



HAL
open science

Interpretable sparse SIR for functional data

Victor Picheny, Rémi Servien, Nathalie Vialaneix

► **To cite this version:**

Victor Picheny, Rémi Servien, Nathalie Vialaneix. Interpretable sparse SIR for functional data. 2016.
hal-01325090v2

HAL Id: hal-01325090

<https://hal.science/hal-01325090v2>

Preprint submitted on 2 Jun 2016 (v2), last revised 1 Mar 2018 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Interpretable sparse SIR for functional data

Victor Picheny

MIAT, Université de Toulouse, INRA, Castanet Tolosan, France,

Rémi Servien

Toxalim, Université de Toulouse, INRA, Toulouse, France

and

Nathalie Villa-Vialaneix

MIAT, Université de Toulouse, INRA, Castanet Tolosan, France

June 2, 2016

Abstract

This work focuses on the issue of variable selection in functional regression. Unlike most work in this framework, our approach does not select isolated points in the definition domain of the predictors, nor does it rely on the expansion of the predictors in a given functional basis. It provides an approach to select full intervals made of consecutive points. This feature improves the interpretability of the estimated coefficients and is desirable in the functional framework for which small shifts are frequent when comparing one predictor (curve) to another. Our method is described in a semiparametric framework based on Sliced Inverse Regression (SIR). SIR is an effective method for dimension reduction of high-dimensional data which computes a linear projection of the predictors in a low-dimensional space, without loss on regression information. We extend the approaches of variable selection developed for multidimensional SIR to select intervals rather than separated evaluation points in the definition domain of the functional predictors. Different and equivalent formulations of SIR are combined in a shrinkage approach with a group-LASSO-like penalty. Finally, a fully automated iterative procedure is also proposed to find the critical (interpretable) intervals. The approach is proved efficient on simulated and real data.

Keywords: functional regression, SIR, lasso, ridge regression, interval selection.

1 Introduction

In numerous applications, data correspond to continuous processes sampled at different evaluation points. Examples of such situations arise in various fields: daily records of meteorological data (temperature, rainfall), financial time series, spectra in chemometrics... The analysis of such data is often referred to as “Functional Data Analysis” (FDA); a comprehensive introduction to FDA can be found in Ramsay and Silverman (1997) or Ferraty and Vieu (2006). From a practical point of view, such data can be viewed as observations of a random variable which takes values in a (infinite dimensional) functional space (*e.g.*, $L^2([0, 1])$, the space of squared integrable functions over $[0, 1]$). The functional random variable is only partially observed at some evaluation points, whose positions can sometimes depend on the observation itself. A challenging problem with FDA lies in its high dimension: the underlying dimension of a functional space is infinite, and even if the digitized version of the curves is considered, the number of evaluation points is typically much larger than the number of observations.

In the present work, we focus on the functional regression model in which a real random variable Y has to be predicted from functional predictors: see for instance Cardot et al. (1999) in a linear framework and Ferraty and Vieu (2004) in a nonparametric framework. Recently, an increasing number of works have focused on variable selection in this framework, in particular in order to determine which parts of the predictors (curves) are relevant to explain Y . Apart from facilitating prediction and solving numerical issues, a desirable feature of variable selection is to enhance the interpretability of the relation between X and Y , which would allow one to reduce the observation range of the predictors to a few influential intervals, or focus on some particular aspects of the curves in order to obtain expected values for Y .

Most of the works that have focused on variable selection in the functional regression problem are related to linear regression: Ferraty et al. (2010) and Aneiros and Vieu (2014) rely on sparse approaches (*e.g.*, LASSO) that select relevant but isolated evaluation points in the predictor curves. However, these approaches have an important drawback in the functional setting: in most practical situations, interesting predictors do not correspond to isolated evaluation points but to sub-intervals of the predictor input range. This is especially true when the predictors are subject to small shifts on the x-axis, due to imprecise measurements. Another frequent approach is to rely on the expansion of the predictors on functional bases: Matsui and Konishi (2011) uses this method, combined with a L^1 regularization in a linear model. James et al. (2009) develop a method that is based on a sparsity constraint on the derivatives of the estimated functional coefficients in a linear models: this leads to obtain functional coefficients that are constant over sub-interval of the predictor input range. A similar method is used in wavelet decompositions of signals to obtain sparse representations in Chen et al. (2015). However, again, these approaches are not easily interpretable because the interpretation must relies on the analysis of the basis components themselves, which must be chosen with care. Fraiman et al. (2016) develop a different point of view and propose a general method which can be applied to various problems (regression, classification, PCA...). This method can select groups of variables (typically intervals) that explain the most a given model. Similarly, for random forests, Gregorutti et al. (2015) propose an adaptation of the importance of the variables which can be used to find groups of variables that are important in a regression model. However, in Fraiman et al. (2016), no specific contiguity constrain is put on groups of variables, which we think is desirable to obtain more relevant groups and Gregorutti et al. (2015) do not provide a method to automatically design relevant groups, which must be given *a priori* to the method.

In the present work, we propose a semi-parametric model that is not purely linear and is interpretable in the functional framework. The method selects intervals in the range of X with an automatic approach. It is based on Sliced Inverse Regression (SIR, Li (1991)), even though it could easily be extended to linear regression. Our choice for SIR is motivated by the fact that the method has already been proven efficient to project the predictors in a small dimensional space which is optimal with respect to the prediction purpose. It has also been extended to the functional regression framework by Ferré and Yao (2003); Ferré and Villa (2006). Since it is based on a prior linear dimension reduction, it can be conveniently penalized by L_1 -type penalty: we use a group-LASSO-like penalty to select the most relevant intervals in the definition domain of the functional predictors. Our approach combines ideas from Li and Nachtsheim (2008); Li and Yin (2008) who introduced sparsity in multivariate SIR in two different ways. The definition of the intervals themselves is based on an iterative procedure that uses the full regularization path of the LASSO. Our contribution is thus twofolds: first, we extend the SIR method to perform variable selection in a way that is efficient and adapted to the interval selection in the functional framework. Second, we provide a fast and automatic procedure to obtain relevant intervals in the definition domain of the predictors without using any prior knowledge.

The paper is organized as follows: Section 2 presents the SIR approach in a multidimensional framework and its adaptations to the high-dimensional and functional frameworks, which are based on regularization and/or sparsity constraints. Section 3 describes our proposal when the range of the predictors are partitioned using a fixed set of intervals. Then, Section 4 describes an automatic procedure to find these intervals and Section 5 provides practical methods to tune the different parameters in a high dimensional framework. Finally, Section 6 evaluates our approach on simulated and real-world datasets.

2 A review on SIR and regularized versions

In this section, we review the standard SIR for multivariate data and its extensions to the high-dimensional setting. Here, (X, Y) denotes a random pair of variables such that X takes values in \mathbb{R}^p and Y is real. We assume given n i.i.d. realizations of (X, Y) , $(\mathbf{x}_i, y_i)_{i=1, \dots, n}$.

2.1 The standard multidimensional case

In standard regression problems when the dimension p of the predictor is large, classical modeling approaches suffer from the curse of dimensionality. This problem might occur even if p is smaller than n . A standard way to overcome this issue is to rely on dimension reduction techniques. This kind of approaches is based on the assumption that there exists a *central space* $\mathcal{S}_{Y|X}$ which is the smallest subspace such that the projection of X on $\mathcal{S}_{Y|X}$ retains all the information on Y contained in the predictor X . More precisely, $\mathcal{S}_{Y|X}$ is assumed of the form $\text{Span}\{\mathbf{a}_1, \dots, \mathbf{a}_d\}$, with $d \ll p$, such that

$$Y = F(\mathbf{a}_1^\top X, \dots, \mathbf{a}_d^\top X, \epsilon), \quad (1)$$

in which $F : \mathbb{R}^{p+1} \rightarrow \mathbb{R}$ is an unknown function and ϵ is an error term independent of X . To estimate this subspace, SIR is one of the most classical approaches when $p < n$: under an appropriate and general enough condition, Li (1991) shows that $\mathbf{a}_1, \dots, \mathbf{a}_d$ can be estimated as the first d Σ -orthonormal eigenvectors of the generalized eigenvalue problem: $\Gamma \mathbf{a} = \lambda \Sigma \mathbf{a}$, in which Σ is the covariance matrix of X and Γ is the covariance matrix of $\mathbb{E}(X|Y)$.

In practice, Σ is replaced by the empirical covariance, $\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i - \bar{X})(\mathbf{x}_i - \bar{X})^T$, and Γ is estimated by “slicing” the observations $(y_i)_i$ as follows. The range of Y is partitioned into H consecutive and non-overlapping slices, denoted hereafter $\mathcal{S}_1, \dots, \mathcal{S}_H$.

An estimate of $\mathbb{E}(X|Y)$ is thus simply obtained by $(\bar{X}_1, \dots, \bar{X}_H)$ in which \bar{X}_h is the average of the observations \mathbf{x}_i in \mathcal{S}_h and \bar{X}_h is associated with the empirical frequency $p_h = \frac{n_h}{n}$ with n_h the number of observations in \mathcal{S}_h . $\hat{\Gamma}$ is thus defined as $\frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i - \bar{X}) (\mathbf{x}_i - \bar{X})^\top$ for \bar{X} the empirical mean of X .

Different equivalent formulations of the SIR have been proposed later, on which regularized and sparse versions of the method are based. Here we briefly describe the regression formulation and the correlation formulation that are used in our method. First, note that the EDR space is often summarized by the $p \times d$ matrix $A = (\mathbf{a}_1, \dots, \mathbf{a}_d)$. Cook (2004); Li and Yin (2008); Bernard-Michel et al. (2008) show that the SIR can be viewed as a regression problem solved by a mean square error minimization: estimates of the EDR space can be obtained by minimizing over $A \in \mathcal{M}_{p \times d}$ and $C = (C_1, \dots, C_H)$, with $C_h \in \mathbb{R}^d$ (for $h = 1, \dots, H$),

$$\mathcal{E}_1(A, C) = \sum_{h=1}^H \hat{p}_h \|(\bar{X}_h - \bar{X}) - \hat{\Sigma} A C_h\|_{\hat{\Sigma}^{-1}}^2, \quad (2)$$

in which $\|\cdot\|_{\hat{\Sigma}^{-1}}^2$ is the norm $\forall u \in \mathbb{R}^p$, $\|u\|_{\hat{\Sigma}^{-1}} = u^\top \hat{\Sigma}^{-1} u$.

An alternative formulation is described in Chen and Li (1998), which rewrites SIR as the following optimization problem:

$$\max_{\mathbf{a}_j, \phi} \text{Cor}(\phi(Y), \mathbf{a}_j^\top X), \quad (3)$$

where ϕ is any function $\mathbb{R} \rightarrow \mathbb{R}$ and $(\mathbf{a}_j)_j$ are Σ -orthonormal. This shows that SIR can be interpreted as a canonical correlation problem. The authors also prove that the solution of ϕ optimizing Equation (3) for a given \mathbf{a}_j is $\phi(y) = \mathbf{a}_j^\top \mathbb{E}(X|Y = y)$, and that \mathbf{a}_j is also obtained as the solution of the mean square error optimization $\min_{\mathbf{a}_j} \mathbb{E}(\phi(Y) - \mathbf{a}_j^\top X)^2$.

However, as explained in Li and Yin (2008); Coudret et al. (2014) among others, in a high dimensional setting ($n < p$), $\hat{\Sigma}$ is singular and the SIR problem is thus ill-posed. The same problem occurs in the functional setting (Dauxois et al., 2001). Solutions to

overcome this difficulty include variable selection (Coudret et al., 2014), ridge regularization or sparsity constraints.

2.2 Regularization approaches for SIR

To avoid the ill-posed problem in the high-dimensional setting, Bernard-Michel et al. (2008) propose an alternative ridge SIR estimator, which minimizes:

$$\mathcal{E}_{r,1}(A, C) = \sum_{h=1}^H \hat{p}_h \left\| (\bar{X}_h - \bar{X}) - \hat{\Sigma} A C_h \right\|_{\hat{\Sigma}^{-1}}^2 + \mu_2 \sum_{h=1}^H \hat{p}_h \|A C_h\|^2. \quad (4)$$

Written as such, the minimization of Equation (4) seems to require the existence of $\hat{\Sigma}^{-1}$, which does not exist in the high dimensional setting case since $\hat{\Sigma}$ is singular. However, Bernard-Michel et al. (2008) show that the estimator of A obtained by minimizing Equation (4) is composed of the first d eigenvectors of $\left(\hat{\Sigma} + \mu_2 \mathbb{I}_p\right)^{-1} \hat{\Gamma}$.

2.3 Sparse SIR

We call sparse estimates of \mathbf{a}_j , vectors of \mathbb{R}^p with only a limited number of non zero entries. Such estimates usually increase the interpretability of the model (here, of the EDR space) by focusing on the most important predictors only. Different approaches can be used to obtain such results. Coudret et al. (2014) propose a variable selection approach based on a re-sampling of the variables and on the fitting of different sub-models. Direct sparse penalty (in the line of the LASSO) can also be used. Up to our knowledge, only two alternatives have been introduced to use such methods, one based on the regression formulation and the other on the correlation formulation of SIR.

Li and Yin (2008) derive a sparse ridge estimator from the work of Ni et al. (2005). More precisely, given (\hat{A}, \hat{C}) the ridge SIR estimates, a shrinkage index vector $\boldsymbol{\alpha} =$

$(\alpha_1, \dots, \alpha_p)^T \in \mathbb{R}^p$ is obtained by minimizing a least square error with L_1 penalty: $\mathcal{E}_{s,1}(\boldsymbol{\alpha}) = \sum_{h=1}^H \hat{p}_h \left\| (\bar{X}_h - \bar{X}) - \widehat{\Sigma} \text{Diag}(\boldsymbol{\alpha}) \hat{A} \hat{C}_h \right\|^2 + \mu_1 \|\boldsymbol{\alpha}\|_{L_1}$, for a given $\mu_1 \in \mathbb{R}^{+*}$. Once the coefficients $\boldsymbol{\alpha}$ have been estimated, the central space is obtained as $\text{Span} \left(\text{Diag}(\hat{\boldsymbol{\alpha}} \hat{A}) \right)$, where $\hat{\boldsymbol{\alpha}}$ is the solution of the minimization of $\mathcal{E}_{s,1}(\boldsymbol{\alpha})$. In this particular setting, the variables selected for all the vectors generating the EDR space are the same and correspond to the non-zero entries of $\boldsymbol{\alpha}$.

An alternative, which does not have this property, is described in Li and Nachtsheim (2008). They use the correlation formulation of the SIR and obtain a sparse estimation of the vectors generating the EDR space in the case where $p < n$: after the standard SIR estimates $\hat{\mathbf{a}}_1, \dots, \hat{\mathbf{a}}_d$ have been computed, they solve d independent minimization problems with sparsity constraints introduced as an L_1 penalty: $\forall j = 1, \dots, d$,

$$\mathcal{E}_{s,2}(\mathbf{a}_j^s) = \sum_{i=1}^n [\mathcal{P}_{\hat{\mathbf{a}}_j}(X|y_i) - (\mathbf{a}_j^s)^T \mathbf{x}_i]^2 + \mu_{1,j} \|\mathbf{a}_j^s\|_{L_1}, \quad (5)$$

in which $\mathcal{P}_{\hat{\mathbf{a}}_j}(X|y_i) = \widehat{\mathbb{E}}(X|Y = y_i)^T \hat{\mathbf{a}}^j$, with $\widehat{\mathbb{E}}(X|Y = y_i) = \bar{X}_h$ for h such that $y_i \in \mathcal{S}_h$ in the case of a sliced estimate of $\widehat{\mathbb{E}}(X|Y)$. In this solution, the sparsity constraints are estimated from the standard solution but independently from each other. Hence, different variables can be selected depending on the direction $j = 1, \dots, d$. The restriction $p < n$ is not a real issue since the solution could be adapted to the case $n < p$ by using for $\hat{\mathbf{a}}_1, \dots, \hat{\mathbf{a}}_d$ the ridge estimates obtained from Equation (4) instead of the standard SIR solution.

The differences between these two approaches can be summarized in two points:

- by using shrinkage coefficients, Li and Yin (2008) constrain the variables selected by the sparse optimization problem to be the same for all dimensions. This makes sense because the vectors \mathbf{a}_j themselves are not relevant: only the space spanned by them is and so there is no need to select different variables for the different estimated

directions. Moreover, this allows to formulate the optimization in a single problem. However, this problem relies on a least square minimization with dependent variables in a high dimensional space \mathbb{R}^p ;

- on the contrary, the approach of Chen and Li (1998) relies on a least square problem based on projections and is thus obtained from d independent optimization problems. The dimension of the dependent variable is reduced (1 instead of p) but the different vectors which span the EDR space are estimated independently and not simultaneously.

In our proposal, we combine both advantages using a single optimization problem based on the correlation formulation of SIR. In this problem, the dimension of the dependent variable is reduced (d instead of p) when compared to the approach of Li and Yin (2008) and it is thus computationally more efficient. Identical sparsity constraints are imposed on all dimensions using a shrinkage approach, but instead of selecting the non-zero variables independently, we adapt the sparsity constraint to the functional setting to avoid selecting isolated measurement points. The next section describes this approach.

3 Interpretable-Sparse SIR

In this section, a functional regression framework is assumed. X is thus a functional random variable, taking value in a (infinite dimensional) Hilbert space. $(x_i, y_i)_{i=1, \dots, n}$ are n i.i.d. realizations of (X, Y) . However, x_i are not perfectly known but observed on a given (deterministic) grid $\tau = \{t_1, \dots, t_p\}$. We denote by $\mathbf{x}_i = (x_i(t_j))_{j=1, \dots, p} \in \mathbb{R}^p$ the i -th observation and by \mathbf{X} the $n \times p$ matrix $(\mathbf{x}_1, \dots, \mathbf{x}_n)^T$. Unlike said so, the notations are derived from the ones introduced in the multidimensional setting (Section 2) by using the observations \mathbf{x}_i as realizations of X .

In the functional setting, it is realistic to select variables \mathbf{x}^j not one by one but in groups, as explained in Gregorutti et al. (2015). More precisely, we assume that the range of the functional predictors X is partitioned into D consistent non-overlapping intervals, τ_1, \dots, τ_D . In the present section, these intervals are supposed to be given *a priori* and we will describe later (in Section 4) a fully automated procedure to obtain them from the data.

Similarly to Li and Yin (2008) and Li and Nachtsheim (2008), we develop a sparse approach using a L^1 penalty set from the correlation formulation and based on the previously obtained ridge estimates. However, unlike Li and Yin (2008); Li and Nachtsheim (2008), we do not impose sparsity constraints on every evaluation point in τ but propose to use a group-Lasso-like constraint with respect to the intervals. This approach is performed through one multiplicative coefficient per interval. These coefficients are estimated using a L^1 penalty to obtain sparsity and are used to shrink the ridge estimates similarly in a given interval. The following two subsections are devoted to the description of the two steps (ridge and sparse) of the method.

3.1 Ridge estimation

First, $\mathcal{E}_{\tau,1}(A, C)$ in Equation (4) is minimized over (A, C) to obtain \hat{A} and \hat{C} , ridge estimates of the SIR. The solution is computed as follow:

\hat{A} calculation As explained in Section 2.2, the estimator of $A \in \mathcal{M}_{p \times d}$ obtained by minimizing Equation (4) is composed of the first d $(\hat{\Sigma} + \mu_2 \mathbb{I}_p)$ -orthonormal eigenvectors of $(\hat{\Sigma} + \mu_2 \mathbb{I}_p)^{-1} \hat{\Gamma}$ associated with the d largest eigenvalues. In practice, the same procedure as the one described in Ferré and Yao (2003); Ferré and Villa (2006) is used: first, orthonormal eigenvectors (denoted hereafter $(\hat{\mathbf{b}}_j)_{j=1, \dots, d}$) of the matrix $(\Sigma + \mu_2 \mathbb{I}_p)^{-1/2} \Gamma (\hat{\Sigma} + \mu_2 \mathbb{I}_p)^{-1/2}$

are computed. Then, \hat{A} is the matrix which columns are equal to $(\hat{\Sigma} + \mu_2 \mathbb{I}_p)^{-1/2} \hat{\mathbf{b}}_j$ for $j = 1, \dots, d$. It is easy to prove that these columns are $(\hat{\Sigma} + \mu_2 \mathbb{I}_p)$ -orthonormal eigenvectors of $(\hat{\Sigma} + \mu_2 \mathbb{I}_p)^{-1} \hat{\Gamma}$.

\hat{C} calculation For a given A , the optimal $\hat{C} = (\hat{C}_1, \dots, \hat{C}_H) \in \mathcal{M}_{d,H}$ is given by $\nabla_{C_h} \mathcal{E}_{r,1}(A, \hat{C}_h) = 0$, which is equivalent to $[A^T \hat{\Sigma} A + \mu_2 A^T A] C_h = A^T (\bar{X}_h - \bar{X})$ that gives $\hat{C}_h = [A^T \hat{\Sigma} A + \mu_2 A^T A]^{-1} A^T (\bar{X}_h - \bar{X}) = A^T (\bar{X}_h - \bar{X})$ because the columns of A are $(\hat{\Sigma} + \mu_2 \mathbb{I}_d)$ -orthonormal.

3.2 Interval-sparse estimation

Once \hat{A} and \hat{C} have been computed, the estimated projections of $(\hat{\mathbb{E}}(X|Y = y_i))_{i=1, \dots, n}$ onto the EDR space are obtained by: $\mathcal{P}_{\hat{A}}(\hat{\mathbb{E}}(X|Y = y_i)) = (\bar{X}_h - \bar{X})^T \hat{A}$, for h such that $y_i \in \mathcal{S}_h$. In the sequel, we will also denote by \mathbf{P}^j (for $j = 1, \dots, d$) the observations of the j -th entries for all observations: $\mathbf{P}^j = (\mathcal{P}_1^j, \dots, \mathcal{P}_n^j)^T \in \mathbb{R}^n$.

In the line of Li and Nachtsheim (2008), we rely on the Multiple Linear Regression (MLR) reformulation of SIR given by Equation (5) to make the estimates \hat{A} sparse. However, as the method requires to solve d independent sparse problems, we suggest to follow the idea of Li and Yin (2008), which introduce sparsity by means of shrinkage coefficients. This approach forces identical sparsity constraints on some intervals τ_k for all dimensions of the EDR space. More precisely, the standard way to shrink $\hat{\mathbf{a}}_j$ is to multiply them by shrinkage coefficients $\boldsymbol{\alpha} \in \mathbb{R}^p$, which are all the same for the p vectors spanning the EDR space and on which a sparsity constraint is imposed by a L_1 penalty. This provides the sparse estimates $\hat{\mathbf{a}}_j^s = \boldsymbol{\alpha} \hat{\mathbf{a}}_j$, which have non zero entries for the same variables than $\boldsymbol{\alpha}$.

Here, we propose to shrink identically all the variables in a same interval τ_k . In this

situation, only D coefficients are needed, $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_D) \in \mathbb{R}^D$. Their estimated values, $\hat{\boldsymbol{\alpha}}$, are used to define the $\hat{\mathbf{a}}_j^s$ of the vectors spanning the EDR space, by:

$$\forall l = 1, \dots, p, \quad \hat{a}_{jl}^s = \hat{\alpha}_k \hat{a}_{jl} \quad \text{for } k \text{ such that } t_j \in \tau_k. \quad (6)$$

To do so, the formulation of Equation (5) is adapted so that i) $\hat{\mathbf{a}}_j^s$ are replaced by their values depending of $\hat{\mathbf{a}}_j$ and $\boldsymbol{\alpha}$, ii) the sparsity constraint is put on $\boldsymbol{\alpha}$ and not directly on $\hat{\mathbf{a}}_j^s$ and iii) $\boldsymbol{\alpha}$ are made identical for all dimensions of the projection $j = 1, \dots, d$ by using a sum over $j = 1, \dots, d$ to perform a trade-off between all dimensions. If $\Lambda(\boldsymbol{\alpha}) = \text{Diag}(\alpha_1 \mathbb{I}_{|\tau_1|}, \dots, \alpha_D \mathbb{I}_{|\tau_D|}) \in \mathcal{M}_{p \times p}$, this leads to solve the following LASSO problem

$$\begin{aligned} \hat{\boldsymbol{\alpha}} &= \arg \min_{\boldsymbol{\alpha} \in \mathbb{R}^D} \sum_{j=1}^d \sum_{i=1}^n \|\mathcal{P}_i^j - (\Lambda(\boldsymbol{\alpha}) \hat{\mathbf{a}}_j)^\top \mathbf{x}_i\|^2 + \mu_1 \|\boldsymbol{\alpha}\|_{L_1} \\ &= \arg \min_{\boldsymbol{\alpha} \in \mathbb{R}^D} \sum_{j=1}^d \|\mathbf{P}^j - (\mathbf{X}\Delta(\hat{\mathbf{a}}_j)) \boldsymbol{\alpha}\|^2 + \mu_1 \|\boldsymbol{\alpha}\|_{L_1}, \end{aligned}$$

with $\Delta(\hat{\mathbf{a}}_j)$ the $(p \times D)$ -matrix such that $\Delta_{kl}(\hat{\mathbf{a}}_j) = \hat{a}_{jl}$ if $t_l \in \tau_k$ and 0 otherwise. This problem can be re-written as:

$$\arg \min_{\boldsymbol{\alpha} \in \mathbb{R}^D} \|\mathbf{P} - \Delta(\mathbf{X}\hat{\mathbf{A}}) \boldsymbol{\alpha}\|^2 + \mu_1 \|\boldsymbol{\alpha}\|_{L_1} \quad (7)$$

with $\mathbf{P} = \begin{pmatrix} \mathbf{P}^1 \\ \vdots \\ \mathbf{P}^d \end{pmatrix}$, a vector of size dn and $\Delta(\mathbf{X}\hat{\mathbf{A}}) = \begin{pmatrix} \mathbf{X}\Delta(\hat{\mathbf{a}}_1) \\ \vdots \\ \mathbf{X}\Delta(\hat{\mathbf{a}}_p) \end{pmatrix}$, a $(dn) \times D$ -matrix.

We finally set $\hat{\mathbf{a}}_j^s = \Lambda(\hat{\boldsymbol{\alpha}}) \hat{\mathbf{a}}_j$, as given in Equation (6).

Once the sparse vectors $(\hat{\mathbf{a}}_j^s)_{j=1, \dots, d}$ have been obtained, an Hilbert-Schmidt orthonormalization approach is used to make them $\hat{\Sigma}$ -orthonormal.

4 An iterative procedure to select the intervals

The previous section described our proposal to detect the subset of relevant intervals among a *fixed*, predefined set of intervals of the predictor input range, $(\tau_k)_{k=1,\dots,D}$. However, choosing *a priori* a proper set of intervals is a challenging task without expert knowledge, and a poor choice (too small, too large, or shifted intervals) may largely hinder interpretability. In the present section, we propose an iterative method to automatically design the intervals, without making any *a priori* choice.

In a closely related framework, Fruth et al. (2015) tackled the problem of designing intervals by combining sensitivity indices, linear regression models and a method called *sequential bifurcation* (Bettonvil, 1995), which allows them to sequentially split in two the most promising intervals (starting from a unique interval covering the entire range of X). Here, we propose the inverse approach: we start with small intervals and merge them sequentially. Our approach is based on the standard sparse SIR (which is used as a starting point) and iteratively performs the most relevant merges in a flexible way (contrary to a splitting approach, we do not need to arbitrary set the splitting positions).

The intervals $(\tau_k)_{k=1,\dots,D}$ are first initialized to a very fine grid, taking for instance $\tau_k = \{t_k\}$ for all $k = 1, \dots, p$ (hence, at the beginning of the procedure, $D = p$). The sparse step described in Section 3.2 is then performed with the *a priori* intervals $(\tau_k)_{k=1,\dots,D}$: the set of solutions of Equation (7), for varying values of the regularization parameter μ_1 , is obtained using a regularization path approach, as described in Friedman et al. (2010). Three elements are retrieved from the path results:

- $(\hat{\boldsymbol{\alpha}}_k^*)_{k=1,\dots,D}$ are the solutions of the sparse problem for the value μ_1^* of μ_1 that minimizes the GCV error;
- $(\hat{\boldsymbol{\alpha}}_k^+)_{k=1,\dots,D}$ and $(\hat{\boldsymbol{\alpha}}_k^-)_{k=1,\dots,D}$ are the first solutions, among the path of solutions, such

that at most (resp. at least) a proportion P of the coefficients are non zero coefficients (resp. are zero coefficients), for a given chosen P , which should be small (0.05 for instance).

Then, the following sets are defined: $\mathcal{D}_1 = \{k : \hat{\alpha}_k^- \neq 0\}$ (called “strong non zeros”) and $\mathcal{D}_2 = \{k : \hat{\alpha}_k^+ = 0\}$ (called “strong zeros”). This step is illustrated in Figure 1.




Figure 1: Example of regularization path with $D = 20$: $(\hat{\alpha}_k)_{k=1,\dots,D}$ are plotted according to different values of the tuning parameter μ_1 . The vertical dotted line represents the optimal value μ_1^* that provides the solutions $(\hat{\alpha}_k^*)_{k=1,\dots,D}$ of the sparse problem. $(\hat{\alpha}_k)_{k \in \mathcal{D}_1}$ and $(\hat{\alpha}_k)_{k \in \mathcal{D}_2}$ are respectively represented in bold and in pointed lines for $P = 0.1$.

Intervals are merged using the following rules:

- “neighbor rule”: consecutive intervals of the same set are merged (τ_k and τ_{k+1} are merged if both k and $k+1$ belong to \mathcal{D}_1 or if they both belong to \mathcal{D}_2) (see a) and b) in Figure 2);
- “squeeze rule”: τ_k, τ_{k+1} and τ_{k+2} are merged if both k and $k+2$ belong to \mathcal{D}_1 while $k+1 \notin \mathcal{D}_2$ (or if both k and $k+2$ belong to \mathcal{D}_2 while $k+1 \notin \mathcal{D}_1$) and $l_k + l_{k+2} > l_{k+1}$ with $l_k = \max \tau_k - \min \tau_k$ (see c) and d) in Figure 2).

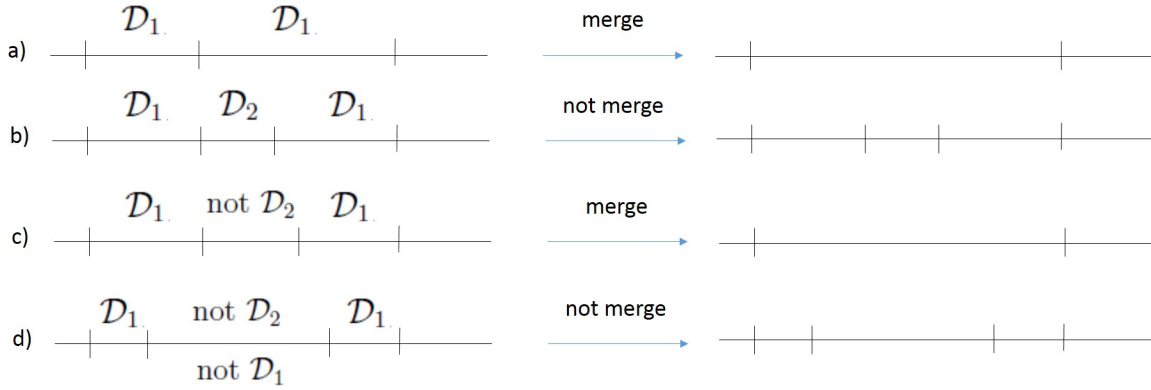


Figure 2: Illustration of the merge procedure for the intervals.

If the current value of P does not yield any fusion between intervals, P is updated by $P \leftarrow P + P_0$ in which P_0 is the initial value of P . The procedure is iterated until all the original intervals have been merged.

The result of the method is a collection of models $(\hat{\alpha}_k^*)_{k=1, \dots, D}$, starting with p intervals and finishing with one. The final selected model is the one that minimizes the CV error. In practice, this often results in a very small number of contiguous intervals which are of the same type (zero or non zero) and are easily interpretable (see Section 6).

Let us remark that the intervals $(\tau_k)_{k=1,\dots,D}$ are not used in the ridge step of Section 3.1, which can thus be performed once, independently of the interval search. The whole procedure is described in Algorithm 1.

Algorithm 1 Overview of the complete procedure

- 1: **Ridge estimation**
 - 2: Choose μ_2 and d according to Section 5
 - 3: Minimize $\mathcal{E}_{r,1}(A, C)$ in Equation (4) to obtain \hat{A} and \hat{C} , ridge estimates of the SIR. See details in Section 3.1
 - 4: **Sparse estimation**
 - 5: Initialize the intervals $(\tau_k)_{k=1,\dots,D}$ to $\tau_k = \{t_k\}$
 - 6: **repeat**
 - 7: Estimate and store $(\hat{\alpha}_k^*)_{k=1,\dots,D}$ the solutions of the sparse problem that minimizes the GCV error
 - 8: Estimate $(\hat{\alpha}_k^+)_{k=1,\dots,D}$ and $(\hat{\alpha}_k^-)_{k=1,\dots,D}$ such that at most (resp. at least) a proportion P of the coefficients are non zero coefficients (resp. are zero coefficients), for a given chosen P (details in Section 4)
 - 9: Update the intervals $(\tau_k)_{k=1,\dots,D}$ according to the "neighbor" and the "squeeze" rules (see Section 4)
 - 10: **until** $\tau_1 \neq [t_1, t_p]$
 - 11: Output : A collection of models $(\hat{\alpha}_k^*)_{k=1,\dots,D}$
 - 12: Select the model $(\hat{\alpha}_k^*)_{k=1,\dots,D}^*$ that minimizes the CV error
 - 13: Active intervals (for interpretation) are consecutive τ_k with non zero coefficients $\hat{\alpha}_k^*$
-

5 Choice of parameters in the high dimensional setting

The method requires to tune four parameters : the number of slices H , the dimension of the central space p , the penalization parameter of the ridge regression μ_2 and of the one of the sparse procedure μ_1 . Two of these parameters, H and μ_1 , are chosen in a standard way (see Section 6 for further details). This section presents a method to jointly choose μ_2 and d , for which no solution has been proposed that is suited to our framework. Two issues are raised to tune these two parameters: i) they depend from each other and ii) the existing methods to tune them are only valid in a low dimensional setting ($p < n$). We propose an iterative method which adapts existing approaches only valid for the low dimension framework and combine them to find an optimal joint choice for μ_2 and d .

5.1 A Cross-Validation (CV) criterion for μ_2

Using the results of Golub et al. (1979), Li and Yin (2008) propose a Generalized Cross-Validation (GCV) criterion to select the regularization parameter μ_2 and Bernard-Michel et al. (2008) explain that this criterion can be applied to their modified estimator, using similar calculations. However, it requires the computation of $\widehat{\Sigma}^{-1/2}$, which is not valid in the high dimensional setting.

We thus used a different strategy, based on L -fold cross-validation (CV), which is also used to select the best dimension of the EDR space, d (see Section 5.2). More precisely, the data are split into L folds, $\mathcal{L}_1, \dots, \mathcal{L}_L$ and a CV error is computed for several values of μ_2 in a given search grid and for a given (large enough d_0). The optimal μ_2 is chosen as the one minimizing the CV error for d_0 .

The CV error is computed based on the original regression problem $\mathcal{E}_1(A, C)$. In the

expression of $\mathcal{E}_1(A, C)$ and for the iteration number $l (\in \{1, \dots, L\})$, A and C_h are replaced by their estimates computed without the observations in fold number l . Then, an error is computed by replacing the values of \hat{p}_h , \bar{X}_h , \bar{X} and $\hat{\Sigma}$ by their empirical estimators for the observations in fold l . The precise expression is given in step 5 of Algorithm 1 (Section 2) in supplementary material.

5.2 Choosing d in a high dimensional setting

The results of CV (*i.e.*, the values of $\mathcal{E}_1(A, C)$ estimated by L -fold CV) are not directly usable for tuning d . The reason is similar to the one developed in Biau et al. (2005); Fromont and Tuleau (2006): different d correspond to different MLR problems which cannot be compared directly using a CV error. In such cases, an additional penalty depending on d is necessary to perform a relevant selection and avoid overfitting due to large d .

Alternatively, a number of works have been dealing with the choice of d in SIR. Many of them are asymptotic methods (Li, 1991; Schott, 1994; Bura and Cook, 2001; Cook and Yin, 2001; Lique and Saracco, 2012) which are not directly applicable in the high dimensional framework. When $n < p$, Zhu et al. (2006); Li and Yin (2008) estimate d using the number of nonzero eigenvalues of Γ , but their approach requires to set a hyper-parameter to which the choice of d is sensitive. Another point of view can be taken from Li (1991) who introduces a quantity, denoted by $R^2(d)$, which is the average of the squared canonical correlation between the space spanned by the columns of $\Sigma^{1/2}A$ and the columns of the space spanned by the columns of $\hat{\Sigma}^{1/2}\hat{A}$. As explained in Ferré (1998), a relevant measure of quality for the choice of a dimension d is $R(d) = d - \mathbb{E} \left[\text{Tr} \left(\Pi_d \hat{\Pi}_d \right) \right]$, in which Π_d is the Σ -orthogonal projector onto the subspace spanned by the columns of A and $\hat{\Pi}_d$ is the $\hat{\Sigma}$ -orthogonal projector onto the space spanned by the columns of \hat{A} . This quantity is equal to $\frac{1}{2} \mathbb{E} \left\| \Pi_d - \hat{\Pi}_d \right\|_F^2$ (in which $\|\cdot\|_F$ is the Frobenius norm; see the proof in Section 1

of supplementary material). In practice, it is often a strictly increasing function of d and we choose the optimal dimension as the largest one before a gap in this increase (“elbow rule”).

However, in the high dimensional setting, the $\widehat{\Sigma}$ -orthogonal projector onto the space spanned by the columns of \widehat{A} is not well defined since the matrix $\widehat{\Sigma}$ is ill-conditioned. We used a similar approach in which this estimate is replaced by its regularized version using the $(\widehat{\Sigma} + \mu_2 \mathbb{I}_p)$ -orthogonal projector onto the space spanned by the columns of \widehat{A} . Finally, $\mathbb{E} \left[\text{Tr} \left(\Pi_d \widehat{\Pi}_d \right) \right]$ is estimated during the same L -fold loop described in Section 5.1: $\widehat{\Pi}_d$ is the $(\widehat{\Sigma} + \mu_2 \mathbb{I}_p)$ -orthogonal projector onto the space spanned by the columns of \widehat{A} . Similarly, for all $l = 1, \dots, L$, we computed the $(\widehat{\Sigma}^{\setminus l} + \mu_2 \mathbb{I}_p)$ -orthogonal projector onto the space spanned by the columns of $\widehat{A}(l)$ in which $\widehat{\Sigma}^{\setminus l}$ and $\widehat{A}(l)$ are computed without the observations in fold number l and averaged the results to obtain an estimate of $\mathbb{E} \left[\text{Tr} \left(\Pi_d \widehat{\Pi}_d \right) \right]$.

5.3 Joint tuning

The estimation of μ_2 and d is jointly performed using a single CV pass in which both parameters are varied. Note that only the number of different values for μ_2 strongly influences the computational time since SIR estimation is only performed once for all values of d , and selecting the first d columns of \widehat{A} for the last computation of the two criteria, the estimation of $\mathcal{E}(A, C)$ and that of $R(d)$. The overall method is described in Section 2 of supplementary material.

6 Experiments

6.1 Simulated data

6.1.1 Model description

To illustrate our approach, we first consider two toy datasets, built as follow. Both are based on the following definition: $Y = \sum_{j=1}^d \log |\langle X, \mathbf{a}_j \rangle|$ with $X(t) = Z(t) + \epsilon$ in which Z is a Gaussian process indexed on $[0, 1]$ with mean $\mu(t) = -5 + 4t - 4t^2$ and the Matern 3/2 covariance function, and ϵ is a centered Gaussian variable independent of Z . The vectors \mathbf{a}_j have a sinusoidal shape, but are nonzero only on specific intervals I_j : $\mathbf{a}_j = \sin\left(\frac{t(2+j)\pi}{2} - \frac{(j-1)\pi}{3}\right) \mathbb{I}_{I_j}(t)$.

From this basis, we consider two models with increasing complexity:

- **(M1):** $d = 1$, $I_1 = [0.2, 0.4]$
- **(M1):** $d = 3$ **and** $I_1 = [0, 0.1]$, $I_2 = [0.5, 0.65]$ **and** $I_3 = [0.65, 0.78]$.

On both cases the datasets consist of $n = 100$ observations of Y , discretized on $p = 200$ and 300 evaluation points, respectively. The number of slices used to estimate the conditional mean $\mathbb{E}(X|Y)$ has been chosen equal to $H = 10$: according to Li (1991); Coudret et al. (2014) among others, the performances of SIR estimates are not sensitive to the choice of H , as long as it is larger than $d + 1$.

The datasets are displayed in Figure 3, with *a priori* intervals provided to test the sparse penalty (see Section 6.1.3 for further details).




Figure 3: Summary of the two simulated datasets: top (**M1**), bottom (**M2**). The left charts display ten samples of X , the colors showing the actual relevant intervals; the middle charts display the functions that span the EDR space with the relevant slices highlighted in color; the right charts display the distribution of Y .

6.1.2 Step 1: Ridge estimation and parameter selection

The method described in Section 3.1 with parameter selection as in Section 5 has been used to obtain the ridge estimates of (\mathbf{a}_j) and to select the parameters μ_2 (in $\mathcal{E}_{r,1}(A, C)$) and d (dimension of the EDR space). Figure 4 shows the evolution of the CV error and of the estimation of $\mathbb{E}(R(d))$ versus (respectively) μ_2 and d among a grid search both for $\mu_2 \in \{10^{-2}, 10^{-1}, \dots, 10^5\}$ and $d \in \{1, 2, \dots, 10\}$. The chosen values for d , given by the “elbow rule” are $d = 1$ for **(M1)** and $d = 2$ for **(M2)** (the true values begin, respectively, $d = 1$ and $d = 3$).





6.1.3 Step 2: Sparse selection and definition of relevant intervals

The approach described in Section 4 is then applied to both models. The value of μ_1 is given by the GCV error as described in Section 4. The algorithm produces a large collection of models with a decreasing number of intervals: a selection of the estimates of \mathbf{a}_1 for **(M1)**, corresponding to those models is shown in Figure 5. Note that the first chart (Figure 5,a) corresponds to the standard sparse penalty in which the constraint is put on isolated evaluation points. Even though most of selected points are found in the relevant interval, the estimated parameter $\hat{\mathbf{a}}_1^s$ has an uneven aspect which does not favor interpretation. On the other side, for a low number of intervals (less than 50), the selected relevant points (those corresponding to non zero coefficients) have a much larger range than the original relevant interval (in red on the figure).

The model selected by minimization of the cross-validation error was found relevant: this approach lead us to choose the model with 147 intervals (plot (b) in Figure 5), which actually correspond to two distinct and consecutive intervals (a first one, which contains only nonzero coefficients and a second one in which no point is selected by the sparse

estimation). This final estimation is very close to the actual direction \mathbf{a}_1 , both in terms of shape and support.


The same method is used for **(M2)**. Given that only the intervals can be interpreted in a relevant way and not the global shape of the estimated directions, we provide a comparison between the true relevant intervals and the estimated ones in Figure 6 (left). The support of each of the two estimates $(\hat{\mathbf{a}}_i)_{i=1,\dots,2}$ is properly retrieved by our approach: in particular, the largest interval in gray is well identified. The other one might suffer from a border effect because the merging procedure has not worked. A very few number of isolated points outside the relevant intervals are also selected. However, compared to the standard sparse approach (right part of Figure 6), the approach is much more efficient to select the relevant intervals.




(M1)

(M2)


Figure 4: Top: CV error versus μ_2 (\log_{10} scale, for $d = 1$ and 2 , respectively) and Bottom: estimation of $\mathbb{E}(R(d))$ versus d (for $\mu_2 = 1$ in both cases), for models (M1) (left) and (M2) (right).




(a) 200 intervals



(b) 147 intervals




(c) 43 intervals




(d) 5 intervals

Figure 5: **(M1)** Values of $\hat{\mathbf{a}}_1^s$ corresponding to four models obtained using the iterative procedure with a different numbers of intervals. (b) is the chosen model and (a) corresponds to a standard sparse estimation with no constraint on intervals.



IS-SIR



standard sparse

Figure 6: Left: comparison between the true intervals and the estimated ones for model (M2). True intervals are represented in the upper side of the figure (in black) and by the gray background. Estimated intervals are represented by the red lines in the bottom of the figure and by the pink background. Right: same representation for the standard sparse approach (penalty is applied to t_j and not to the intervals).

6.2 Tecator dataset

Additionally, we tested the approach with the well-known Tecator dataset (Borggaard and Thodberg, 1992), which consists of spectrometric data from the food industry. This dataset is a standard benchmark for functional data analysis. It contains 215 observations of near infrared absorbency spectra of a meat sample recorded on a Tecator Infratec Food and Feed Analyzer. Each spectrum was sampled at 100 wavelengths uniformly spaced in the range 850–1050 nm. The composition of each meat sample was determined by analytic chemistry, among which we focus on the percentage of fat content. The data is displayed in Figure 7: the left chart displays the original spectra whereas the right chart displays the first order derivatives (obtained by simple finite differences). The fat content is represented in both graphics by the color level and, as is already well known with this dataset, the derivative is a good predictor of this quantity: these derivatives were thus used as predictors (X) to explain the fat content (Y).

6.2.1 Parameter estimation and relevant intervals

We first applied the method once on the entire dataset to check the relevance of the estimated EDR space and corresponding intervals in the range 850–1050 nm. Using the ridge estimation and the method described in Section 5, we set $\mu_2 = 10^{-4}$ and $d = 2$. We follow the approach described in Section 4 to find the number of intervals only not using the “squeeze rule” which was found not so efficient in this situation, because of the small number of original intervals ($p = 100$). The optimal number of intervals was found equal to 61 and the corresponding estimates $\hat{\mathbf{a}}_1^s$ and $\hat{\mathbf{a}}_2^s$ are given in Figure 8 (left), in which a gray background has been superimposed to highlight the active concatenated intervals (that correspond to only 5 consecutive sets of points). Even though the two last intervals

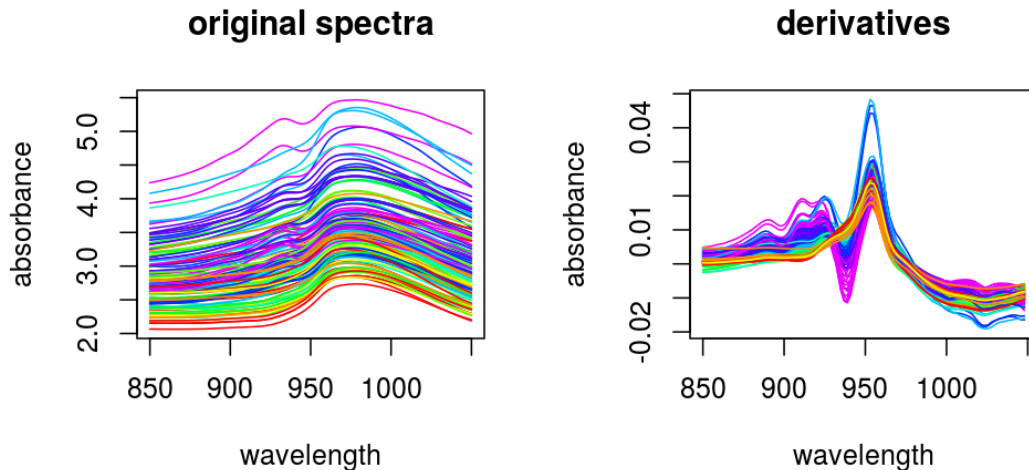


Figure 7: *Tecator*. 215 near infrared spectra from the “Tecator” dataset (left) and corresponding first order derivatives (right). The color level indicates the percentage of fat content.

are not relevant, the method automatically selects two intervals that allows an accurate ordering of Y . In particular, the one located between 934–945 indeed discriminates Y (fat content) almost perfectly, even for the smallest value of the fat content. Compared to what was found with a standard sparse approach (Figure 8, right) which only selects three isolated points, the approach is more relevant in a functional context, especially if small translations occur between discriminant features in the predicting curves.

6.2.2 Application in regression

We finally evaluated the relevance of the approach in a regression setting. Following the simulation setting described in Hernández et al. (2015), we split the data into a training and test sets with 150 observations for the training. This separation of the data was

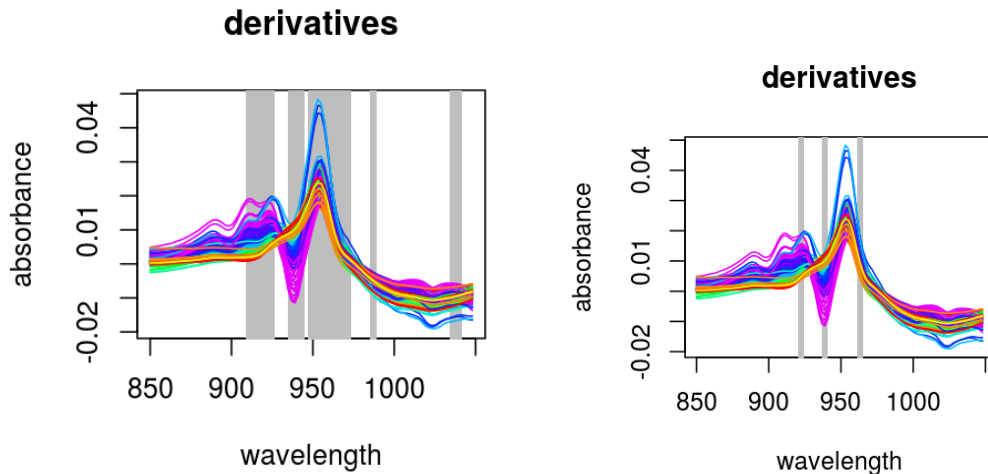


Figure 8: *Tecator*. Left: original predictors (first order derivatives) with a gray background superimposed to highlight the active concatenated intervals found by our procedure. Right: same figure for the standard sparse approach (no constraint on intervals).

performed 100 times randomly. For each training data set, the EDR space is estimated the projection retrieved, and a Support Vector Machine (SVM, package **e1071**, Meyer et al. (2015)) was used to fit the link function F of Equation (1). The mean square error is then computed on the test set. We found an averaged value equal to 10.88, which is smaller than the results obtained by the Nadaraya Watson kernel estimate results reported in Hernández et al. (2015). Even if some methods achieve better performance on this data set (Hernández et al. (2015) reported an average MSE of 2.41 for their non parametric approach), our method has the advantage of being easily interpretable because it extracts a few components which are themselves composed of a small number of relevant intervals.

7 Conclusion

Variable selection is an important issue in high dimensional problems, which is especially relevant in functional regression problems in which predictors are functions often finely digitized. In this framework, selecting intervals in the definition domain of the predictors leads to improved interpretability of the coefficients and to estimates which are more robust to small shifts. We have proposed an approach based on a group-Lasso-like penalty, which is able to select such intervals. The method has been described in a semi-parametric framework, whereas it could be easily extended to a linear regression framework for instance. Moreover, our proposal also incorporate an automatic method to design the intervals themselves. Experiments on simulated and real data show that the selected intervals are indeed relevant and improves the interpretability of the estimated coefficients when compared to standard LASSO. Perspective of developments would extend the approach to multiple functional predictors, allowing to design common or separated interval selections for the different predictors, according to their meaning.

Supplementary material

The supplementary material contains proof for equivalent formulations of $R(d)$ and the complete algorithm to jointly tune μ_2 and d .

References

- Aneiros, G. and Vieu, P. (2014). Variable in infinite-dimensional problems. *Statistics and Probability Letters*, 94:12–20.
- Bernard-Michel, C., Gardes, L., and Girard, S. (2008). A note on sliced inverse regression with regularizations. *Biometrics*, 64(3):982–986.

- Bettonvil, B. (1995). Factor screening by sequential bifurcation. *Communications in Statistics, Simulation and Computation*, 24(1):165–185.
- Biau, G., Bunea, F., and Wegkamp, M. (2005). Functional classification in Hilbert spaces. *IEEE Transactions on Information Theory*, 51:2163–2172.
- Borggaard, C. and Thodberg, H. (1992). Optimal minimal neural interpretation of spectra. *Analytical Chemistry*, 64(5):545–551.
- Bura, A. and Cook, R. (2001). Extending sliced inverse regression: the weighted chi-squared test. *Journal of the American Statistical Association*, 96(455):996–1003.
- Cardot, H., Ferraty, F., and Sarda, P. (1999). Functional linear model. *Statistics and Probability Letters*, 45(1):11–22.
- Chen, C. and Li, K. (1998). Can SIR be as popular as multiple linear regression? *Statistica Sinica*, 8:289–316.
- Chen, S., Donoho, D., and Saunders, M. (2015). Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61.
- Cook, R. (2004). Testing predictor contributions in sufficient dimension reduction. *Annals of Statistics*, 32(3):1061–1092.
- Cook, R. and Yin, X. (2001). Dimension reduction and visualization in discriminant analysis. *Australian & New-Zealand Journal of Statistics*, 43(2):147–199.
- Coudret, R., Liquet, B., and Saracco, J. (2014). Comparison of sliced inverse regression approaches for undetermined cases. *Journal de la Société Française de Statistique*, 155(2):72–96.
- Dauvois, J., Ferré, L., and Yao, A. (2001). Un modèle semi-paramétrique pour variable aléatoire hilbertienne. *Comptes Rendus Mathématique. Académie des Sciences. Paris*, 327(I):947–952.
- Ferraty, F., Hall, P., and Vieu, P. (2010). Most-predictive design points for functional data predictors. *Biometrika*, 97(4):807–824.

- Ferraty, F. and Vieu, P. (2004). Nonparametric models for functional data, with application in regression, time series prediction and curves discrimination. *Journal of Nonparametric Statistics*, 16(1-2):111–125.
- Ferraty, F. and Vieu, P. (2006). *NonParametric Functional Data Analysis*. Springer.
- Ferré, L. (1998). Determining the dimension in sliced inverse regression and related methods. *Journal of the American Statistical Association*, 93(441):132–140.
- Ferré, L. and Villa, N. (2006). Multi-layer perceptron with functional inputs: an inverse regression approach. *Scandinavian Journal of Statistics*, 33(4):807–823.
- Ferré, L. and Yao, A. (2003). Functional sliced inverse regression analysis. *Statistics*, 37(6):475–488.
- Fraiman, R., Gimenez, Y., and Svarc, M. (2016). Feature selection for functional data. *Journal of Multivariate Analysis*, 146:191–208.
- Friedman, J., Hastie, T., and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1–22.
- Fromont, M. and Tuleau, C. (2006). Functional classification with margin conditions. In Lugosi, G. and Simon, H., editors, *Proceedings of the 19th Annual Conference on Learning Theory (COLT 2006)*, volume 4005 of *Lecture Notes in Computer Science*, pages 94–108, Pittsburgh, PA, USA. Springer (Berlin/Heidelberg).
- Fruth, J., Roustant, O., and Kuhnt, S. (2015). Sequential designs for sensitivity analysis of functional inputs in computer experiments. *Reliability Engineering & System Safety*, 134:260–267.
- Golub, T., Slonim, D., and Wahba, G. (1979). Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21(2):215–223.
- Gregorutti, B., Michel, B., and Saint-Pierre, P. (2015). Grouped variable importance with random forests and application to multiple functional data analysis. *Computational Statistics and Data Analysis*, 90:15–35.
- Hernández, N., Biscay, R., Villa-Vialaneix, N., and Talavera, I. (2015). A non parametric approach for calibration with functional data. *Statistica Sinica*, 25:1547–1566.

- James, G., Wang, J., and Zhu, J. (2009). Functional linear regression that's interpretable. *Annals of Statistics*, 37(5A):2083–2108.
- Li, K. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414):316–342.
- Li, L. and Nachtsheim, C. (2008). Sparse sliced inverse regression. *Technometrics*, 48(4):503–510.
- Li, L. and Yin, X. (2008). Sliced inverse regression with regularizations. *Biometrics*, 64(1):124–131.
- Liquet, B. and Saracco, J. (2012). A graphical tool for selecting the number of slices and the dimension of the model in SIR and SAVE approaches. *Computational Statistics*, 27(1):103–125.
- Matsui, H. and Konishi, S. (2011). Variable selection for functional regression models via the l_1 regularization. *Computational Statistics and Data Analysis*, 55(12):3304–3310.
- Meyer, D., Dimitriadou, E., Hornik, K., Weingessel, A., and Leisch, F. (2015). *e1071: Misc Functions of the Department of Statistics, Probability Theory Group (Formerly: E1071), TU Wien*. R package version 1.6-7.
- Ni, L., Cook, D., and Tsai, C. (2005). A note on shrinkage sliced inverse regression. *Biometrika*, 92(1):242–247.
- Ramsay, J. and Silverman, B. (1997). *Functional Data Analysis*. Springer Verlag, New York.
- Schott, J. (1994). Determining the dimensionality in sliced inverse regression. *Journal of the American Statistical Association*, 89(425):141–148.
- Zhu, L., Miao, B., and Peng, H. (2006). On sliced inverse regression with high-dimensional covariates. *Journal of the American Statistical Association*, 101(474):360–643.