

Robust Function-on-Function Regression

Harjit Hullait

STOR-i Centre for Doctoral Training, Lancaster University

and

David S. Leslie

Department of Mathematics and Statistics, Lancaster University

and

Nicos G. Pavlidis

Department of Management Science, Lancaster University

and

Steve King

Rolls Royce PLC

July 20, 2020

Abstract

Functional linear regression is a widely used approach to model functional responses with respect to functional inputs. However, classical functional linear regression models can be severely affected by outliers. We therefore introduce a Fisher-consistent robust functional linear regression model that is able to effectively fit data

in the presence of outliers. The model is built using robust functional principal component and least squares regression estimators. The performance of the functional linear regression model depends on the number of principal components used. We therefore introduce a consistent robust model selection procedure to choose the number of principal components. Our robust functional linear regression model can be used alongside an outlier detection procedure to effectively identify abnormal functional responses. A simulation study shows our method is able to effectively capture the regression behaviour in the presence of outliers, and is able to find the outliers with high accuracy. We demonstrate the usefulness of our method on jet engine sensor data. We identify outliers that would not be found if the functional responses were modelled independently of the functional input, or using non-robust methods.

Keywords: Robust Functional Data Analysis, Robust Model Selection, Outlier Detection.

1 Introduction

Functional Linear Regression (FLR) in the function-on-function case (Ramsay and Dalzell, 1991) is a widely used technique for modelling functional responses with respect to functional inputs. However classical FLR models can be severely affected by outliers as we will demonstrate via a simulation study in Section 6. We therefore develop a Robust FLR (RFLR) model, which is able to effectively fit the data in the presence of outliers. The model is built using the robust functional principal component model by Bali et al. (2011) and the Multivariate Least Trimmed Squares (MLTS) estimator by Agulló et al. (2008).

Our study of FLR is motivated by the need to identify unusual jet engine behaviour using sensor data captured during a *pass-off* test (the last test prior to the deployment of a jet engine). Jet engines are highly complex machines consisting of tens of thousands of parts, and although the physical behaviour of individual components can be modelled by systems of partial or ordinary differential equations, such a description is far from adequate to model the entire engine (Houstis et al., 2002). During a pass-off test an engineer that controls the speed of the engine performs a number of manoeuvres, which can be defined as different accelerations and decelerations starting and ending at a set idle speed. Although each manoeuvre is characterised by a distinct engine speed profile, its execution differs each time it is performed. Throughout the duration of the test sensor measurements for numerous engine parameters, including speed, pressure, temperature and vibration at different parts on the engine, are recorded at very high frequency. The volume of data generated in each test is very large and only a small proportion of this is manually inspected to identify potential engine issues. Currently there exists no engine simulation model to predict the evolution of these parameters during the manoeuvres performed by the controller. Instead engineers use their judgement to determine whether the sensor measurements deviate sufficiently from their expectations to indicate a potential engine issue.

An automated approach to identifying pass-off test segments that require closer inspection will enable a larger proportion of the pass-off data to be processed, and can thus improve the early identification of engine issues. However, to achieve this goal it is first necessary to develop accurate models of the expected behaviour of a healthy engine during different manoeuvres. The available data consist of sensor measurements captured during pass-off tests. For each individual engine we do not know whether it had issues or not, but we do know that the dataset contains tests for engines that had issues.

In this work we develop RFLR to model engine temperature parameters during the Vibration Survey (VS), a key pass-off test manoeuvre during which the engine is gradually accelerated to a certain speed and then slowly decelerated. Both engine speed and temperature are continuous quantities that are recorded at very high-frequencies. Therefore the central assumption underlying functional data analysis, namely that the observed time series constitute discrete measurements from a smooth functional process is realistic in our application. We use FLR due to the documented ability of this methodology to capture complex dependency structures (Morris, 2015). Ramsay and Silverman (2005) have illustrated that the FLR model simultaneously uses information across functions to identify their commonalities (called *replication*), and information about the expected underlying structural relationships within each function (called *regularisation*). However, the flexibility to use information from “future” time points implies that the models developed do not truly approximate the underlying mechanical process.

For our analysis we will use 199 VS datasets, which include information about the Turbine Pressure Ratio (TPR) (the human controlled parameter that measures engine speed) and various temperature features including the Turbine Gas Temperature (TGT). Figure 1 provides envelope plots of TPR and TGT for all the VS manoeuvres. To anonymise the data we have transformed the time index onto the interval $[0, 1]$ and the sensor measurements to the range $[0, 100]$. The VS manoeuvres are performed by a human controller, which causes variability in TPR as depicted in Figure 1(a). This naturally affects the

TGT curves and can mask unusual behaviour arising from an engine issue. In fact outlier detection methods applied directly on the TGT curves (as shown in Figure 1(b)) fail to identify meaningful outliers, and instead identify exclusively curves produced by unusual TPR profiles. The green TGT curve in Figure 1(b) is such an example. Direct outlier detection methods identify this curve as an outlier but there are no issues with this engine. This curve only appears to be an outlier due to the relatively uncommon TPR schedule during this manoeuvre (see Figure 1(a)).

We therefore require a method of detecting outliers in the presence of the controller induced variability. In particular, we expect that the relationship between engine speed (TPR) and engine temperature for different VS manoeuvres should be the same irrespective of the way the manoeuvre is performed. For example given a certain engine acceleration we would expect a certain temperature response. If however the response differs from expectation this could be indicative of an engine issue. In Section 6 we show that classical FLR can be severely affected when the dataset contains outliers, which in our case correspond to data from VS manoeuvres of engines with issues. The RFLR model we propose aims to overcome this limitation. Applied on the data in Figure 1, our model identifies the TGT curves with purple and red colour as outliers. Notice that both curves do not appear to diverge from the general pattern in the envelope plot in Figure 1(b), and are only identified as outliers when the functional input from the corresponding TPR curves is taken into account. Both VS manoeuvres correspond to engines for which issues were identified during the test.

The paper is organised as follows. In Section 2 we outline the classical FLR model. In Section 3, we will outline robust Functional Data Analysis (FDA) techniques to obtain a robust FLR model. We also introduce a robust model selection procedure. In Section 4 we prove consistency results for the robust FLR model and the robust model selection procedure. In Section 5, we describe an outlier detection method, which acts on the residuals of the robust FLR model. In Section 6 we perform a simulation study to illustrate the




Figure 1: Envelope plots of Turbine Pressure Ratio (TPR) and Turbine Gas Temperature (TGT) time series in grey. Sample 33 (green) is identified as an outlier when the relationship between TPR and TGT is not accounted for, but is not associated with an engine issue. Samples 106 (purple) and 77 (red) are not identified as outliers by direct outlier detection methods, but are linked with engine faults. These are correctly identified using RFLR.

model fitting and outlier detection capabilities of the robust model. In Section 7 we apply the robust model on the engine data and highlight unusual observations that can not be detected by using outlier detection methods directly on the temperature curves. The paper ends with concluding remarks in Section 8.

2 Classical Functional Data Analysis

In this section we give a brief summary of the FDA tools that we will later apply in our model. In the following sections we will use the vector space $L^2(I)$ which is the Hilbert space of square integrable functions on the compact interval I with the inner product $\langle f, g \rangle = \int_I f(t)g(t)dt$ for functions $f, g \in L^2(I)$.

We will define $X(t), Y(t)$ to be univariate stochastic processes defined on I , with mean functions $\mu^X(t)$ and $\mu^Y(t)$, and covariance functions $C_X(s, t) = \text{cov}\{X(s), X(t)\}$ and $C_Y(s, t) = \text{cov}\{Y(s), Y(t)\}$ for all $s, t \in I$. We shall define $\mathbf{x}(t) = [x_1(t), \dots, x_n(t)]$ and $\mathbf{y}(t) = [y_1(t), \dots, y_n(t)]$ to be n independent and identically distributed realisations of $X(t)$

and $Y(t)$ respectively.

In practice we observe $x_i(t)$ and $y_i(t)$ at discrete time points. We shall assume for simplicity of exposition that observations are made at equally spaced time points t_1, \dots, t_T . We will outline Functional Linear Regression and Functional Principal Component Analysis with respect to the underlying functions $x(t), y(t)$. In Section 2.3 we need to use the discretely observed data to define a suitable model selection criterion.

2.1 Functional Linear Regression

In this section we will introduce the classical FLR model (Ramsay and Dalzell, 1991). In FLR we model the relationship between predictor $x_i(t)$ and response $y_i(t)$ as:

$$y_i(t) = \alpha(t) + \int_I x_i(s)\beta(s, t)ds + \epsilon_i(t), \quad (2.1)$$

where $\alpha(t)$ is the intercept function, $\beta(s, t)$ is the regression function and $\epsilon_i(t)$ is the error process. For a fixed t , we can think of $\beta(s, t)$ as the relative weight placed on $x_i(s)$ to predict $y_i(t)$. As in Chiou et al. (2016) we will assume the mean functions $\mu^X(t) = 0$ and $\mu^Y(t) = 0$ which thereby means $\alpha(t) = 0$. This is a reasonable assumption as in practice we can calculate the mean functions $\mu^X(t)$ and $\mu^Y(t)$ efficiently for dense data and then pre-process the data by subtracting $\mu^X(t)$ and $\mu^Y(t)$ from the observed curves.

FLR in the function-on-function case can be modelled parametrically (Yao et al., 2005; Chiou et al., 2016) or nonparametrically (Ferraty et al., 2012; Ivanescu et al., 2015; Scheipl et al., 2015). We use a parametric approach which models the regression matrix in terms of pre-defined basis functions. We obtain the approximations $\hat{x}_i(t)$ and $\hat{y}_i(t)$ in terms of (M, K) pre-chosen orthonormal basis functions $\phi_m^X(t), \phi_k^Y(t)$ respectively:

$$\hat{x}_i(t) = \sum_{m=1}^M z_{im}\phi_m^X(t) \quad \text{and} \quad \hat{y}_i(t) = \sum_{k=1}^K w_{ik}\phi_k^Y(t).$$

where $z_{im}, w_{ik} \in \mathbb{R}$. We define $\boldsymbol{\phi}^{\mathbf{X}}(t) = [\phi_1^{\mathbf{X}}(t), \dots, \phi_M^{\mathbf{X}}(t)]$, $\boldsymbol{\phi}^{\mathbf{Y}}(s) = [\phi_1^{\mathbf{Y}}(s), \dots, \phi_K^{\mathbf{Y}}(s)]$, $\mathbf{z}_i = [z_{i1}, \dots, z_{iM}]$ and $\mathbf{w}_i = [w_{i1}, \dots, w_{iK}]$. We then model the regression surface using a double basis expansion (Ramsay and Silverman, 2005):

$$\beta(s, t) = \sum_{m=1}^M \sum_{k=1}^K b_{mk} \phi_m^{\mathbf{X}}(s) \phi_k^{\mathbf{Y}}(t) = \boldsymbol{\phi}^{\mathbf{X}}(s)^T \mathbf{B} \boldsymbol{\phi}^{\mathbf{Y}}(t), \quad (2.2)$$

for an $M \times K$ regression matrix \mathbf{B} . Using the above in Eq. (2.1); expressing $\epsilon_i(t) = \mathbf{q}_i^T \boldsymbol{\phi}^{\mathbf{Y}}(t)$, and using the orthonormality of $\boldsymbol{\phi}^{\mathbf{X}}(t)$ and $\boldsymbol{\phi}^{\mathbf{Y}}(t)$, gives:

$$\begin{aligned} \hat{y}_i(t) &= \int_I \hat{x}_i(s) \beta(s, t) ds + \epsilon_i(t), \\ \hat{y}_i(t) &= \int_I [\mathbf{z}_i^T \boldsymbol{\phi}^{\mathbf{X}}(s)] [\boldsymbol{\phi}^{\mathbf{X}}(s)^T \mathbf{B} \boldsymbol{\phi}^{\mathbf{Y}}(t)] ds + \epsilon_i(t), \\ \hat{y}_i(t) &= \mathbf{z}_i^T \left\{ \int_I \boldsymbol{\phi}^{\mathbf{X}}(s) \boldsymbol{\phi}^{\mathbf{X}}(s)^T ds \right\} \mathbf{B} \boldsymbol{\phi}^{\mathbf{Y}}(t) + \epsilon_i(t), \\ \hat{y}_i(t) &= \mathbf{z}_i^T \mathbf{B} \boldsymbol{\phi}^{\mathbf{Y}}(t) + \epsilon_i(t), \\ \mathbf{w}_i^T \boldsymbol{\phi}^{\mathbf{Y}}(t) &= \mathbf{z}_i^T \mathbf{B} \boldsymbol{\phi}^{\mathbf{Y}}(t) + \mathbf{q}_i^T \boldsymbol{\phi}^{\mathbf{Y}}(t), \\ \mathbf{w}_i^T \int_I \boldsymbol{\phi}^{\mathbf{Y}}(t) \boldsymbol{\phi}^{\mathbf{Y}}(t)^T dt &= \mathbf{z}_i^T \mathbf{B} \int_I \boldsymbol{\phi}^{\mathbf{Y}}(t) \boldsymbol{\phi}^{\mathbf{Y}}(t)^T dt + \mathbf{q}_i^T \int_I \boldsymbol{\phi}^{\mathbf{Y}}(t) \boldsymbol{\phi}^{\mathbf{Y}}(t)^T dt, \\ \mathbf{w}_i &= \mathbf{z}_i^T \mathbf{B} + \mathbf{q}_i^T. \end{aligned} \quad (2.3)$$

This parametrisation of the residual function is also used by Chiou et al. (2016). We can then estimate \mathbf{B} using standard multivariate regression methods typically assuming Gaussian \mathbf{q}_i .

2.2 Functional Principal Component Analysis

In this section we describe Functional Principal Component Analysis (FPCA), which we will use to estimate the basis functions $\boldsymbol{\phi}^{\mathbf{X}}(t)$ and $\boldsymbol{\phi}^{\mathbf{Y}}(t)$ for $x_i(t)$ and $y_i(t)$, respectively. These basis functions give effective, low-dimensional representations and will be used in the Functional Linear Regression model described in Section 2.1.

Functional Principal Component Analysis (FPCA) is a method of finding dominant modes of variance for functional data. These dominant modes of variance are called the Functional Principal Components (FPCs). FPCA is also used as a dimensionality reduction tool, as a set of observed curves can be effectively approximated by a linear combination of a small set of FPCs.

Let $X(t)$ be a mean-centred stochastic process with covariance function $C_X(s, t)$, with eigenfunctions $\phi_m^X(t)$ and eigenvalues λ_m^X with $m \in \mathbb{N}$, where the eigenfunctions are ordered in decreasing order of the corresponding eigenvalues. The Karhunen-Loève theorem (Shang, 2014) states

$$X(t) = \sum_{m=1}^{\infty} \xi_m^X \phi_m^X(t),$$

where each ξ_m^X is a random variable. Let $x_i(t)$, $i = 1, \dots, n$ be n independent realisations of $X(t)$. We can write $x_i(t) = \sum_{m=1}^{\infty} z_{im} \phi_m^X(t)$ where the principal component scores, z_{im} , constitute realisations of ξ_m^X . Since the eigenfunctions are orthonormal the scores can be defined as $z_{im} = \int_I x_i(t) \phi_m^X(t) dt$. In practise a large proportion of the variability in $x_i(t)$ can be captured by using a finite number M , of the first eigenfunctions. We can define the M -truncation as

$$\hat{x}_i^M(t) = \sum_{m=1}^M z_{im} \phi_m^X(t). \quad (2.4)$$

This minimises the reconstruction error, subject to the constraint that M basis functions are used:

$$\frac{1}{n} \sum_{i=1}^n \|x_i - \hat{x}_i^M\|^2 = \frac{1}{n} \sum_{i=1}^n \int_I [x_i(t) - \hat{x}_i^M(t)]^2 dt. \quad (2.5)$$

To choose M we will use an information criterion discussed in the next subsection. An analogous procedure is used to determine the K eigenfunctions $\phi_k^Y(t)$ for $y(t)$.

In Figure 2(a) we depict the first four eigenfunctions for the TGT curves shown in Figure 1(b). Note that the eigenfunctions do not resemble the shape of the TGT curves because, as described above, they are estimated after the data is preprocessed by subtract-

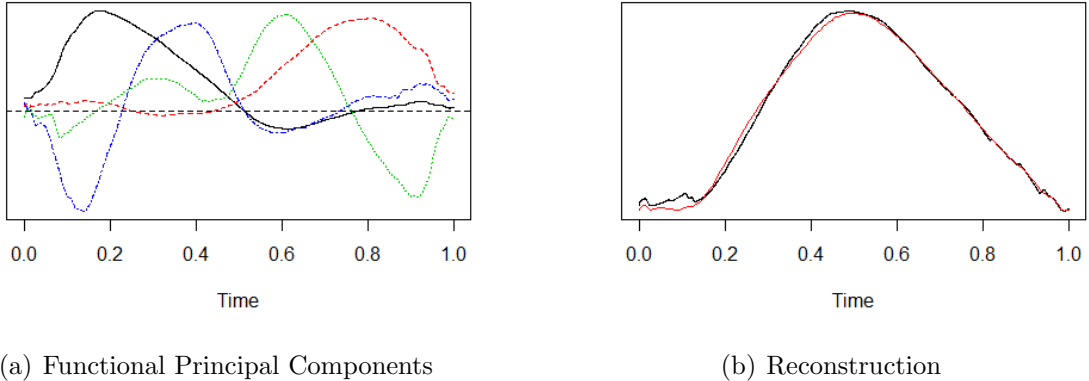


Figure 2: Plot of the first four eigenfunctions for Turbine Gas Temperature (TGT) data and a plot of a TGT curve (black) and its reconstruction using the first four eigenfunctions (red).

ing the mean function, $\alpha(t)$ in Eq. (2.1). In Figure 2(b) we depict the reconstruction of a single TGT curve using the first four eigenfunctions (four-truncation). The reconstruction is formed by taking the mean function plus a weighted sum of the eigenfunctions. We can see that using only four eigenfunctions produces an accurate approximation of the shape of this TGT curve except at the start where the discrepancy is noticeable.

2.3 Bayesian Information Criterion for FLR

In the FLR model described in Section 2.1 we need to choose terms M and K . Typically M and K are chosen independently (Yao et al., 2005), however the estimation of $\beta(s, t)$ also depends on M and K and this should be incorporated into the estimation of these terms. In this section we formulate a Bayesian Information Criterion (BIC) to determine the basis size M and K , similarly to Matsui (2020).

A component of the BIC is the log likelihood, often expressed as a squared error term. It is tempting to use the squared error resulting from Equation (2.3). However the objective is to fit the data \mathbf{y}_i so we should use a likelihood of this data instead of a squared error term of basis coefficients.

We have a set of models $J = \{(M, K) \mid M = 1, \dots, M_{\max}, K = 1, \dots, K_{\max}\}$, where M_{\max} and K_{\max} are pre-set maximum number of FPCs that will be considered in the model. Let vector $\vec{\mathbf{y}}_i$ be the values of $y_i(t)$ evaluated at discrete time points: $\vec{\mathbf{y}}_i = [y_i(t_1), \dots, y_i(t_T)]$. Let $\mathbf{z}_i^{(M)}$ be the first M principal scores of $x_i(t)$ with respect to the FPCs $\phi^{\mathbf{X}}(t)$ and let $\phi^{(\mathbf{K})}$ be the matrix with (k, i) entry $\phi_k^{\mathbf{Y}}(t_i)$. We assume there exists a true model (M_0, K_0) with associated $M_0 \times K_0$ matrix \mathbf{B}^{M_0, K_0} such that

$$\vec{\mathbf{y}}_i = (\mathbf{z}_i^{(M_0)})^T \mathbf{B}^{M_0, K_0} \phi^{(K_0)} + \boldsymbol{\epsilon}_i, \quad (2.6)$$

where the error $\boldsymbol{\epsilon}_i = [\epsilon_i(t_1), \dots, \epsilon_i(t_T)]$ is assumed for simplicity to be sampled from $N(0, v^2 \mathbf{I}_T)$, where \mathbf{I}_T is the identity matrix of size T . The assumption of independent observations in Eq. (2.8) is not tenable in the general case. However in our application this simplification is realistic because in accord with the practice followed by engineers, we subsample the high frequency sensor data at every 200 time points.

For Model (M, K) we define the parameters $\boldsymbol{\theta}^{M, K} = (\mathbf{B}^{M, K}, v^{M, K})$ and the prediction $\hat{\mathbf{y}}_i^{M, K} = (\mathbf{z}_i^{(M)})^T \mathbf{B}^{M, K} \phi^{(K)}$. We want to identify this true model (M_0, K_0) , which we can use to obtain consistent estimates of $\boldsymbol{\theta}^{M_0, K_0}$. For each Model (M, K) we define the likelihood for sample i as

$$f(\vec{\mathbf{y}}_i | \boldsymbol{\theta}^{M, K}) = \frac{1}{(2\pi)^{\frac{T}{2}} (v^{M, K})^T} \exp \left\{ -\frac{[\vec{\mathbf{y}}_i - \hat{\mathbf{y}}_i^{M, K}]^T [\vec{\mathbf{y}}_i - \hat{\mathbf{y}}_i^{M, K}]}{2(v^{M, K})^2} \right\}, \quad (2.7)$$

and the log-likelihood $l(\boldsymbol{\theta}^{M, K}) = \sum_{i=1}^n \log(f(\vec{\mathbf{y}}_i | \boldsymbol{\theta}^{M, K}))$. As in Eilers and Marx (1996)

$$BIC_n(M, K) = -2l(\hat{\boldsymbol{\theta}}^{M, K}) + \omega(M, K) \log(n), \quad (2.8)$$

where $\hat{\boldsymbol{\theta}}^{M, K}$ is the maximum likelihood estimator and the penalty $\omega(M, K) = MK + 1$, in which MK is the number of free parameters in the model and the 1 comes from v . We will denote $(M^*, K^*)_n = \arg \min_{(M, K) \in J} BIC_n(M, K)$, which depends on the sample size n .

To summarise, we estimate the FPCs for X and Y and solve the FLR model for different models (M, K) . We then choose model $(M^*, K^*)_n$ that minimises the BIC criterion. The robust equivalent of this procedure is given in Algorithm 1.

3 Robust Functional Linear Regression

In Section 2 we have defined the FLR model and have outlined the use of FPCA bases to estimate parameters of the model. In this section we will introduce robust versions of the FDA techniques outlined in Section 2. This will allow us to fit a model that captures the expected relationship of the predictor and response functions even in the presence of outliers. We shall also propose a robust BIC procedure for model selection.

We will replace classical FPCA with robust FPCA estimates by Bali et al. (2011) which ensure that outliers do not unduly affect the FPCA estimates. Note that FPCA minimises the residual error given in (2.5). To obtain robust FPCA estimates Bali et al. (2011) minimise a robust scale estimator, using a projection pursuit approach, which iteratively performs a weighted least squares till the estimators stabilise.

Analogous to (2.4), the robust FPCs $\tilde{\phi}_m^X(t)$, with $m = 1, \dots, M$, and $\tilde{\phi}_k^Y(t)$ with $k = 1, \dots, K$, are orthonormal functions such that

$$\tilde{x}_i(t) = \sum_{m=1}^M \tilde{z}_{im} \tilde{\phi}_m^X(t), \quad \tilde{y}_i(t) = \sum_{k=1}^K \tilde{w}_{ik} \tilde{\phi}_k^Y(t), \quad (3.1)$$

are good approximations for $x_i(t)$ and $y_i(t)$. We define $\tilde{y}_i(t) = \tilde{\mathbf{w}}_i^T \tilde{\boldsymbol{\phi}}^Y(t)$ and assume as in (2.3) that $\boldsymbol{\epsilon}_i = \tilde{\mathbf{q}}_i^T \tilde{\boldsymbol{\phi}}^Y(s)$. We can now write the robust counterpart of (2.3) as

$$\tilde{\mathbf{w}}_i^T = \tilde{\mathbf{z}}_i^T \tilde{\mathbf{B}} + \tilde{\mathbf{q}}_i^T. \quad (3.2)$$

To obtain a robust estimate of the regression matrix $\tilde{\mathbf{B}}$, we will use the Multivariate Least Trimmed Squares (MLTS) estimator by Agulló et al. (2008), to mitigate the affect

of outliers with respect to the regression relationship.

Let $[\cdot] : \mathbb{R} \rightarrow \mathbb{Z}$ denote the function that rounds its input to the nearest integer. Then, for $\alpha \in [0, 1]$ we define the set $\mathcal{S} = \{S \subset \{1, \dots, n\}, |S| = [\alpha n]\}$. The objective of MLTS is to find a subset S such that

$$S = \arg \min_{S \in \mathcal{S}} \sum_{i \in S} \|\tilde{\mathbf{w}}_i^T - \tilde{\mathbf{z}}_i^T \tilde{\mathbf{B}}\|^2.$$

This is robust as outliers will not be in the subset by definition so shall not affect the model estimation. According to the engineers' assessment between 1 and 5% of the data correspond to engines with issues. To ensure that no data from such engines are included in the calculation of the robust estimators we used a subset of size $[0.8n]$.

3.1 Robust Bayesian Information Criterion for FLR

The BIC model selection method is known to be non-robust (Machado, 1993). In particular outliers can significantly affect the loglikelihood estimation. We therefore outline a robust BIC (RBIC) model, which, similar to MLTS, maximises over a subset of samples S . RBIC can therefore give good model selection performance in the presence of outliers.

Denote as $\tilde{\boldsymbol{\theta}}^{M,K} = (\tilde{\mathbf{B}}^{M,K}, \tilde{v}^{M,K})$ the robust estimated parameters for model (M, K) , where $\tilde{\mathbf{B}}^{M,K}$ is the robust estimator of the regression matrix, obtained by solving Eq. (3.2); and $\tilde{v}^{M,K}$ is a robust trimmed estimator of the variance in Eq. (2.7). The robust prediction of \mathbf{y}_i by model (M, K) is given by $\tilde{\mathbf{y}}_i^{M,K} = (\tilde{\mathbf{z}}_i^{(M)})^T \tilde{\mathbf{B}}^{M,K} \tilde{\boldsymbol{\Phi}}^{(K)}$, where $\tilde{\boldsymbol{\Phi}}^{(K)} = [\tilde{\boldsymbol{\phi}}_1^Y, \dots, \tilde{\boldsymbol{\phi}}_K^Y]$ are the first K robust FPCs used in Eq. (3.1), and estimated through the method of Bali et al. (2011). We define the trimmed log-likelihood for model (M, K) and set S as

$$\tilde{l}(\tilde{\boldsymbol{\theta}}^{M,K}, S) = \sum_{i \in S} \left(\frac{[\tilde{\mathbf{y}}_i - \tilde{\mathbf{y}}_i^{M,K}]^T [\tilde{\mathbf{y}}_i - \tilde{\mathbf{y}}_i^{M,K}]}{(\tilde{v}^{M,K})^2} \right) + rT \log(2\pi) + 2rT \log(\tilde{v}^{M,K}). \quad (3.3)$$

We will define $S^{M,K} = \arg \min_{S \in \mathcal{S}} \tilde{l}(\tilde{\boldsymbol{\theta}}^{M,K}, S)$, where $\mathcal{S} = \{S \subset \{1, \dots, n\}, |S| = \lceil \alpha n \rceil\}$ where $\alpha = 0.8$. Then

$$\begin{aligned} RBIC_n(M, K) &= -2 \min_{S \in \mathcal{S}} \tilde{l}(\tilde{\boldsymbol{\theta}}^{M,K}, S) + \omega(M, K) \log(\lceil \alpha n \rceil) \\ &= -2 \tilde{l}(\tilde{\boldsymbol{\theta}}^{M,K}, S^{M,K}) + \omega(M, K) \log(\lceil \alpha n \rceil) \end{aligned} \quad (3.4)$$

We will denote $(\tilde{M}, \tilde{K})_n = \arg \min_{(M,K) \in J} RBIC_n(M, K)$, and assume that there is a unique minimiser, as is standard in the BIC literature (Schwarz, 1978).

In Algorithm 1 we outline the calculation of the RFLR model, which incorporates the RBIC procedure. In the algorithm we estimate the model for different values of (M, K) and choose the model with the minimum RBIC value. We consider $M = 1, \dots, M_{\max}$ and $l = 1, \dots, K_{\max}$ where M_{\max}, K_{\max} are chosen to ensure that 99.99% of the variance in the raw data is captured.

Data: Let $(\mathbf{x}_i, \mathbf{y}_i)$ be mean-corrected time series of length T for $i = 1, \dots, n$.

1. Estimate $\{\tilde{\phi}_1^X(t), \dots, \tilde{\phi}_{M_{\max}}^X(t)\}, \{\tilde{\phi}_1^Y(t), \dots, \tilde{\phi}_{K_{\max}}^Y(t)\}$ (Bali et al., 2011).

for $M = 1, \dots, M_{\max}$ **do**

for $K = 1, \dots, K_{\max}$ **do**

Estimate the regression matrix $\mathbf{B}^{M,K}$ using MLTS (Agulló et al., 2008).

Obtain the $RBIC_n(M, K)$ value using (3.4)

end for

end for

2. Select model $(\tilde{M}, \tilde{K})_n$.

return Regression matrix $\tilde{\mathbf{B}}$ from model $(\tilde{M}, \tilde{K})_n$ and $\{\tilde{\phi}_1^X(t), \dots, \tilde{\phi}_{\tilde{M}}^X(t)\}, \{\tilde{\phi}_1^Y(t), \dots, \tilde{\phi}_{\tilde{K}}^Y(t)\}$.

Algorithm 1: Robust FLR procedure

4 Consistency Results

In Section 3 we proposed a Robust FLR model for the function-on-function problem. In this section we outline our asymptotic consistency results. The proofs are contained in the Supplementary Material. Asymptotic consistency is a minimum requirement for a

good model. As Silverman (1996) argues asymptotic consistency results should be seen as large-sample approximation results, which can aid our intuition about a method and are practically relevant when reasonably large amounts of data are available. Silverman (1996) also argues that such results are particularly important for functional data analysis because in certain contexts “obvious” procedures are not consistent and may therefore produce misleading insights.

Definition 1. Let X_1, X_2, \dots, X_n be a sequence of real-valued random variables. An estimator $T_n := T(X_1, X_2, \dots, X_n)$ of a parameter θ is said to be (asymptotically) **consistent** if for all $\varepsilon > 0$

$$\lim_{n \rightarrow \infty} P(|T_n - \theta| > \varepsilon) = 0.$$

Definition 2. Let X_1, X_2, \dots, X_n be a sequence of real-valued random variables with an associated cumulative distribution function F_θ , which depends on an unknown parameter θ . Let the estimator $T_n := T(F_n)$ of a parameter θ , be a function of the empirical distribution function F_n . We say this estimator is **Fisher-consistent** for the parameter θ if

$$T(F_\theta) = \theta$$

Remark 1. Fisher consistency is equivalent to (asymptotic) consistency if the empirical distribution function F_n converges pointwise to the true distribution function F_θ . This can be shown to be the case for iid real multivariate random variables using the Glivenko-Cantelli theorem (Pollard, 2012).

Lemma 4.1. Assume C1-C6 (given in Section 1 of the supplemental material) hold then the robust regression function $\tilde{\beta}(s, t)$ is Fisher-consistent.

Corollary 4.1. If $\{x_1(t), y_1(t)\}, \dots, \{x_n(t), y_n(t)\}$ are iid samples with cumulative distribution function (F_X, F_Y) . Then, assuming C1-C6 (given in Section 1 of the supplemental

material) hold, $\tilde{\beta}(s, t)$ is consistent.

Note that $x_i(t)$ and $y_i(t)$ are defined on a finite number of eigenfunctions, so are defined by finite score vectors. Therefore Corollary 4.1 follows from Lemma 4.1 and Remark 1, which states almost sure convergence of the empirical distribution for iid multivariate random variables. In this case Fisher-consistency is equivalent to consistency.

Theorem 4.1. *Given Assumptions 1 and 2 (given in Section 1 of the supplemental material) hold, and there exists a true model (M_0, K_0) , then $(\tilde{M}, \tilde{K})_n$ is a consistent estimator of (M_0, K_0) .*

5 Outlier Detection

There is a rich literature on outlier detection methods for functional data. This includes methods based on outlyingness measures such as Arribas-Gil and Romo (2014), and Dai and Genton (2018), and methods relying on functional depth such as the thresholding approach by Febrero-Bande et al. (2008) and the functional boxplot by Sun and Genton (2011). These methods are typically applied directly to functional representations of the response time-series, which in our case would be those of the engine temperature parameters. As the discussion around Figure 1 in the introduction indicated outliers identified through this approach can be meaningless, because the dependency between engine speed and temperature is ignored. Our approach relies on the premise that once this dependency is accounted for through an RFLR model the residual curve $r_i(t) = y_i(t) - \tilde{y}_i(t)$ of an engine that has issues will differ from other residual curves, and thus constitute an outlier. We therefore apply the outlier detection method of Febrero-Bande et al. (2008) to residual curves from RFLR. We next outline the notion of depth and the motivation for selecting this approach.

To order univariate data one can directly use order statistics, however ordering multivariate data is not trivial. This motivates depth functions, defined as $D : \mathbb{R}^d \rightarrow \mathbb{R}$, that

map a multivariate point to a univariate depth value (Tukey, 1975). A key property of a depth function is that a point x_1 that is closer to the “centre” of a data cloud than a point x_2 will be assigned a larger depth value. Points sufficiently far from the centre can be considered as outliers, and can be identified by their low depth values. The notion of depth has been extended to functional data (Nieto-Reyes and Battey, 2016) and is referred to as functional depth. The functional depth based outlier detection method of Febrero-Bande et al. (2008) allows us to detect not only residual curves with large integrated squared error, but also *shape outliers* (Hubert et al., 2015). The latter are curves that do not appear unusual if viewed at each time point but are abnormal across the entire trajectory. Shape outliers can be missed if we use methods that rely solely on integrated square error.

There are a number of depth functions for functional data that can be chosen. We will use the h -modal depth (Cuevas et al., 2007) to rank samples r_i , as it works well in practice. For a given kernel G_h (typically Gaussian with bandwidth h), the h -modal depth of r_i with respect to $\mathbf{r} = (r_1, \dots, r_n)$ is given by:

$$D(r_i|\mathbf{r}, h) = \frac{1}{n} \sum_{j=1}^n G\left(\frac{\|r_i - r_j\|}{h}\right). \quad (5.1)$$

The h -modal depth has two useful properties. First, it uses a distance metric therefore the depth values will be proportional to the distance from the “centre”. Second, it doesn’t assume a unique centre and therefore it is appropriate when there are more than one “mode” in the data cloud.

In the algorithm we need to choose the bandwidth h and a threshold C to identify outliers. The bandwidth h is taken to be the 15th percentile of the empirical distribution of $\{\|r_i - r_j\|, i, j = 1, \dots, n\}$ (Febrero-Bande et al., 2008). The threshold C is chosen such that $P(D(r_i|\mathbf{r}, h) \leq C) = \delta$, where δ is a pre-chosen percentile. To estimate the threshold C they use a bootstrapping approach, which estimates a value of C for different random sets of samples and then aggregates these estimates. We describe the outlier detection

algorithm in Algorithm 2.

Data: Centred curves $\{x_i(t), y_i(t)\}$ for $i = 1, \dots, n$ and percentile δ .

1. Use Algorithm 1 to obtain $\tilde{\phi}_k^Y(t)$, \tilde{z}_m and $\tilde{\mathbf{B}}$.
2. Calculate residual curves $r_i(t)$.
3. Estimate bandwidth h .
4. For each $r_i(t)$ calculate $D(r_i|\mathbf{r}, h)$.
5. Estimate C for given percentile δ .
6. If $D(r_i|\mathbf{r}, h) < C$ sample i is an outlier.

Algorithm 2: Outlier Detection using robust FLR.

6 Simulation Study

In this section we will provide a simulation study to investigate the finite sample properties of RBIC and robust FLR (RFLR) in comparison to BIC and classical FLR (CFLR). Code to reproduce the experimental results reported in this section can be found at the GitHub repository <https://github.com/hullait/RobustFLR>. In the simulation study we will generate data using a FLR process and corrupt a certain number of samples, which will be the outliers. The outliers have been designed to be undetectable, if the response curves are considered independently of the predictor curves. Therefore standard functional data outlier detection algorithms such as those discussed in Section 5 will perform poorly.

The main motivation for the RFLR model is to obtain good model fitting in the presence of outliers. In this simulation study we compare the fitting error (FE) given in (6.1), for the non-outlier samples using the robust model, which uses RFLR and RBIC with the classical approach using CFLR and BIC. We define the indicator variable $u_i = 1$ if sample i is an outlier and zero otherwise. Letting $\hat{y}_i(t)$ be the estimation of $y_i(t)$ and given that proportion a of the samples have been contaminated then FE is given by:

$$\text{FE} = \frac{1}{(1-a)n} \sum_{i=1}^n (1-u_i) \|y_i - \hat{y}_i\|^2. \quad (6.1)$$

Next we compare the outlier detection capabilities of robust and classical approaches us-

ing the receiver operating characteristic (ROC) curve to determine the sensitivity/specificity trade-off for different thresholds. If we have perfect outlier detection for all thresholds then the area under the curve (AUC) of the ROC curve would be one. We can therefore use the AUC value as a measure of outlier detection accuracy regardless of threshold.

FPCA is performed by taking the principal components of a cubic B-spline representation of each of the predictor and response curves (Ramsay and Silverman, 2005). The robust FPCA approach outlined in Section 3 is performed using the CR algorithm proposed by Croux and Ruiz-Gazen (1996) on the same B-spline coefficients. The MLTS estimator is calculated using the heuristic given by Agulló et al. (2008) using different trimming proportions $(1 - \alpha)$ for $\alpha \in [0, 1]$.

6.1 Scenarios

We will generate samples $x(t)$ using a FPCA based model with mean function $\mu_X(t) = -10(t - 0.5)^2 + 2$ for $t \in [0, 1]$ and eigenfunctions:

$$\phi_1^X = \sqrt{2} \sin(\pi t), \quad \phi_2^X = \sqrt{2} \sin(7\pi t), \quad \phi_3^X = \sqrt{2} \cos(7\pi t).$$

The principal scores are sampled from Gaussian distributions with mean 0 and variances 40, 10 and 1 for the eigenfunctions respectively. Note that we do not create any outliers in the FPCA decompositions of the predictor curves. We generate 400 predictor curves $x_1(t), \dots, x_{400}(t)$, which are observed at $T = 500$ equidistant points in the interval $[0, 1]$.

The samples $y(t)$ will have eigenfunctions:

$$\phi_1^Y = \sqrt{2} \sin(12\pi t), \quad \phi_2^Y = \sqrt{2} \sin(5\pi t), \quad \phi_3^Y = \sqrt{2} \cos(2\pi t),$$

and mean function $\mu_Y(t) = 60 \exp(-(t - 1)^2)$. We will generate $\beta(s, t) = \boldsymbol{\phi}^{\mathbf{X}}(s)^T \mathbf{B} \boldsymbol{\phi}^{\mathbf{Y}}(t)$ where \mathbf{B} will have random entries between $[-3, 3]$. We generate non-outlier curves:

$$y_i(t) = \mu_Y(t) + \int_I \beta(s, t)(x_i(s) - \mu_X(s))ds + \epsilon_i(t),$$

where the residual function $\epsilon_i(t) = \mathbf{q}_i^T \boldsymbol{\phi}^Y(t) + d_i$ where \mathbf{q}_i^T and d_i are sampled iid from $N(0, 0.1)$. We will consider three cases when the proportion of outliers are $a = 0.1, 0.2$ and 0.3 .

In **Scenario 1** outliers will be generated by replacing \mathbf{B} with $\mathbf{B}_1 = \mathbf{B} + \mathbf{R}$ where \mathbf{R} has random entries sampled from $N(0, 0.5)$ giving $\beta_1(s, t) = \boldsymbol{\phi}^X(s)^T \mathbf{B}_1 \boldsymbol{\phi}^Y(t)$. Outliers $y'_i(t)$ are given by

$$y'_i(t) = \mu_Y(t) + \int_I \beta_1(s, t)(x_i(s) - \mu_X(s))ds + \epsilon_i(t).$$

In **Scenario 2** we generate outliers by adding a random B-spline function $p(t)$ defined on an interval of length $1/10$. Letting $\beta_2(s, t) = \boldsymbol{\phi}^X(s)^T \mathbf{B}_2 [\boldsymbol{\phi}^Y(t), p(t)]$, for 3×4 matrix $\mathbf{B}_2 = [\mathbf{B}, \mathbf{l}]$ for $\mathbf{l} \sim N(2, 1)$, then the outliers $y''_i(t)$ are given by

$$y''_i(t) = \mu_Y(t) + \int_I \beta_2(s, t)(x_i(s) - \mu_X(s))ds + \epsilon_i(t).$$

Note that the outliers in Scenario 1 affect the regression function across the entire interval whereas the outliers in Scenario 2 only affect a small interval of the curves.

In Figure 3 we have a plot of the predictor curves $x_i(t)$ and response curves $y_i(t)$ with outliers from Scenario 1 and Scenario 2. The figure shows the outliers are masked by the variability in the curves and therefore cannot be identified using standard outlier detection algorithms. To make the outliers clearer we have plotted the residuals of the response curves using the true regression function and mean functions. In the bottom row of Figure 3 we can see that the outliers in Scenario 2 are localised to a fixed interval whereas in Scenario 1 the outliers affect the response curve at all time points.

The RFLR model depends on the proportion of trimming α . To investigate the effect

of the trimming we will consider trimming proportions $\alpha = 0.1, 0.2$ and 0.3 . We shall also investigate the performance using BIC and RBIC with fixed trimmed sample size of $r = [0.8n]$.

We sample 400 predictor and response curve datasets and generate classical and robust models to calculate the average FE (6.1). In Tables 1 and 2 we present the results for Scenario 1 and 2 respectively. The CFLR model gives a smaller FE value in the case of no-outliers $a = 0$, however the robust model still gives good model fits. If we compare the FE using BIC and RBIC, we can see that BIC gives better model choices when $a = 0$. This is due to BIC using all the data and in particular using samples in the tails of the distribution. In the presence of outliers the robust model outperforms the classical model, and as expected the difference in FE increases as the number of outliers increases. We should also note that RBIC is giving better model choices than BIC when outliers are present. Next, we can see using trimming proportion $\alpha = 0.1$ we obtain significantly large FE values when $a = 0.3$. However the FE values for $\alpha = 0.2$ and 0.3 are very similar in the case of $a = 0.3$. The outliers generated can have different sizes, therefore in the $\alpha = 0.2$ robust model only small outliers are present, which only affect the model fitting slightly .

In Figure 4 we have two ROC curves generated for one of the repetitions in Scenario 1 and 2 in which we have contaminated 20% of the samples. In both scenarios the robust model outperforms the classical model. We also deploy the approach of Febrero-Bande et al. (2008) to the response curves, disregarding the predictor curves (henceforth called the Direct approach). The ROC curves show that the robust and classical models are more effective than the Direct method in identifying the outliers in Scenario 1 and 2. By only using the specificity and sensitivity for a fixed threshold a lot of information is being lost, therefore a better comparison would be the area under the curve (AUC). Using the AUC metric we can understand the model outlier detection capabilities overall, in particular how well are the outliers separated from the other samples. We have taken the average AUC values over the 100 iterations performed for Scenario 1, which are shown in Table

Table 1: Average fitting errors (FE) for 100 replications for Scenario 1, using classic FPCA and robust FPCA with different amount of trimming in the MLTS estimator and using models selected by BIC and RBIC.

	Trim	Model	a=0	a=0.1	a=0.2	a=0.3
Classic	$\alpha = 0.0$	BIC	5.326	18.441	48.771	101.320
Robust	$\alpha = 0.1$	BIC	8.283	14.166	21.118	33.907
	$\alpha = 0.1$	RBIC	9.285	9.179	10.674	28.393
	$\alpha = 0.2$	BIC	8.288	14.178	15.750	16.623
	$\alpha = 0.2$	RBIC	9.292	9.207	9.535	13.436
	$\alpha = 0.3$	BIC	8.294	14.199	15.815	16.518
	$\alpha = 0.3$	RBIC	9.301	9.214	9.544	12.334

Table 2: Average fitting errors (FE) for 100 replications for Scenario 2, using classic FPCA and robust FPCA with different amount of trimming in the MLTS estimator and using models selected by BIC and RBIC.

	Trim	Model	a=0	a=0.1	a=0.2	a=0.3
Classic	$\alpha = 0.0$	BIC	5.326	17.252	48.906	85.063
Robust	$\alpha = 0.1$	BIC	8.283	15.242	21.524	28.758
	$\alpha = 0.1$	RBIC	9.285	9.074	9.919	18.546
	$\alpha = 0.2$	BIC	8.288	16.745	20.652	21.928
	$\alpha = 0.2$	RBIC	9.292	9.191	8.997	13.628
	$\alpha = 0.3$	BIC	8.294	16.808	20.695	21.750
	$\alpha = 0.3$	RBIC	9.301	9.233	9.018	11.439

3. We have considered the average AUC values for trimming levels $\alpha = 0.1, 0.2$ and 0.3 . The robust models give larger AUC values than the classical model. However the different trimming levels do not seem to have a significant effect on the AUC values. In Scenario 2 we have the results in Table 4. The same patterns appear as in Scenario 1 except the AUC values are notably smaller. This is to be expected given the outliers in Scenario 2 are defined on a small time interval.




Figure 3: *Left:* Plots of the predictor curves $x_i(t)$, response curves $y_i^{(1)}(t)$ and residuals curves $r_i^{(1)}(t)$ for Scenario 1. *Right:* Plots of the predictor curves $x_i(t)$, response curves $y_i^{(2)}(t)$ and residuals curves $r_i^{(2)}(t)$ for Scenario 2. The residual curves are generated using the true regression function and mean functions. In each scenario there are 5 outliers each in a distinctive colour. The predictors curves $x_i(t)$ are identical for both scenarios, and the response curves look very similar due to mean and functional components masking the outliers. However the residuals are clearly distinctive.

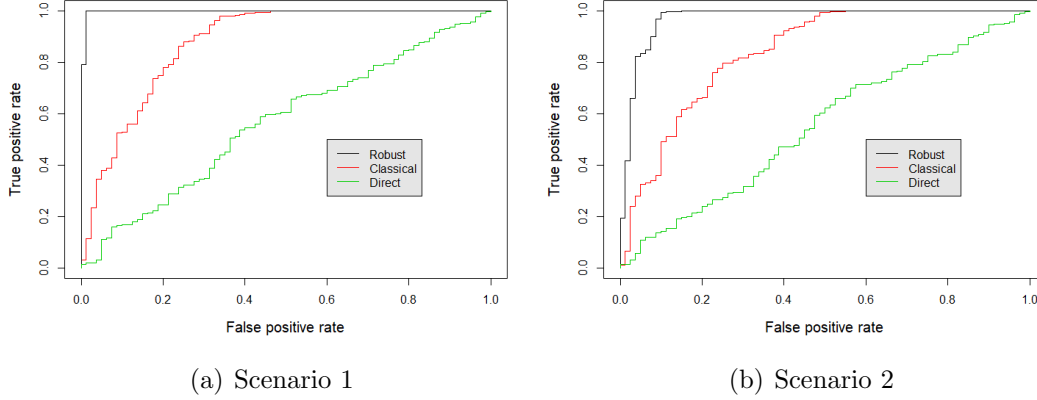


Figure 4: ROC curve for one instance of Scenario 1 and 2 with the proportion of outlier $a=0.2$ and proportion trimmed $\alpha = 0.2$.

Table 3: Average AUC values over 100 replications for Scenario 1, using proportion of outliers $a=0.1, 0.2$ and 0.3 . Using Direct compared to classic FPCA with BIC, and using robust FPCA with RBIC and trimming levels $\alpha = 0.1, 0.2$ and 0.3 .

		a=0.1	a=0.2	a=0.3
Direct	-	0.532	0.538	0.550
Classic	$\alpha = 0.0$	0.960	0.898	0.797
Robust	$\alpha = 0.1$	0.995	0.991	0.953
	$\alpha = 0.2$	0.996	0.996	0.987
	$\alpha = 0.3$	0.996	0.996	0.990

Table 4: Average AUC values over 100 replications for Scenario 2, using proportion of outliers $a=0.1, 0.2$ and 0.3 . Using Direct compared to classic FPCA with BIC, and using robust FPCA with RBIC and trimming levels $\alpha = 0.1, 0.2$ and 0.3 .

		a=0.1	a=0.2	a=0.3
Direct	-	0.512	0.548	0.554
Classic	$\alpha = 0.0$	0.922	0.838	0.734
Robust	$\alpha = 0.1$	0.985	0.964	0.932
	$\alpha = 0.2$	0.980	0.980	0.966
	$\alpha = 0.3$	0.980	0.980	0.968

7 Jet Engine data

The jet engine dataset contains sensor measurements taken during 199 Vibration Survey (VS) manoeuvres. These manoeuvres have been taken from multiple engine tests. To minimise the variability in the data we have ensured that the tests were performed on the same engine type and at the same test facility as environmental factors can affect engine parameters during the test. The VS manoeuvre has a distinctive engine speed profile with a slow acceleration followed by a slow deceleration, as shown in Figure 1(a). The TPR is the engineer-controlled measure of engine thrust. As shown in Figure 1 engine temperature parameters like the TGT follow a broadly similar pattern to TPR, during the acceleration phase. As expected, the temperature parameters reach their maximum value with some delay compared to TPR and although they decrease during the deceleration phase they are considerably higher than their initial values at the end of the manoeuvre. This highlights the trajectory-dependent behaviour that we seek to model.

We do not have labels for whether any of the individual engines have outliers but we do have log books from the engine test, from which we can obtain insights into the Vibration Survey manoeuvres which our method flags as outliers. There are a number of temperature features measured within an engine including the TGT, discussed previously. In addition we have four other temperature readings T25 (the temperature reading at sensor station 25), T30 (the temperature reading at sensor station 30), TCAR (Temperature of Casing at Rear) and TCAF (Temperature of Casing at Front). All the temperature features are shown in Figure 7. The TCAR is particularly interesting as it has two distinct curve behaviours. The VS manoeuvres time series are of similar length. To standardise we have fitted a B-spline basis of 400 basis functions (with equally spaced knots) to each to ensure the time series are well approximated. We have experimented with different number of basis functions in the range between 200 and 500 and our results were not affected by this choice. A lower dimensional fit can be achieved with a smarter placement of knots, but

we avoided this option to minimise the degree of manual tuning required. We have taken 1000 equally spaced points on the B-spline representations to be our inputs $x_i(t)$ and $y_i(t)$. Note that if all time series are of equal length we could use standard PCA. However, as shown by Ramsay and Silverman (2005) this can lead to noisy estimates of the principal components, which can negatively affect the subsequent model fitting.

We first estimate an RFLR model for each temperature parameter, through the procedure outlined in Algorithm 1. As specified in this algorithm, the number of eigenfunctions to represent the predictor and response curves, M and K respectively, is determined through RBIC. In all cases M and K range between three and five.

We can then apply the outlier detection algorithm described in Algorithm 2. We compare these outliers with those detected on the temperature curves directly and using CFLR and BIC in Algorithm 2. We can look at the residuals curves to verify whether the outliers appear to be abnormal. In particular we want to show that using functional regression we are able to determine outliers that would otherwise be missed by investigating the temperature curves directly.

Using the depth based outlier detection (Direct) (Febrero-Bande et al., 2008) directly on the temperature curves (with a default threshold of $\delta = 0.01$), we obtain the outliers in Table 5. We can see that the outliers in the TPR are the same as the outliers in the temperature features. This suggests the outliers being identified are arising from the controller induced variability. We therefore need to model the dependency between the control feature (TPR) and the temperature features.

We applied the outlier detection algorithm given in Algorithm 2 using CFLR and BIC with threshold $\delta = 0.01$. The outliers identified are given in Table 5. The residuals curves are shown in Figures 5 and 6, with the outliers coloured in blue. It is not clear from this plot that the outliers are truly different from the other data.


Lastly we applied Algorithm 2 using RFLR and RBIC with threshold $\delta = 0.01$. The outlier samples are given in Table 5 for each temperature feature. In Figures 5 and 6 we

Temp	Direct	CFLR	RFLR
TPR	33, 106, 167	-	-
T25	33, 106, 167	24, 182	24, 70, 106
T30	33, 106, 167	24, 182, 192	24, 44, 70, 106, 196
TGT	33, 106, 167	119, 153	44, 70, 106, 117
TCAR	33, 106	36, 91, 106	70, 106
TCAF	33, 167	65, 167, 170, 171	24, 70, 106


Table 5: Outliers detected for temperature features (Temp) using outlier detection on the temperature features directly (Direct), and the outliers found using CFLR and RFLR.

have the residual curves using RFLR. We can see that the RFLR model fits the majority of the temperature curves well. The outliers that are picked up clearly look abnormal, with significant deviations from the general behaviour. The RFLR model is therefore able to identify interesting behaviour, which may otherwise have been undetected. Engineers have informed us that Sample 24 comes from an engine in which they detected damaged hardware. All the other outliers in the RFLR column of Table 5 were also noted to come from engines that displayed odd behaviour during the Pass-Off test. This is not the case for the outliers reported in the CFLR column.


In Figure 7 we have a plot of the temperature parameters with the outliers identified using the curves directly in green, those using the RFLR model in red and those detected by both in purple. We can see that the outliers from the RFLR model do not necessarily appear as abnormal if we look at the temperature curves directly. Sample 106 is identified as an outlier by multiple temperature features and also when the depth based outlier detection is used on the temperature curves directly. Comparing the outliers identified using a classical approach, we can see Sample 24 is identified as an outlier multiple times using the classical and robust approaches. However most of the outliers from the classical approaches differ with the outliers identified using the robust approach. We can also see that the outliers using the RFLR are significantly more distinctive than the outliers using CFLR.




(a) Residuals of T25 through CFLR




(b) Residuals of T25 through RFLR




(c) Residuals of T30 through CFLR



(d) Residuals of T30 through RFLR




(e) Residuals of TGT through of CFLR




(f) Residuals of TGT through RFLR


Figure 5: Plots of the residual curves for temperature at stations 25 and 30, T25 and T30, respectively, and turbine gas temperature (TGT). Non-outliers are in the grey region while outliers are depicted in blue and red for the CFLR and RFLR models respectively. Figures on the left column are obtained through CFLR while those on the right through RFLR.




(a) Residuals of TCAR through CFLR



(b) Residuals of TCAR through RFLR




(c) Residuals of TCAF through CFLR




(d) Residuals of TCAF through RFLR


Figure 6: Plots of residuals curves for temperature of casing at rear and at Front, TCAR and TCAF respectively. Non-outliers are in the grey region and outliers are depicted in blue and red for the CFLR and RFLR models respectively. Figures on the left column are obtained through CFLR while those on the right through RFLR.




(a) Turbine Pressure Ratio




(b) Temperature at station 25




(c) Temperature at station 30



(d) Turbine Gas Temperature




(e) Temperature of Casing at Rear




(f) Temperature of Casing At Front

Figure 7: Plots of temperature time series with with non-outliers in grey region and outliers using robust FLR in red; those using the curves directly in green and those for both in purple.



(a) T25



(b) TGT

Figure 8: Plot of the regression functions $\tilde{\beta}(s, t)$ for T25 and TGT parameters.

The RFLR model produces a regression function $\tilde{\beta}(s, t)$, which models the relationship between the engine speed and engine temperature. In Figure 8 we depict two indicative regression surfaces for the temperature parameters T25 and TGT. Note that the number of eigenfunctions used to represent the response and the predictor curves, K and M respectively, affects the estimated $\tilde{\beta}(s, t)$. For both regression functions in Figure 8 large in absolute value entries appear on or close to the diagonal, which corresponds to $s = t$. However, not all the weight is on the diagonal, and in particular it is commonly the case that $\tilde{\beta}(s, t)$ is non-zero for $s > t$. This suggests that using information from the current and past time-points only ($s \leq t$) can substantially limit the accuracy of an FLR model. FLR models that are restricted to not use information from future time-points are called *historical FLR* models (Malfait and Ramsay, 2003). We developed such models and indeed found that they fit the data less well. A further limitation of this approach is that no robust historical FLR models exist and developing one is challenging. One major difficulty is that the basis functions required for the historical FLR model are not necessarily orthogonal.

8 Conclusion

There exist a number of functional regression models for functional inputs and responses, however these methods are not robust to outliers. We have introduced a robust FLR model that is able to produce good model fits in the presence of outliers. Alongside the robust FLR model we have also introduced a robust model selection procedure and proven the consistency of the robust FLR and model selection procedure. Using a simulation study we have shown the need for a robust approach to obtain good models in the presence of outliers. The robust FLR model is also effective in identifying global and localised outliers. Finally using jet engine sensor data as a motivating application for robust FLR we have identified unusual temperature behaviour. In particular the outliers identified in the jet engine sensor data would not have been detected if we modelled the response variables independently of human controlled driving variable.

SUPPLEMENTARY MATERIAL

Supplementary paper: A PDF containing proofs of the results in Section 4.

RobFLR: R code used for simulation study available at [github/hullait/RobFLR](https://github.com/hullait/RobFLR)

References

- Agulló, J., C. Croux, and S. Van Aelst (2008). The multivariate least-trimmed squares estimator. *Journal of Multivariate Analysis* 99(3), 311–338.
- Arribas-Gil, A. and J. Romo (2014). Shape outlier detection and visualization for functional data: the outliergram. *Biostatistics* 15 4, 603–19.
- Bali, J. L., G. Boente, D. E. Tyler, and J.-L. Wang (2011). Robust functional principal components: A projection-pursuit approach. *Annals of Statistics* 39(6), 2852–2882.

- Chiou, J.-M., Y.-F. Yang, and Y.-T. Chen (2016). Multivariate functional linear regression and prediction. *Journal of Multivariate Analysis* 146, 301 – 312.
- Croux, C. and A. Ruiz-Gazen (1996). A fast algorithm for robust principal components based on projection pursuit. In A. Prat (Ed.), *COMPSTAT*, pp. 211–216. Physica-Verlag HD.
- Cuevas, A., M. Febrero, and R. Fraiman (2007). Robust estimation and classification for functional data via projection-based depth notions. *Computational Statistics* 22(3), 481–496.
- Dai, W. and M. G. Genton (2018). Multivariate functional data visualization and outlier detection. *Journal of Computational and Graphical Statistics* 27(4), 923–934.
- Eilers, P. H. C. and B. D. Marx (1996). Flexible smoothing with b -splines and penalties. *Statistical Science* 11(2), 89–121.
- Febrero-Bande, M., P. Galeano, and W. González-Manteiga (2008). Outlier detection in functional data by depth measures, with application to identify abnormal nox levels. *Environmetrics* 19(2), 331 – 345.
- Ferraty, F., I. V. Keilegom, and P. Vieu (2012). Regression when both response and predictor are functions. *Journal of Multivariate Analysis* 109, 10 – 28.
- Houstis, E., A. Catlin, P. Tsompanopoulou, D. Gottfried, G. Balakrishnan, K. Su, and J. Rice (2002). Gasturbnlab: a multidisciplinary problem solving environment for gas turbine engine design on a network of nonhomogeneous machines. *Journal of Computational and Applied Mathematics* 149(1), 83 – 100.
- Hubert, M., P. J. Rousseeuw, and P. Segaert (2015). Multivariate functional outlier detection. *Statistical Methods & Applications* 24(2), 177–202.

- Ivanescu, A. E., A.-M. Staicu, F. Scheipl, and S. Greven (2015). Penalized function-on-function regression. *Computational Statistics* 30(2), 539–568.
- Machado, J. A. F. (1993). Robust model selection and M-estimation. *Econometric Theory* 9(3), 478–493.
- Malfait, N. and J. O. Ramsay (2003). The historical functional linear model. *The Canadian Journal of Statistics / La Revue Canadienne de Statistique* 31(2), 115–128.
- Matsui, H. (2020). Quadratic regression for functional response models. *Econometrics and Statistics* 13, 125–136.
- Morris, J. S. (2015). Functional Regression. *Annual Review of Statistics and Its Application* 2, 321–359.
- Nieto-Reyes, A. and H. Battey (2016, 02). A topologically valid definition of depth for functional data. *Statistical Science* 31(1), 61–79.
- Pollard, D. (2012). *Convergence of stochastic processes*. Springer Science & Business Media.
- Ramsay, J. and B. W. Silverman (2005). *Functional Data Analysis (Springer Series in Statistics)*.
- Ramsay, J. O. and C. J. Dalzell (1991). Some tools for functional data analysis. *Journal of the Royal Statistical Society. Series B (Methodological)* 53(3), 539–572.
- Scheipl, F., A.-M. Staicu, and S. Greven (2015). Functional additive mixed models. *Journal of Computational and Graphical Statistics* 24(2), 477–501.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics* 6(2), 461–464.
- Shang, H. L. (2014). A survey of functional principal component analysis. *AStA Advances in Statistical Analysis* 98(2), 121–142.

- Silverman, B. W. (1996). Smoothed functional principal components analysis by choice of norm. *The Annals of Statistics*. *24*(1), 1–24.
- Sun, Y. and M. G. Genton (2011). Functional boxplots. *Journal of Computational and Graphical Statistics* *20*(2), 316–334.
- Tukey, J. W. (1975). Mathematics and the picturing of data. In R. D. James (Ed.), *International Congress of Mathematicians*, Volume 2, pp. 523–532.
- Yao, F., H.-G. Müller, and J.-L. Wang (2005). Functional linear regression analysis for longitudinal data. *The Annals of Statistics* *33*(6), 2873–2903.