



HAL
open science

Combining clustering of variables and feature selection using random forests: the CoV/VSURF procedure

Marie Chavent, Robin Genuer, Jerome Saracco

► **To cite this version:**

Marie Chavent, Robin Genuer, Jerome Saracco. Combining clustering of variables and feature selection using random forests: the CoV/VSURF procedure. 2016. hal-01345840v2

HAL Id: hal-01345840

<https://hal.science/hal-01345840v2>

Preprint submitted on 23 Aug 2016 (v2), last revised 6 Nov 2018 (v3)

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.

Combining clustering of variables and feature selection using random forests: the CoV/VSURF procedure

Marie CHAVENT^{1,2}, Robin GENUER^{3,4}, Jérôme SARACCO^{1,2}

¹ Institut de Mathématiques de Bordeaux, UMR CNRS 5251, Université de Bordeaux,
351 cours de la libération, 33405 Talence Cedex, France.

² INRIA Bordeaux Sud-Ouest, CQFD team, France.

³ ISPED, INSERM U-1219, Université de Bordeaux,
146 rue Léo Saignat, 33076 Bordeaux, France.

⁴ INRIA Bordeaux Sud-Ouest, SISTM team, France.

e-mail : `Marie.Chavent`, `Jerome.Saracco@math.u-bordeaux.fr`, `Robin.Genuer@isped.u-bordeaux2.fr`

Abstract

High-dimensional data classification is a challenging problem. A standard approach to tackle this problem is to perform variables selection, e.g. using stepwise procedures or LASSO approaches. Another standard way is to perform dimension reduction, e.g. by Principal Component Analysis (PCA) or Partial Least Square (PLS) procedures. The approach proposed in this paper combines both dimension reduction and variables selection. First, a procedure of clustering of variables (**CoV**) is used to built groups of correlated variables in order to reduce the redundancy of information. This dimension reduction step relies on the R package `ClustOfVar` which can deal with both numerical and categorical variables. Secondly, the most relevant synthetic variables (which are numerical variables summarizing the groups obtained in the first step) are selected with a procedure of variable selection using random forests (**VSURF**), implemented in the R package `VSURF`. Numerical performances of the proposed methodology called **CoV/VSURF** are compared with direct applications of **VSURF** or random forests (**RF**) on the original p variables. Improvements obtained with the **CoV/VSURF** procedure are illustrated on two simulated mixed datasets (cases $n > p$ and $n \ll p$) and on a real proteomic dataset.

***Keywords:** High-dimensional data, Supervised classification, Clustering of variables, Variable selection, Random forests, Mixed data*

1 Introduction

High-dimensional data classification is a challenging problem. This kind of data, such as genomics or proteomics data, are difficult to handle because the number p of variables largely exceeds the number n of

available observations. Furthermore, there are often two objectives when leading such an analysis: prediction (to be able to predict classes associated to new observations) and feature selection (to manage to extract the most interesting variables). Typically, in a medical context, the first goal could be to succeed in predicting if a patient will or will not respond well to a treatment, given his gene expression profile, whereas the second aim could be to determine which part of the genome is responsible to the good or bad response to the treatment. We stress that these two objectives are related, because it could be easier to perform prediction if useless variables have already been eliminated.

A classical way of addressing such issues is to use a variable selection technique (see e.g. Guyon and Elisseeff, 2003). The hope is that the method will be able to select the most interesting variables, while preserving good prediction performances. Moreover, high-dimensional data often come with many highly correlated variables and succeeding in selecting all variables in a group of correlated variables can be very difficult. Even if getting all these variables does not always appear to be interesting for prediction purpose, it can be useful for interpretation purpose, depending on the application.

In this article, we propose a two steps methodology for dimension reduction in the context of supervised classification which can also be applied in the context high-dimensional regression.

First, one eliminates redundancy using clustering of variables, based on the R package (R Core Team, 2016) `ClustOfVar` (Chavent et al., 2012a). This clustering approach, denoted by **CoV** hereafter, allows to deal with both numerical and categorical variables. The clustering of variables groups together highly correlated variables and provides, for each group (cluster), a synthetic variable which is a numerical variable summarizing the variables within a cluster. The main advantage of this step is to eliminate redundancy and to keep together all the variables of a cluster during the rest of the analysis. Moreover it reduces the dimension of the data by replacing the p original variables by K synthetic variables (where K denotes the selected number of clusters). Note that this clustering of variables approach does not require to define a priori groups of variables as in group lasso or sparse group lasso approaches (see e.g. Yuan and Lin, 2006). In addition, contrary to principal component analysis (PCA) (Jolliffe, 2002), the reduction of dimension provides K synthetic variables which are only constructed with the variables within a cluster and not using all the p original variables as for the synthetic variables in PCA (the principal components). Hence, an original variable takes action in the construction of a unique synthetic variable, which is nicer for interpretation purpose.

Secondly, the most important synthetic variables obtained in the first step are selected using a procedure based on random forests (**RF**), introduced in Genuer et al. (2010) and implemented in the R package **VSURF** (see Genuer et al., 2015). This variable selection procedure, denoted **VSURF** hereafter, is applied to the reduced dataset consisting of the n observations described with the K synthetic variables. Thus a list

of selected synthetic variables (i.e. a list of clusters of variables) is obtained and the prediction for new observations can be done with a predictor built on these selected synthetic variables.

Combining dimension reduction and variable selection has already been proposed by Chun and Keles (2010) and Lê Cao et al. (2011) via sparse partial least squares regression (PLS). However this approach is purely linear both in the construction of the synthetic variables (the PLS components) and in the link between the dependent variable and the selected components. In our methodology combining **CoV** and **VSURF**, the synthetic variables are linear combinations of the variables within a group, but the **VSURF** step is purely non parametric.

The rest of the paper is organized as follows. Section 2 gives an overview of both the clustering of variables method and the feature selection procedure. The methodology combining **CoV** and **VSURF** is then described in Section 2.3. In Section 3, the numerical performances of this methodology are compared, via a simulation study, with a straightforward applications of **VSURF** or **RF** (on the original p variables) and with application of **RF** on all the synthetic variables obtained in the first step. Two cases are taken into consideration: $n > p$ and $n \ll p$. In the corresponding simulated datasets, groups of correlated variables have been generated with various sizes of groups (small, moderate, large), various types of variables within a group (numerical, categorical, mixed). Moreover some groups are linked to the response variable (informative group), others are not. Noise variables are also present in the dataset. Finally, a real proteomic dataset is analyzed in Section 4, before giving concluding remarks in Section 5.

2 Description of the methodology

First, the two underlying methods, **CoV** for unsupervised dimension reduction and **VSURF** for both variable selection and prediction, are presented respectively in Sections 2.1 and 2.2. Then, the proposed methodology combining **CoV** and **VSURF**, named **CoV/VSURF** hereafter, is described in Section 2.3.

Let us consider a p -dimensional explanatory variable $X = (X^1, \dots, X^j, \dots, X^p)'$ and a univariate response variable Y , which takes its values in $\{1, \dots, L\}$. Let n be the number of observations of these variables. More precisely, let us consider a set of p_1 numerical variables measured on the n observations and denoted by $\{\mathbf{x}^1, \dots, \mathbf{x}^{p_1}\}$ and a set of p_2 categorical variables denoted by $\{\tilde{\mathbf{x}}^1, \dots, \tilde{\mathbf{x}}^{p_2}\}$ with $p_1 + p_2 = p$. Let \mathbf{X} be the corresponding data matrix of dimension $n \times p$. The i -th row of \mathbf{X} is denoted \mathbf{x}_i . Let \mathbf{y} be the vector of the n observations of the response variable.

2.1 Clustering of variables

The objective of clustering of variables is to sort variables into homogeneous clusters, that is to construct clusters of variables which are strongly related to each other and thus provide similar information. The idea

is to summarize all the variables belonging to a cluster by a synthetic numerical variable which is the most “linked” to all the variables within this cluster. In this section, we focus on a method devoted to the clustering of variables with no restriction on the type (numerical or categorical) of the variables. The clustering approach relies on an homogeneity criterion of a cluster and then on an homogeneity criterion of a partition. This criterion is based on the **PCAmix** method, a simple SVD (Singular Value Decomposition) version of the method PCAMIX introduced first by Kiers (1991). **PCAmix** is a principal component analysis method for a mixture of categorical and numerical variables which includes the ordinary principal component analysis (PCA) and multiple correspondence analysis (MCA) as special cases, see Chavent et al. (2012b) for details. More precisely, the ascendant hierarchical clustering algorithm implemented in the package `ClustOfVar` is used in the proposed methodology and presented hereafter, but the k -means type algorithm could also have been used. For seek of simplicity, the name **CoV** is used in the rest of the paper to refer to the hierarchical clustering method implemented in the package `ClustOfVar`.

Note that like PCA, **CoV** is a dimension reduction method but, contrary to PCA, it can be an helpful tool for variable selection. Indeed, each synthetic variable of **CoV** is a linear combination of a subset of variables (the variables within the corresponding cluster) whereas the principal components in PCA are linear combination of all the original variables. Selecting synthetic variables of **CoV** means selecting of subsets of original variables, which is not the case when selecting principal components in PCA.

Synthetic variable of a cluster C_k . This variable is defined as the numerical variable $\mathbf{f}^k \in \mathbb{R}^n$ which is the “most linked” to all the variables in C_k :

$$\mathbf{f}^k = \arg \max_{\mathbf{u} \in \mathbb{R}^n} \left\{ \sum_{\mathbf{x}^j \in C_k} r_{\mathbf{u}, \mathbf{x}^j}^2 + \sum_{\tilde{\mathbf{x}}^j \in C_k} \eta_{\mathbf{u}|\tilde{\mathbf{x}}^j}^2 \right\},$$

where $r_{\mathbf{u}, \mathbf{x}^j}^2 \in [0, 1]$ is the squared Pearson correlation between the numerical variables \mathbf{u} and \mathbf{x}^j , and $\eta_{\mathbf{u}|\tilde{\mathbf{x}}^j}^2 \in [0, 1]$ is the correlation ratio between \mathbf{u} and $\tilde{\mathbf{x}}^j$ (which measures the part of the variance of \mathbf{u} explained by the levels of $\tilde{\mathbf{x}}^j$). It has been shown that:

- \mathbf{f}^k is the first principal component of the method **PCAmix** applied to the variables in C_k .
- \mathbf{f}^k is as a linear combination of the numerical variables and of the dummy variables of the levels of the categorical variables in C_k . This linear combination given in (7) can be used to predict the value (score) of a new observation on the synthetic variable of C_k .
- the empirical variance of \mathbf{f}^k is equal to: $V(\mathbf{f}^k) = \sum_{\mathbf{x}^j \in C_k} r_{\mathbf{x}^j, \mathbf{f}^k}^2 + \sum_{\tilde{\mathbf{x}}^j \in C_k} \eta_{\mathbf{f}^k|\tilde{\mathbf{x}}^j}^2$.

Details on the **PCAmix** method and on the prediction of principal component scores can be found in the

Appendix. Note that the **PCAmix** algorithm is implemented in the R package **PCAmixdata** (Chavent et al., 2014).

Homogeneity H of a cluster C_k . This criterion H measures the adequacy between the variables within the cluster and its associated synthetic variable \mathbf{f}^k . It is defined as follows:

$$H(C_k) = \sum_{\mathbf{x}^j \in C_k} r_{\mathbf{x}^j, \mathbf{f}^k}^2 + \sum_{\tilde{\mathbf{x}}^j \in C_k} \eta_{\mathbf{f}^k | \tilde{\mathbf{x}}^j}^2 = \lambda_1^k, \quad (1)$$

where λ_1^k denotes the first eigenvalue of **PCAmix** applied to the cluster C_k .

Clearly the first term (based on the squared Pearson correlation r^2) quantifies the link between the numerical variables in C_k and \mathbf{f}^k , independently of the sign of the relationship. The second term (based on the correlation ratio η^2) measures the link between the categorical variables in C_k and \mathbf{f}^k . The homogeneity of a cluster is then maximized when all the numerical variables are perfectly correlated (or anti-correlated) to \mathbf{f}^k and when all the correlation ratios of the categorical variables are equal to 1. Hence all the variables in the cluster C_k bring the same information which is summarized by the corresponding synthetic variable \mathbf{f}^k .

Homogeneity \mathcal{H} of a partition $P_K = \{C_1, \dots, C_K\}$. The homogeneity criterion \mathcal{H} is defined as the sum of the homogeneities of its clusters:

$$\mathcal{H}(P_K) = \sum_{k=1}^K H(C_k) = \lambda_1^1 + \dots + \lambda_1^K. \quad (2)$$

A hierarchical clustering algorithm. The objective of this algorithm is to find a partition of a set of the p available (numerical and/or categorical) variables. This partition must be such that the variables within a cluster are strongly related to each other in the sense of the homogeneity criterion introduced previously. More specifically, the aim is to find a partition P_K which maximizes the homogeneity function \mathcal{H} defined in (2). To this end, a hierarchical clustering algorithm can be used and is described hereafter.

This algorithm builds a set of p nested partitions of variables as follows.

- **Step $l = 0$: Initialization.** Start with the partition in p clusters (i.e. one variable per cluster).
- **Step $l = 1, \dots, p - 2$: Aggregation of two clusters.** The objective is to aggregate two clusters of the partition in $p - l + 1$ clusters to get a new partition in $p - l$ clusters. To this end, we have to choose the two clusters A and B which provide the smallest dissimilarity $d(A, B)$ defined as:

$$d(A, B) = H(A) + H(B) - H(A \cup B) = \lambda_A^1 + \lambda_B^1 - \lambda_{A \cup B}^1. \quad (3)$$

This dissimilarity (aggregation measure) quantifies the lost of homogeneity observed when the two clusters A and B merge. Based on this measure of aggregation, the new partition in $p - l$ clusters

maximizes \mathcal{H} among all the partitions in $p - l$ clusters obtained by aggregation of two clusters of the partition in $p - l + 1$ clusters.

- **Step $l = p - 1$: Stop.** The final partition in one cluster (i.e. containing all the p variables) is obtained.

In the package `ClustOfVar`, the function `hclustvar` allows the user to build a hierarchy of the p variables according to the described hierarchical clustering algorithm. Moreover, the function `plot.hclustvar` gives the dendrogram of this hierarchy in which the height of the cluster $A \cup B$ is defined as $d(A, B)$. The function `cutreevar` cuts this dendrogram according to the number K of cluster given by the user and provides of the corresponding partition. Finally, the function `predict.clustvar` gives the scores of a new observation on the associated K synthetic variables.

Note that a partitioning clustering algorithm (in the spirit of k -means approach) is also available in the package `ClustOfVar`, as well as a bootstrap procedure to evaluate the stability of the p nested partitions of the dendrogram obtained with `hclustvar`.

2.2 Variable selection using random forests

In this section, we focus on the **VSURF** procedure. This method is based on random forests (Breiman, 2001), which provide a non-parametric predictor, with very good performance in prediction in lots of applied situations including high-dimensional data (see e.g. Verikas et al., 2011). Furthermore, the **VSURF** procedure is fully automatic (i.e. it does not need a pre-specified number of variables to select). Finally, **VSURF** can be applied for both supervised classification and regression problems. For the sake of simplicity, we focus on the supervised classification case, where the response variable is categorical.

Random Forests. A Random Forests (RF in the sequel) predictor is obtained by aggregating a collection of randomly perturbed decision trees. The randomness comes at two levels: the individuals level with a preliminary bootstrap sample draw, and the variables level with random variables subsamples draws before splitting a node of a tree. The main idea is to build diverse trees, which is mandatory to gain with the aggregation.

Let us denote by $\hat{h}_1, \dots, \hat{h}_q$ an ensemble of binary tree predictors. Each tree is a piece-wise constant function, which associates a class label for every input vector. A binary tree is made of internal nodes (which are split in two children nodes) and leaves (also called terminal nodes). The collection of the leaves forms a partition of the input space. The first node (which contains all data) is called the root of the tree. A node is pure if it contains observations belonging to the same class. Finally each internal node t is split according to a splitting variable X^{j_t} and a splitting value v_t .

Let us detail RF methodology in the following algorithm, which is implemented in the `randomForest` R package (Liaw and Wiener, 2002).

Input: a learning sample of size n and an input vector X .

Goal: predict the class label associated to X .

- **For each** $l = 1, \dots, q$:

- **Bootstrap sample.** Draw a bootstrap sample of the learning set, by randomly choosing n observations among the n available, with replacement.

- **Tree construction.**

- * Initialize the tree by putting all observations of the bootstrap sample in the root node.

- * **While it exists an impure node among current terminal nodes,**

- for each impure terminal node, randomly choose a subset of `mtry` variables among the p variables without replacement, seek the best split of the node only among splits involving the selected variables, and split the node into two children nodes according to the best split (see details below).

- **Tree prediction.**

- Let X go down the tree and note the leaf it falls into. Return $\hat{h}_l(X)$, the class label of the observations of the learning set belonging to this leaf.

- **Aggregation.**

Return the majority class among trees predictions: $\hat{f}(X) = \operatorname{argmax}_{c \in \{1, \dots, L\}} \sum_{l=1}^q \mathbf{1}_{\hat{h}_l(X)=c}$

We now precise what “looking for the best split” means. For supervised classification, the Gini index is used for heterogeneity measure of a node, in terms of class labels. It is defined, for a node t , as follows:

$$\text{Gini}(t) = \sum_{c \in \{1, \dots, L\}} p_c(t)(1 - p_c(t))$$

where $p_c(t)$ is the proportion of observations of class c in node t . In the 2-class problem, it simplifies to $2p_0(t)p_1(t)$. Hence, the best split of a node t is the one minimizing $\text{Gini}(t)$ over all possible splits.

For a numerical variable, possible splits are of the form $\{X^j \leq v\}$ with $v \in [\min(\mathbf{x}^j), \max(\mathbf{x}^j)[$. This split means that observations with a j -th variable value not larger than v are sent to the left child node and the others to the right one. For a categorical variable, possible splits are of the form $\{X^j \in \mathcal{M}_j\}$ where \mathcal{M}_j is a subset of the j -th variable levels (except empty set and total set).

Note, that compared to Classification And Regression Trees (CART) introduced in Breiman et al. (1984), the search for the best split of a node is only made among the randomly selected variables instead of among all the p available variables. Furthermore, trees are fully developed (it remains either several observations with the same class label, or only one observation in a leaf) and are not pruned.

OOB error and variable importance. During a RF run, an estimation of the prediction error and a measure of variable importance can be computed.

Out-Of-Bag (OOB) error of a RF \hat{f} is defined as follows:

$$\text{OOBerror}(\hat{f}) = \frac{1}{n} \text{Card} \{i \in \{1, \dots, n\} \mid y_i \neq \hat{y}_i\}$$

where \hat{y}_i is the majority class label among predictions of the trees \hat{h}_l for which y_i is OOB, that is for which y_i was not chosen in the bootstrap sample used to build \hat{h}_l .

Variable Importance (VI) also uses the OOB samples (all observations of the learning set not included in a bootstrap sample) to compute a measure of the link between Y and a variable X^j . It is defined as follows:

$$\text{VI}(X^j) = \frac{1}{q} \sum_{l=1}^q \left(\text{errperm}_j(\hat{h}_l) - \text{err}(\hat{h}_l) \right) \quad (4)$$

where

$$\text{err}(\hat{h}_l) = \frac{\text{Card} \left\{ i \mid (\mathbf{x}_i, y_i) \text{ is OOB for } \hat{h}_l \text{ and } y_i \neq \hat{h}_l(\mathbf{x}_i) \right\}}{\text{Card} \left\{ i \mid (\mathbf{x}_i, y_i) \text{ is OOB for } \hat{h}_l \right\}}$$

and

$$\text{errperm}_j(\hat{h}_l) = \frac{\text{Card} \left\{ i \mid (\mathbf{x}_i, y_i) \text{ is OOB for } \hat{h}_l \text{ and } y_i \neq \hat{h}_l(\mathbf{x}_i^{\text{perm}_j}) \right\}}{\text{Card} \left\{ i \mid (\mathbf{x}_i, y_i) \text{ is OOB for } \hat{h}_l \right\}}$$

where \mathbf{x}_i is the i -th row of the data matrix \mathbf{X} and $\mathbf{x}_i^{\text{perm}_j}$ is the i -th row of the data matrix obtained by randomly permuting the values \mathbf{x}_i^j of the j -th column, for which i is OOB for \hat{h}_l .

The idea is that the more the mean error of a tree on its OOB sample increases when the link between X^j and Y is broken, the more important the variable is.

Variable Selection Using Random Forests. The **VSURF** procedure works with three steps.

- The first one begins by sorting variables based on random forests VI, and eliminates useless variables by an adaptive thresholding. The threshold is set to the estimation of the standard deviation (over multiple RF runs) of the VI of an unimportant variable.

- The second one starts with the previously kept variables and performs an ascendant variable introduction strategy, which builds embedded RF models. The model which attains the minimum OOB error rate is then selected, and the variables set on which it is based is called the interpretation set.
- The third step consists in eliminating the redundancy of interpretation variables and leads to a smaller variables set called the prediction set. It consists in a step-wise ascendant strategy, which at each step verifies that the next variable to introduce helps to decrease enough the OOB error rate.

More details on those three steps can be found in Genuer et al. (2015).

2.3 The CoV/VSURF procedure

We now describe the proposed methodology, which combines **CoV** and **VSURF**, in the following algorithm.

Input: a dataset (\mathbf{X}, \mathbf{y}) and a new observation \mathbf{x} of X .

Goal: select groups of informative variables and predict the class label of \mathbf{x} .

(a) **Groups of informative variable selection:**

1. Apply **CoV** on \mathbf{X} to obtain a hierarchy (a tree) of variables.
2. For each $K = 2, \dots, p$, cut **CoV** tree in K clusters, train a RF with the K synthetic variables $\mathbf{f}^1, \dots, \mathbf{f}^K$ as predictors and \mathbf{y} as output variable and compute its OOB error rate.
3. Choose the optimal number K^* of clusters, which leads to the minimum OOB error rate.
Cut **CoV** tree in K^* clusters.
4. Perform **VSURF** with the K^* synthetic variables $\mathbf{f}^1, \dots, \mathbf{f}^{K^*}$ as predictors and \mathbf{y} as output variable. Denote by $m \leq K^*$ the number of selected informative synthetic variables (corresponding to the interpretation set of **VSURF**).

(b) **Prediction of a new observation \mathbf{x} :**

1. Train a random forest, \hat{f} , on the dataset consisting of the m selected synthetic variables and \mathbf{y} .
2. Compute the scores of \mathbf{x} on the m selected synthetic variables and predict its class label using \hat{f} .

This method is illustrated on simulated examples and a real proteomic dataset in the following sections. Finally, we mention that we have implemented the method **CoV/VSURF** in an R package, which is under development¹.

¹<https://github.com/robingenuer/CoVVSURF>

3 Simulation study

We first describe the simulated model used to generate the data. We then consider a large n (resp. a small n) simulated dataset to evaluate the proposed methodology in comparison with alternative approaches.

3.1 Simulated model

We consider a binary response variable Y and a p -dimensional explanatory variable X . We model the conditional probability $P(Y = 1|X = x) =: p(x)$ as a function of x using the well known logistic regression model formally defined as

$$\log\left(\frac{p(x)}{1-p(x)}\right) = x'\beta - 9, \quad (5)$$

where $\beta \in \mathbb{R}^p$. Note that the term “-9” in (5) allows us to center the index $X'\beta$ (for the considered choice of parameters) and then to obtain equibalanced response for Y .

Let us now specify how we construct the p -dimensional explanatory variable X (with $p = 120$) in order to get 10 groups of variables (components X^j of X) of several types (numerical, categorical or mixed) such that some groups of variables are informative or not. Table 1 provides a brief description of the 10 groups of these $p = 120$ variables, details are given below.

Type of variables	Size of group	Informative group	Components of β	Names of variables
Numerical	Small (3)	Yes	(1, 2, 3)	NumS j with $j = 1, \dots, 3$
Numerical	Moderate (12)	No	(0, ..., 0)	NumM j with $j = 1, \dots, 12$
Numerical	Large (15)	Yes	$\frac{1}{5}(\underbrace{1, \dots, 1}_5, \underbrace{2, \dots, 2}_5, \underbrace{3, \dots, 3}_5)$	NumL j with $j = 1, \dots, 15$
Categorical	Small (3)	Yes	(1, 2, 3)	CategS j with $j = 1, \dots, 3$
Categorical	Moderate (12)	No	(0, ..., 0)	CategM j with $j = 1, \dots, 12$
Categorical	Large (15)	Yes	$\frac{1}{5}(\underbrace{1, \dots, 1}_5, \underbrace{2, \dots, 2}_5, \underbrace{3, \dots, 3}_5)$	CategL j with $j = 1, \dots, 15$
Mixed	Small (3)	Yes	(1, 2, 3)	MixedS j with $j = 1, \dots, 3$
Mixed	Moderate (12)	No	(0, ..., 0)	MixedM j with $j = 1, \dots, 12$
Mixed	Large (15)	Yes	$\frac{1}{5}(\underbrace{1, \dots, 1}_5, \underbrace{2, \dots, 2}_5, \underbrace{3, \dots, 3}_5)$	MixedL j with $j = 1, \dots, 15$
Numerical	Big (30)	No	(0, ..., 0)	Noise j with $j = 1, \dots, 30$

Table 1: Overview of the 10 groups of variables in the simulation study.

Let us first introduce the underlying explanatory variable Z following the multivariate Gaussian distribution $\mathcal{N}_p(\mu, \Sigma)$ where $p = 120$, $\mu = 0_p$ and the covariance matrix Σ is a block-diagonal matrix in order to get 10 groups of variables. These groups are independent from each other. We also introduce the parameter $\rho \in [-1, 1]$ which allows to control the link between the variables within a group. We then define the

covariance matrix of Z :

$$\Sigma = \text{diag} \left(\Sigma_{3,\rho}, \Sigma_{15,\rho}, \Sigma_{12,\rho}, \Sigma_{3,\rho}, \Sigma_{15,\rho}, \Sigma_{12,\rho}, \Sigma_{3,\rho}, \Sigma_{15,\rho}, \Sigma_{12,\rho}, \sigma^2 I_{30} \right),$$

where

$$\Sigma_{s,\rho} = \begin{pmatrix} 1 & \rho & \cdots & \rho \\ \rho & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \rho \\ \rho & \cdots & \rho & 1 \end{pmatrix}$$

is a $s \times s$ matrix, $\sigma^2 > 0$ and I_{30} is the identity matrix of dimension 30. When the value of ρ is close to 1, we obtain nine groups of highly correlated numerical variables (3 groups of 3 variables, 3 groups of 15 variables, 3 groups of 12 variables) and one group of 30 variables which are independent. In the following, we set $\rho = 0.9$.

For the nine groups we are interested in studying the general case where the user faces up to mixed data, with both numerical and/or categorical variables within a group. Therefore we generate, for each size of group (3, 12 and 15), a group of numerical variables, a group of categorical variables and a group of mixed variables. The construction of the categorical variables from the numerical ones is described below in the data generating process. Moreover, for each type of variables, we choose the components of β such that the 3 small (size equal to 3) groups and the 3 large (size equal to 15) groups are informative, whereas the 3 moderate (size equal to 12) ones are non-informative (i.e. the corresponding components of β are null). The last group of 30 variables is non-informative and is formed of non-structured (independent) variables.

Let us now specify the data generating process. We generate n mutually independent random p -dimensional vectors $\mathbf{z}_1, \dots, \mathbf{z}_n$ from the Gaussian distribution $\mathcal{N}_p(\mu, \Sigma)$. From the \mathbf{z}_i 's, we first construct the output variable values $y_i \in \{0, 1\}$ using the logistic regression model (5).

Since the \mathbf{z}_i 's original data are numerical, we have to binarize some numerical variables Z^j to obtain categorical ones by thresholding at their median value: if the variable value is less than the median, the value of the corresponding binarized variable X^j is 0, and this value is 1 otherwise. We proceed as follows to obtain the entire \mathbf{x}_i 's data:

- the first 30 variables are unchanged (i.e. $X^j := Z^j$) and provide the 3 groups of numerical variables;
- the following 30 variables X^j are binarized and provide the 3 groups of categorical variables;
- the next 30 variables X^j provide the 3 groups of mixed variables. More precisely, a third of the variables are binarized in each group:
 - in the small group of 3 correlated variables, only the last one is binarized,
 - in the moderate group of 12 correlated variables, only the last 4 are binarized,

- in the large group of 15 correlated variables, only the last 5 are binarized;
- the last 30 independent variables are unchanged (i.e. $X^j := Z^j$).

To sum up, we have a binary response variable and nine groups of correlated variables: one small and one large informative groups of numerical variables, one moderate non-informative group of numerical variables, and this structure is repeated for the three groups of categorical variables and the three groups of mixed variables. In addition, we have 30 non-informative and non-structured (independent) variables.

3.2 A large n simulated learning dataset

Application of CoV/VSURF to a learning dataset. From the previous model, a learning sample of $n = 600$ observations of the $p = 120$ explanatory variables and of the binary response variable Y is generated.

First, the method **CoV** is applied to the 600×120 matrix of explanatory variables. The dendrogram of the hierarchy of the 120 variables (also called the **CoV** tree) is given in Figure 1. This dendrogram suggests a partition in 9 clusters.

However, in our methodology, the number of clusters is not chosen according to the shape of the dendrogram but according to the prediction of the binary response variable Y . Indeed, for each value of K between 2 and 120, we cut the **CoV** tree, build a random forest on the K synthetic variables of the corresponding K -clusters partition and compute the RF OOB error rate. This procedure is illustrated Figure 2. The (nearly) optimal value we get for this learning dataset is $K^* = 9$. This partition in $K^* = 9$ clusters retrieves almost the complete structure of the data. We recover 8 of the 9 groups of correlated variables, while all noise variables are pulled together in the last cluster with the last group of correlated variables (the small one). However, the synthetic variable of this last cluster does not really take the noise variables into account, whose loadings (coefficients in the linear combination) are very low contrary to the loadings of the variables of the small group of correlated variables. Since this small group of numerical variables is informative, it means that this synthetic variable is also informative.

Next, we apply **VSURF** on the $K^* = 9$ synthetic variables of the previously chosen partition. **VSURF** selects 6 synthetic variables, corresponding to the 6 informative groups of variables. From the interpretation point of view, we succeed in selecting all informative variables, with in addition the clustering structure. For comparison, **VSURF** directly applied on the 120 original variables selects 39 variables among the 54 informative ones, with at least one per informative group (the large group of categorical variables being the least recovered).

Prediction performances on a test dataset. In addition to the learning dataset, a test sample of $n = 600$ observations is now generated. We focus on the prediction performances of the proposed methodology




Figure 1: **CoV** tree of the $p = 120$ variables based on the learning sample of size $n = 600$. Heights correspond to dissimilarities between two aggregated clusters. In red, the partition in $K^* = 9$ clusters.




Figure 2: Random Forests OOB error rate according to the number of clusters of the partitions obtained by cutting the **CoV** tree for the simulated learning dataset of $n = 600$ observations. The dashed red vertical line corresponds to $K^* = 9$ clusters.

CoV/VSURF on this test sample and do comparisons with 3 other methodologies:

- **VSURF**: variable selection using random forests is applied on the original 120 variables,
- **RF**: random forests are applied on the original 120 variables,
- **CoV/RF**: random forests are applied on the K^* synthetic variables obtained by **CoV** in a first step.

Note that the test sample of size $n = 600$ is fixed all along this simulation study. We obtain in Figure 3 the boxplots of test error rates.




Figure 3: Comparison of test error rates with learning datasets of size $n = 600$. **CoV/VSURF** and **CoV/RF** correspond to **VSURF** and random forests (RF) applied on the K^* synthetic variables given by **CoV**. **VSURF** and **RF** refer to methods applied on the original 120 variables. (a) Test error rates of 100 forests trained on one learning dataset. (b) Averaged (over 100 forests) test error rates computed on 50 learning datasets.

On the left-hand side of Figure 3, boxplots correspond to 100 runs of random forests. More precisely,

- For **CoV/VSURF** and **CoV/RF**, the clustering of variables is performed once and the optimal number K^* of clusters is chosen automatically using OOB error rate. For **Cov/RF**, 100 forests are trained on the K^* synthetic variables, while for **Cov/VSURF**, 100 forests are trained on the $m < K^*$ synthetic variables selected by **VSURF** (which is then performed also a single time).
- For **VSURF**, 100 forests are trained on the variables selected with one run of **VSURF** applied on the $p = 120$ original variables.
- For **RF**, 100 forests are trained on the original variables.

The results are better here with **CoV/VSURF** and **CoV/RF** i.e. when **VSURF** and **RF** are applied after a **CoV** step. In addition, the approach **CoV/VSURF** is slightly better than **CoV/RF** in this setting. We note also that the variability due to the random nature of the **RF** is relatively small here and very similar with the four approaches.

On the right-hand side of Figure 3, 50 learning datasets are generated in order to take into account the variability of the results due to the random nature of the simulation procedure. The four methodologies explained above for a single learning dataset are applied on each dataset. Boxplots correspond then to the 50 averaged (over 100 forests) test error rates. The results are again better here with **CoV/VSURF** and **CoV/RF** and the improvement obtained by using **CoV** in a first step before the classification step is confirmed here, despite the variability due to the random simulation procedure.

To sum up, applying **CoV** in a first step allows an efficient data dimension reduction which leads — at least for this example — to a gain in prediction both with **RF** and **VSURF**. But applying **VSURF** in a second step permits furthermore to select informative groups of linked variables without loss in prediction. In the next section, the exact same experiment is repeated but in a more challenging situation where only $n = 60$ observations are generated, hence $n \ll p$.

3.3 A “small n , large p ” simulated learning dataset

We simulate now a learning sample of $n = 60$ observations of the $p = 120$ explanatory variables and of the binary response variable Y .

First, the method **CoV** is applied to the 60×120 matrix of explanatory variables. The dendrogram of the hierarchy is represented in Figure 4. This tree suggests to retain the partition in $K = 9$ clusters which matches with the underlying partition in 9 groups of correlated variables (i.e. the informative and non informative structured groups of variables) while the 30 noise variables are pulled together either with the small group of numerical variables or with the small group of mixed variables.

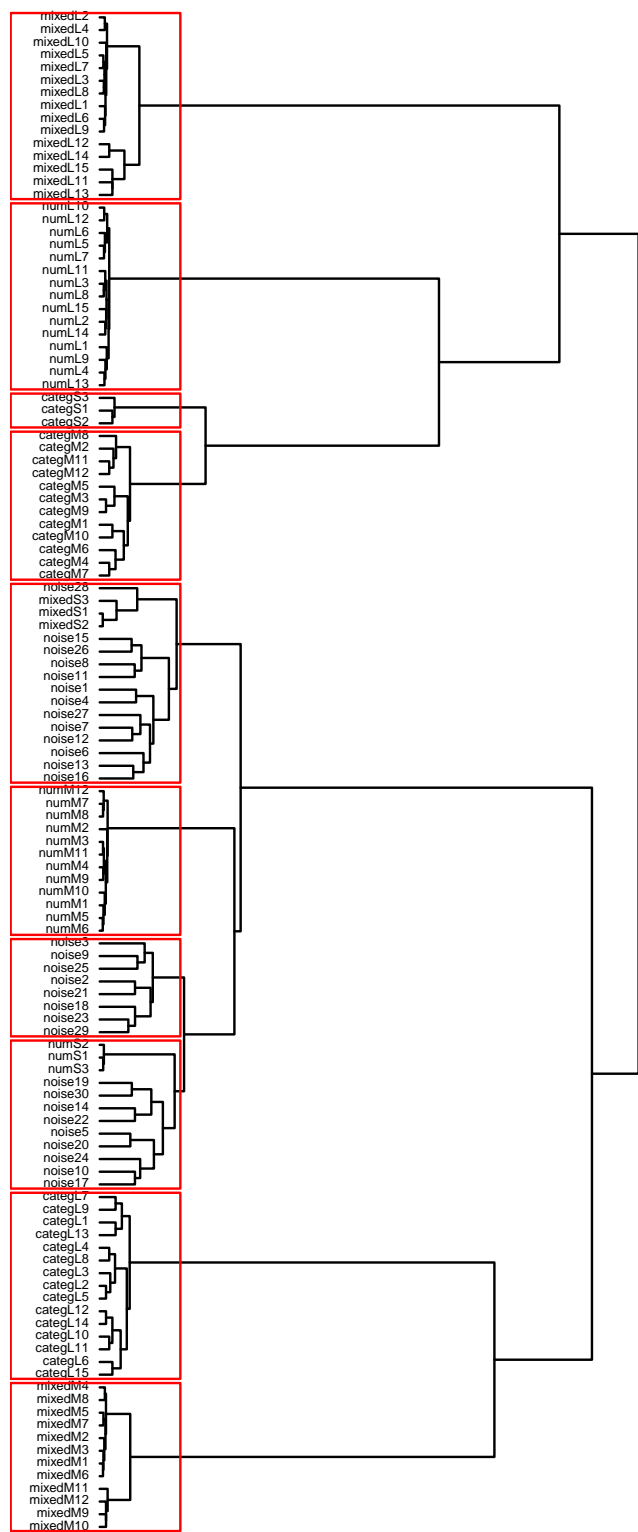


Figure 4: **CoV** tree of the $p = 120$ variables based on the learning sample of size $n = 60$. Heights correspond to dissimilarities between two aggregated clusters. In red, the partition in $K^* = 10$ clusters.

However, as illustrated in Figure 5, our methodology chooses to cut the dendrogram in $K^* = 10$ clusters. The structure in 9 groups of correlated variables is conserved and the 10-th cluster only contains 8 noise variables. **VSURF** applied on the $K^* = 10$ associated synthetic variables selects 4 of them. The two groups of informative numerical variables and the two groups of informative mixed variables are retained. In this more difficult context where $n \ll p$, the **CoV/VSURF** methodology misses the informative groups of categorical variables. This can be explained by the fact that those variables are binary, hence less explanatory than the numerical ones. **VSURF** applied on the 120 original variables only selects 8 variables, with at least one per group selected by our approach, and still no informative categorical variable are retained.

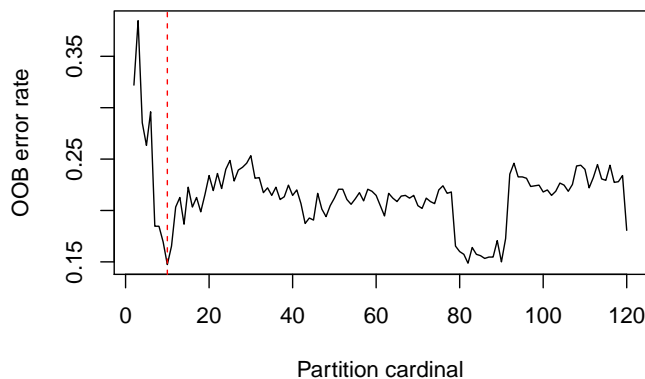


Figure 5: Random Forests OOB error rate according to **CoV** partition cardinal for the learning dataset of $n = 60$ observations. The dashed red vertical line corresponds to $K^* = 10$ clusters.

We focus now on the prediction performances of the proposed methodology **CoV/VSURF**. The test sample used here is the same as in the previous section and Figure 6 is obtained in the exact same way than in the previous $n = 600$ case. This Figure shows that the variability due to random forests (left-hand side) remains relatively small compared to the variability due to random generation of the 50 learning dataset (right-hand side). However, in this more difficult situation, all methodologies are fairly comparable in prediction, with a slight advantage to **Cov/RF**. But we remind again that the proposed methodology **CoV/VSURF** has the advantage to select groups of informative variables, without increasing too much — at least for this example — the prediction error rate.

4 Proteomic data application

In this section, we illustrate our methodology **CoV/VSURF** with an application on a real proteomic dataset. This dataset comes from a clinical study sponsored by University Hospital in Bordeaux. It involved $n = 44$ patients with a rectum cancer who undertook a treatment of chemotherapy and radiotherapy, before a surgery intervention. Some patients responded favorably to the treatment and hence had a smaller tumor at t_1 , the




Figure 6: Comparison of test error rates with learning datasets of size $n = 60$. **CoV/VSURF** and **CoV/RF** correspond to **VSURF** and random forests (RF) applied on the K^* synthetic variables given by **CoV**. **VSURF** and **RF** refer to methods applied on the original 120 variables. (a) Test error rates of 100 forests trained on one learning dataset. (b) Averaged (over 100 forests) test error rates computed on 50 learning datasets.

time of surgery, compared to t_0 , a time just before the beginning of the treatment, and some patients did not.

The main goal of this study was to predict if the patient will be a good treatment responder or not, using proteomic information, measured at t_0 . This is indeed a crucial information, because it is useless to give the treatment to a bad responder, for which other alternatives should be tried. Furthermore, oncologists are interested in knowing the proteins which discriminates the most the two kinds of treatment response.

Let us describe more precisely the data. We have $p = 4786$ numerical explanatory variables which are measures of protein abundances. Measurements are done on peptides, i.e. a “part” of a protein (one protein is made of several peptides). So, the total number of peptides is $p = 4786$, whereas total number of proteins is 868. The a priori knowledge of what peptide is part of which protein is not used in our approach. However this information is used for interpretation purpose.

The proposed **CoV/VSURF** methodology, described in Section 2.3, is applied to this dataset. Since the number p of variables is relatively large, the number of clusters of peptides is chosen among 2 and 2000. The resulting random forests OOB error rates are displayed in Figure 7 and the (nearly) optimal number of clusters is $K^* = 68$.

The **VSURF** method, applied on those $K^* = 68$ synthetic variables, selects 4 of them. These 4 synthetic variables are sorted by decreasing order of their variable importance (VI, defined in 4). The corresponding 4 groups of peptides gather respectively 37, 61, 20, 25 peptides, that is a total of 143 peptides. Those 143 peptides come from 73 different proteins. The most homogeneous one is the 4th with 22 peptides (over 25) coming from the same protein, *PZP* (Pregnancy Zone Protein). The 3rd one contains 18 peptides (over 20) coming from two different proteins, *IGHD* (Ig Delta Chain) protein and *THRB* (Prothrombin) protein. The




Figure 7: Random forests OOB error rate according to **CoV** partition cardinal (number of peptides clusters) for the proteomic data. The two dashed red vertical lines correspond respectively to partitions in $K^* = 68$ and $K = 142$ clusters. For visibility, the x-axis is truncated around $K = 500$ clusters (the error rate remains quite stable until $K = 2000$).

second cluster is the more heterogeneous with 61 peptides coming from 50 different proteins. The first one contains 13 peptides (over 37) coming from *A2GL* (Alpha 2 Glycoprotein), 6 from *CRP* (C-reactive protein).

This cluster of 37 peptides (which has the highest VI) is responsible of the large decrease of OOB error rates in Figure 7: it appears when the number of synthetic variables goes from 67 to 68 i.e. when the OOB error rate decreases from around 40% to 20%. In addition, when this particular cluster is split into two new clusters (when the partition cardinal K reaches 142), the OOB error increases significantly from 20% to 30%. This confirms the importance of this cluster of peptides to predict treatment response.

Note that, among the peptides selected by **CoV/VSURF**, we find peptides which come from *APOA1* (Apolipoprotein A1), *HPT* (Haptoglobin), *TRFE* (Serotransferrin) and *PGRP2* (N-Acetylmuramoyl L-Alanine Amidase). Those proteins make sense to oncologists in this context, who are currently investigating in more depth the relation between those proteins and the response to the treatment.

Note also that, when **VSURF** is applied alone on the original data, only 35 peptides are selected. So, **VSURF** gives a more sparse variable selection, but with no group structure.

To get an estimation of classification error rates for this real dataset, we perform an external leave-one-out cross-validation, to fairly estimate prediction performance of our approach (see Ambroise and McLachlan, 2002). This means that, for each patient in this study, we perform the entire methodology **CoV/VSURF** on the dataset containing all other patients, before predicting him. As before, we compare the performance with the three approaches **CoV/RF**, **VSURF** and **RF**. Figure 8 provides the results. According to these boxplots, there is almost no difference between **CoV/VSURF**, **CoV/RF** and **VSURF** (recall that since $n = 42$, one more good classification of an individual allows to gain around 0.02 in error rate). The main advantage of **CoV/VSURF** is again that it selects groups of informative peptides without increasing prediction error.




Figure 8: Comparison of leave-one out error rates for proteomic data. **CoV/VSURF** and **CoV/RF** correspond to **VSURF** and random forests (RF) applied on the $K^* = 68$ synthetic variables given by **CoV**. **VSURF** and **RF** refer to methods applied on the original $p = 4786$ peptides.

5 Concluding remarks

In this paper, we proposed a two steps methodology, named **CoV/VSURF**, to tackle the problem of high-dimensional supervised classification. The **CoV** step constructs "unidimensional" groups of variables by clustering of variables. Latent variables of the groups are then selected in the **VSURF** step with an automatic procedure based on random forests. We emphasized the advantage of using clustering of variables (**CoV**) rather than PCA in the dimension reduction step and **VSURF** rather than **RF** in the prediction step, to select groups of informative variables. Moreover the simulation study has shown good numerical performances of this methodology compared to straightforward applications of **VSURF** or **RF**.

This methodology can be applied with mixed data type (i.e. mixture of numerical and categorical variables) and it also applies in the context of high-dimensional regression. Based on the existing R packages `ClustOfVar` and `VSURF`, we have implemented **CoV/VSURF** in an R package under development.

Acknowledgments.

The authors would like to thank members of the USMR unit of Bordeaux University Hospital (especially H el ene Savel and Paul Perez) for allowing access to the clinical study data and for insightful discussions.

Appendix: the PCAmix algorithm.

This PCAmix algorithm is applied on a $n \times (p_1 + p_2)$ data matrix \mathbf{X} where the first p_1 columns correspond to p_1 numerical variables measured on the n observations and the following p_2 columns correspond to p_2 categorical variables. It works in three main steps.

1. PREPROCESSING STEP.

- (a) Build the numerical data matrix $\mathbf{Z} = (\mathbf{Z}_1|\mathbf{Z}_2)$ of dimension $n \times (p_1 + m)$ where \mathbf{Z}_1 is the standardized version of the first p_1 columns of \mathbf{X} corresponding to the p_1 numerical variables, \mathbf{Z}_2 is the centered indicator matrix of the last p_2 columns of \mathbf{X} corresponding to the p_2 categorical variables and m denotes the total number of levels.
- (b) Build the diagonal matrix \mathbf{N} of the weights of the rows. The n rows are weighted by $\frac{1}{n}$, that is $N = \text{diag}(\frac{1}{n}, i = 1, \dots, n)$.
- (c) Build the diagonal matrix \mathbf{M} of the weights of the columns.
 - The p_1 first columns are weighted by 1.
 - The m last columns are weighted by $\frac{n}{n_s}$, with n_s the number of observations with category s .

2. GSVD STEP.

The GSVD (Generalized Value Decomposition) of \mathbf{Z} with the diagonal metrics of the weights \mathbf{N} and \mathbf{M} gives the decomposition

$$\mathbf{Z} = \mathbf{U}\mathbf{\Lambda}\mathbf{V}^t \quad (6)$$

where

- $\mathbf{\Lambda} = \text{diag}(\sqrt{\lambda_1}, \dots, \sqrt{\lambda_r})$ is the $r \times r$ diagonal matrix of the singular values of $\mathbf{Z}\mathbf{M}\mathbf{Z}^t\mathbf{N}$ and $\mathbf{Z}^t\mathbf{N}\mathbf{Z}\mathbf{M}$, and r denotes the rank of \mathbf{Z} ;
- \mathbf{U} is the $n \times r$ matrix of the first r eigenvectors of $\mathbf{Z}\mathbf{M}\mathbf{Z}^t\mathbf{N}$ such that $\mathbf{U}^t\mathbf{N}\mathbf{U} = \mathbb{I}_r$;
- \mathbf{V} is the $p \times r$ matrix of the first r eigenvectors of $\mathbf{Z}^t\mathbf{N}\mathbf{Z}\mathbf{M}$ such that $\mathbf{V}^t\mathbf{M}\mathbf{V} = \mathbb{I}_r$.

3. SCORES PROCESSING STEP.

- (a) The set of factor scores for rows, called principal components, is computed as:

$$\mathbf{F} = \mathbf{U}\mathbf{\Lambda}.$$

- (b) The set of factor scores for columns, called loadings, is computed as:

$$\mathbf{A} = \mathbf{M}\mathbf{V}\mathbf{\Lambda} = \begin{pmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{pmatrix} \begin{matrix} \} p_1 \\ \} m \end{matrix}$$

where \mathbf{A}_1 (resp. \mathbf{A}_2) contains the scores of the p_1 numerical variables (resp. the scores of the m levels).

Principal components prediction. The α th principal component \mathbf{f}_α is as a linear combination of the columns of \mathbf{X}_1 formed by the first p_1 columns of \mathbf{X} and of the m columns of \mathbf{G} where \mathbf{G} is the indicator matrix of the p_2 categorical variables (the last p_2 columns of \mathbf{X}). Let us denote by $\check{\mathbf{X}} = [\dots, \check{\mathbf{x}}_j, \dots]$ the corresponding $n \times (p_1 + m)$ matrix. Then, \mathbf{f}_α writes as follows:

$$\mathbf{f}_\alpha = \beta_0 + \sum_{j=1}^{p_1+m} \beta_j \check{\mathbf{x}}_j \quad (7)$$

with $\beta_0 = -\sum_{k=1}^{p_1} v_{k\alpha} \frac{\bar{\mathbf{x}}_k}{s_k} - \sum_{k=p_1+1}^{p_1+m} v_{k\alpha}$, $\beta_j = v_{j\alpha} \frac{1}{s_j}$ for $j = 1, \dots, p_1$ and $\beta_j = v_{j\alpha} \frac{n}{n_j}$ for $j = p_1 + 1, \dots, p_1 + m$. These coefficients can then be used for instance to predict the scores of a new observation on the α th principal component.

References

- Ambrose, C. and McLachlan, G. J. (2002). Selection bias in gene extraction on the basis of microarray gene-expression data. *Proceedings of the national academy of sciences*, 99(10):6562–6566.
- Breiman, L. (2001). Random forests. *Machine learning*, 45(1):5–32.
- Breiman, L., Friedman, J., Olshen, R., and Stone, C. (1984). *Classification and Regression Trees*. Chapman & Hall, New York.
- Chavent, M., Kuentz-Simonet, V., Labenne, A., and Saracco, J. (2014). Multivariate analysis of mixed data: The PCAmixdata R package. *arXiv preprint arXiv:1411.4911*.
- Chavent, M., Kuentz-Simonet, V., Liquet, B., and Saracco, J. (2012a). Clustofvar: An R package for the clustering of variables. *Journal of Statistical Software*, 50(13):1–16.
- Chavent, M., Kuentz-Simonet, V., and Saracco, J. (2012b). Orthogonal rotation in PCAMIX. *Advances in Data Analysis and Classification*, 6(2):131–146.
- Chun, H. and Keles, S. (2010). Sparse partial least squares regression for simultaneous dimension reduction and variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(1):3–25.
- Genuer, R., Poggi, J.-M., and Tuleau-Malot, C. (2010). Variable selection using random forests. *Pattern Recognition Letters*, 31(14):2225–2236.
- Genuer, R., Poggi, J.-M., and Tuleau-Malot, C. (2015). VSURF: An R package for variable selection using random forests. *The R Journal*, 7(2):19–33.

- Guyon, I. and Elisseeff, A. (2003). An introduction to variable and feature selection. *The Journal of Machine Learning Research*, 3:1157–1182.
- Jolliffe, I. (2002). *Principal component analysis*. Wiley Online Library.
- Kiers, H. A. (1991). Simple structure in component analysis techniques for mixtures of qualitative and quantitative variables. *Psychometrika*, 56(2):197–212.
- Lê Cao, K.-A., Boitard, S., and Besse, P. (2011). Sparse pls discriminant analysis: biologically relevant feature selection and graphical displays for multiclass problems. *BMC bioinformatics*, 12(1):1.
- Liaw, A. and Wiener, M. (2002). Classification and regression by randomForest. *R News*, 2(3):18–22.
- R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Verikas, A., Gelzinis, A., and Bacauskiene, M. (2011). Mining data with random forests: A survey and results of new tests. *Pattern Recognition*, 44(2):330–349.
- Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67.