of objectivity being based on the analyst's personal decision. On the contrary, the robust approach produces parameter estimates and associated tests and confidence intervals more robust, to the presence of outliers. For a general perspective on this topic in the classical setting see Huber (2004); Hampel et al. (2011); Maronna et al. (2019). By looking at Fig. 9.2, the marked curve could be safely deleted if the diagnostic approach is considered. However, as we will see in the subsequent analysis, it is not always easy to label an observation as outlier, especially when complex process dynamics and lack of measurable covariates make the search for root causes a difficult task.

In the very last years, several works have explored robust estimation for functional data. Fraiman and Muniz (2001) defined trimmed means for functional data based on a functional depth defined as an integral of the univariate depths for each domain value. In order to obtain robust estimates of the center of a functional distribution, Cuesta-Albertos and Fraiman (2006) extended the notion of impartial trimming to a functional data framework. Other location estimators based on depth functions for functional data were proposed by Cuesta-Albertos and Nieto-Reyes (2008); Cuevas and Fraiman (2009); López-Pintado and Romo (2009, 2011). The above methods are all extensions of the classical linear combination types estimators (i.e., L -estimator) (Maronna et al., 2019) to the functional setting. More recently, Sinova et al. (2018) extended the notion of maximum likelihood type estimators (i.e., M -estimators) to the functional data setting. M -estimators (Huber et al., 1964) are less influenced by outliers than the standard least-squares/maximum likelihood estimators, because they are based on loss functions that increase less rapidly than the usual square loss. This estimator has been used by Kalogridis and Van Aelst (2019) to robustly estimate the functional linear model.

The FANOVA methods are not necessarily robust to outliers as they rely both on the functional mean and on quadratic forms, which are known to be highly sensitive to outlying observations. In the classical setting, robust ANOVA methods have been proposed by Schrader and Mc Kean (1977); Schrader and Hettmansperger (1980), where Huber's M -estimates are adapted to be used in both a modified F -statistic and a likelihood ratio type test. However, to the best of our knowledge, no robust ANOVA has been introduced so far in the functional setting.

In this chapter, we propose a robust nonparametric functional ANOVA method (RoFANOVA) that is able to test differences among group functional means. It is based on a functional generalization of the test statistic proposed by Schrader and Mc Kean (1977) included in a permutational framework (Good, 2013; Pesarin and Salmaso, 2010). Applications of nonparametric methods in FDA can be found in Ramsay and Silverman (2005); Corain et al. (2014); Pini and Vantini (2017); Pini et al. (2017). Moreover, to obtain the test statistic, we introduce a functional extension of the normalized median absolute deviation (NMAD) estimator, referred to as functional normalized median absolute deviation (FuNMAD) estimator, as well as an equivariant version of the functional M -estimator proposed by Sinova et al. (2018). An extensive Monte Carlo simulation study is presented to quantify the performance of the RoFANOVA with respect to FANOVA tests already present in the literature, both in one-way and two-way designs. The application of the proposed approach to the real case study in additive manufacturing also highlights that the RoFANOVA is more effective than other methods in identifying interaction effects that are relevant to get deeper insights about the functional response variable of interest.

The chapter is organized as follows. In Section 9.2, the robust functional analysis of variance is introduced together with the functional normalized median absolute deviation and the scale equivariant functional M -estimator. Section 9.3 presents a Monte Carlo simulation study that compares the RoFANOVA with competing methods both in one-way and two-way designs. Then, in Section 9.4 the RoFANOVA is applied to the real-case study devoted to the study of the spatter behaviour in the L-PBF process. Conclusions are provided in Section 9.5. All computations and plots have been created by using R software (R Core Team, 2020a).

9.2 The Robust Functional Analysis of Variance

The Scale Equivariant Functional M -estimator and the Functional Normalized Median Absolute Deviation Estimator

This section introduces the equivariant functional M -estimator and the functional normalized median absolute deviation estimators. Let us consider the random element X with value in $L^2(\mathcal{T})$, the Hilbert space of square integrable functions defined on the compact set $\mathcal{T} \subset \mathbb{R}^p$, with the usual norm $\|f\| = \left(\int_{\mathcal{T}} f^2(t) dt \right)^{1/2}$, for $f \in L^2(\mathcal{T})$, having mean function $\mu(t) = E[X(t)]$ and covariance function $\gamma(s, t) = \text{Cov}[X(s), X(t)]$, for $s, t \in \mathcal{T}$. Moreover, let $\mathbf{X} = (X_1, \dots, X_n)^T$ be a vector whose elements X_i are independent realizations of X . Recently, Sinova et al. (2018) proposed a functional M -estimators of location defined as

$$\hat{\mu}_s = \underset{y \in L^2(\mathcal{T})}{\operatorname{argmin}} \sum_{i=1}^n \rho(\|X_i - y\|), \quad (9.1)$$

where $\rho : \mathbb{R}^+ \rightarrow \mathbb{R}$ is the *loss function* which is continuous, non-decreasing and satisfies $\rho(0) = 0$. As shown by Sinova et al. (2018), each version of $\hat{\mu}_s$ is well-defined and enjoys good theoretical properties, e.g., it has maximal breakdown value and is strong consistent under suitable model assumptions. Unfortunately, these estimators are not scale equivariant. This means that if all X_i are equally scaled, the resulting robust estimator is not necessarily equally scaled in analogy with the multivariate case (Maronna et al., 2019). Following Maronna et al. (2019), we propose a scale equivariant M -estimator of location defined as

$$\hat{\mu} = \underset{y \in L^2(\mathcal{T})}{\operatorname{argmin}} \sum_{i=1}^n \rho\left(\left\| \frac{X_i - y}{\sigma} \right\| \right), \quad (9.2)$$

where $\sigma(t) = \sqrt{\gamma(t, t)}$, for $t \in \mathcal{T}$. If σ was known, the problem would have been reduced to the case of a L^2 random element with $\sigma = 1$. Since σ is rarely known, it can be suitably substituted by a robust scale estimator $\hat{\sigma}$. In this regard, we introduce the FuNMAD estimator defined as

$$\text{FuNMAD}(\mathbf{X}) = \frac{1}{c} \text{Med}\left(|\mathbf{X} - \hat{\mu}_{s,med}|\right), \quad (9.3)$$

with $c = 0.6745$ and where $\hat{\mu}_{s,med}$, the functional generalization of the median, is the solution of the optimization problem in equation (9.1) with $\rho^{med}(\cdot) = |\cdot|$; $|\mathbf{X} - \hat{\mu}_{s,med}| = \left(|X_1 - \hat{\mu}_{s,med}|, \dots, |X_n - \hat{\mu}_{s,med}|\right)^T$ and $\text{Med}(\cdot)$ transforms a vector of functions to a function of pointwise medians. The constant $c = 0.6745$ is used to make FuNMAD an asymptotically pointwise consistent estimator of σ as shown in the supplementary materials.

Because the minimization problem in equation (9.1) has not a closed-form solution, in general, Sinova et al. (2018) proposed a standard iteratively re-weighted least-squares algorithm to approximate $\hat{\mu}_s$. The algorithm is specifically modified to approximate $\hat{\mu}$ in equation (9.2) where σ is approximated by $\hat{\sigma} = \text{FuNMAD}(\mathbf{X})$, and can be summarized in the following steps.

Step 1. Select initial weight vector $\mathbf{w}^{(0)} = (w_1^{(0)}, \dots, w_n^{(0)}) \in \mathbb{R}^n$ such that $w_i^{(0)} \geq 0$ and $\sum_{i=1}^n w_i^{(0)} = 1$.

Step 2. Generate a sequence $\{\hat{\mu}^{(k)}\}_{k \in \mathbb{N}}$ iterating the following procedure:

$$\hat{\mu}^{(k)} = \sum_{i=1}^n w_i^{(k-1)} X_i, \quad w_i^{(k)} = \frac{\psi\left(\left\|\frac{X_i - \hat{\mu}^{(k)}}{\sigma}\right\|\right)}{\sum_{i=1}^n \psi\left(\left\|\frac{X_i - \hat{\mu}^{(k)}}{\sigma}\right\|\right)},$$

where $\psi = \rho'$ is the first derivative of ρ .

Step 3. Terminate the algorithm when

$$\frac{|J(\hat{\mu}^{(k)}) - J(\hat{\mu}^{(k-1)})|}{J(\hat{\mu}^{(k-1)})} < \varepsilon,$$

where $J(h) = \sum_{i=1}^n \rho\left(\left\|\frac{X_i - h}{\sigma}\right\|\right)$ and for a tolerance $\varepsilon > 0$.

The initial weight vector can be chosen as $w_i^{(0)} = \frac{\psi\left(\left\|\frac{X_i - \hat{\mu}^{(0)}}{\sigma}\right\|\right)}{\sum_{i=1}^n \psi\left(\left\|\frac{X_i - \hat{\mu}^{(0)}}{\sigma}\right\|\right)}$ where $\hat{\mu}^{(0)}$ is an initial estimate of $\hat{\mu}$.

The loss function ρ in equation (9.2) defines the properties of the resulting estimator $\hat{\mu}$. For instance, the *Huber's family* of loss functions (Huber et al., 1964), which generates monotone functional M -estimators of location, is given by

$$\rho_a^{HU}(x) = \begin{cases} x^2/2 & \text{if } 0 \leq x \leq a \\ a(x - a/2) & \text{if } a < x, \end{cases}$$

with tuning parameter $a > 0$. It gives less importance to large errors compared to the standard least-squares loss function $\rho^{sq}(x) = x^2$. Functional M -estimators arise from the *bisquare* or *Tukey's biweight family* of loss functions (Beaton and Tukey, 1974) defined by

$$\rho_a^{BI}(x) = \begin{cases} a^2/6 \left[1 - \left(1 - (x/a)^2\right)^3\right] & \text{if } 0 \leq x \leq a \\ a^2/6 & \text{if } a < x, \end{cases}$$

with tuning parameter $a > 0$. M -estimators obtained by using ρ_a^{BI} are re-descending, that is values of $x > a$ give the same contribution to the loss, regardless of their distance from a .




Figure 9.3. The loss functions ρ_a^{HU} (HUB), ρ_a^{BI} (BIS), $\rho_{a,b,c}^{HA}$ (HAM), ρ_a^{OP} (OPT), with tuning constants chosen to achieve 95% asymptotic efficiency, and, ρ^{sqr} (SQR) and ρ^{med} (MED).

Another very used family of loss functions is the *Hampel's one* (Hampel, 1974), defined by

$$\rho_{a,b,c}^{HA}(x) = \begin{cases} x^2/2 & \text{if } 0 \leq x < a \\ a(x - a/2) & \text{if } a \leq x < b \\ \frac{a(x-c)^2}{2(b-c)} + a(b+c-a)/2 & \text{if } b \leq x < c \\ a(b+c-a)/2 & \text{if } c \leq x, \end{cases}$$

with tuning parameter $a, b, c > 0$. M -estimators obtained by using $\rho_{a,b,c}^{HA}$ are redescending as well. Finally, the *optimal family* of loss functions (Maronna et al., 2019) is defined by

$$\rho_a^{OP}(x) = \int_0^x \left(-\frac{\Phi'(|x|) + a}{\Phi(|x|)} \right)_+ dx,$$

where Φ is the standard normal density, $a > 0$ is a tuning parameter and $(t)_+$ denotes the positive part of t . The tuning parameters used in ρ_a^{HU} , ρ_a^{BI} , $\rho_{a,b,c}^{HA}$ and ρ_a^{OP} are chosen in order to ensure given asymptotic efficiency with respect to the normal distribution (Maronna et al., 2019). The loss functions ρ_a^{HU} , ρ_a^{BI} , $\rho_{a,b,c}^{HA}$ and ρ_a^{OP} with tuning constants chosen to achieve 95% asymptotic efficiency, along with ρ^{sqr} and ρ^{med} , are displayed in Fig. 9.3.

The Proposed Robust Method for the Functional Analysis of Variance

The aim of this section is to describe the proposed RoFANOVA for the multiway functional ANOVA design. Without loss of generality, and for ease of notation, we will focus on the two-way functional ANOVA design with an interaction, but the extension to more complex designs is straightforward. To introduce the two-way functional ANOVA design with an interaction, let us consider a functional response X , which is a random element with value in $L^2(\mathcal{T})$, where $\mathcal{T} \subset \mathbb{R}^p$, and is possibly affected by two factors, say A and B (with I and J levels, respectively). In this model, X will be expressed as the sum of two main effects and an interaction between them, plus a random error. Our aim is to test the statistical significance of the main effects and of the interaction term. Let X_{ijk} , for $k = 1, \dots, n_{ij}$,

denote the realizations of X at level i of the factor A, $i = 1, \dots, I$, and level j of the factor B, $j = 1, \dots, J$. Then, the two-way functional ANOVA model to be tested is

$$X_{ijk}(t) = m(t) + f_i(t) + g_j(t) + h_{ij}(t) + \varepsilon_{ijk}(t) \quad t \in \mathcal{T}, \quad (9.4)$$

where m is the functional grand mean, which describes the overall shape of the process, f_i and g_j are the functional main effects and h_{ij} is the interaction term, all these terms being in $L^2(\mathcal{T})$. The functional errors ε_{ijk} are assumed to be independent and identically distributed zero-mean random functions with covariance function γ . They are not required to be Gaussian. In order to make the model identifiable, we will assume, that, $\sum_{i=1}^I \sum_{j=1}^J n_{ij} f_i(t) = \sum_{j=1}^J \sum_{i=1}^I n_{ij} g_j(t) = \sum_{i=1}^I \sum_{j=1}^J n_{ij} h_{ij}(t) = 0$. To test the significance of the coefficients in model (9.4), (that is, to extend the classical ANOVA test to functional data), we consider the following null and alternative hypotheses

$$H_{0,A} : f_1 = \dots = f_I = \mathbf{0}, \quad H_{1,A} : (H_{0,A})^C, \quad (9.5)$$

$$H_{0,B} : g_1 = \dots = g_J = \mathbf{0}, \quad H_{1,B} : (H_{0,B})^C, \quad (9.6)$$

$$H_{0,AB} : h_{11} = \dots = h_{IJ} = \mathbf{0}, \quad H_{1,AB} : (H_{0,AB})^C, \quad (9.7)$$

where $\mathbf{0}$ is a function almost everywhere equal to zero. The hypotheses $H_{0,A}$ against $H_{1,A}$ and $H_{0,B}$ against $H_{1,B}$ are for the effects of the main factors A and B, respectively, whereas, the hypothesis $H_{0,AB}$ against $H_{1,AB}$ is for the interaction term.

Each test is carried out through a nonparametric permutational approach. In this regard, we introduce a test statistic that is a functional extension of the robust F-statistic proposed by Schrader and Mc Kean (1977). In that paper, the authors considered a robust version of the classical F -test statistic, defined as the fraction of the drop in residual sum of squares between the full model (i.e., the model when H_0 is false) and the reduced model (i.e., the model when H_0 is true), and the standard deviation of the error distribution, where all the quantities are estimated by using the least-squares approach. They modified the F -test statistic by considering a residual sum of dispersions (identified by a specific loss function as the ones described in Section 9.2) instead of the residual sum of squares, and a robust estimate, instead of the least-squares estimate, of the standard deviation of the error distribution.

In detail, to test the hypothesis (9.5), we propose to use the following test statistic

$$F_A = (I - 1)^{-1} \left(\sum_{i=1}^I \sum_{j=1}^J \sum_{k=1}^{n_{ij}} \rho \left(\left\| \frac{X_{ijk} - \bar{X}_r - \bar{X}_{r,ij} + \bar{X}_{r,i}}{\hat{\sigma}_{r,e}} \right\| \right) - \sum_{i=1}^I \sum_{j=1}^J \sum_{k=1}^{n_{ij}} \rho \left(\left\| \frac{X_{ijk} - \bar{X}_{r,ij}}{\hat{\sigma}_{r,e}} \right\| \right) \right),$$

where ρ is a given loss function. The test statistic F_A is the mean difference between the standardized residual sum of dispersions under the reduced model and the full model (analogously to Schrader and Mc Kean (1977) in the classical setting), where \bar{X}_r , $\bar{X}_{r,i}$, and $\bar{X}_{r,ij}$ are scale equivariant functional M -estimators (Section 9.2) of the functional grand mean m , and of the group means of $\{X_{ijk}\}_{k=1, \dots, n_{ij}, i=1, \dots, I}$ and $\{X_{ijk}\}_{k=1, \dots, n_{ij}}$, respectively, and, $\hat{\sigma}_{r,e}$ is a robust estimate of the functional standard deviation of the error distribution.

In detail, \bar{X}_r , $\bar{X}_{r,i\cdot}$, $\bar{X}_{r,ij}$ and $\hat{\sigma}_{r,e}$ are defined as

$$\begin{aligned}\bar{X}_r &= \operatorname{argmin}_{y \in L^2(\mathcal{T})} \sum_{i=1}^I \sum_{j=1}^J \sum_{k=1}^{n_{ij}} \rho \left(\left\| \frac{X_{ijk} - y}{\hat{\sigma}_r} \right\| \right), & \hat{\sigma}_r &= \text{FuNMAD}(\{X_{ijk}\}_{k=1, \dots, n_{ij}, i=1, \dots, I, j=1, \dots, J}), \\ \bar{X}_{r,i\cdot} &= \operatorname{argmin}_{y \in L^2(\mathcal{T})} \sum_{j=1}^J \sum_{k=1}^{n_{ij}} \rho \left(\left\| \frac{X_{ijk} - y}{\hat{\sigma}_{r,i\cdot}} \right\| \right), & \hat{\sigma}_{r,i\cdot} &= \text{FuNMAD}(\{X_{ijk}\}_{k=1, \dots, n_{ij}, j=1, \dots, J}), \\ \bar{X}_{r,ij} &= \operatorname{argmin}_{y \in L^2(\mathcal{T})} \sum_{k=1}^{n_{ij}} \rho \left(\left\| \frac{X_{ijk} - y}{\hat{\sigma}_{r,ij}} \right\| \right), & \hat{\sigma}_{r,ij} &= \text{FuNMAD}(\{X_{ijk}\}_{k=1, \dots, n_{ij}}), \\ \hat{\sigma}_{r,e} &= \frac{1}{0.6745} \operatorname{Med}(|\{X_{ijk} - \bar{X}_{r,ij}\}_{k=1, \dots, n_{ij}, i=1, \dots, I, j=1, \dots, J}|).\end{aligned}$$

Intuitively, F_A is a measure of the discrepancy between the residuals, obtained by using robust statistics, of the model under $H_{0,A}$ and under $H_{1,A}$. In this latter case, F_A assumes high values. Analogously, to test the hypotheses (9.6) and (9.7), we define

$$\begin{aligned}F_B &= (J-1)^{-1} \left(\sum_{i=1}^I \sum_{j=1}^J \sum_{k=1}^{n_{ij}} \rho \left(\left\| \frac{X_{ijk} - \bar{X}_r - \bar{X}_{r,ij} + \bar{X}_{r,j\cdot}}{\hat{\sigma}_{r,e}} \right\| \right) - \sum_{i=1}^I \sum_{j=1}^J \sum_{k=1}^{n_{ij}} \rho \left(\left\| \frac{X_{ijk} - \bar{X}_{r,ij}}{\hat{\sigma}_{r,e}} \right\| \right) \right), \\ F_{AB} &= ((I-1)(J-1))^{-1} \left(\sum_{i=1}^I \sum_{j=1}^J \sum_{k=1}^{n_{ij}} \rho \left(\left\| \frac{X_{ijk} - \bar{X}_{r,i\cdot} - \bar{X}_{r,j\cdot} + \bar{X}_r}{\hat{\sigma}_{r,e}} \right\| \right) - \sum_{i=1}^I \sum_{j=1}^J \sum_{k=1}^{n_{ij}} \rho \left(\left\| \frac{X_{ijk} - \bar{X}_{r,i\cdot}}{\hat{\sigma}_{r,e}} \right\| \right) \right),\end{aligned}$$

where

$$\bar{X}_{r,j\cdot} = \operatorname{argmin}_{y \in L^2(\mathcal{T})} \sum_{i=1}^I \sum_{k=1}^{n_{ij}} \rho \left(\left\| \frac{X_{ijk} - y}{\hat{\sigma}_{r,j\cdot}} \right\| \right), \quad \hat{\sigma}_{r,j\cdot} = \text{FuNMAD}(\{X_{ijk}\}_{k=1, \dots, n_{ij}, i=1, \dots, I}).$$

Different versions of the proposed test statistics may emerge by the choice of the loss function ρ as defined in Section 9.2, and, note that, to estimate $\bar{X}_{r,ij}$, $\hat{\sigma}_{r,ij} = \hat{\sigma}_{r,e}$ could be used as well.

Another element to choose in a permutation test is the method to approximate the distribution of the considered statistic under the null hypothesis. In our case, we selected the Manly's scheme (Gonzalez and Manly, 1998; Manly, 2006) that consists of simply permuting the raw data without restrictions. Although other schemes could be used, the Manly's one has demonstrated to have good performance in spite of its simplicity, especially when the sample size for given factor levels is small, see Gonzalez and Manly (1998) and Anderson (2001) for further details.

Therefore, let F be the chosen statistic (resp. F_A or F_B or F_{AB}) to test, at level α , H_0 against H_1 (resp. $H_{0,A}$ against $H_{1,A}$, or $H_{0,B}$ against $H_{1,B}$, or $H_{0,AB}$ against $H_{1,AB}$). Then, the proposed permutation test is composed by the following steps.

- Step 1. Compute the observed value of the test statistic F_{obs} , by considering the original sample $\{X_{ijk}\}_{k=1, \dots, n_{ij}, i=1, \dots, I, j=1, \dots, J}$.
- Step 2. Randomly permute the data, among the Factor A and Factor B combinations, B times, and for each permuted sample compute the value F_1^*, \dots, F_B^* of the statistic F .
- Step 3. Compute the approximated p-value as

$$p = \frac{1}{B} \sum_{i=1}^B I(F_i^* \geq F_{obs}),$$

where $I(E)$ takes values 1 or 0 depending on whether E is true or false.

Step 4. Accept H_0 if $p > \alpha$, otherwise reject H_0 .

This test is an approximate (asymptotically exact) level α test for H_0 against H_1 (Anderson, 2001). The larger the number of permutations B the lower the approximation error. We suggest to select the number of permutations B equal or larger than 1000 (Good, 2013).

9.3 Simulation Study

In this section, the performance in terms of empirical size and power of the proposed method are assessed by means of an extensive Monte Carlo simulation study. In particular, two scenarios are investigated:

Scenario 1 A one-way FANOVA model (i.e., model (9.4) with $m = 0$, $g_1 = \dots = g_J = 0$ and $h_{11} = \dots = h_{IJ} = 0$) is considered (Section 9.3).

Scenario 2 A two-way FANOVA model (i.e., model (9.4)) is considered (Section 9.3).

In each scenario, the FANOVA model is contaminated by different type of outlying curves. To do so, we use the same contamination models as in previous works on robust FDA (Fraiman and Muniz, 2001; López-Pintado and Romo, 2009; Sinova et al., 2018). All the details about the data generation process are provided in the supplementary materials.

One-way functional analysis of variance

The proposed simulation study framework for one-way FANOVA has been inspired by Cuevas et al. (2004); Górecki and Smaga (2015). Three different model M1, M2 and M3, with 3 level main effect f_i , $i = 1, 2, 3$, are considered, where the curves are defined for $\mathcal{T} = [0, 1]$. Model M1 corresponds to a situation where $H_0: f_1 = f_2 = f_3$ is true; in this case the empirical size is studied. M2 and M3 provides examples, with H_0 false, of monotone functions with different increasing patterns, where for the latter f_i are quite separated, whereas for the former differences are less apparent; in this case the empirical power is studied. Moreover, to simulate different type of outlying curves, seven contamination model C0-6 are considered. The model C0 is representative of no contamination; C1-4 are magnitude contaminations, i.e., generate curves far from the center, where C1 (C3) and C2 (C4) are symmetric and partial trajectories contamination models, that are independent (dependent) of the level of the main effect. Models C5-6 are shape contamination models (López-Pintado and Romo, 2009; Sinova et al., 2018).

In all the cases considered, the response curves are independent realizations of a Gaussian process with covariance function $\gamma(s, t) = \sigma^2 e^{(-|s-t|^{10^{-5}})}$ and are observed through 25 evenly spread discrete points with σ equal to $\sigma_1 = 1/25$, $\sigma_2 = 1.8/25$, $\sigma_3 = 2.6/25$, $\sigma_4 = 3.4/25$, $\sigma_5 = 4.2/25$, $\sigma_6 = 5/25$ (Cuevas et al., 2004). We expect the higher σ the worse the performance in terms of both empirical size and power. Fig. 9.4 shows the realizations of the response curve for the three models M1, M2, and M3 in presence of no contamination (C0) for $\sigma = \sigma_1$. Five implementations of the RoFANOVA method introduced in Section 9.2 are considered, which are defined by different choices of the loss function, that is the RoFANOVA with median loss ρ^{med} , referred to as RoFANOVA-MED, Huber loss ρ_a^{HU} , referred to as RoFANOVA-HUB, bisquare loss ρ_a^{BI} , referred to as RoFANOVA-BIS, Hampel loss $\rho_{a,b,c}^{HA}$,




Figure 9.4. The response curve $Y(t)$ realizations for the three models M1, M2, and M3 in presence of no contamination (C0) for $\sigma = \sigma_1$ in Scenario 1.

referred to as RoFANOVA-HAM, and, optimal loss ρ_a^{OP} , referred to as RoFANOVA-OPT. The tuning constants are chosen to achieve 95% asymptotic efficiency, the number of permutations B are set equal to 1000 and the functional 0.8% deepest curve following the FM criteria (Febrero-Bande and Oviedo de la Fuente, 2012) is chosen as starting value to compute the robust equivariant functional M -estimators (Section 9.2). The proposed tests are compared with some non robust methods already appeared in the literature. In particular, we consider the method proposed by Górecki and Smaga (2015), referred to as FP, which is a permutation test based on a basis function representation of the response function, the method proposed by Zhang and Liang (2014), referred to as GPF, based on a globalized version of the pointwise F -test, the method proposed by Zhang et al. (2007), referred to as L^2B , a L^2 -norm-based test with the bias-reduced method to estimate the unknown parameters, and, finally, the method proposed by Zhang (2011), referred to as FB, an F -type test based the bias-reduced estimation method. All these methods are implemented with the default settings of the R package `fdANOVA` (Gorecki and Smaga, 2018). In addition, the method proposed by Cuesta-Albertos and Febrero-Bande (2010), based on randomly chosen one-dimensional projections, with both the Bonferroni (referred to as TRPbon) and the false discovery rate (referred to as TRPfdr) corrections, is considered. The TRPbon and TRPfdr are run with 30 random projections through the R package `fda.usc` (Febrero-Bande and Oviedo de la Fuente, 2012).

For each triplet (Ml, Cm, σ_n) , $l = 1, \dots, 3$, $m = 0, \dots, 6$, $n = 1, \dots, 6$, the five proposed and the seven competitor methods are applied $N = 500$ times to the generated functional sample to test $H_0: f_1 = f_2 = f_3$ against $H_1: (H_0)^C$ at level $\alpha = 0.05$. Then, for each case, the empirical sizes (for model M1) and powers (for models M2 and M3) of the tests were computed as the proportional number of rejections out of the N replications with standard deviation equals at most to 0.0224 (corresponding to the case of probability of rejection equals to 0.5).

Fig. 9.5 displays the results for model M1, that is the empirical size of the eleven tests as a function of σ_n , $n = 1, \dots, 6$, for different contamination models (C0-6). In this case, the tests provide satisfactory results in controlling the level α , i.e., the empirical size is approximately less than or equal to 0.05, in case of no contamination (C0), symmetric magnitude contamination (C1-2) and shape contamination both symmetric (C5) and asymmetric (C6). On the contrary, for asymmetric magnitude contamination (C3-4), only the RoFANOVA tests based on redescending loss functions, that is RoFANOVA-BIS, RoFANOVA-HAM and RoFANOVA-OPT, are able to control the level α by assuring an empirical size approximately less or equal than 0.05. This is expected, because redescending estimators give no weight to observations far from the center (Maronna et al., 2019). The estimators used in the RoFANOVA-MED and RoFANOVA-HUB tests do not have this property and,

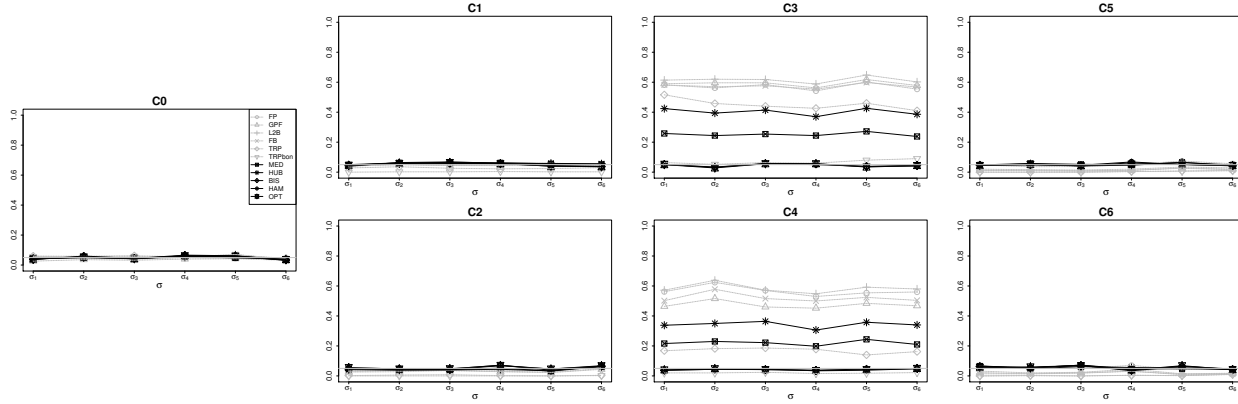


Figure 9.5. Empirical size of all tests for H_0 against H_1 (at level $\alpha = 0.05$) as a function of σ_n , $n = 1, \dots, 6$, for different contamination models (C0-6) in model M1 of Scenario 1. The proposed and competitor tests are displayed as black and grey lines, respectively.

thus, they suffer from the presence of contaminations depending on the level of the main factor. Note that, among the competitors, the TRPbon approximately controls the level for contamination model C4, while is slightly affected by outliers in model C5. This comes from the Bonferroni correction property of being conservative for high dimensional multiple comparisons (Lehmann and Romano, 2006).

In Fig. 9.6, the results, in terms of empirical power, for model M2 are shown. These tend to get worse as σ_n increases. In case of no contamination (C0), the FP test achieves the highest empirical power even though all the RoFANOVA tests have comparable results. For contamination model C1-6, it is extremely clear the proposed RoFANOVA tests outperform all the competitors. In particular, among the RoFANOVA tests, those based on redescending functional M -estimators, i.e., RoFANOVA-BIS, RoFANOVA-HAM and RoFANOVA-OPT, are the best ones. Note that, for model C3-4 only the RoFANOVA-BIS, RoFANOVA-HAM and RoFANOVA-OPT tests and the TRPbon (for C4) test should be considered because the other methods are not able to successfully control the level α (see Fig. 9.5).

Fig. 9.7 shows the empirical power for model M2. Also in this case, when σ_n increases, the empirical powers of the tests generally become smaller. The results are similar to those for model M2, even though the empirical power tend to be higher due to the more apparent separation of the main effect. Again, the proposed RoFANOVA test outperforms the competitors when there is contamination (C1-6) and have satisfactory power in case of no contamination (C0). The best results are achieved by the RoFANOVA-BIS, RoFANOVA-HAM and RoFANOVA-OPT tests.

Two-way functional analysis of variance

In this section, the two-way FANOVA model (9.4) is considered. The simulation design is inspired by Cuesta-Albertos and Febrero-Bande (2010).

As for Scenario 1, let $\mathcal{T} = [0, 1]$; then, the functional response depends on a grand mean m , 2 level main effects f_i and g_i , and interaction term h_{ij} through two parameter a and b with values in $\{0, 0.05, 0.10, 0.25, 0.50\}$. Here, the higher the values of a (b) the more far away f_i (g_i) is from the grand mean m . Thus, the empirical power should be an increasing function of a and b , respectively. The empirical size is studied for $a = 0$ or $b = 0$. The grand

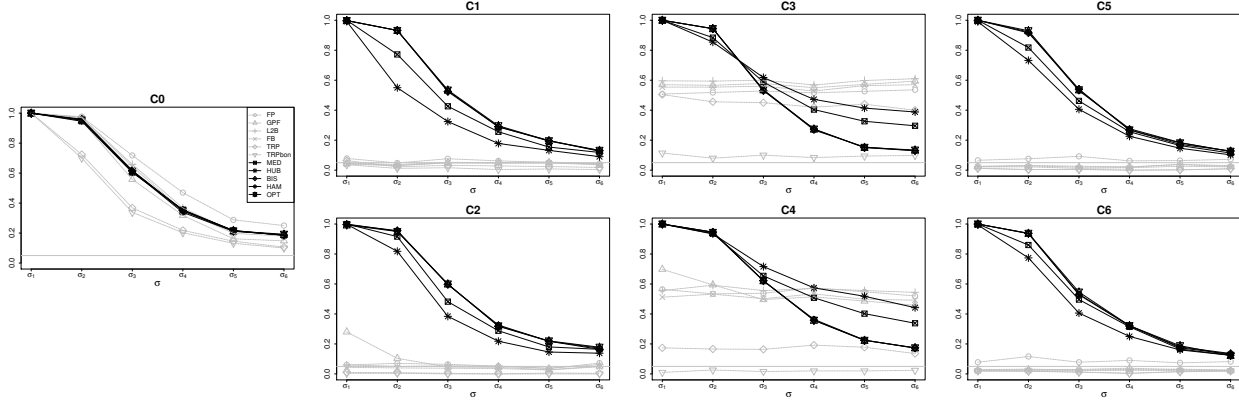


Figure 9.6. Empirical power of all tests for H_0 against H_1 (at level $\alpha = 0.05$) as a function of σ_n , $n = 1, \dots, 6$, for different contamination models (C0-6) in model M2 of Scenario 1. The proposed and competitor tests are displayed as black and grey lines, respectively.

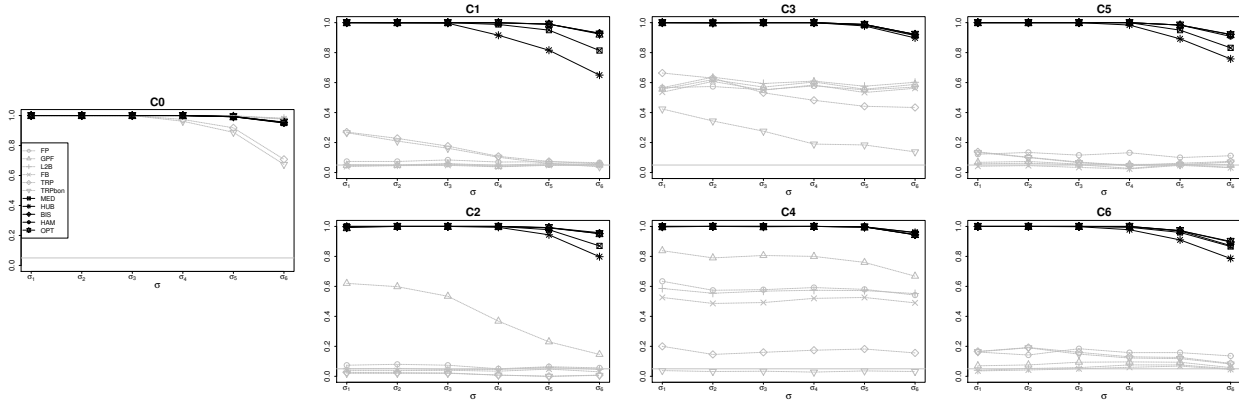


Figure 9.7. Empirical power of all tests for H_0 against H_1 (at level $\alpha = 0.05$) as a function of σ_n , $n = 1, \dots, 6$, for different contamination models (C0-6) in model M3 of Scenario 1. The proposed and competitor tests are displayed as black and grey lines, respectively.

mean m , as well as the 2 level main effects f_i and g_i are shown in Figure 9.8.

As in Scenario 1, seven contamination models C0-6 are considered, and the response curves are independent realizations of a Gaussian process with covariance function $\gamma(s, t) = \sigma^2 e^{-|s-t|10^{-5}}$. Data are observed through 25 evenly spread discrete points with $\sigma = 0.3$. Fig. 9.8 shows the realizations of the response curve for $a = b = 0$, $a = 0.5$ and $b = 0$, and, $a = 0$ and $b = 0.5$ in presence of no contamination (C0).

Also in this scenario, we consider the five versions of the proposed method, i.e., RoFANOVA-MED, RoFANOVA-HUB, RoFANOVA-BIS, RoFANOVA-HAM, and RoFANOVA-OPT, with tuning parameters chosen as in Scenario 1. As competitors, we consider (i) the permutation version of the method proposed by Zhang (2011), referred to as FNDP, which is permutation test based on a F -type statistic, and (ii) the global version of the method proposed by Pini and Vantini (2017), which is the two-way extension of the method of Zhang and Liang (2014), referred to as TGPF. Both for the FNDP and TGPF




Figure 9.8. The response curve realizations for $a = b = 0$, $a = 0.5$ and $b = 0$, and, $a = 0$ and $b = 0.5$ in presence of no contamination (C0) in Scenario 2.

methods, the distribution of the test statistic is approximated by using a Manly's scheme (Manly, 2006) with 1000 random permutations. Moreover, also the TRPbon and TRPfdm method (Section 9.3) are considered with 30 random projections.

For each triplet (Cm, a, b) , $m = 0, \dots, 6$, $a, b \in \{0, 0.05, 0.10, 0.25, 0.50\}$, the five proposed and the four competitor methods are applied $N = 500$ times to the generated functional sample to test $H_{0,A}$, $H_{0,B}$ and $H_{0,AB}$ against $H_{1,A}$, $H_{1,B}$ and $H_{1,AB}$ resp., at level $\alpha = 0.05$. Then, for each triplet and each test, the empirical sizes (when $a = b = 0$ for $H_{0,A}$, $H_{0,B}$ against $H_{1,A}$, $H_{1,B}$ and $a < 0.25$ for $H_{0,AB}$ against $H_{1,AB}$) and powers (when $a \neq 0$ or $b \neq 0$ for $H_{0,A}$, $H_{0,B}$ against $H_{1,A}$, $H_{1,B}$ and $a \geq 0.25$ for $H_{0,AB}$ against $H_{1,AB}$) of the tests were computed as the proportion of rejections out of N replications (also in this case, with maximum standard deviation equals to 0.0224).

For the sake of brevity, we summarize the results for cases that are statistically equivalent. For instance, when analyzing the null hypothesis $H_{0,A}$ (resp. $H_{0,B}$), for each value of a (resp. b) the five values corresponding to $b = \{0, 0.05, 0.10, 0.25, 0.50\}$ (resp. $a = \{0, 0.05, 0.10, 0.25, 0.50\}$) are summarized through their median. Similarly when analyzing $H_{0,AB}$ the values corresponding to $a < 0.25$ are substituted by their median for each value of b .

Fig. 9.9 shows the empirical size ($a = 0$) and the empirical power ($a \neq 0$) of all tests for $H_{0,A}$ against $H_{1,A}$ as a function of a . When a increases the performance of all the methods to reject $H_{0,A}$ enhance. In terms of empirical size (i.e., when $a = 0$), the results are quite satisfactory for all the methods in case of no contamination (C0), symmetric magnitude contamination (C1-2) and shape contamination both symmetric (C5) and asymmetric (C6). However, in case of asymmetric magnitude contamination (C3-4), only the RoFANOVA-BIS, RoFANOVA-HAM and RoFANOVA-OPT tests are able to control the level α , being approximately less or equal to 0.05. This behavior is analogous to Scenario 1 (Section 9.3). In terms of empirical power ($a \neq 0$), the proposed RoFANOVA test have comparable performance, when there are no outliers (C0), whereas they are far better than the competitors for the contamination models C1-6. Note that for asymmetric magnitude contamination (C3-4), only the RoFANOVA-BIS, RoFANOVA-HAM and RoFANOVA-OPT tests should be considered, being the only ones able to control the level α .

In Fig. 9.10, the empirical size ($b = 0$) and the empirical power ($b \neq 0$) of all tests for $H_{0,B}$ against $H_{1,B}$ (at level $\alpha = 0.05$) as a function of b are displayed. Also in this case, the proposed tests outperform the competitors in terms of power for contamination models C1-6 by simultaneously having comparable performance in presence of no contamination (C0). Moreover, all the tests are able to approximately control the level α , even for the contamination models C3-4 (differently from Scenario 1). This is expected, because, in this case, the asymmetry in the contamination affects only the main effect f_i by leaving

9. ROBUST FUNCTIONAL ANOVA WITH APPLICATION TO ADDITIVE MANUFACTURING

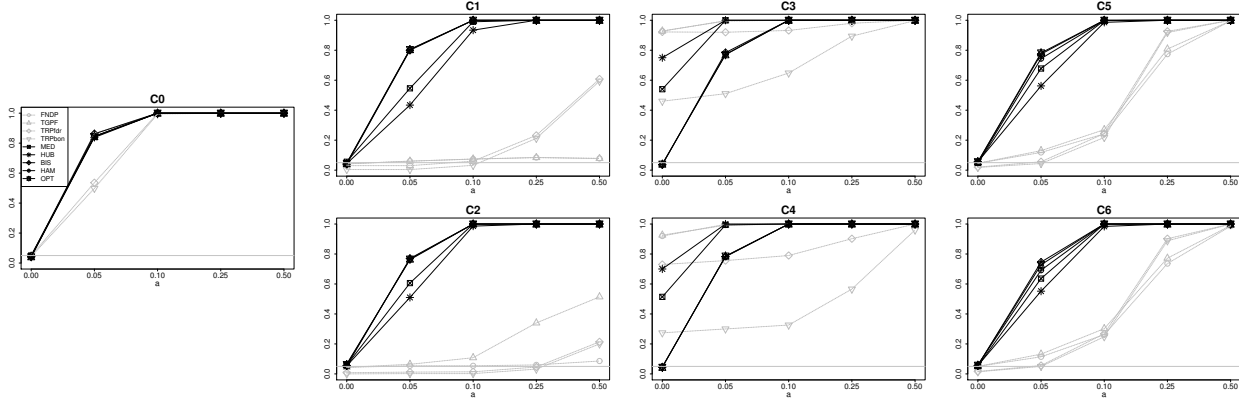


Figure 9.9. Empirical power and empirical size ($a = 0$) of all tests for $H_{0,A}$ against $H_{1,A}$ (at level $\alpha = 0.05$) as a function of a , for different contamination models (C0-6) in Scenario 2. The proposed and competitor tests are displayed as black and grey lines, respectively.

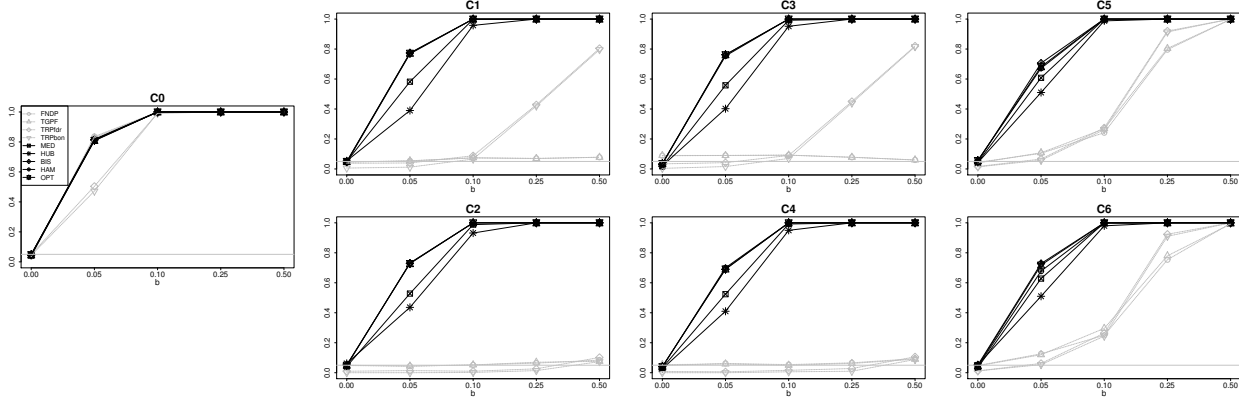


Figure 9.10. Empirical size ($b = 0$) and empirical power ($b \neq 0$) of all tests for $H_{0,B}$ against $H_{1,B}$ (at level $\alpha = 0.05$) as a function of b , for different contamination models (C0-6) in Scenario 2. The proposed and competitor tests are displayed as black and grey lines, respectively.

g_i unaffected. Among the proposed tests, the RoFANOVA-BIS, RoFANOVA-HAM and RoFANOVA-OPT ones tend to perform better than the ones based on monotonic functional M -estimator, i.e., the RoFANOVA-MED and RoFANOVA-HUB tests.

Fig. 9.11 shows the empirical size ($a \leq 0.10$ and $b = 0$) and empirical power ($a = 0.25, 0.50$ and $b \neq 0$) of all tests for $H_{0,AB}$ against $H_{1,AB}$ (at level $\alpha = 0.05$) as a function of b , for different contamination models (C0-6). In terms of empirical size ($a \leq 0.10$ and $b = 0$), all the test are able to approximately control the level α , except for the FNDP and TGPF tests for model C3 and for RoFANOVA-MED and RoFANOVA-HUB one for model C4 at $b = 0.50$. For $a = 0.25, 0.50$ and $b \neq 0$, the empirical powers of the proposed tests is much higher than those of the competitors for all the contamination model C1-6. Moreover, in case of no contamination (C0), the power of the RoFANOVA tests is comparable to that of the competitors for $a = 0.25, 0.50$. Also in this case, among the RoFANOVA tests, the RoFANOVA-BIS, RoFANOVA-HAM and RoFANOVA-OPT are the best ones.

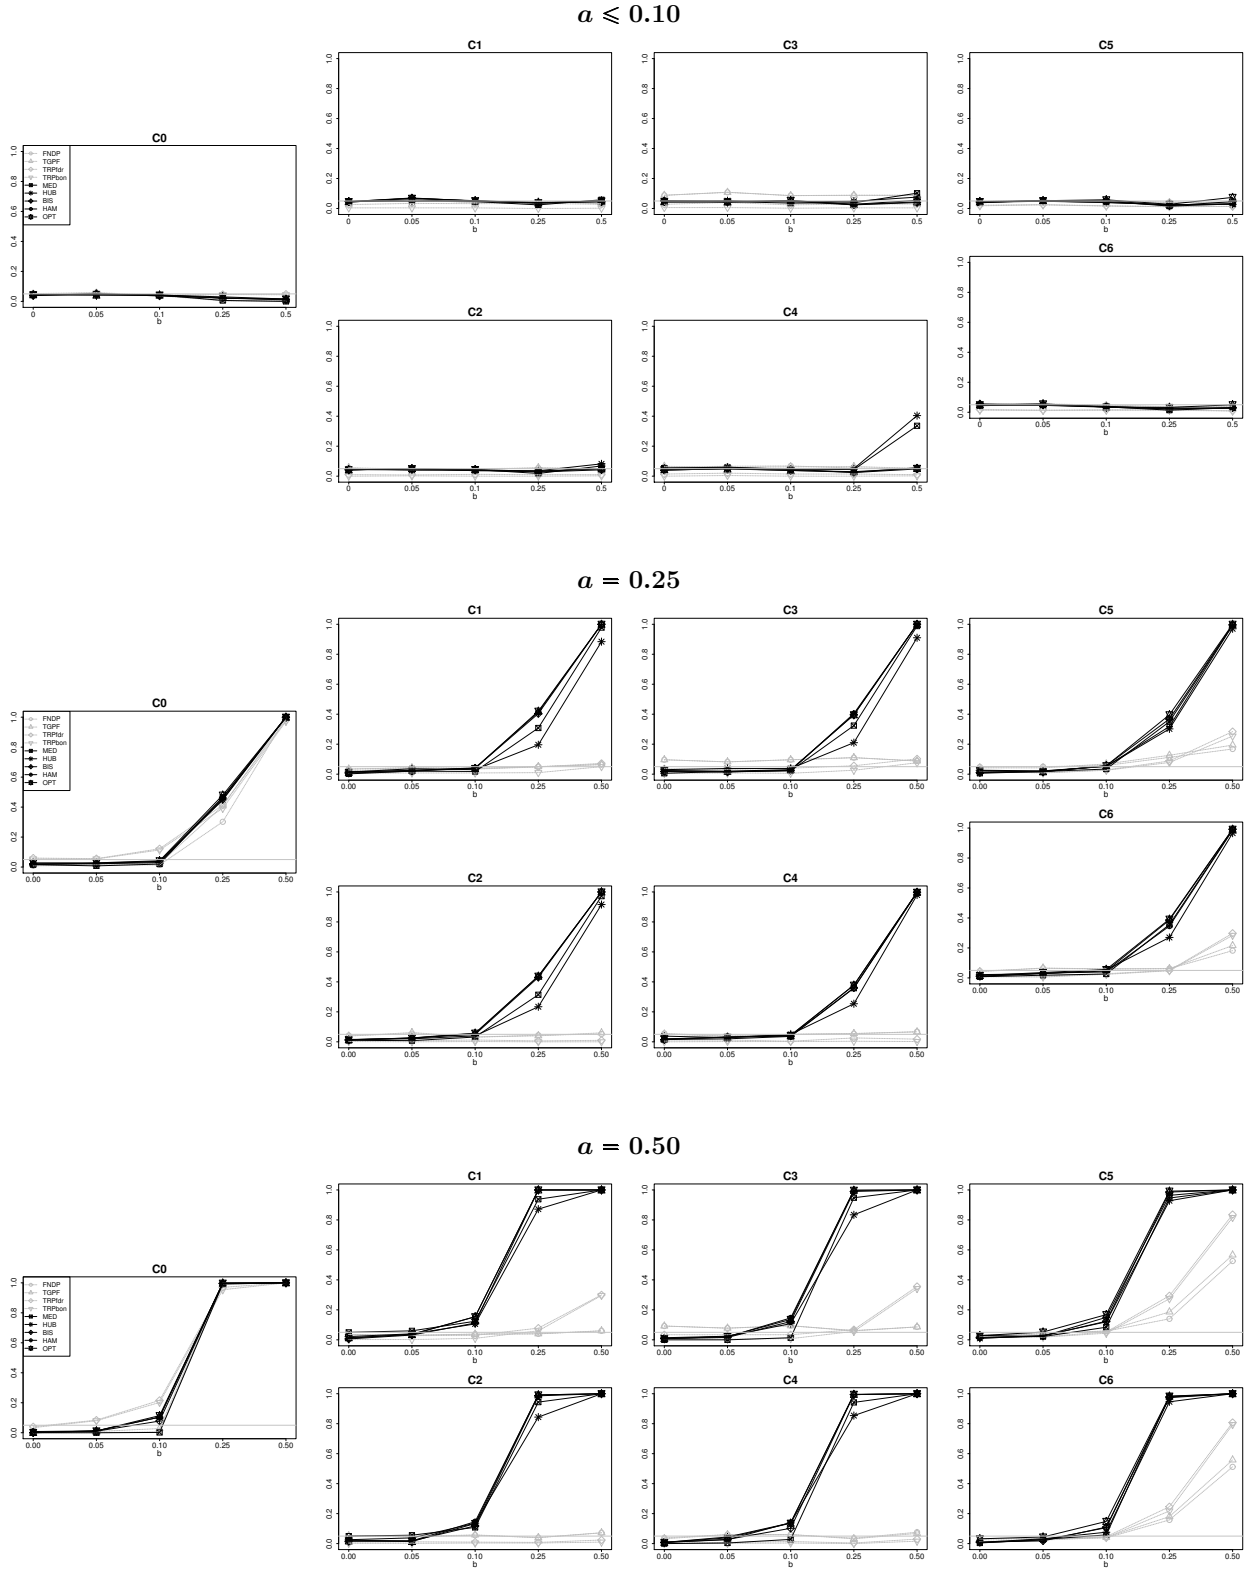


Figure 9.11. Empirical size ($a \leq 0.10$ and $b = 0$) and empirical power ($a = 0.25, 0.50$ and $b \neq 0$) of all tests against $H_{0,AB}$ against $H_{1,AB}$ (at level $\alpha = 0.05$) as a function of b , for different contamination models (C0-6) in Scenario 2. The proposed and competitor tests are displayed as black and grey lines, respectively.

9.4 Real Case-study: Analysis of Variance of Applied to the Analysis of Spatter Behaviour in Laser Powder Bed Fusion

To demonstrate the potential of the proposed approach, this Section presents the real case-study in additive manufacturing. In L-PBF, spatters are process by-products that can be ejected either by the melt pool in the form of hot and liquid droplets or by the powder bed regions surrounding the melt pool (Young et al., 2020; Ly et al., 2017; Bidare et al., 2018). In the latter case, spatters consist of powder particles entrained by convective motions above and around the melt pool. In case they pass through the hot material vaporization region above the melt pool, they are quickly heated up and become hot particles moving away from the powder bed. For more details about the spatter generation mechanism, the reader is referred to (Young et al., 2020; Ly et al., 2017; Bidare et al., 2018) and the literature cited therein. The analysis of L-PBF process by-products has gathered an increasing interest in the last years because they can be driver of relevant information about the process state and the final quality of the manufactured part (Yang et al., 2020; Tan et al., 2020; Yin et al., 2020; Andani et al., 2017; Repossini et al., 2017; Ly et al., 2017; Bidare et al., 2018). Studying the effect of controllable process factors and other operating conditions on the spatter behaviour allows getting a deeper comprehension of underlying physical phenomena. Such knowledge may be used to tune the process conditions to enhance the quality and mechanical performances of the products, or to design in-line and real-time process monitoring methodologies that, by looking at the behaviour of process by-products, allow the detection of possible anomalies and unstable process states (Colosimo and Grasso, 2020).

Hot spatters ejected as a consequence of the laser-material interaction can be observed by means of high-speed cameras installed into the L-PBF machine or placed outside its viewports. The mainstream literature devoted to spatter analysis and monitoring in L-PBF relies on video image processing methods to compute synthetic indexes that capture salient aspects of the spatter behaviour, e.g., the number of ejected spatters in each video frame, their size, speed, travelled distance from the melt pool, etc. (Grasso et al., 2017; Everton et al., 2016). In the real case study presented in this Section, instead of treating synthetic descriptors of the spatter ejections as univariate or multivariate variables, they are translated into a functional form by means of the spatter intensity function introduced in Section 1. Such function captures the spatial spread of ejected spatter and can be estimated for each manufactured layer and for each test treatment. Section 4.1 presents the main experimental settings, whereas the results of the analysis and the comparison against benchmark methods are reported in Section 4.2.

Experimental Setting and Data Pre-processing

The case study involves the production of specimens of size 5 x 5 x 12 mm via L-PBF of 18Ni(300) maraging steel powder, a steel alloy commonly used for tooling applications, with average particle size between 25 and 35 200μ . An industrial L-PBF system, namely a Renishaw AM250, was used, with a high-speed camera in the visible range placed outside the front viewport of the machine as shown in Fig. 9.12, left panel. Videos were recorded during the production of six layers with a sampling rate of 1000 fps (frames per second) and a spatial resolution of about 200μ m/pixel. Specimens were placed as shown in Fig. 9.12, right panel, and produced by varying the energy density provided by the laser to the material as shown in Table 9.1.

The laser was displaced by a scanner along a predefined path consisting of parallel scan

9.4. Real Case-study: Analysis of Variance of Applied to the Analysis of Spatter Behaviour in Laser Powder Bed Fusion




Figure 9.12. Setup of the high-speed camera in front of the Renishaw AM250 machine's viewport (left panel) and placement of manufactured specimens in the build area, within the camera's field of view (right panel): numbers shown in the specimens correspond to the energy density level, from 1 to 6, applied during the process.

Table 9.1. Process parameters and corresponding energy density levels.

Energy density level	Laser exposure time t (μ s)	Distance between exposed points along laser scan track dp (μ m)	Distance between parallel laser scan tracks dh (μ m)	Laser power P (W)	Powder bed thickness z (μ m)	Energy density F (kJ/cm ³)
1	39	65	80	200	50	30
2	85	85	80	200	50	50
3	104	65	80	200	50	80
4	125	62.5	80	200	50	100
5	115	50	80	200	50	115
6	104	40	80	200	50	130

lines, whose orientation changed layer by layer, with a default rotation of about 67° every layer. Along each scan line, the laser melts the material with a pulsed mode, i.e., by exposing points equispaced apart of a quantity d along each scan line with a point exposure duration, t . The energy density was varied by varying t and d . The energy density is known to be a factor of primary importance to determine the quality of the process, as insufficient or excessive energy densities may produce defects, like internal pores, and other deviations from the expected quality (Grasso et al., 2017; Everton et al., 2016; Mani et al., 2015). Within the build chamber, where the L-PBF process takes place, a laminar flow of inert gas, called shielding gas, is used to prevent ejected spatters from falling on the build area with consequent potential contamination effects, and vaporized material from depositing on the laser window leading to possible attenuation of the laser beam (Anwar and Pham, 2018).

The functional response variable, i.e. the so-called spatter intensity, was estimated by applying the video image pre-processing method presented in Repossini et al. (2017). Each acquired frame consists of a dark background and different hot (bright) areas in foreground. Those bright areas correspond either to spatters (being visible as particles thanks to the very low video image integration time) and the region of the powder bed heated up by the laser, also known as laser-heated zone. A perspective correction was first applied to acquired videos to compensate the image distortion due to the inclination of the camera above the build area. Then, the resulting grey-scale frames were converted into a binary format by means of a thresholding algorithm to isolate and identify each single connected component (i.e., spatters and the laser-heated zone). A criterion based on the size and location of each connected component was used to filter out the laser heated zone and focus the analysis on all other regions of interest, i.e., the spatters. Eventually, the centroid of each spatter

9. ROBUST FUNCTIONAL ANOVA WITH APPLICATION TO ADDITIVE MANUFACTURING

Table 9.2. Location of analysed layers along the specimen build direction (distance from the baseplate) and orientation of the laser scan direction relative to the shielding gas flow in each layer.

Analysed layer	Layer height along the build direction (mm)	Laser scan angle relative to the shielding gas flow
1	31	10°
2	56	40°
3	83	80°
4	110	85°
5	137	90°
6	163	30°

in the frame was computed and used to determine the spatial coordinates, (s, t) , of each detected spatter. Additional details about the video image pre-processing steps can be found in Repossini et al. (2017). In order to spatially map the amount of spatters ejected during the production of each specimen in each layer, three additional pre-processing operations were performed. First, the location of spatters was referred to a spatial domain centered in the center of the scanned area of each specimen, to allow comparing the functional response variables for specimens produced in different locations. Second, the spatial domain was discretized into 60 by 80 adjacent squared cells, in order to count the number of spatters ejected in each layer within each cell. Based on these pre-processing steps, the spatial spread of the spatters, in each layer and for each specimen, could be summarized into the function $Y_{i,j,k}(s, t)$ defined on the bi-dimensional domain $\mathcal{T} = [0, 1] \times [0, 1]$, where indices $i = 1 \dots, 6$, $j = 1, \dots, 6$ and $k = 1, \dots, n_{ij}$ indicate the energy density level, the layer, and the number of replicates (specimens) for each treatment. It is worth noticing that a possible layer effect could be related to two different aspects. The first regards the orientation of the scan lines that is varied layer by layer. Indeed, it is known that the orientation of the laser scan with respect to the shielding gas flow direction affects the spatter trajectory and hence the way in which they displace above the powder bed (Bidare et al., 2018). The relative orientation of the laser scan direction with respect to the shielding gas flow for each layer is shown in Table 9.2. The second regards a variation of the material thermal properties as more and more layers are produced. In the present case study, this second effect is assumed to be negligible with respect to the first one, due to the small size and simple shape of the specimens. The spatter intensity function $Y_{i,j,k}(s, t)$ is a smoothed version of the actual amount of spatters counted in every location of the spatial domain. The number of replicates n_{ij} is fixed and equal to 3, as three specimens were produced for each energy density level, except for $i = 6$ and $j = 1$ where $n_{ij} = 2$, due to a missing data point (a delamination occurred in initial layers prevented from producing one of the three specimens with the lowest energy density level). The $Y_{i,j,k}$ are obtained by means of a smoothing phase based on tensor product bases of cubic splines with second derivative penalty as marginal smooths. The marginal basis dimensions, set equal to 30, and the smoothing parameter were chosen by using restricted maximum likelihood (REML) (Wood, 2017). The smoothing phase was performed by using the R package `mgcv` (Wood, 2017). Then, in order to reduce phase variability, a registration phase was performed (Ramsay and Silverman, 2005). It consists in the shifting of each $Y_{i,j,k}$ along the s and t axes to minimize the L^2 distance with respect to the reference curve, which

9.4. Real Case-study: Analysis of Variance of Applied to the Analysis of Spatter Behaviour in Laser Powder Bed Fusion




Figure 9.13. The functional observations $Y_{i,j,k}$ for $t = 0.75$, (a) and (c), and $s = 0.5$, (b) and (d), in the real-case study, for different fluency levels ((a) and (b)) and different scan strategies ((c) and (d)).

was chosen such that the mean of pairwise distances among the aligned curves is minimum. The functional observations $Y_{i,j,k}$, $i = 1, \dots, 6$, $j = 1, \dots, 6$ and $k = 1, \dots, n_{ij}$, for $t = 0.75$ and $s = 0.5$ are represented in Fig. 9.13 at different energy density levels and in different layers. The graphical representation of cross-sections of the spatter intensity function in Fig. 9.13 was adopted to aid the superimposition and direct comparison of functional patterns corresponding to different experimental treatments.

Results

The spatter intensity functions $Y_{i,j,k}$ ($i = 1, \dots, 6$, $j = 1, \dots, 6$ and $k = 1, \dots, n_{ij}$) are modeled according to (9.4), where f_i is the energy density functional effect, g_i is the layer functional effect, h_{ij} is the interaction term between the energy density and the layer.

The aim of the analysis is therefore to test the energy density effect $H_{0,Flu} = H_{0,A}$ (9.5), the layer effect $H_{0,Lay} = H_{0,B}$ (9.6) (mainly related to the layer by layer variation of the laser scan direction) and their interaction effect $H_{0,FluLay} = H_{0,AB}$ against the alternatives $H_{1,Flu} = H_{1,A}$ (9.5), $H_{1,Lay} = H_{1,B}$ (9.6) and $H_{1,FluLay} = H_{1,AB}$. As already introduced in Section 9.1, some outliers could be present. In particular, Fig. 9.14 shows (a) the residuals of the fitted model for $t = 0.75$ (the approximate t value of the spatter intensity peak), obtained by using the RoFANOVA-BIS test as implemented in Section 9.3, and (b) the boxplot of their L^1 norms, defined as $\|f\|_1 = \int_{\mathcal{T}} |f(t)| dt$, for $f \in L^2(\mathcal{T})$. Because $\mathcal{T} = [0, 1] \times [0, 1]$, the L^1 norm can be interpreted as the average value of the function over its domain. It is clear from Fig. 9.14 that some outliers are present in this real case study. However, except from few points that have residuals far from the bulk of the data, there are some points that could not be easily labeled as outliers. As mentioned in the introduction, the L-PBF process is characterized by complex dynamics and many transient and local phenomena that not




Figure 9.14. (a) Residuals of the fitted model for $t = 0.75$, obtained by using the RoFANOVA-BIS test as implemented in Section 9.3, and (b) boxplot of their L^1 norms.

Table 9.3. p-values of all robust tests for $H_{0,Flu}$, $H_{0,Lay}$ and $H_{0,FluLay}$ against $H_{1,Flu}$, $H_{1,Lay}$ and $H_{1,FluLay}$.

	RoFANOVA-MED	RoFANOVA-HUB	RoFANOVA-BIS	RoFANOVA-HAM	RoFANOVA-OPT
$H_{0,FluLay}$	0.00	0.01	0.00	0.00	0.00
$H_{0,Flu}$	0.00	0.00	0.00	0.00	0.00
$H_{0,Lay}$	0.00	0.00	0.00	0.00	0.00

only affect the natural variability of the measured quantities, but could lead also to outlying pattern. Determining whether an experimental point is an outlier and identifying its root causes can be a difficult task, which makes the diagnostic approach hardly applicable in the absence of additional data and information.

Therefore, we applied the RoFANOVA test described in Section 9.3, i.e., the RoFANOVA-MED, RoFANOVA-HUB, RoFANOVA-BIS, RoFANOVA-HAM, and, RoFANOVA-OPT tests, specifically adapted for bi-dimensional functional data. As in the Monte Carlo simulation study (9.3) the tuning constants are chosen to achieve 95% asymptotic efficiency, the number of permutations B are set equal to 1000. In this case, the functional sample mean is used as starting value to compute the robust equivariant functional M -estimators (Section 9.2). The results are shown in Table 9.3. All the tests agree in considering significant the interaction between the energy density and the layer.

When an the interaction effect is present, it is well-known that an interpretation of the main effects becomes less straightforward than if the interaction is not significant (Miller Jr, 1997) because, the layer effect upon the spatter intensity will differ depending on the energy density level. In this case, the best way to interpret the results is through the interaction plot (Montgomery, 2017), which graphically represents the response means at different factor levels. Fig. 9.15 shows an interaction plot adapted to deal with bi-dimensional data. In particular, the L^1 norms of the group means corresponding to the RoFANOVA-BIS test are plotted as a function of the energy density level and the layer. In this case, if an interaction is present, then, the trace of the average response across the levels of one factor plotted separately for each level of the other factor will not be parallel (Montgomery, 2017).

Fig. 9.15 shows that, as the energy density increases, the spatter intensity tends to increase as well. This is in agreement with the fact that a higher energy density generates a larger and hotter melt pool with more intense convective and recoil motions, which translate into a more intense spatter ejection (Yang et al., 2020; Repossini et al., 2017; Bidare et al., 2018). More interestingly, Fig. 9.15 shows different patterns corresponding to different layers.

Indeed, in layers 1, 2, and 6, there is an approximately linear trend of the spatter intensity with respect to the energy density. These three levels were characterized by very similar laser scan directions, with a low angle relative to the shielding gas flow (between 10° and 40°). Since most spatters are ejected forward or backward along the scan direction, when the scan direction is parallel (or little angled) to the gas flow, more powder bed particles are pushed along the laser path, increasing the occurrence of particles being heated up by the hot metal vapour emission and being then ejected as hot spatters. Under these conditions, increasing the energy density increases the intensity of convective motions that entrap the powder particles into the hot vapour emission and hence the spatter intensity (Bidare et al., 2018).

A different influence of the energy density on the spatter intensity was observed in layers 3, 4 and 5. In these layers, the laser scan direction was almost perpendicular to the shielding gas flow direction, i.e., with angles in the range 80° to 90° . Under these conditions, particles are dragged away from the scan path, reducing the amount of particles entrained in the hot vapour plume above the melt pool, and hence reducing the overall spatter intensity with respect to layers where lower scan angles were applied (Bidare et al., 2018). In addition, the analysis reveals that when the laser scan direction was about perpendicular to the gas flow, there was a range of intermediate energy densities (from level 3 to level 5) at which the influence of the energy density itself on the spatter intensity was reduced or even inverted. This can be interpreted as follows. When the laser scan direction is parallel to the gas flow, an increase of the energy density causes an increase of convective motions and metal vapour emissions that result also in higher spatter intensity. When the laser scan direction is perpendicular to the gas flow, an increase of the energy density still causes an increase of convective motions and metal vapour emissions, but such vapour emission has little effect on the spatter intensity, which makes the influence of the energy density mainly evident at very low or very high energy density levels only. Such interaction between the energy density and the laser scan direction on the spatter intensity was explored in very few studies in the literature, but it is particularly relevant to understand the underlying by-product behaviour and to design either process optimization or process monitoring tools that rely on the in-line observation of such by-products.

Finally, we cannot affirm with a sufficient confidence that the spatter intensity is affected by the layer (i.e., by the laser scan direction) as well, because it is not clear if the differences among layers are results of interactions only or there is also a systematic laser scan direction effect.

Even if the use of the RoFANOVA tests is recommended because the results in the Monte Carlo simulation study (Section 9.3), for the sake of completeness, the bi-dimensional version of the FNDP and TGPF test have been applied. For the latter, the Manly's scheme (Manly, 2006) with 1000 random permutations are used to approximate the test statistic distribution. The additional results are shown in Table 9.4. By comparing the proposed tests (Table 9.3) with the competitors, they disagree in considering significant the interaction between the energy density and the layer. In particular, the FNDP and the TGPF tests accept the null hypothesis of no interaction (i.e., large p-values). Therefore, analogously to the Monte Carlo simulation results suggests in case of two-way FANOVA design (Section 9.3), the FNDP and the TGPF tests have not as statistical power as needed to detect a significant interaction among the main factor. In this regard, it appears that the proposed tests spot a technologically relevant interaction that the FNDP and the TGPF tests fail in finding significant.




Figure 9.15. Interaction plot the energy density level and the layer (i.e., the laser scan direction) in the real case-study.

Table 9.4. p-values of the FNDP and TGPF tests for $H_{0,Flu}$, $H_{0,Lay}$ and $H_{0,FluLay}$ against $H_{1,Flu}$, $H_{1,Lay}$ and $H_{1,FluLay}$.

	FNDP	TGPF
$H_{0,FluLay}$	0.72	0.23
$H_{0,Flu}$	0.00	0.00
$H_{0,Lay}$	0.00	0.00

9.5 Conclusions

In this chapter, we have proposed the RoFANOVA test for the functional analysis of variance problem. In particular, the proposed method has been designed to be robust against functional outliers, which are increasingly common in complex problems and, as it is well known, can severely bias the analyses. Robustness comes from the use of robust test statistics based on the functional equivariant M -estimator and the functional normalized median absolute deviation, which are the extensions of the classical M -estimator and normalized median absolute deviation to functional data. The test statistic is, then, incorporated in a permutation test, in order to solve the FANOVA problem in a non parametric fashion. However, the proposed approach is very general, as different versions of the proposed test emerge by the choices of particular loss functions, which are specifically designed to reduce the abnormal observation weights in the computation of the test statistic in comparisons with the standard least-squares loss function, and, it is designed for both one-dimensional and bi-dimensional functional data. To the best of the authors' knowledge, this is the first example of a robust method for the FANOVA problem in the literature, where attention has been mainly focused on non robust methods.

The performance of the proposed method have been investigated by means of an extensive Monte Carlo simulation study, where the proposed RoFANOVA tests have been compared with other methods already present in the literature. The results showed that the proposed

tests outperform all the competitors in presence of outlier contamination in terms of both empirical size and empirical power. Moreover, the loss of power between the RoFANOVA tests and the competitors is negligible when no outlier contamination is present.

Eventually, the proposed method was applied to a motivating real case-study in additive manufacturing. The aim of the analysis was to study the effects of the energy density and laser scan direction on the bi-dimensional distribution of the spatter intensity, i.e., a function mapping the amount of spatters ejected within the bi-dimensional field of view of the machine vision system used for in-line measurements. Apart from the known influence of the energy density on the spatter intensity, in agreement with previous study, the RoFANOVA test revealed a statistically significant interaction between the energy density and the laser scan direction relative to the shielding gas flow. The knowledge of such interaction is of particular importance to interpret and understand the underlying process by-product dynamics and to design method that use such by-products as proxies of the process quality and stability. The statistical significance of the interaction between the two factors was not identified by other non robust tests, which confirm the effectiveness of the proposed approach for an application where complex process dynamics may lead to outlying patterns that contaminate the experimental dataset. However, the validity of the proposed approach is completely general, and not limited to the case study here presented and even not limited to manufacturing applications.

In future research, the effects of heteroscedasticity on the RoFANOVA tests should be investigated in order to be able to deal with a wider variety of settings. In addition, some efforts should be made to extend the proposed robust method to more complex FANOVA designs.

9.6 Supplementary Materials

Derivation of the Constant c in the FuNMAD Expression

Following Theorem 3.4 of Sinova et al. (2018), $\hat{\mu}_{s,med}$ is a strongly consistent estimator of $\tilde{\mu}_{s,med} = \operatorname{argmin}_{y \in L^2(\mathcal{T})} \mathbb{E} [||X_i - y||]$. Moreover, let assume that X is a Gaussian random process, than, by Proposition 3.2 of Sinova et al. (2018), $\tilde{\mu}_{s,med} = \mu$, where μ is the mean function of the random element X . Therefore, from the population version of Equation (9.3) by applying for each $t \in \mathcal{T}$ the definition of univariate population median, we have asymptotically

$$0.5 = \Pr [|X(t) - \mu(t)| < c \text{FuNMAD}(t)] = \Pr \left[|Z| < c \frac{\text{FuNMAD}(t)}{\sigma(t)} \right],$$

where Z is a standard normal random variable. Therefore, we must have

$$\Phi \left(c \frac{\text{FuNMAD}(t)}{\sigma(t)} \right) - \Phi \left(-c \frac{\text{FuNMAD}(t)}{\sigma(t)} \right) = 0.5,$$

where Φ is the cumulative distribution function of the standard normal distribution. Noticing that

$$\Phi \left(-c \frac{\text{FuNMAD}(t)}{\sigma(t)} \right) = 1 - \Phi \left(c \frac{\text{FuNMAD}(t)}{\sigma(t)} \right),$$

then,

$$c \frac{\text{FuNMAD}(t)}{\sigma(t)} = \Phi^{-1}(3/4) = 0.6745, \quad (9.8)$$

where Φ^{-1} is the quantile function for the standard normal distribution.

Therefore, if $c = 0.6745$ then FuNMAD is an asymptotically pointwise consistent estimator of σ .

Details on Data Generation

In this section, the data generation process for Scenario 1 and Scenario 2 of the simulation study is described. For Scenario 1, let $\mathcal{T} = [0, 1]$, then, three different model with the following 3 level main effect f_i are considered

$$\text{M1 } f_i(t) = t(1-t) \text{ for } t \in [0, 1] \text{ and } i = 1, 2, 3,$$

$$\text{M2 } f_i(t) = t^i(1-t)^{6-i} \text{ for } t \in [0, 1] \text{ and } i = 1, 2, 3,$$

$$\text{M3 } f_i(t) = t^{i/5}(1-t)^{6-i/5} \text{ for } t \in [0, 1] \text{ and } i = 1, 2, 3.$$

Case M1 corresponds to a situation where $H_0: f_1(t) = f_2(t) = f_3(t)$, for $t \in [0, 1]$, is true; M2 and M3 provides examples, with H_0 false, of monotone functions with different increasing patterns, where for the latter f_i are quite separated, whereas for the former differences are less apparent. To simulate different type of outlying curves, let B and U two independent random variables following a Bernoulli (with parameter p) and a discrete uniform on $\{-1, 1\}$ distributions, respectively, and T a random number generated from a uniform distribution on $(0, 0.75)$. Then, the following contamination models C_i are considered

$$\text{C0 } C_i(t) = 0 \text{ for } t \in [0, 1] \text{ and } i = 1, 2, 3,$$

$$\text{C1 } C_i(t) = BUM \text{ for } t \in [0, 1] \text{ and } i = 1, 2, 3,$$

$$\text{C2 } C_i(t) = \begin{cases} BUM & \text{if } t \geq T \\ 0 & \text{if } t < T, \end{cases} \text{ for } t \in [0, 1] \text{ and } i = 1, 2, 3,$$

$$\text{C3 } C_i(t) = (-1)^i BM \text{ for } t \in [0, 1] \text{ and } i = 1, 2, 3,$$

$$\text{C4 } C_i(t) = \begin{cases} (-1)^i BM & \text{if } t \geq T \\ 0 & \text{if } t < T, \end{cases} \text{ for } t \in [0, 1] \text{ and } i = 1, 2, 3,$$

with contamination size constant $M = 25$ and $p = 0.1$. The model C0 is representative of no contamination; C1 (C3) and C2 (C4) are symmetric and partial trajectories contamination models, that are independent (dependent) of the level i of the main effect. Then, the curves X_{ik} are generated, for $i = 1, 2, 3$ and $k = 1, \dots, 20$, as

$$X_{ik}(t) = f_i(t) + C_i(t) + \varepsilon_{ik}(t) \quad t \in [0, 1],$$

where the errors ε_{ik} are independent Gaussian processes with mean zero and covariance function $\gamma(s, t) = \sigma^2 e^{(-|s-t|10^{-5})}$. Models C1-4 are magnitude contaminations, i.e., generate curves far from the center, in what follows we consider two shape contamination models (López-Pintado and Romo, 2009; Sinova et al., 2018) that are both independent and dependent of the level i of the main effect. In this setting, the curves X_{ik} are generated, for $i = 1, 2, 3$ and $k = 1, \dots, 20$, as

$$X_{ik}(t) = (1 - B)Y_{ik}(t) + BZ_{ik}(t) \quad t \in [0, 1],$$

with

$$Y_{ik}(t) = f_i(t) + \varepsilon_{ik}(t), \quad Z_{ik}(t) = f_i(t) + \varepsilon_{ik,c}(t) \quad t \in [0, 1],$$

where $\varepsilon_{ik,c}$ are independent Gaussian processes with mean zero and covariance function $\gamma_{i,c}(s, t) = \sigma^2 e^{(-|s-t|k_{\gamma_c,i}10^{-5})}$. The following choices for $k_{\gamma_c,i}$ are considered

C5 $k_{\gamma_c,i} = 10^2$ for $i = 1, 2, 3$,

C6 $k_{\gamma_c,i} = 10^{2+i}$ for $i = 1, 2, 3$.

In all the cases considered, the curves X_{ik} are observed through 25 evenly spread discrete points and σ is equal to $\sigma_1 = 1/25$, $\sigma_2 = 1.8/25$, $\sigma_3 = 2.6/25$, $\sigma_4 = 3.4/25$, $\sigma_5 = 4.2/25$, $\sigma_6 = 5/25$.

For Scenario 2, let consider $\mathcal{T} = [0, 1]$, then, the functional response depends on the grand mean m , 2 level main effects f_i and g_i , and interaction term h_{ij} through two parameter a and b as follows

- $m(t) = t(1-t)$ for $t \in [0, 1]$,
- $f_i(t) = a(-1)^i |\sin(4\pi t)|$ for $t \in [0, 1]$ and $i = 1, 2$,
- $g_j(t) = b(-1)^j I(t > 0.5)$ for $t \in [0, 1]$ and $j = 1, 2$,
- $h_{ij}(t) = -f_i(t)g_j(t)I(a \geq 0.25)$ for $t \in [0, 1]$ and $i = 1, 2$, $j = 1, 2$,

with $a, b \in \{0, 0.05, 0.10, 0.25, 0.50\}$. For the contamination models C0-4 (Section 9.3), the curves X_{ijk} are generated, for $i = 1, 2$, $j = 1, 2$ and $k = 1, \dots, 20$, as

$$X_{ijk}(t) = m(t) + f_i(t) + g_j(t) + h_{ij}(t) + C_i(t) + \varepsilon_{ijk}(t) \quad t \in [0, 1],$$

where the errors ε_{ijk} are independent Gaussian processes with mean zero and covariance function $\gamma(s, t) = \sigma^2 e^{(-|s-t|10^{-5})}$. The curves X_{ijk} are observed through 25 evenly spread discrete points.

Whereas, for the contamination models C5-6 (Section 9.3), the curves X_{ijk} are generated, for $i = 1, 2$, $j = 1, 2$ and $k = 1, \dots, 20$, as

$$X_{ijk}(t) = (1-B)Y_{ijk}(t) + BZ_{ijk}(t) \quad t \in [0, 1],$$

with

$$Y_{ijk}(t) = m(t) + f_i(t) + g_j(t) + h_{ij}(t) + \varepsilon_{ijk}(t), \quad Z_{ijk}(t) = m(t) + f_i(t) + g_j(t) + h_{ij}(t) + \varepsilon_{ijk,c}(t),$$

for $t \in [0, 1]$, where $\varepsilon_{ijk,c}$ are independent Gaussian process with mean zero and covariance function $\gamma_{ij,c}(s, t) = \sigma^2 e^{(-|s-t|k_{\gamma_c,i}10^{-5})}$ with, as for Scenario 1, $k_{\gamma_c,i} = 10^2$ for C5 and $k_{\gamma_c,i} = 10^{2+i}$ for C6. The random variable B follows a Bernoulli (with parameter $p = 0.1$) distribution. In this case, the curves X_{ik} are observed through 25 evenly spread discrete points with $\sigma = 0.3$.

Chapter 10

Functional Real-time Monitoring Control Chart

This chapter sketches the main elements of a new method for real-time monitoring of functional data, which is subject of an ongoing research with Murat Kulahci and Max Spooner of the DTU Compute, Technical University of Denmark, and Biagio Palumbo and Antonio Lepore of the Univeristy of Naples.

Recent improvements in data acquisition technologies have produced data-rich environments in every field. Often such data are big, but more often they are complex and high dimensional. Particularly relevant is the case where data are apt to be modelled as functions defined on multidimensional domain, which are referred to as *functional data*. Functional data analysis (FDA) deals with the analysis and theory of functional data (Ramsay and Silverman, 2005; Kokoszka and Reimherr, 2017). A typical problem in industrial applications deals with evaluating the stability over time of some quality characteristics of interest. This problem is addressed by statistical process control (SPC) methods (Montgomery, 2007). Profile monitoring (Noorossana et al., 2012) is the suite of SPC methods that deal with quality characteristics that are functional data.

Most of the profile monitoring methods consider the case where the aim of the analysis is to assess the stability of the functional quality characteristic in its entirety (Centofanti et al., 2020b; Capezza et al., 2020). However, in some applications, the interest relies in understanding if the process is working properly before its completion. In other words, in these applications, practitioners want to understand, while the process is still running, if only normal sources of variation (i.e., common causes) apply to the process, i.e., the process is in-control (IC), or assignable sources of variations (i.e., special causes) act on it, that is the process is out-of-control (OC). This kind of control is referred to as functional real-time monitoring (FRTM). Although the idea of FRTM is totally new in the profile monitoring literature, a similar problem is faced by SPC methods for monitoring batch processes. Indeed, in batch processes, quantities of raw materials are subjected to a sequence of steps and conditions over a finite duration to transform them into the final product. Thus, the data produced by a batch process has a three-way structure where for each batch, a set of variables are measured over time throughout the batch's duration. In this setting, the aim is real-time monitoring the batch in order to assess if the process is either IC or OC

(Spooner and Kulahci, 2018; Kassidas et al., 1998).

The first issue to face to perform FRTM is related to the need of registration or alignment of the partially observed functional data. Indeed, during the monitoring phase, we need to real-time compare, i.e., while the process is running, the observed values of the quality characteristic with some reference observations. Unfortunately, often this phase is cumbersome either because functions have different length or because phase variability is present. In functional data analysis the problem of registration is well-known in the literature (Marron et al., 2015) and refers to lateral displacements in curve features, referred to as phase variation, as opposed to amplitude variation in curve height. In general, the main reason for separating phase and amplitude variation is to better preserve the structure of the observed data, since a separate modelling of amplitude and phase variability will be more natural, parsimonious and efficient (Srivastava et al., 2011).

When a reasonable comparison is identified through the alignment step, the aligned partially observed functional data could be seen as complete on the partial domain and, thus, standard profile monitoring methods could be applied at each time point. However, this approach considers the registration process as a pre-processing step aimed to identify the right reference point, and, thus, discards all the information related to the phase variability. This could be risky because all the OC conditions related to variation in phase would be overlooked. In applications, where this behaviour is unwanted, the phase component should be taken into account. The mainstream literature devoted to profile monitoring does not deal with registration, as it is deemed an unnecessary task or simply treated as a pre-processing step to be neglected when profile monitoring is applied. The first work, in the direction of combining curve registration algorithms within the profile monitoring framework is Grasso et al. (2016), where a novel approach to jointly monitor the stability over time of both the registered profiles and the warping functions used to align them is proposed.

On this line of research, the present work aims to present a new method, referred to as functional real-time monitoring control chart (FRTMCC), that is able to real-time monitor a functional quality characteristic. It is based on the idea of real-time alignment and simultaneously monitoring of phase and amplitude variations. The FRTMCC apply iteratively at each time point a procedure consisting of three main steps: i) align the partially observed functional data to the reference observations through a registration procedure; ii) perform a dimensionality reduction through a modification of the functional principal component analysis (FPCA) (Happ, 2018; Ramsay and Silverman, 2005) specifically designed to take into account the phase variability; iii) monitor the resulting coefficients. The first step is performed through the functional dynamic time warping (FDTW) (Wang et al., 1997), which is the functional extension of the well-known dynamic time warping (DTW) designed to align two signals with different dynamics (Sakoe and Chiba, 1978; Itakura, 1975). The idea of the FDTW is to align two functional data by minimizing a distance as a function of the warping function which is a nonlinear transformation mapping the system to the clock time. In particular, a modification of the FDTW is considered to take into account settings with partial matches, which is referred to as open-end/open-begin FDTW. The output of this phase is the aligned and the warping functions. Because functions are intrinsically infinite dimensional, the second step is related to dimensionality reduction. The most used method to perform dimensionality reduction of functional data is the FPCA. However, standard FPCA is not designed to take into account the constrained nature of the warping functions, which are by definition strictly monotonic. Indeed, if standard FPCA is applied to the warping function, then the reconstructed function is not guaranteed to satisfy the monotonic constraint. Therefore, a modification of the FPCA is considered where the

warping function is transformed through an isometric isomorphism from the Bayes space of the warping functions to the space of square integrable function (Happ et al., 2019). In this way the constrained warping function is substituted by an unconstrained one in the space of square integrable functions. The last step concerns with the identification of the appropriate statistics to be monitored. In particular, we consider the profile monitoring approach introduced by Woodall et al. (2004) and, then, used in Noorossana et al. (2012); Grasso et al. (2016); Pini et al. (2017), which is based on the simultaneous application of the *Hotelling's* T^2 and the squared prediction error (*SPE*) control charts.

Conclusion and Future Developments

This thesis stands as one of the first attempt to use functional data analysis (FDA) with the specific aim of solving problems in the industry. Indeed, a suite of methods based on this object oriented data analysis (OODA) paradigm have been presented to address a whole variety of industrial applications. The advantages of these approaches have been extensively discussed, that undoubtedly reveal the utility of FDA methods in the industrial setting.

In particular, in Part I, FDA methods have been proposed to monitor CO₂ emissions during the navigation phase of a roll-on/roll-off passenger cruise ship (Chapter 1) and to deal with the case where measurements of functional covariates are available along with the functional quality characteristic (Chapter 2 and Chapter 3). Chapter 4 represents one of the first contribution in the direction of integrating FDA methods, and in particular functional regression methods, into the manufacturing field, whereas, Chapter 5 shows the potentiality and the practical applicability of functional clustering methods to find out homogeneous groups of dynamic resistance curve in a real-case study in the automotive industry.

Again stimulated by real applications, Part II of this thesis presents new methodological contributions. The focus in this part is on the notion of *interpretability*. Indeed, new interpretable FDA methods are presented. In particular, in Chapter 6, a new interpretable estimator, named S-LASSO, is proposed to estimate the coefficient function of the function-on-function linear regression model. Chapter 7 presents the adaptive smoothing spline (AdaSS) estimator for the same model that is able to adapt to the true coefficient function over regions of large curvature and not to be undersmoothed over the remaining part of the domain. Regarding the problem of clustering functional data, Chapter 8 introduces the sparse and smooth functional clustering (SaS-Funclust) that is able to classify a sample of curves into homogeneous groups while jointly detecting the most informative portions of domain. Eventually, Part III presents two ongoing researches. In particular, in Chapter 9 a new robust method, referred to as Ro-FANOVA, for the functional analysis of variance is presented, whereas Chapter 10 contains a sketch of a new method for functional real-time monitoring.

The possibilities for FDA and more in general OODA methods in the industrial field are countless. Indeed, the more and more complex industrial environment will always require new methods able to deal with these complexity. FDA deals with complex objects that are functions, but what about other types of complex objects such as shapes, network, and so on? Moreover, nowadays datasets of complex objects are of either moderate or small size, but in the future, huge datasets of complex objects will be available as well. Most of the methods now present in the literature are not designed for these types of datasets, how could we deal with this issue? Still, how could the dependence structure among complex objects be used to provide more insightful analysis results? In this thesis, only linear regression has been investigated, could non-linear relationships among complex objects be fruitfully used in

industrial applications? And, in this increasingly complex industrial setting, how could we guarantee that the analysis results are interpretable as well as robust against the presence of anomalous observations?

So, the questions are many, but the answers are still few. The voyage of methods for complex objects in the industry field is just at the beginning.

Bibliography

- Abraham, C., P.-A. Cornillon, E. Matzner-Løber, and N. Molinari (2003). Unsupervised curve clustering using b-splines. *Scandinavian journal of statistics* 30(3), 581–595.
- Abrahamsen, P. and N. Regnesentral (1997). *A Review of Gaussian Random Fields and Correlation Functions*. Norsk Regnesentral/Norwegian Computing Center.
- Abramovich, F. and D. M. Steinberg (1996). Improved inference in nonparametric regression using lk-smoothing splines. *Journal of Statistical Planning and Inference* 49(3), 327–341.
- Abramowitz, M. and I. A. Stegun (1964). *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*. Courier Corporation.
- Adams, D. W., C. D. Summerville, B. M. Voss, J. Jeswiet, and M. C. Doolan (2017). Correlating variations in the dynamic resistance signature to weld strength in resistance spot welding using principal component analysis. *Journal of Manufacturing Science and Engineering* 139(4).
- Andani, M. T., R. Dehghani, M. R. Karamooz-Ravari, R. Mirzaeifar, and J. Ni (2017). Spatter formation in selective laser melting process using multi-laser technology. *Materials & Design* 131, 460–469.
- Anderson, M. J. (2001). Permutation tests for univariate or multivariate analysis of variance and regression. *Canadian journal of fisheries and aquatic sciences* 58(3), 626–639.
- Andrew, G. and J. Gao (2007). Scalable training of l1-regularized log-linear models. In *Proceedings of the 24th international conference on Machine learning*, pp. 33–40. ACM.
- Anwar, A. B. and Q.-C. Pham (2018). Study of the spatter distribution on the powder bed during selective laser melting. *Additive Manufacturing* 22, 86–97.
- ASTM B822–17 (2017). *Standard test method for particle size distribution of metal powders and related compounds by light scattering*. ASTM, American Society for Testing and Materials.
- ASTM E8/E8M–16a (2016). *Standard Test Methods for Tension Testing of Metallic Materials*. ASTM, American Society for Testing and Materials.
- Bäck, T., D. B. Fogel, and Z. Michalewicz (1997). *Handbook of evolutionary computation*. CRC Press.

- Barclift, M., S. Joshi, T. Simpson, and C. Dickman (2016). Cost modeling and depreciation for reused powder feedstock in powder bed fusion additive manufacturing. In *Proceedings of the 27th annual international solid freeform fabrication symposium—an additive manufacturing conference. Austin, TX, USA*.
- Beaton, A. E. and J. W. Tukey (1974). The fitting of power series, meaning polynomials, illustrated on band-spectroscopic data. *Technometrics* 16(2), 147–185.
- Bergstra, J., R. Bardenet, Y. Bengio, and B. Kégl (2011). Algorithms for hyper-parameter optimization. In *25th annual conference on neural information processing systems (NIPS 2011)*, Volume 24. Neural Information Processing Systems Foundation.
- Bergstra, J. and Y. Bengio (2012). Random search for hyper-parameter optimization. *Journal of Machine Learning Research* 13(Feb), 281–305.
- Besse, P. C. and H. Cardot (1996). Approximation spline de la prévision d’un processus fonctionnel autorégressif d’ordre 1. *Canadian Journal of Statistics* 24(4), 467–487.
- Bidare, P., I. Bitharas, R. Ward, M. Attallah, and A. J. Moore (2018). Fluid and particle dynamics in laser powder bed fusion. *Acta Materialia* 142, 107–120.
- Bocchetti, D., A. Lepore, B. Palumbo, and L. Vitiello (2015). A statistical approach to ship fuel consumption monitoring. *Journal of Ship Research* 59(3), 162–171.
- Bosq, D. (2012). *Linear Processes in Function Spaces: Theory and Applications*. Lecture Notes in Statistics. Springer New York.
- Bouveyron, C. and J. Jacques (2011). Model-based clustering of time series in group-specific functional subspaces. *Advances in Data Analysis and Classification* 5(4), 281–300.
- Brumback, B. A. and J. A. Rice (1998). Smoothing spline models for the analysis of nested and crossed samples of curves. *Journal of the American Statistical Association* 93(443), 961–976.
- Canale, A. and S. Vantini (2016). Constrained functional time series: Applications to the italian gas market. *International Journal of Forecasting* 32(4), 1340–1351.
- Candes, E., T. Tao, et al. (2007). The dantzig selector: Statistical estimation when p is much larger than n. *The Annals of Statistics* 35(6), 2313–2351.
- Capezza, C., F. Centofanti, A. Lepore, and B. Palumbo (2020a). funclustrsw: functional clustering for resistance spot-welding data. <https://github.com/unina-sfere/funclustrSW>.
- Capezza, C., F. Centofanti, A. Lepore, and B. Palumbo (2020b). Functional clustering methods for resistance spot welding process data in the automotive industry. *arXiv preprint arXiv:2007.09128*.
- Capezza, C., A. Lepore, A. Menafoglio, B. Palumbo, and S. Vantini (2019). Control charts for monitoring ship operating conditions and CO₂ emissions based on scalar-on-function regression. *manuscript - MOX report*.

- Capezza, C., A. Lepore, A. Menafoglio, B. Palumbo, and S. Vantini (2020). Control charts for monitoring ship operating conditions and CO₂ emissions based on scalar-on-function regression. *Applied Stochastic Models in Business and Industry*.
- Cardot, H., F. Ferraty, and P. Sarda (2003). Spline estimators for the functional linear model. *Statistica Sinica* 13, 571–591.
- Centofanti, F., M. Fontana, A. Lepore, and S. Vantini (2020). Smooth lasso estimator for the function-on-function linear regression model. *arXiv preprint arXiv:2007.00529*.
- Centofanti, F., A. Lepore, A. Menafoglio, B. Palumbo, and S. Vantini (2020a). Adaptive smoothing spline estimator for the function-on-function linear regression model. *arXiv preprint arXiv:2011.12036*.
- Centofanti, F., A. Lepore, A. Menafoglio, B. Palumbo, and S. Vantini (2020b). Functional regression control chart. *Technometrics*, 1–14.
- Charpentier, A. and E. Flachaire (2015). Log-transform kernel density estimation of income distribution. *L'Actualité économique* 91(1-2), 141–159.
- Charrad, M., N. Ghazzali, V. Boiteau, and A. Niknafs (2012). Nbclust package: finding the relevant number of clusters in a dataset. *J. Stat. Softw.*
- Chen, H., P. T. Reiss, and T. Tarpey (2014). Optimally weighted l2 distance for functional data. *Biometrics* 70(3), 516–525.
- Chiou, J.-M., Y.-T. Chen, and Y.-F. Yang (2014). Multivariate functional principal component analysis: A normalization approach. *Statistica Sinica* 24, 1571–1596.
- Chiou, J.-M. and P.-L. Li (2007). Functional clustering and identifying substructures of longitudinal data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69(4), 679–699.
- Chiou, J.-M. and H.-G. Müller (2009). Modeling hazard rates as functional data for the analysis of cohort lifetables and mortality forecasting. *Journal of the American Statistical Association* 104(486), 572–585.
- Chiou, J.-M., Y.-F. Yang, and Y.-T. Chen (2016). Multivariate functional linear regression and prediction. *Journal of Multivariate Analysis* 146, 301–312.
- Chou, Y.-M., R. L. Mason, and J. C. Young (2001). The control chart for individual observations from a multivariate non-normal distribution. *Communications in statistics-Theory and methods* 30(8-9), 1937–1949.
- Colosimo, B. M. and M. Grasso (2020). On-machine measurement, monitoring and control. *Precision Metal Additive Manufacturing*, 102.
- Colosimo, B. M., Q. Huang, T. Dasgupta, and F. Tsung (2018). Opportunities and challenges of quality engineering for additive manufacturing. *Journal of Quality Technology* 50(3), 233–252.
- Colosimo, B. M. and M. Pacella (2007). On the use of principal component analysis to identify systematic patterns in roundness profiles. *Quality and reliability engineering international* 23(6), 707–725.

- Colosimo, B. M. and M. Pacella (2010). A comparison study of control charts for statistical monitoring of functional data. *International Journal of Production Research* 48(6), 1575–1601.
- Corain, L., V. B. Melas, A. Pepelyshev, and L. Salmaso (2014). New insights on permutation approach for hypothesis testing on functional data. *Advances in Data Analysis and Classification* 8(3), 339–356.
- Cordova, L., M. Campos, and T. Tinga (2019). Revealing the effects of powder reuse for selective laser melting by powder characterization. *JOM* 71(3), 1062–1072.
- Cuesta-Albertos, J. and M. Febrero-Bande (2010). A simple multiway anova for functional data. *Test* 19(3), 537–557.
- Cuesta-Albertos, J. A. and R. Fraiman (2006). Impartial trimmed means for functional data. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science* 72, 121.
- Cuesta-Albertos, J. A. and R. Fraiman (2007). Impartial trimmed k-means for functional data. *Computational Statistics & Data Analysis* 51(10), 4864–4877.
- Cuesta-Albertos, J. A. and A. Nieto-Reyes (2008). The random tukey depth. *Computational Statistics & Data Analysis* 52(11), 4979–4988.
- Cuevas, A. (2014). A partial overview of the theory of statistics with functional data. *Journal of Statistical Planning and Inference* 147, 1–23.
- Cuevas, A., M. Febrero, and R. Fraiman (2004). An anova test for functional data. *Computational statistics & data analysis* 47(1), 111–122.
- Cuevas, A., M. Febrero, and R. Fraiman (2006). On the use of the bootstrap for estimating functions with functional data. *Computational statistics & data analysis* 51(2), 1063–1074.
- Cuevas, A. and R. Fraiman (2009). On depth measures and dual statistics. a methodology for dealing with general data. *Journal of Multivariate Analysis* 100(4), 753–766.
- Davis, J. R. (2003). *Metals Handbook, Desk Edition 2nd Edition I*, Volume 464. Materials Park, Ohio: ASM International.
- Davison, A. C. and D. V. Hinkley (1997). *Bootstrap methods and their application*. Cambridge university press.
- De Boor, C., C. De Boor, E.-U. Mathématicien, C. De Boor, and C. De Boor (1978). *A practical guide to splines*, Volume 27. springer-verlag New York.
- De Castro, B. F., S. Guillas, and W. G. Manteiga (2005). Functional samples and bootstrap for predicting sulfur dioxide levels. *Technometrics* 47(2), 212–222.
- Del Re, F., V. Contaldi, A. Astarita, B. Palumbo, A. Squillace, P. Corrado, and P. Di Petta (2018). Statistical approach for assessing the effect of powder reuse on the final quality of alsil0mg parts produced by laser powder bed fusion additive manufacturing. *The International Journal of Advanced Manufacturing Technology* 97(5–8), 1–10.

- Delaigle, A., P. Hall, and T. Pham (2019). Clustering functional data into groups by using projections. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 81(2), 271–304.
- Dickinson, D., J. Franklin, A. Stanya, et al. (1980). Characterization of spot welding behavior by dynamic electrical parameter monitoring. *Welding Journal* 59(6), 170.
- Dunn, J. C. (1973). A fuzzy relative of the isodata process and its use in detecting compact well-separated clusters. *Journal of Cybernetics* 3(3), 32–57.
- Dunson, D. B. (2018). Statistics in the big data era: Failures of the machine. *Statistics & Probability Letters* 136, 4–9.
- Efron, B. et al. (1979). Bootstrap methods: Another look at the jackknife. *The Annals of Statistics* 7(1), 1–26.
- Efron, B., T. Hastie, I. Johnstone, R. Tibshirani, et al. (2004). Least angle regression. *The Annals of statistics* 32(2), 407–499.
- Efron, B. and R. Tibshirani (1986). Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science* 1(1), 54–75.
- El-Banna, M., D. Filev, and R. B. Chinnam (2008). Online qualitative nugget classification by using a linear vector quantization neural network for resistance spot welding. *The International Journal of Advanced Manufacturing Technology* 36(3-4), 237–248.
- Erto, P., A. Lepore, B. Palumbo, and L. Vitiello (2015). A procedure for predicting and controlling the ship fuel consumption: Its implementation and test. *Quality and Reliability Engineering International* 31(7), 1177–1184.
- Eubank, R. L. (1999). *Nonparametric regression and spline smoothing*. CRC press.
- Everitt, B. S., S. Landau, M. Leese, and D. Stahl (2011). *Cluster analysis*. John Wiley & Sons.
- Everton, S. K., M. Hirsch, P. Stravroulakis, R. K. Leach, and A. T. Clare (2016). Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing. *Materials & Design* 95, 431–445.
- Fan, J. and R. Li (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association* 96(456), 1348–1360.
- Fan, J., H. Peng, et al. (2004). Nonconcave penalized likelihood with a diverging number of parameters. *The Annals of Statistics* 32(3), 928–961.
- Fan, Y., N. Foutz, G. M. James, W. Jank, et al. (2014). Functional response additive model estimation with online virtual stock markets. *The Annals of Applied Statistics* 8(4), 2435–2460.
- Faraway, J. J. (1997). Regression analysis for a functional response. *Technometrics* 39(3), 254–261.
- Febrero-Bande, M. and M. Oviedo de la Fuente (2012). Statistical computing in functional data analysis: The R package fda.usc. *Journal of Statistical Software* 51(4), 1–28.

- Ferraty, F. and P. Vieu (2006). *Nonparametric functional data analysis: theory and practice*. Springer Science & Business Media.
- Floriello, D. and V. Vitelli (2017). Sparse clustering of functional data. *Journal of Multivariate Analysis* 154, 1–18.
- Fraiman, R. and G. Muniz (2001). Trimmed means for functional data. *Test* 10(2), 419–440.
- Friedman, J., T. Hastie, and R. Tibshirani (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software* 33(1), 1.
- Friedman, J. H. and J. J. Meulman (2004). Clustering objects on subsets of attributes (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 66(4), 815–849.
- Gardan, J. (2016). Additive manufacturing technologies: state of the art and trends. *International Journal of Production Research* 54(10), 3118–3132.
- Giacofci, M., S. Lambert-Lacroix, G. Marot, and F. Picard (2012). *curvclust: Curve clustering*. R package version 0.0.1.
- Giacofci, M., S. Lambert-Lacroix, G. Marot, and F. Picard (2013). Wavelet-based clustering for mixed-effects functional models in high dimension. *Biometrics* 69(1), 31–40.
- Gibson, I., D. Rosen, B. Stucker, and M. Khorasani (2014). *Additive manufacturing technologies*, Volume 17. Springer.
- Goldsmith, J., J. Bobb, C. M. Crainiceanu, B. Caffo, and D. Reich (2011). Penalized functional regression. *Journal of Computational and Graphical Statistics* 20(4), 830–851.
- Gonzalez, L. and B. F. Manly (1998). Analysis of variance by randomization with small data sets. *Environmetrics: The official journal of the International Environmetrics Society* 9(1), 53–65.
- González-Manteiga, W. and A. Martínez-Calvo (2011). Bootstrap in functional linear regression. *Journal of Statistical Planning and Inference* 141(1), 453–461.
- Good, P. (2013). *Permutation tests: a practical guide to resampling methods for testing hypotheses*. Springer Science & Business Media.
- Górecki, T. and Ł. Smaga (2015). A comparison of tests for the one-way anova problem for functional data. *Computational Statistics* 30(4), 987–1010.
- Gorecki, T. and L. Smaga (2018). fdanova: Analysis of variance for univariate and multivariate functional data. R package version 0.1.2.
- Grasso, M., B. M. Colosimo, and F. Tsung (2017). A phase i multi-modelling approach for profile monitoring of signal data. *International Journal of Production Research* 55(15), 4354–4377.
- Grasso, M., A. Menafoglio, B. M. Colosimo, and P. Secchi (2016). Using curve-registration information for profile monitoring. *Journal of Quality Technology* 48(2), 99.

- Green, P. J. and B. W. Silverman (1993). *Nonparametric regression and generalized linear models: a roughness penalty approach*. Chapman and Hall/CRC.
- Grenander, U. (1950). Stochastic processes and statistical inference. *Arkiv för matematik* 1(3), 195–277.
- Gu, C. (2013). *Smoothing spline ANOVA models*, Volume 297. Springer Science & Business Media.
- Gu, D., W. Meiners, K. Wissenbach, and R. Poprawe (2012). Laser additive manufacturing of metallic components: materials, processes and mechanisms. *International materials reviews* 57(3), 133–164.
- Guo, J., E. Levina, G. Michailidis, and J. Zhu (2010). Pairwise variable selection for high-dimensional model-based clustering. *Biometrics* 66(3), 793–804.
- Guo, W. (2002). Inference in smoothing spline analysis of variance. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64(4), 887–898.
- Hague, R., S. Mansour, and N. Saleh (2004). Material and design considerations for rapid manufacturing. *International Journal of Production Research* 42(22), 4691–4708.
- Hall, P., J. L. Horowitz, et al. (2007). Methodology and convergence rates for functional linear regression. *The Annals of Statistics* 35(1), 70–91.
- Hall, P. and M. Hosseini-Nasab (2006). On properties of functional principal components analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68(1), 109–126.
- Hampel, F. R. (1974). The influence curve and its role in robust estimation. *Journal of the american statistical association* 69(346), 383–393.
- Hampel, F. R., E. M. Ronchetti, P. J. Rousseeuw, and W. A. Stahel (2011). *Robust statistics: the approach based on influence functions*, Volume 196. John Wiley & Sons.
- Happ, C. (2018). *MFPCA: Multivariate Functional Principal Component Analysis for Data Observed on Different Dimensional Domains*. R package version 1.3.
- Happ, C. and S. Greven (2018). Multivariate functional principal component analysis for data observed on different (dimensional) domains. *Journal of the American Statistical Association* 113(522), 649–659.
- Happ, C., F. Scheipl, A.-A. Gabriel, and S. Greven (2019). A general framework for multivariate functional principal component analysis of amplitude and phase variation. *Stat* 8(1), e220.
- Hastie, T., R. Tibshirani, and J. Friedman (2009a). *The elements of statistical learning: data mining, inference, and prediction*. Springer series in statistics New York, NY, USA.
- Hastie, T., R. Tibshirani, and J. Friedman (2009b). *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media.
- Hawkins, D. M. (1991). Multivariate quality control based on regression-adjusted variables. *Technometrics* 33(1), 61–75.

- Hawkins, D. M. (1993). Regression adjustment for variables in multivariate quality control. *Journal of Quality Technology* 25, 170–182.
- Herlihy, M. and N. Shavit (2011). *The art of multiprocessor programming*. Morgan Kaufmann.
- Herzog, D., V. Seyda, E. Wycisk, and C. Emmelmann (2016). Additive manufacturing of metals. *Acta Materialia* 117, 371–392.
- Horváth, L. and P. Kokoszka (2012). *Inference for functional data with applications*, Volume 200. Springer Science & Business Media.
- Hsing, T. and R. Eubank (2015). *Theoretical foundations of functional data analysis, with an introduction to linear operators*. John Wiley & Sons.
- Huber, P. J. (2004). *Robust statistics*, Volume 523. John Wiley & Sons.
- Huber, P. J. et al. (1964). Robust estimation of a location parameter. *The Annals of Mathematical Statistics* 35(1), 73–101.
- Hubert, L. and P. Arabie (1985). Comparing partitions. *Journal of classification* 2(1), 193–218.
- Hubert, M., P. J. Rousseeuw, and P. Segaeert (2015). Multivariate functional outlier detection. *Statistical Methods & Applications* 24(2), 177–202.
- Hullait, H., D. S. Leslie, N. G. Pavlidis, and S. King (2020). Robust function-on-function regression. *Technometrics*, 1–14.
- Hunter, D. R. and R. Li (2005). Variable selection using mm algorithms. *Annals of statistics* 33(4), 1617.
- Ieva, F., A. M. Paganoni, D. Pigoli, and V. Vitelli (2013). Multivariate functional clustering for the morphological analysis of electrocardiograph curves. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 62(3), 401–418.
- Ighodaro, O., E. Biro, and Y. Zhou (2016, 12). Study and applications of dynamic resistance profiles during resistance spot welding of coated hot-stamping steels. *Metallurgical and Materials Transactions A* 48, 1–14.
- IMO (2012a). Air Pollution and Greenhouse Gas (GHG) Emissions from International Shipping, MARPOL Annex 6. London, U.K.
- IMO (2012b). Guidelines on the method of calculation of the attained Energy Efficiency Design Index (EEDI) for new ships, MEPC.212 Annex 8. London, U.K.
- IMO (2014). 2014 Guidelines on survey and certification of the Energy Efficiency Design Index (EEDI), London, U.K.
- ISO 6892:2016 (2016). *Metallic Materials - Tensile Testing. Part 1: Method of Test at Room Temperature*. ISO, International Organization for Standardization.
- Itakura, F. (1975). Minimum prediction residual principle applied to speech recognition. *IEEE Transactions on acoustics, speech, and signal processing* 23(1), 67–72.

- Ivanescu, A. E., A.-M. Staicu, F. Scheipl, and S. Greven (2015). Penalized function-on-function regression. *Computational Statistics* 30(2), 539–568.
- Jacob, G., G. Jacob, C. U. Brown, M. A. Donmez, S. S. Watson, and J. Slotwinski (2017). *Effects of powder recycling on stainless steel powder and built material properties in metal powder bed fusion processes*. US Department of Commerce, National Institute of Standards and Technology.
- Jacques, J. and C. Preda (2013). Funclust: A curves clustering method using functional random variables density approximation. *Neurocomputing* 112, 164–171.
- Jacques, J. and C. Preda (2014). Functional data clustering: a survey. *Advances in Data Analysis and Classification* 8(3), 231–255.
- Jaderberg, M., V. Dalibard, S. Osindero, W. M. Czarnecki, J. Donahue, A. Razavi, O. Vinyals, T. Green, I. Dunning, and K. Simonyan (2017). Population based training of neural networks. *arXiv preprint arXiv:1711.09846*.
- James, G. M. (2002). Generalized linear models with functional predictors. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64(3), 411–432.
- James, G. M. and P. Radchenko (2009). A generalized dantzig selector with shrinkage tuning. *Biometrika* 96(2), 323–337.
- James, G. M. and C. A. Sugar (2003). Clustering for sparsely sampled functional data. *Journal of the American Statistical Association* 98(462), 397–408.
- James, G. M., J. Wang, J. Zhu, et al. (2009). Functional linear regression that’s interpretable. *The Annals of Statistics* 37(5A), 2083–2108.
- Jia, J. and B. Yu (2010). On model selection consistency of the elastic net when $p \geq n$. *Statistica Sinica*, 595–611.
- Jin, J. and J. Shi (1999). Feature-preserving data compression of stamping tonnage information using wavelets. *Technometrics* 41(4), 327–339.
- Johnson, R. A., D. W. Wichern, et al. (2002). *Applied multivariate statistical analysis*, Volume 5. Prentice hall Upper Saddle River, NJ.
- Jolliffe, I. (2011). *Principal component analysis*. Springer.
- Kalogridis, I. and S. Van Aelst (2019). Robust functional regression based on principal components. *Journal of Multivariate Analysis* 173, 393–415.
- Kassidas, A., J. F. MacGregor, and P. A. Taylor (1998). Synchronization of batch trajectories using dynamic time warping. *AIChE Journal* 44(4), 864–875.
- Khairallah, S. A., A. T. Anderson, A. Rubenchik, and W. E. King (2016). Laser powder-bed fusion additive manufacturing: Physics of complex melt flow and formation mechanisms of pores, spatter, and denudation zones. *Acta Materialia* 108(15), 36–45.
- Khajavi, S. H., G. Deng, J. Holmström, P. Puukko, and J. Partanen (2018). Selective laser melting raw material commoditization: impact on comparative competitiveness of additive manufacturing. *International Journal of Production Research* 56(14), 1–23.

- Kokoszka, P. and M. Reimherr (2017). *Introduction to functional data analysis*. CRC Press.
- Lackenby, H. (1978). *ITTC Dictionary of Ship Hydrodynamics*. Royal institution of naval architects.
- Ledoux, M. and M. Talagrand (2013). *Probability in Banach Spaces: isoperimetry and processes*. Springer Science & Business Media.
- Lehmann, E. L. and J. P. Romano (2006). *Testing statistical hypotheses*. Springer Science & Business Media.
- Lepore, A., B. Palumbo, and C. Capezza (2018). Analysis of profiles for monitoring of modern ship performance via partial least squares methods. *Quality and Reliability Engineering International* 34(7), 1424–1436.
- Lepore, A., M. S. Reis, B. Palumbo, R. Rendall, and C. Capezza (2017). A comparison of advanced regression techniques for predicting ship CO₂ emissions. *Quality and Reliability Engineering International* 33(6), 1281–1292.
- Li, Y., T. Hsing, et al. (2007). On rates of convergence in functional linear regression. *Journal of Multivariate Analysis* 98(9), 1782–1804.
- Lin, Z., J. Cao, L. Wang, and H. Wang (2017). Locally sparse estimator for functional linear regression models. *Journal of Computational and Graphical Statistics* 26(2), 306–318.
- López-Pintado, S. and J. Romo (2009). On the concept of depth for functional data. *Journal of the American Statistical Association* 104(486), 718–734.
- López-Pintado, S. and J. Romo (2011). A half-region depth for functional data. *Computational Statistics & Data Analysis* 55(4), 1679–1695.
- Lowry, C. A. and D. C. Montgomery (1995). A review of multivariate control charts. *IIE transactions* 27(6), 800–810.
- Luo, R. and X. Qi (2017). Function-on-function linear regression by signal compression. *Journal of the American Statistical Association* 112(518), 690–705.
- Luo, R. and X. Qi (2019). Interaction model and model selection for function-on-function regression. *Journal of Computational and Graphical Statistics* 28(2), 1–14.
- Ly, S., A. M. Rubenchik, S. A. Khairallah, G. Guss, and M. J. Matthews (2017). Metal vapor micro-jet controls material redistribution in laser powder bed fusion additive manufacturing. *Scientific reports* 7(1), 1–12.
- Mandel, B. (1969). The regression control chart. *Journal of Quality Technology* 1(1), 1–9.
- Mani, M., S. Feng, B. Lane, A. Donmez, S. Moylan, and R. Fesperman (2015). Measurement science needs for real-time control of additive manufacturing powder bed fusion processes.
- Mani, M., B. M. Lane, M. A. Donmez, S. C. Feng, and S. P. Moylan (2017). A review on measurement science needs for real-time control of additive manufacturing metal powder bed fusion processes. *International Journal of Production Research* 55(5), 1400–1418.

- Manly, B. F. (2006). *Randomization, bootstrap and Monte Carlo methods in biology*, Volume 70. CRC press.
- Manufacturers' Alliance, R. W. (2003). Resistance welding manual.
- Maronna, R. A., R. D. Martin, V. J. Yohai, and M. Salibián-Barrera (2019). *Robust statistics: theory and methods (with R)*. John Wiley & Sons.
- Marron, J. S. and A. M. Alonso (2014). Overview of object oriented data analysis. *Biometrical Journal* 56(5), 732–753.
- Marron, J. S., J. O. Ramsay, L. M. Sangalli, and A. Srivastava (2015). Functional data analysis of amplitude and phase variation. *Statistical Science*, 468–484.
- Martín, Ó., M. Pereda, J. I. Santos, and J. M. Galán (2014). Assessment of resistance spot welding quality based on ultrasonic testing and tree-based techniques. *Journal of Materials Processing Technology* 214(11), 2478–2487.
- Matsui, H., S. Kawano, and S. Konishi (2009). Regularized functional regression modeling for functional response and predictors. *Journal of Math-for-industry* 1(3), 17–25.
- Maugis, C., G. Celeux, and M.-L. Martin-Magniette (2009). Variable selection for clustering with gaussian mixture models. *Biometrics* 65(3), 701–709.
- McLachlan, G. J. and D. Peel (2004). *Finite mixture models*. John Wiley & Sons.
- Menafoglio, A., M. Grasso, P. Secchi, and B. M. Colosimo (2018). Profile monitoring of probability density functions via simplicial functional PCA with application to image data. *Technometrics* 60(4), 497–510.
- Meng, X.-L. and D. B. Rubin (1993). Maximum likelihood estimation via the ecm algorithm: A general framework. *Biometrika* 80(2), 267–278.
- Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence* 267, 1–38.
- Miller Jr, R. G. (1997). *Beyond ANOVA: basics of applied statistics*. CRC press.
- Montgomery, D. C. (2007). *Introduction to statistical quality control*. John Wiley & Sons.
- Montgomery, D. C. (2017). *Design and analysis of experiments*. John Wiley & Sons.
- Morris, J. S. (2015). Functional regression. *Annual Review of Statistics and Its Application* 2, 321–359.
- Müller, H.-G. and U. Stadtmüller (2005). Generalized functional linear models. *the Annals of Statistics* 33(2), 774–805.
- Nocedal, J. and S. Wright (2006). *Numerical optimization*. Springer Science & Business Media.
- Noorossana, R., A. Saghaei, and A. Amiri (2012). *Statistical analysis of profile monitoring*. John Wiley & Sons.
- Ohm, G. S. (1827). *Die galvanische Kette, mathematisch bearbeitet*. TH Riemann.

- Palumbo, B., F. Centofanti, and F. Del Re (2020). Function-on-function regression for assessing production quality in industrial manufacturing. *Quality and Reliability Engineering International* 36(8), 2738–2753.
- Pan, W. and X. Shen (2007). Penalized model-based clustering with application to variable selection. *Journal of Machine Learning Research* 8(May), 1145–1164.
- Pesarin, F. and L. Salmaso (2010). *Permutation tests for complex data: theory, applications and software*. John Wiley & Sons.
- Petrovic, V., J. Vicente Haro Gonzalez, O. Jordá Ferrando, J. Delgado Gordillo, J. Ramón Blasco Puchades, and L. Portolés Griñan (2011). Additive layered manufacturing: sectors of industrial application shown through case studies. *International Journal of Production Research* 49(4), 1061–1079.
- Pini, A. and S. Vantini (2017). Interval-wise testing for functional data. *Journal of Nonparametric Statistics* 29(2), 407–424.
- Pini, A., S. Vantini, B. M. Colosimo, and M. Grasso (2017). Domain-selective functional analysis of variance for supervised statistical profile monitoring of signal data. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*.
- Pintore, A., P. Speckman, and C. C. Holmes (2006). Spatially adaptive smoothing splines. *Biometrika* 93(1), 113–125.
- Prchal, L. and P. Sarda (2007). Spline estimator for the functional linear regression with functional response. *Preprint*.
- Qi, X. and R. Luo (2018). Function-on-function regression with thousands of predictive curves. *Journal of Multivariate Analysis* 163, 51–66.
- Qi, X. and R. Luo (2019). Nonlinear function on function additive model with multiple predictor curves. *Statistica Sinica* 29, 719–739.
- R Core Team (2020a). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- R Core Team (2020b). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Raftery, A. E. and N. Dean (2006). Variable selection for model-based clustering. *Journal of the American Statistical Association* 101(473), 168–178.
- Ramberg, Walter and Osgood, William R (1943). *Description of stress-strain curves by three parameters*. National Advisory Committee for Aeronautics.
- Ramsay, J. (1982). When the data are functions. *Psychometrika* 47(4), 379–396.
- Ramsay, J. O. and C. Dalzell (1991). Some tools for functional data analysis. *Journal of the Royal Statistical Society. Series B (Methodological)*, 539–572.
- Ramsay, J. O., S. Graves, and G. Hooker (2020). *fda: Functional Data Analysis*. R package version 5.1.5.

- Ramsay, J. O., G. Hooker, and S. Graves (2009). *Functional data analysis with R and MATLAB*. Springer Science & Business Media.
- Ramsay, J. O. and B. W. Silverman (2005). *Functional data analysis*. Wiley Online Library.
- Rao, C. R. (1958). Some statistical methods for comparison of growth curves. *Biometrics* 14(1), 1–17.
- Raelison, R., A. Fuentes, P. Rogeon, P. Carre, T. Loulou, D. Carron, and F. Dechalotte (2012). Contact conditions on nugget development during resistance spot welding of zn coated steel sheets using rounded tip electrodes. *Journal of Materials Processing Technology* 212(8), 1663–1669.
- Ray, S. and B. Mallick (2006). Functional clustering by bayesian wavelet methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68(2), 305–332.
- Reis, M. S., R. Rendall, B. Palumbo, A. Lepore, and C. Capezza (2020). Predicting ships’ CO₂ emissions using feature-oriented methods. *Applied Stochastic Models in Business and Industry* 36(1), 110–123.
- Reiss, P. T. and R. Todd Ogden (2009). Smoothing parameter selection for a class of semiparametric linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 71(2), 505–523.
- Repossini, G., V. Laguzza, M. Grasso, and B. M. Colosimo (2017). On the use of spatter signature for in-situ monitoring of laser powder bed fusion. *Additive Manufacturing* 16, 35–48.
- Rigon, T. (2019). An enriched mixture model for functional clustering. *arXiv preprint arXiv:1907.02493*.
- Rodríguez, A., D. B. Dunson, and A. E. Gelfand (2009). Bayesian nonparametric functional data analysis through density estimation. *Biometrika* 96(1), 149–162.
- Rossi, F., B. Conan-Guez, and A. El Golli (2004). Clustering functional data with the som algorithm. In *ESANN*, pp. 305–312.
- Rousseeuw, P. J. (1987). Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics* 20, 53–65.
- Ruppert, D. and R. J. Carroll (2000). Theory & methods: Spatially-adaptive penalties for spline fitting. *Australian & New Zealand Journal of Statistics* 42(2), 205–223.
- Sakoe, H. and S. Chiba (1978). Dynamic programming algorithm optimization for spoken word recognition. *IEEE transactions on acoustics, speech, and signal processing* 26(1), 43–49.
- Sames, W. J., F. A. List, S. Pannala, R. R. Dehoff, and S. S. Babu (2016). The metallurgy and processing science of metal additive manufacturing. *International Materials Reviews* 61(5), 315–360.
- Scheipl, F. and S. Greven (2016). Identifiability in penalized function-on-function regression models. *Electronic Journal of Statistics* 10(1), 495–526.

- Scheipl, F., A.-M. Staicu, and S. Greven (2015). Functional additive mixed models. *Journal of Computational and Graphical Statistics* 24(2), 477–501.
- Schmutz, A. and J. J. . C. Bouveyron (2019). *funHDDC: Univariate and Multivariate Model-Based Clustering in Group-Specific Functional Subspaces*. R package version 2.3.0.
- Schrader, R. M. and T. P. Hettmansperger (1980). Robust analysis of variance based upon a likelihood ratio criterion. *Biometrika* 67(1), 93–101.
- Schrader, R. M. and J. W. Mc Kean (1977). Robust analysis of variance. *Communications in Statistics-Theory and Methods* 6(9), 879–894.
- Schumaker, L. (2007). *Spline functions: basic theory*. Cambridge University Press.
- Secchi, P. (2018). On the role of statistics in the era of big data: A call for a debate. *Statistics & Probability Letters* 136, 10–14.
- Serban, N. and L. Wasserman (2005). Cats: clustering after transformation and smoothing. *Journal of the American Statistical Association* 100(471), 990–999.
- Shang, H. L. (2014). A survey of functional principal component analysis. *AStA Advances in Statistical Analysis* 98(2), 121–142.
- Shen, Q. and J. Faraway (2004). An f test for linear models with functional responses. *Statistica Sinica*, 1239–1257.
- Shu, L., F. Tsung, and K.-L. Tsui (2004, 7). Run-length performance of regression control charts with estimated parameters. *Journal of Quality Technology* 36(3), 280–292.
- Silverman, B. (1986). *Density Estimation for Statistics and Data Analysis*. Chapman and Hall.
- Silverman, B. W. (2018). *Density estimation for statistics and data analysis*. Routledge.
- Sinova, B., G. Gonzalez-Rodriguez, S. Van Aelst, et al. (2018). M-estimators of location for functional data. *Bernoulli* 24(3), 2328–2357.
- Slotwinski, J. A. and E. J. Garboczi (2015). Metrology needs for metal additive manufacturing powders. *Jom* 67(3), 538–543.
- Spooner, M. and M. Kulahci (2018). Monitoring batch processes with dynamic time warping and k-nearest neighbours. *Chemometrics and Intelligent Laboratory Systems* 183, 102–112.
- Srivastava, A., W. Wu, S. Kurttek, E. Klassen, and J. S. Marron (2011). Registration of functional data using fisher-rao metric. *arXiv preprint arXiv:1103.3817*.
- Stein, M. (1999). *Interpolation of Spatial Data: Some Theory for Kriging*. Springer New York.
- Storlie, C. B., H. D. Bondell, and B. J. Reich (2010). A locally adaptive penalty for estimation of functions with varying roughness. *Journal of Computational and Graphical Statistics* 19(3), 569–589.

- Sun, X., P. Du, X. Wang, and P. Ma (2018). Optimal penalized function-on-function regression under a reproducing kernel hilbert space framework. *Journal of the American Statistical Association*, 1–11.
- Sutton, A. T., C. S. Kriewall, M. C. Leu, and J. W. Newkirk (2017). Powder characterisation techniques and effects of powder characteristics on part properties in powder-bed fusion processes. *Virtual and Physical Prototyping* 12(1), 3–29.
- Tan, Z., Q. Fang, H. Li, S. Liu, W. Zhu, and D. Yang (2020). Neural network based image segmentation for spatter extraction during laser-based powder bed fusion processing. *Optics & Laser Technology* 130, 106347.
- Tarpey, T. and K. K. Kinader (2003). Clustering functional data. *Journal of classification* 20(1), 093–114.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* 58(1), 267–288.
- Tibshirani, R., G. Walther, and T. Hastie (2001). Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63(2), 411–423.
- Tuddenham, R. D. (1954). Physical growth of california boys and girls from birth to eighteen years. *University of California publications in child development* 1, 183–364.
- Vitelli, V. (2019). A novel framework for joint sparse clustering and alignment of functional data. *arXiv preprint arXiv:1912.00687*.
- Wade, M. R. and W. H. Woodall (1993). A review and analysis of cause-selecting control charts. *Journal of quality technology* 25(3), 161–169.
- Wahba, G. (1990). *Spline models for observational data*, Volume 59. Siam.
- Walnut, D. F. (2013). *An introduction to wavelet analysis*. Springer Science & Business Media.
- Wang, H., J. Marron, et al. (2007). Object oriented data analysis: Sets of trees. *The Annals of Statistics* 35(5), 1849–1873.
- Wang, K., T. Gasser, et al. (1997). Alignment of curves by dynamic time warping. *The annals of Statistics* 25(3), 1251–1276.
- Wang, S. and J. Zhu (2008). Variable selection for model-based high-dimensional clustering and its application to microarray data. *Biometrics* 64(2), 440–448.
- Wang, X., P. Du, and J. Shen (2013). Smoothing splines with varying smoothing parameter. *Biometrika* 100(4), 955–970.
- Ward Jr, J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American statistical association* 58(301), 236–244.
- Witten, D. M. and R. Tibshirani (2010). A framework for feature selection in clustering. *Journal of the American Statistical Association* 105(490), 713–726.

- Wood, S. N. (2017). *Generalized additive models: an introduction with R*. CRC press.
- Woodall, W. H., D. J. Spitzner, D. C. Montgomery, and S. Gupta (2004). Using control charts to monitor process and product quality profiles. *Journal of Quality Technology* 36(3), 309.
- Xie, B., W. Pan, and X. Shen (2008). Variable selection in penalized model-based clustering via regularization on grouped parameters. *Biometrics* 64(3), 921–930.
- Xu, L. D., E. L. Xu, and L. Li (2018). Industry 4.0: state of the art and future trends. *International Journal of Production Research* 56(8), 2941–2962.
- Yang, L. and Y. Hong (2017). Adaptive penalized splines for data smoothing. *Computational Statistics & Data Analysis* 108, 70–83.
- Yang, L., L. Lo, S. Ding, and T. Özel (2020). Monitoring and detection of meltpool and spatter regions in laser powder bed fusion of super alloy inconel 625. *Progress in Additive Manufacturing* 5(4), 367–378.
- Yao, F. and H.-G. Müller (2010). Functional quadratic regression. *Biometrika* 97(1), 49–64.
- Yao, F., H.-G. Müller, and J.-L. Wang (2005a). Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association* 100(470), 577–590.
- Yao, F., H.-G. Müller, and J.-L. Wang (2005b). Functional linear regression analysis for longitudinal data. *The Annals of Statistics*, 2873–2903.
- Yin, J., D. Wang, L. Yang, H. Wei, P. Dong, L. Ke, G. Wang, H. Zhu, and X. Zeng (2020). Correlation between forming quality and spatter dynamics in laser powder bed fusion. *Additive Manufacturing* 31, 100958.
- Young, Z. A., Q. Guo, N. D. Parab, C. Zhao, M. Qu, L. I. Escano, K. Fezzaa, W. Everhart, T. Sun, and L. Chen (2020). Types of spatter and their features and formation mechanisms in laser powder bed fusion additive manufacturing process. *Additive Manufacturing* 36, 101438.
- Zhang, H. and J. Senkara (2011). *Resistance welding: fundamentals and applications*. CRC press.
- Zhang, J.-T. (2011). Statistical inferences for linear models with functional responses. *Statistica Sinica*, 1431–1451.
- Zhang, J.-T. (2013). *Analysis of variance for functional data*. CRC Press.
- Zhang, J.-T., J. Chen, et al. (2007). Statistical inferences for functional data. *The Annals of Statistics* 35(3), 1052–1079.
- Zhang, J.-T. and X. Liang (2014). One-way anova for functional data via globalizing the pointwise f-test. *Scandinavian Journal of Statistics* 41(1), 51–71.
- Zhao, X., Y. Zhang, and G. Chen (2006). Research for ultrasonic fast-identification of the stick-weld defect. In *2006 IEEE Instrumentation and Measurement Technology Conference Proceedings*, pp. 81–85. IEEE.

- Zhou, J., N.-Y. Wang, and N. Wang (2013). Functional linear model with zero-value coefficient function at sub-regions. *Statistica Sinica* 23(1), 25.
- Zhou, K. and L. Cai (2013, 10). Online nugget diameter control system for resistance spot welding. *The International Journal of Advanced Manufacturing Technology* 68.
- Zhou, K. and L. Cai (2014). Study on effect of electrode force on resistance spot welding process. *Journal of applied physics* 116(8), 084902.
- Zhou, M. and T. Goh (2016). Effects of model accuracy on residual control charts. *Quality and Reliability Engineering International* 32(5), 1785–1794.
- Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American statistical association* 101(476), 1418–1429.
- Zou, H. and H. H. Zhang (2009). On the adaptive elastic-net with a diverging number of parameters. *Annals of statistics* 37(4), 1733.