



HAL
open science

Joint rank and variable selection for parsimonious estimation in high-dimension finite mixture regression model.

Emilie Devijver

► **To cite this version:**

Emilie Devijver. Joint rank and variable selection for parsimonious estimation in high-dimension finite mixture regression model.. 2015. hal-01099296v1

HAL Id: hal-01099296

<https://hal.science/hal-01099296v1>

Preprint submitted on 2 Jan 2015 (v1), last revised 14 Feb 2017 (v2)

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.

JOINT RANK AND VARIABLE SELECTION FOR PARSIMONIOUS ESTIMATION IN HIGH-DIMENSION FINITE MIXTURE REGRESSION MODEL.

EMILIE DEVIJVER

ABSTRACT. We study a dimension reduction method for finite mixture of multivariate response regression models in high dimension. Both the number of responses and of predictors may exceed the sample size. We consider jointly predictor selection and rank reduction for obtaining lower-dimensional approximations of parameter matrices. This methodology was already developed in [8]. In this paper, we prove that these estimators are adaptive to the unknown matrix sparsity. More precisely, we exhibit a penalty for which the model selected by the penalized likelihood satisfies an oracle inequality. We support our theoretical result with simulation study and data analysis.

1. INTRODUCTION

The multivariate response regression model

$$Y = \beta X + E$$

postulates a linear relationship between Y , the $q \times n$ matrix containing q responses for n subjects, and X , the $p \times n$ matrix on p predictor variables. The term E is an $q \times n$ matrix with independent columns, $E_i \sim N_q(0, \Sigma)$ for all $i \in \{1, \dots, n\}$. The unknown $q \times p$ coefficient matrix β needs to be estimate. In a more general way, we could use finite mixture of linear models, which model the relationship between response and predictors, arising from different subpopulations: if the variable Y , conditionally to X , belongs to the cluster k , there exists β_k and Σ_k such that $Y = \beta_k X + E$, with $E \sim N_q(0, \Sigma_k)$.

If we use this model to deal with high-dimensional data, the number of variables can be quickly much larger than the sample size, because and predictors and response variables could be high dimensional. To solve this problem, we will have to reduce the parameter dimension.

One way to solve the dimension problem is to select relevant variables, in order to reduce the number of unknowns. Indeed, all the information should not be interesting for the clustering. In a density estimation way, we could cite Pan and Shen, in [16], who focus on mean variable selection, Zhou and Pan, in [20], who use the Lasso estimator to regularize Gaussian mixture model with general covariance matrices, Sun and Wang, in [18], who propose to regularize the k-means algorithm to deal with high-dimensional data, Guo et al, in [10], who propose a pairwise variable selection method.

In a regression framework, we could use the Lasso estimator, introduced by Tibshirani in [19], which is a sparse estimator, by penalizing the maximum likelihood estimator by the ℓ_1 -norm, which achieves the sparsity, as the ℓ_0 penalty, but leads also to a convex optimization. Because we work with the multivariate linear model, to deal with the matrix structure, we could prefer the group Lasso, variables grouped by columns, which selects columns rather than coefficients. This estimator was introduced by Zhou in [21] in the general case. If we select $|J|$ columns among the p possible, we have to estimate $|J|q$ coefficients rather than pq for A , which could be smaller than n if $|J|$ is smaller enough.

Another estimator which reduces the dimension, known in the linear model, is the low rank estimator: introduced by Izenman in [11], and more used the last decades, with among others Bunea et al. in [4] and Giraud in [9], the regression matrix could be estimated by matrix of rank r , $r < p \wedge q$. Then, we have to estimate $r(p + q - r)$ coefficients, which could be smaller than n .

In this paper, we have chosen to mix these two estimators to provide a sparse and low rank estimator in mixture models. This method was introduced by Bunea et al. in [4], in the case of linear model and known noise covariance matrix. They present different ways, more or less computational, with more or less good results in theory.

They get an oracle inequality, which say that, among a model collection, they are able to choose an estimator with good rank and good active variables. For this model, Ma et al. in [12] get a minimax lower bound, which precise that they attain nearly optimal rates of convergence adaptively for square Schatten norm losses.

In this paper, we consider finite mixture of K linear models in high-dimension. This model is studied in details by Städler et al. article for real response variable in [17], and by Devijver for multivariate response variable in [8]. We will estimate β_k for all $k \in \{1, \dots, K\}$ by a column sparse and low rank estimator. The Lasso estimator is used to select variables, whereas we refit the estimation by a low rank estimator, restricted on relevant variables. We propose a procedure which is based on a modeling that recasts variable selection, rank selection, and clustering problems into a model selection problem. This procedure is developed in [8], with methodology, computational issues, simulations and data analysis. In this paper, we focus on theoretical point of view, and developed simulations and data analysis for the low rank issue. Our procedure constructs a model collection, with models more or less sparse, and with rank vector more or less small. Among this collection, we have to select a model. We use the slope heuristic, which is a non-asymptotic criterion. In a theoretical way, in this paper, we get an oracle inequality for the collection constructed by our procedure, which makes a comparison performance between our model and the oracle for a specified penalty.

This result is an extension of the work of Bunea et al in [4], to mixture models and with covariance matrices $(\Sigma_k)_{1 \leq k \leq K}$ unknown. They ensure that mixing sparse estimator and low rank matrix could be interesting. Indeed, whereas we have to estimate $p \times q$ coefficients in each cluster for the regression matrix, we get only $r(|J| + q - r)$ unknown variables, which could be smaller than the number of observations n if $|J|$ and r are small. Even if the oracle inequality we get in this paper is an extension of Bunea et al. result, we use a really different way to prove it. Considering the model collection constructed, we want to select a model as good as possible. For that, we use the slope heuristic, which constructs a penalty, proportional to the dimension, and select the model minimizing the penalized log-likelihood. Theoretically, we construct also a penalty, proportional to the dimension (up to a logarithm term). We provide an oracle inequality which compare, up to a constant, the Jensen-Kullback-Leibler divergence of our model and the true model to the Kullback-Leibler divergence between the oracle and the true model. In estimation term, we do as well as possible. This oracle inequality is deduced from a general model selection theorem for maximum likelihood estimator of Massart in [13]. Controlling the bracketing entropy of models, we could prove the result. Remark that we work in a regression framework, then we rather use an extension of this theorem proved in Cohen and Le Pennec article [5], and that our model collection is random, constructed by the Lasso, then we rather use an extension of this theorem proved in [7]. To illustrate this procedure, in a computational way, we validate it on simulated dataset, and benchmark dataset. If the data have a low rank structure, we could easily find it with our methodology.

This paper is organized as follows. In the Section 2, we describe the finite mixture regression model used in this procedure, and the main step of the procedure. In the Section 3, we present the main result of this paper, which is an oracle inequality for the procedure proposed. Finally, in Section 4, we illustrate the procedure on simulated and benchmark dataset. Proof details of the oracle inequality are given in Appendix.

2. THE MODEL AND THE MODEL COLLECTION

We introduce our procedure of estimation by sparse and low rank matrix in the linear model, in Section 2.1, and extend it in Section 2.2 in mixture models.

2.1. Linear model. We consider the observations $((x_i, y_i)_{i=1, \dots, n})$ which realized random variables (X, Y) , satisfying the linear model

$$Y = \beta X + \epsilon$$

where $Y \in \mathbb{R}^q$ are the responses, $X \in \mathbb{R}^p$ are the regressors, $\beta \in \mathbb{R}^{p \times q}$ is an unknown matrix, and $\epsilon \in \mathbb{R}^q$ are random errors, $\epsilon \sim N_q(0, \Sigma)$, with $\Sigma \in \mathbb{S}_q^{++}$ a symmetric positive definite matrix. We will work in high-dimension, then $p \times q$ could be larger than the number of observations n .

We will construct an estimator which is sparse and low rank for β to cope with the high-dimension issue. Moreover, to reduce the covariance matrix dimension, we compute a diagonal estimator of Σ . The procedure we will propose could be explained into two steps. First, we estimate the active columns of β thanks to the Lasso estimator, for $\lambda > 0$, using the estimator

$$\hat{\beta}_\lambda^{\text{Lasso}} = \underset{\beta \in \mathbb{R}^{p \times q}}{\operatorname{argmin}} \{ \|Y - \beta X\|_2^2 + \lambda \|\beta\|_1 \};$$

where $\|\beta\|_1 = \sum_{j=1}^p \sum_{z=1}^q |\beta_{j,z}|$. We assume that the variance is unknown. From this approach, we could rescale the parameters by $\Sigma^{-1} = P^t P$ the Cholesky decomposition, and $\phi = P^{-1}\beta$, and then get

the estimates by, T_q denoting the set of triangular matrices of size q , for $\lambda > 0$,

$$(1) \quad (\hat{\phi}_\lambda^{\text{Lasso}}, \hat{P}_\lambda^{\text{Lasso}}) = \underset{(\phi, P) \in \mathbb{R}^{p \times q} \times T_q}{\operatorname{argmin}} \left\{ \|PY - \phi X\|_2^2 + \lambda \|\phi\|_1 \right\}.$$

This approach was considered in Städler et al. in [17], in scalar response case. This reparametrization is done in order to get a scale-invariant estimator and a convex minimization problem.

For $\lambda > 0$, from the Lasso estimator of $\hat{\phi}_\lambda^{\text{Lasso}}$, we could deduce the relevant columns.

Restricted to these relevant columns, in the second step of the procedure, we compute a low rank estimator of β , saying of rank at most r . Indeed, as explained in Giraud in [9], we restrict the maximum likelihood estimator to have a rank at most r , keeping only the r biggest singular values in the corresponding decomposition. We get an explicit formula.

This two steps procedure leads to an estimator of β which is sparse and has a low rank. We have also reduced the dimension into two ways. We refit the covariance matrix estimator by the maximum likelihood estimator.

This estimator is studied in Bunea et al. in [4], in method 3. Let extend it in mixture models.

2.2. Mixture model. We observe n independent couples $(\mathbf{x}, \mathbf{y}) = ((x_i, y_i)_{1 \leq i \leq n})$ of random variables (X, Y) , with $Y \in \mathbb{R}^q$ and $X \in \mathbb{R}^p$. The conditional density is assumed to be a multivariate Gaussian mixture regression model. If the observation i belongs to the cluster k , we assume that there exists $\beta_k \in \mathbb{R}^{p \times q}$, and $\Sigma_k \in \mathbb{S}_q^{++}$ such that $y_i = \beta_r x_i + \epsilon_i$ where $\epsilon_i \sim N_q(0, \Sigma_r)$.

Thus, the random response variable $Y \in \mathbb{R}^q$ depends on a set of explanatory variables, written $X \in \mathbb{R}^p$, through a mixture of linear regression-type model. Give more precisions on the assumptions.

- The variables Y_i are independent conditionally to X_i , for all $i = 1, \dots, n$;
- we let $Y_i | X_i = x_i \sim s_\xi(y | x_i) dy$, with

$$(2) \quad s_\xi(y | x) = \sum_{k=1}^K \frac{\pi_k}{(2\pi)^{\frac{q}{2}} \det(\Sigma_k)^{1/2}} \exp \left(-\frac{(y - \beta_k x)^t \Sigma_k^{-1} (y - \beta_k x)}{2} \right)$$

$$\xi = (\pi_1, \dots, \pi_K, \beta_1, \dots, \beta_K, \Sigma_1, \dots, \Sigma_K) \in (\Pi_K \times (\mathbb{R}^{q \times p})^K \times (\mathbb{S}_q^{++})^K)$$

$$\Pi_K = \left\{ (\pi_1, \dots, \pi_K); \pi_k > 0 \text{ for } k \in \{1, \dots, K\} \text{ and } \sum_{k=1}^K \pi_k = 1 \right\}$$

\mathbb{S}_q^{++} is the set of symmetric positive definite matrices on \mathbb{R}^q .

We want to estimate the conditional density function s_ξ from the observations. For all $k \in \{1, \dots, K\}$, β_k is the matrix of regression coefficients, and Σ_k is the covariance matrix in the mixture component k . The π_k s are the mixture proportions. In fact, for all $k \in \{1, \dots, K\}$, for all $z \in \{1, \dots, q\}$, $[\beta_k^t x]_z = \sum_{j=1}^p [\beta_k]_{j,z} x_j$ is the z th component of the mean of the mixture component k for the conditional density $s_\xi(\cdot | x)$.

We could introduce, for all $k \in \{1, \dots, K\}$,

$$\Sigma_k^{-1} = P_k^t P_k$$

$$\phi_k = P_k^{-1} \beta_k.$$

We could define an extension of the Lasso estimator,

$$(3) \quad (\hat{\phi}_\lambda^{\text{Lasso}}, \hat{P}_\lambda^{\text{Lasso}}) = \underset{(\phi, P) \in \mathbb{R}^{p \times q} \times T_q}{\operatorname{argmin}} \left\{ \|PY - \phi X\|_2^2 + \lambda \sum_{k=1}^K \pi_k \|\phi_k\|_1 \right\}.$$

Remark that the penalty take into account the mixture weight. We could also define a low rank estimator $(\hat{\beta}_\lambda^{\text{LR}}, \hat{P}_\lambda^{\text{LR}})$ restricted to relevant variables detected by the Lasso estimator, for the regularization parameter λ .

In a computational way, we will use two generalized EM algorithms, in order to deal with high dimensional data and get a sparse and low rank estimator. Give some details about those algorithms.

Initially, the EM algorithm was introduced by Dempster in [6]. It alternates two steps, an expectation step to cluster data, and a maximization step to update estimation. In our procedure, we want first to know which columns are relevant, then we want to compute (3). This algorithm was used and explained in [8]. It is a generalization of the EM algorithm, for the Lasso estimator, and in a regression context. From the estimate $\hat{\phi}_\lambda^{\text{Lasso}}$, we could deduce which columns are relevant. The second algorithm we use lead to determine ϕ_k on relevant columns, for all $k \in \{1, \dots, K\}$, with rank R_k . We alternate two steps,

E-step and M-step, until relative convergence of the parameters and of the likelihood. We restrict the dataset to relevant columns, and construct an estimator of size $|J| \times q$ rather than $p \times q$.

- E-step: compute for $k \in \{1, \dots, K\}$, $i \in \{1, \dots, n\}$, the expected value of the log likelihood function,

$$\begin{aligned}\hat{\gamma}_{i,k} &= E_{\theta^{(\text{ite})}}(Z_{i,k}|Y) \\ &= \frac{\varphi_k}{\sum_{l=1}^K \varphi_l}\end{aligned}$$

where

$$\varphi_l = \pi_l^{(\text{ite})} \det P_l^{(\text{ite})} \exp^{-\frac{1}{2} \left(P_l^{(\text{ite})} Y_i - X_i \Phi_l^{(\text{ite})} \right)^t \left(P_l^{(\text{ite})} Y_i - X_i \Phi_l^{(\text{ite})} \right)}$$

- M-step:
 - To get estimation in linear model, we assign each observation in its estimated cluster, by the MAP principle. We could apply this thanks to the E-step, which compute the a posteriori probability. Therefore, we say that y_i comes from component number $\underset{k \in \{1, \dots, K\}}{\operatorname{argmax}} \hat{\gamma}_{i,k}$.
 - Then, we can define $\tilde{\beta}_k^{(\text{ite})} = (\mathbf{x}_{|k}^t \mathbf{x}_{|k})^{-1} \mathbf{x}_{|k}^t \mathbf{y}_{|k}$, in which $x_{|k}$ and $y_{|k}$ are the sample restriction to the cluster k . We decompose $\tilde{\beta}_k^{(\text{ite})}$ in singular values such that $\tilde{\beta}_k^{(\text{ite})} = USV^t$ with $S = \operatorname{diag}(s_1, \dots, s_q)$ and $s_1 \geq s_2 \geq \dots \geq s_q$ the singular values. Then, the estimator $\hat{\beta}_k^{(\text{ite})}$ is defined by $\hat{\beta}_k^{(\text{ite})} = US_r V^t$ with $S_r = \operatorname{diag}(s_1, \dots, s_r, 0, \dots, 0)$.

To select the regularization parameter for the Lasso estimator and the ranks, we could use criterion as BIC or cross-validation. In practice, we prefer to construct a model collection with various active columns and various ranks, and select one as final step. To get various active columns, we construct a data-driven grid of regularization parameters, coming from EM algorithm formula. See [8] for more details. To get various ranks, we estimate parameters for different values of ranks, belonging to $\{r_{\min}, \dots, r_{\max}\}$.

From this procedure, we construct a model with K clusters, $|J|$ relevant columns and matrix of regression coefficients of ranks $R \in \mathbb{N}^K$, as described by the next model $S_{(K,J,R)}$.

$$(4) \quad S_{(K,J,R)} = \{y \in \mathbb{R}^q | x \in \mathbb{R}^p \mapsto s_{\xi^{(K,J,R)}}(y|x)\}$$

where

$$\begin{aligned}s_{\xi^{(K,J,R)}}(y|x) &= \sum_{k=1}^K \frac{\pi_k \det(P_k)}{(2\pi)^{q/2}} \exp\left(-\frac{1}{2}(y - \beta_k^{R(k)} x_i^{[J]})^t \Sigma_k^{-1} (y - \beta_k^{R(k)} x_i^{[J]})\right); \\ \xi^{(K,J,R)} &= (\pi_1, \dots, \pi_K, \beta_1^{R(1)}, \dots, \beta_K^{R(K)}, \Sigma_1, \dots, \Sigma_K) \in \Xi_K \\ \Xi_K &= \Pi_K \times \Psi_{(K,J,R)} \times (\mathbb{S}_q^{++})^K; \\ \Psi_{(K,J,R)} &= \{(\beta_1^{R(1)}, \dots, \beta_k^{R(K)}) \in (\mathbb{R}^{q \times |J|})^K | \operatorname{Rank}(\beta_k) = R(k)\}.\end{aligned}$$

Varying $K \in \mathcal{K} \subset \mathbb{N}^*$, $J \in \mathcal{J} \subset \mathcal{P}([1, p])$, and $R \in \mathcal{R} \subset \{1, \dots, p \wedge q\}^K$, we get a model collection with various number of components, active columns and matrix of regression coefficients.

Among this model collection, during the last step, a model has to be selected. As in Maugis and Michel in [14], and in Maugis and Meynet in [15], among others, a non asymptotic penalized criterion is used. The slope heuristic was introduced by Birgé and Massart in [3], and developed in practice by Maugis and Michel in [2] with the Capushe package. To use it in our context, we have to extend the theoretical result to determine the penalty shape in the high-dimensional context, with a random model collection, in a regression framework. The main result is described in the next section, whereas proof details are given in appendix.

3. ORACLE INEQUALITY

In a theoretical point of view, we want to ensure that the slope heuristic which construct a penalty will select a good model. We follow the approach developed by Birgé and Massart in [3] which consists of defining a non asymptotic penalized criterion, leading to an oracle inequality. In the context of regression, Cohen and Le Pennec, in [5], and Devijver in [7], propose a general model selection theorem for maximum likelihood estimation. The result we get is a theoretic penalty, for which the model selected is as good as the best one, according to the Kullback-Leibler loss.

3.1. Framework and model collection. Among the model collection constructed by the procedure developed in section 2.2, with various rank and various sparsity, we want to select an estimator which is close to the truth. The oracle is by definition the model belonging to the collection which minimizes the contrast with the true model. In practice, we do not have access to the true model, then we can not know the oracle. Nevertheless, the goal of the model selection step of our procedure is to be nearest to the oracle. In this section, we present an oracle inequality, which mean that if we have penalized the log-likelihood in a good way, we will select a model which is as good as the oracle, according to the Kullback-Leibler loss.

We consider the model collection defined by (4).

Because we work in high dimension, p could be big, and it will be time-consuming to test all the parts of $\{1, \dots, p\}$. We construct a sub-collection denoted by \mathcal{J}^L , which is constructed by the Lasso, which is also random. This step is explained in more details in [8].

Moreover, to get the oracle inequality, we assume that the parameters are bounded:

$$\begin{aligned} S_{(K,J,R)}^{\mathcal{B}} = & \{s_{(K,J,R)} \in S_{(K,J,R)} \mid \text{for all } k \in \{1, \dots, K\}, \\ & \Sigma_k = \text{diag}([\Sigma_k]_{1,1}, \dots, [\Sigma_k]_{q,q}), \\ & \text{for all } z \in \{1, \dots, q\}, a_{\Sigma} \leq [\Sigma_k]_{z,z} \leq A_{\Sigma}, \\ & \text{for all } k \in [1, K], \beta_k^{R(k)} = \sum_{l=1}^{R(k)} [\sigma_k]_l [u_k^t]_{\cdot,l} [v_k]_{l,\cdot}, \\ & \text{for all } l \in \{1, \dots, R(k)\}, [\sigma_k]_l < A_{\sigma}\}. \end{aligned}$$

Remark that this decomposition of β_k is the singular value decomposition, with $(\sigma_l)_{1 \leq l \leq R(k)}$ the singular values, and u_k and v_k unit vectors, for $k \in \{1, \dots, K\}$.

We also assume that covariates belong to an hypercube: without restrictions, we could assume that $X \in [0, 1]^p$.

Fixing \mathcal{K} the possible number of components, \mathcal{J}^L the active column set constructed by the Lasso, and \mathcal{R} the possible vector of ranks, we get a model collection

$$(5) \quad \bigcup_{K \in \mathcal{K}} \bigcup_{J \in \mathcal{J}} \bigcup_{R \in \mathcal{R}} S_{(K,J,R)}^{\mathcal{B}}.$$

3.2. Notations. Before enunciate the main theorem which leads to the oracle inequality, for the model collection (5) we need to define some metrics used to compare the conditional densities. First, the Kullback-Leibler divergence is defined by

$$KL_{\lambda}(s, t) = \int \log\left(\frac{s}{t}\right) s d\lambda$$

for s and t two densities, $s d\lambda$ absolutely continuous with respect to $t d\lambda$. To deal with regression data, for fixed covariates (x_1, \dots, x_n) , we define

$$KL_{\lambda}^{\otimes n}(s, t) = E \left(\frac{1}{n} \sum_{i=1}^n KL_{\lambda}(s(\cdot|x_i), t(\cdot|x_i)) \right)$$

for s and t two densities, $s d\lambda$ absolutely continuous with respect to $t d\lambda$.

We also define the Jensen-Kullback-Leibler divergence, first introduced in Cohen and Le Pennec [5], by

$$JKL_{\lambda, \rho}(s, t) = \frac{1}{\rho} KL_{\lambda}(s, (1 - \rho)s + \rho t)$$

for $\rho \in]0, 1[$, s and t two densities, $s d\lambda$ absolutely continuous with respect to $t d\lambda$. The tensorized one is defined by

$$JKL_{\lambda, \rho}^{\otimes n}(s, t) = E \left(\frac{1}{n} \sum_{i=1}^n JKL_{\lambda, \rho}(s(\cdot|x_i), t(\cdot|x_i)) \right).$$

Note that those divergences are not metrics, not satisfying the triangular inequality and no symmetric, but are also wild used in statistics to compare two densities.

3.3. Oracle inequality. Let state the main theorem.

Theorem 3.1. Assume that we observe $((x_i, y_i)_{1 \leq i \leq n})$ with unknown conditional density s_0 . Let $\mathcal{M} = \mathcal{K} \times \mathcal{J} \times \mathcal{R}$ and $\mathcal{M}^L = \mathcal{K} \times \mathcal{J}^L \times \mathcal{R}$, where \mathcal{J}^L is constructed by the Lasso estimator. Let $\bar{s}_{(K,J,R)} \in S_{(K,J,R)}^{\mathcal{B}}$ such that, for $\delta_{KL} > 0$,

$$KL_{\lambda}^{\otimes n}(s_0, \bar{s}_{(K,J,R)}) \leq \inf_{t \in S_{(K,J,R)}^{\mathcal{B}}} KL_{\lambda}^{\otimes n}(s_0, t) + \frac{\delta_{KL}}{n}$$

and such that there exists $\tau > 0$ such that

$$(6) \quad \bar{s}_{(K,J,R)} \geq e^{-\tau} s_0.$$

Consider the collection $\{\hat{s}_{(K,J,R)}\}_{(K,J,R) \in \mathcal{M}}$ of rank constrained log-likelihood minimizer in $S_{(K,J,R)}$, satisfying

$$\hat{s}_{(K,J,R)} = \operatorname{argmin}_{s_{(K,J,R)} \in S_{(K,J,R)}^{\mathcal{B}}} \left\{ -\frac{1}{n} \sum_{i=1}^n \log(s_{(K,J,R)}(y_i|x_i)) \right\}.$$

Denote by $D_{(K,J,R)}$ the dimension of the model $S_{(K,J,R)}^{\mathcal{B}}$. Let $\operatorname{pen} : \mathcal{M} \rightarrow \mathbb{R}_+$ defined by, for all $(K, J, R) \in \mathcal{M}$,

$$\begin{aligned} \operatorname{pen}(K, J, R) \geq & \kappa \frac{D_{(K,J,R)}}{n} \{ [2B^2(A_{\beta}, A_{\Sigma}, a_{\sigma}, q) \\ & - \log \left(\frac{D_{(K,J,R)}}{n} B^2(A_{\beta}, A_{\Sigma}, a_{\sigma}, q) \wedge 1 \right) \\ & + \log \left(\frac{4epq}{D_{(K,J,R)} - q^2 \wedge pq} + R \right) \} \end{aligned}$$

for $\kappa > 0$ an absolute constant.

Then, the estimator $\hat{s}_{(\hat{K}, \hat{J}, \hat{R})}$, with

$$(\hat{K}, \hat{J}, \hat{R}) = \operatorname{argmin}_{(K,J,R) \in \mathcal{M}^L} \left\{ -\frac{1}{n} \sum_{i=1}^n \log(\hat{s}_{(K,J,R)}(y_i|x_i)) + \operatorname{pen}(K, J, R) \right\},$$

$$\sum_{i=1}^n -\log(\hat{s}_{(\hat{K}, \hat{J}, \hat{R})}(y_i|x_i)) + \operatorname{pen}(\hat{K}, \hat{J}, \hat{R})$$

$$\leq \inf_{(K,J,R) \in \mathcal{M}} \left(\sum_{i=1}^n -\log(\hat{s}_{(K,J,R)}(y_i|x_i)) + \operatorname{pen}(K, J, R) \right) + \eta'$$

satisfies

$$(7) \quad \begin{aligned} & E(JKL_{\rho, \lambda}^{\otimes n}(s_0, \hat{s}_{(\hat{K}, \hat{J}, \hat{R})})) \\ & \leq CE \left(\inf_{(K,J,R) \in \mathcal{M}^L} \left(\inf_{t \in S_{(K,J,R)}} KL_{\lambda}^{\otimes n}(s_0, t) + \frac{\operatorname{pen}(K, J, R)}{n} \right) + \frac{\Sigma^2}{n} \right) \end{aligned}$$

for $C > 0$.

The proof of the theorem 3.1 is given in Section 5. Note that condition (6) leads to control the random model collection. The mixture parameters are bounded in order to construct brackets over $\mathcal{S}_{(K,J,R)}$, and thus to upper bound the entropy number. The inequality (7) we obtain is not exactly an oracle inequality, since the Jensen-Kullback-Leibler risk is upper bounded by the Kullback-Leibler bias.

Because we are looking on a small random sub-collection of models, our estimator $\hat{s}_{(K,J,R)}$ is attainable in practice. Moreover, it is a non-asymptotic result, which allows us to study cases for which p increases with n .

Note that we use the Jensen-Kullback-Leibler divergence rather than the Kullback-Leibler divergence, because it is bounded. This boundedness turns out to be crucial to control the loss of the penalized maximum likelihood estimator under mild assumptions on the complexity of the model and their collection.

We could compare our bound with the one of Bunea et al, in [4], who computed a procedure similar to ours, in a linear model. According to consistent group selection for the group Lasso, they get adaptivity of the estimator to an optimal rate, and their estimators perform the bias variance trade-off among

TABLE 1. Performances of our procedure. Mean number $\{K, \hat{R}, M, FA, ARI\}$ of the Kullback-Leibler divergence between the model selected and the true model, the estimated rank of the model selected in each cluster, the missed variables, the false relevant variables, and the ARI, over 20 simulations.

	KL	\hat{R}	M	FA	ARI
$p > n$	19.03	[2.8,3]	0	20	0.95
$p < n$	3.28	[3,3]	0	0.6	0.99

all reduced rank estimators. Nevertheless, their results are obtained according to some assumption, for instance the mutual coherence on $X^t X$, which postulates that the off-diagonal elements have to be small. Some assumptions on the design are required, whereas our result just need to deal with bounded parameters and bounded covariates.

4. NUMERICAL STUDIES

We will illustrate our procedure with simulations and real datasets, to highlight advantages of our method. We adapt the simulations part of Bunea et al. article, [4]. Indeed, we work in the same way, to get sparse and low rank estimator. Nevertheless, we add the mixture to be consistent with our clustering method, and to have more flexibility.

4.1. Simulations. To illustrate our procedure, we use simulations adapted from the article of Bunea [4], extended for mixture models.

The design matrix X has iid rows X_i from a multivariate normal distribution $N(0, \Sigma)$ with $\Sigma = \rho I$, $\rho > 0$. We consider a mixture with 2 components. Depending on the cluster, the coefficient matrix β_k has the form

$$\beta_k = \begin{bmatrix} b_k B^0 & b_k B^1 \\ 0 & 0 \end{bmatrix}$$

for $k \in \{1, 2\}$, with B^0 a $J \times R(k)$ matrix and B^1 a $R(k) \times q$ matrix. All entries in B^0 and B^1 are iid $N(0, 1)$. The noise matrix E has independent $N(0, 1)$ entries. Let E_i denote its i th row.

The proportion vector π is defined by $\pi = [\frac{1}{2}, \frac{1}{2}]$, all clusters having the same probability.

Each row Y_i in Y is then generated as, if the observation i belongs to the cluster k , $Y_i = \beta_k X_i + E_i$, for $1 \leq i \leq n$. This setup contains many noise features, but the relevant one lie in a low-dimensional subspace. We report two settings:

- $p > n$: $n = 50, |J| = 6, p = 100, q = 10, R = [3, 3], \rho = 0.1, b = [3, -3]$.
- $p < n$: $n = 200, |J| = 6, p = 10, q = 10, R = [3, 3], \rho = 0.01, b = [3, -3]$.

The current setups show that variable selection, without taking the rank information into consideration, may be suboptimal, even if the correlations between predictors are low. Each model was simulated 20 times, and table 1 summarizes our findings. We evaluate the prediction accuracy of each estimator $\hat{\beta}$ by the mean squared error (MSE) using a test sample at each run. We also report the median rank estimate (denoted by \hat{R}) over all runs, rates of non included true variables (denoted by M for misses) and the rates of incorrectly included variables (FA for false actives). Ideally, we are looking for a model with low MSE , low M and low FA .

We can draw the following conclusions from table 1. When we work in low-dimensional framework, we get very good results. Even if we could use any estimator, because we do not have dimension problem, with our procedure we get the matrix structure involved by the model. We get almost exact clustering, and the Kullback-Leibler divergence between the model we construct and the true model is really low. In case of high-dimensional data, when p is larger than n , we get also good results. We find the good structure, selecting the relevant variables (our model will have false relevant variables, but no missed variables), and selecting the true ranks. We could remark that false relevant variables have low values. Comparing to another procedure which will not reduce the rank, we will perform the dimension reduction.

4.2. Application. In this section, we apply our procedure to real data set. The Norwegian paper quality data were obtained from a controlled experiment that was carried out at a paper factory in Norway to uncover the effect of three control variables X_1, X_2, X_3 on the quality of the paper which was measured by 13 response variables. Each of the control variables X_i takes values in $\{-1, 0, 1\}$. To account for possible interactions and nonlinear effects, second order terms were added to the set of

predictors, yielding $X_1, X_2, X_3, X_1^2, X_2^2, X_3^2, X_1X_2, X_1X_3, X_2X_3$, and the intercept term. There were 29 observations with no missing values made on all response and predictor variables. The Box Behnken design of the experiment and the resulting data are described in Aldrin [1] and Izenman [11]. Moreover, Bunea et al in [4] also study this dataset. We always center the responses and the predictors. The dataset clearly indicates that dimension reduction is possible, making it a typical application for reduced rank regression methods. Moreover, our method will exhibit different classes among this sample.

We construct a model collection varying the number of clusters in $\mathcal{K} = \{2, \dots, 5\}$. We select a model with 2 classes. We select all variables except X_1X_2 and X_2X_3 , which is consistent with comments of Bunea et al. In the two classes, we get two mean matrices, with ranks equal to 2 and 4. One cluster describes the mean compartment (with rank equal to 2), whereas the other cluster contains values more different.

5. APPENDICES

In those appendices, we present the details of the proof of the theorem 3.1. This derives from a general model selection theorem, enunciated in section 5.1, and proved in the paper [7]. Then, the proof of the theorem 3.1 could be summarized by satisfying assumptions 1, 2 and 3 described in section 5.1.

5.1. A general oracle inequality for model selection. Model selection appears with the AIC criterion and BIC criterion. In a non-asymptotic way, a theory was developed by Birgé and Massart in [3]. With some assumptions that we will detail here, we get an oracle inequality for the maximum likelihood estimator among a model collection. Le Pennec and Cohen, in [5], generalize this theorem in regression framework. We have to use a generalization of this theorem because we consider a random collection of models.

Let us consider a model collection $(S_m)_{m \in \mathcal{M}}$.

Before enunciate the general theorem, begin by talking about the assumptions. First, we impose a structural assumption. It is a bracketing entropy condition on the model S_m with respect to the Hellinger divergence

$$d_H^{2\otimes n}(s, t) = E \left[\frac{1}{n} \sum_{i=1}^n d_H^2(s(\cdot|x_i), t(\cdot|x_i)) \right].$$

A bracket $[t^-, t^+]$ is a pair of functions such that for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$, $t^-(y, x) \leq s(y|x) \leq t^+(y, x)$. The bracketing entropy $H_{[\cdot]}(\delta, S, d_H^{\otimes n})$ of a set S is defined as the logarithm of the minimum number of brackets $[t^-, t^+]$ of width $d_H^{\otimes n}(t^-, t^+)$ smaller than δ such that every functions of S belong to one of these brackets.

Assumption 1 (H_m). *There is a non-decreasing function ϕ_m such that $\delta \mapsto \frac{1}{\delta} \phi_m(\delta)$ is non-increasing on $(0, +\infty)$ and for every $\sigma \in \mathbb{R}^+$ and every $s_m \in S_m$,*

$$\int_0^\sigma \sqrt{H_{[\cdot]}(\delta, S_m(s_m, \sigma), d_H^{\otimes n})} d\delta \leq \phi_m(\sigma);$$

where $S_m(s_m, \sigma) = \{t \in S_m, d_H^{\otimes n}(t, s_m) \leq \sigma\}$. The model complexity \mathcal{D}_m is then defined as $n\sigma_m^2$ with σ_m the unique root of

$$(8) \quad \frac{1}{\sigma} \phi_m(\sigma) = \sqrt{n}\sigma.$$

Denote that the model complexity depends on the bracketing entropies not of the global models S_m but of the ones of smaller localized sets. This is a weaker assumption.

For technical reason, a separability assumption is also required.

Assumption 2 (Sep_m). *There exists a countable subset S'_m of S_m and a set \mathcal{Y}'_m with $\lambda(\mathcal{Y} \setminus \mathcal{Y}'_m) = 0$, for λ the Lebesgue measure, such that for every $t \in S_m$, there exists a sequence $(t_k)_{k \geq 1}$ of elements of S'_m such that for every x and every $y \in \mathcal{Y}'_m$, $\log(t_k(y|x))$ goes to $\log(t(y|x))$ as k goes to infinity.*

We also need an information theory type assumption on our model collection. We assume the existence of a Kraft-type inequality for the collection.

Assumption 3 (K). *There is a family $(x_m)_{m \in \mathcal{M}}$ of non-negative numbers such that*

$$\sum_{m \in \mathcal{M}} e^{-x_m} \leq \Sigma < +\infty.$$

Then, we could write our main global theorem to get an oracle inequality in regression framework, with a random collection of models.

Theorem 5.1. *Assume we observe $((x_i, y_i)_{1 \leq i \leq n})$ with unknown conditional density s_0 . Let $\mathcal{S} = (S_m)_{m \in \mathcal{M}}$ be at most countable collection of conditional density sets. Let assumption (K) holds while assumptions (H_m) and (Sep_m) hold for every models $S_m \in \mathcal{S}$. Let $\delta_{KL} > 0$, and $\bar{s}_m \in S_m$ such that*

$$KL_{\lambda}^{\otimes n}(s_0, \bar{s}_m) \leq \inf_{t \in S_m} KL_{\lambda}^{\otimes n}(s_0, t) + \frac{\delta_{KL}}{n};$$

and let $\tau > 0$ such that

$$(9) \quad \bar{s}_m \geq e^{-\tau} s_0.$$

Introduce $(S_m)_{m \in \hat{\mathcal{M}}}$ some random sub-collection of $(S_m)_{m \in \mathcal{M}}$. Consider the collection $(\hat{s}_m)_{m \in \hat{\mathcal{M}}}$ of η -log-likelihood minimizer in S_m , satisfying

$$\sum_{i=1}^n -\log(\hat{s}_m(y_i|x_i)) \leq \inf_{s_m \in S_m} \left(\sum_{i=1}^n -\log(s_m(y_i|x_i)) \right) + \eta.$$

Then, for any $\rho \in (0, 1)$ and any $C_1 > 1$, there are two constants κ_0 and C_2 depending only on ρ and C_1 such that, as soon as for every index $m \in \mathcal{M}$,

$$(10) \quad pen(m) \geq \kappa(\mathcal{D}_m + (1 \vee \tau)x_m)$$

with $\kappa > \kappa_0$, and where the model complexity \mathcal{D}_m is defined in (8), the penalized likelihood estimate $\hat{s}_{\hat{m}}$ with $\hat{m} \in \hat{\mathcal{M}}$ such that

$$\sum_{i=1}^n -\log(\hat{s}_{\hat{m}}(y_i|x_i)) + pen(\hat{m}) \leq \inf_{m \in \hat{\mathcal{M}}} \left(\sum_{i=1}^n -\log(\hat{s}_m(y_i|x_i)) + pen(m) \right) + \eta'$$

satisfies

$$(11) \quad E(JKL_{\rho, \lambda}^{\otimes n}(s_0, \hat{s}_{\hat{m}})) \leq C_1 E \left(\inf_{m \in \hat{\mathcal{M}}} \inf_{t \in S_m} KL_{\lambda}^{\otimes n}(s_0, t) + 2 \frac{pen(m)}{n} \right)$$

$$(12) \quad + C_2 (1 \vee \tau) \frac{\Sigma^2}{n} + \frac{\eta' + \eta}{n}.$$

Remark 5.2. *We get that, among a random model collection, we are able to choose one which is as good as the oracle, up to a constant C_1 , and some additive terms being around $\frac{1}{n}$. This result is non-asymptotic, and gives a theoretic penalty to select this model.*

Remark 5.3. *The proof of this theorem is detailed in [7]. Nevertheless, we could give the main ideas to understand the assumptions. From assumptions 1 and 2, we could use maximal inequalities which lead to, except on a set of probability less than $e^{-x_{m'} - x}$, for all x , a control of the ratio of the centered empirical process of $\log(\hat{s}_{m'})$ over the Hellinger distance between s_0 and $\hat{s}_{m'}$, this control being around $\frac{1}{n}$. Thanks to Bernstein inequality, satisfied according to the inequality (9), and thanks to the assumption 3, we get the oracle inequality.*

Now, to prove theorem (3.1), we have to satisfy assumptions (1), (3), and (2).

5.2. Assumption (H_m) .

5.2.1. *Decomposition.* As done in Cohen and Le Pennec [5], we could decompose the entropy by

$$H_{[\cdot]}(\epsilon, \mathcal{S}_{(K, J, R)}^B, d_H^{\otimes n}) \leq H_{[\cdot]}(\epsilon, \Pi_K, d_H^{\otimes n}) + kH_{[\cdot]}(\epsilon, \mathcal{F}_{(J, R)}, d_H^{\otimes n})$$

where

$$\mathcal{S}_{(K,J,R)}^{\beta} = \left\{ \begin{array}{l} y \in \mathbb{R}^q | x \in \mathbb{R}^p \mapsto s_{\theta}(y|x) = \sum_{k=1}^K \pi_k \Phi(y | (\beta_k^{R(k)} x)_{|J}, \Sigma_k) \\ \theta = \left\{ \pi_1, \dots, \pi_K, \beta_1^{R(1)}, \dots, \beta_K^{R(K)}, \Sigma_1, \dots, \Sigma_K \right\} \in \Theta_K \\ \Theta_k = \Pi_K \times \tilde{\Psi}_{(K,J,R)} \times ([a_{\Sigma}, A_{\Sigma}]_{+*}^q)^K \end{array} \right\}$$

$$\Psi_{(K,J,R)} = \{(\beta_1^{R(1)}, \dots, \beta_K^{R(K)}) \in [-A_{\beta}, A_{\beta}]^K | \text{Rank}(\beta_k) = R(k)\}$$

$$\Pi_k = \left\{ (\pi_1, \dots, \pi_K) \in (0, 1)^K; \sum_{k=1}^K \pi_k = 1 \right\}$$

$$\mathcal{F}_{(J,R)} = \left\{ \Phi(\cdot | (\beta^R X)_{|J}, \Sigma); \beta \in [a_{\beta}, A_{\beta}]^{|J|}, \text{rank}(\beta^R) = R, \right. \\ \left. \Sigma = \text{diag}(\Sigma_1^2, \dots, \Sigma_q^2) \in [a_{\Sigma}^2, A_{\Sigma}^2]^q \right\}$$

Φ the Gaussian density.

For the proportions, we get that

$$H_{[\cdot]}(\epsilon, \Pi_K, d_H^{\otimes n}) \leq \log \left(K(2\pi e)^{K/2} \left(\frac{3}{\epsilon} \right)^{K-1} \right).$$

Looking after the Gaussian entropy.

5.2.2. *For the Gaussian.* We want to bound the integrated entropy. For that, first we have to construct some brackets to recover S_m . Fix $f \in S_m$. We are looking for functions t^- and t^+ such that $t^- \leq f \leq t^+$. Because f is a Gaussian, t^- and t^+ are dilatations of Gaussians. We then have to determine the mean, the variance and the dilatation coefficient of t^- and t^+ . We need the both following lemmas to construct these brackets.

Lemma 5.4. *Let $\Phi(\cdot | \mu_1, \Sigma_1)$ and $\Phi(\cdot | \mu_2, \Sigma_2)$ be two Gaussian densities. If their variance matrices are assumed to be diagonal, with $\Sigma_a = \text{diag}([S_a^2]_1, \dots, [S_a^2]_q)$ for $a \in \{1, 2\}$, such that $[S_2]_z^2 > [S_1]_z^2 > 0$ for all $z \in \{1, \dots, q\}$, then for all $x \in \mathbb{R}^q$,*

$$\frac{\Phi(x | \mu_1, \Sigma_1)}{\Phi(x | \mu_2, \Sigma_2)} \leq \prod_{z=1}^q \frac{\sqrt{[\Sigma_2]_z}}{\sqrt{[\Sigma_1]_z}} \exp^{\frac{1}{2}(\mu_1 - \mu_2)^t \text{diag}\left(\frac{1}{[\Sigma_2]_1 - [\Sigma_1]_1}, \dots, \frac{1}{[\Sigma_2]_q - [\Sigma_1]_q}\right)(\mu_1 - \mu_2)}.$$

Lemma 5.5. *The Hellinger distance of two Gaussian densities with diagonal variance matrices is given by the following expression:*

$$d_H^2(\Phi(\cdot | \mu_1, \Sigma_1), \Phi(\cdot | \mu_2, \Sigma_2)) = 2 - 2 \left(\prod_{q_1=1}^q \frac{2[\Sigma_1^2]_{q_1} [\Sigma_2^2]_{q_1}}{[\Sigma_1^2]_{q_1} + [\Sigma_2^2]_{q_1}} \right)^{1/2} \\ \times \exp \left\{ -\frac{1}{4}(\mu_1 - \mu_2)^t \text{diag} \left(\left(\frac{1}{[\Sigma_1^2]_{q_1} + [\Sigma_2^2]_{q_1}} \right)_{q_1 \in \{1, \dots, q\}} \right) (\mu_1 - \mu_2) \right\}$$

To get an ϵ -bracket for the densities, we have to construct a δ -net for the variance and the mean, δ to be precised later.

• **Step 1: construction of a net for the variance**

Let $\epsilon \in]0, 1]$, and $\delta = \frac{\epsilon}{\sqrt{2q}}$. Let $b_j^2 = (1 + \delta)^{1 - \frac{j}{2}} A_{\Sigma}^2$. For $2 \leq j \leq N$, we have $[a_{\Sigma}, A_{\Sigma}] = [b_N, b_{N-1}] \cup \dots \cup [b_3, b_2]$, where N is chosen to recover everything. We want that

$$a_{\Sigma}^2 = (1 + \delta)^{1 - N/2} A_{\Sigma}^2 \\ \Leftrightarrow 2 \log \frac{a_{\Sigma}}{A_{\Sigma}} = \left(1 - \frac{N}{2} \right) \log(1 + \delta) \\ \Leftrightarrow N = \frac{4 \log\left(\frac{A_{\Sigma}}{a_{\Sigma}} \sqrt{1 + \delta}\right)}{\log(1 + \delta)}.$$

We want N to be an integer, then $N = \left\lceil \frac{4 \log\left(\frac{A_{\Sigma}}{a_{\Sigma}} \sqrt{1 + \delta}\right)}{\log(1 + \delta)} \right\rceil$. We get a regular net for the variance.

We could let $B = \text{diag}(b_{i(1)}^2, \dots, b_{i(q)}^2)$, close to Σ (and deterministic, independent of the values of Σ), where i is a permutation such that $b_{i(z)+1} \leq \Sigma_{z,z} \leq b_{i(z)}$ for all $z \in \{1, \dots, q\}$.

• **Step 2: construction of a net for the mean vectors**

We use the singular decomposition of β , $\beta = \sum_{l=1}^R \sigma_l u_l^t v_l$, with $(\sigma_l)_{1 \leq l \leq R}$ the singular values, and $(u_l)_{1 \leq l \leq R}$ and $(v_l)_{1 \leq l \leq R}$ unit vectors. Those vectors are also bounded.

We are looking for t^- and t^+ such that $d_H(t^-, t^+) \leq \epsilon$, and $t^- \leq f \leq t^+$. We will use a dilatation of a Gaussian to construct such an ϵ -bracket of Φ .

We let

$$\begin{aligned} t^-(x, y) &= (1 + \delta)^{-(pqr+3q/4)} \Phi(y | \nu_{J,R} x, (1 + \delta)^{-1/4} B^{l+1}) \\ t^+(x, y) &= (1 + \delta)^{pqr+3q/4} \Phi(y | \nu_{J,R} x, (1 + \delta) B^l) \end{aligned}$$

where B^l and B^{l+1} are constructed such that, for all $z \in \{1, \dots, q\}$, $B_{z,z}^{l+1} \leq \Sigma_{z,z} \leq B_{z,z}^l$ (see step 1).

The means $\nu_{J,R} \in \mathbb{R}^{p \times q}$ will be specified later. Just remark that J is the set of the relevant columns, and R the rank of $\nu_{J,R}$: we will decompose $\nu_{J,R} = \sum_{l=1}^R \tilde{\sigma}_l \tilde{u}_l^t \tilde{v}_l$, $\tilde{u} \in \mathbb{R}^{|J| \times R}$, and $\tilde{v} \in \mathbb{R}^{q \times R}$.

We get

$$t^-(x, y) \leq f(y|x) \leq t^+(x, y)$$

if we have

$$\|\beta x - \nu_{J,R} x\|_2^2 \leq pqr \frac{\delta^2}{2} a_\Sigma^2 (1 - 2^{-1/4}).$$

Remark that $\|\beta x - \nu_{J,R} x\|_2^2 \leq p \|\beta - \nu_{J,R}\|_2^2 \|x\|_\infty$. We need then

$$(13) \quad \|\beta - \nu_{J,R}\|_2^2 \leq pqr \frac{\delta^2}{2} a_\Sigma^2 (1 - 2^{-1/4})$$

According to [7], $d_H(t^-, t^+) \leq 2(p^2 q R + 3q/4)^2 \delta^2$, then, with

$$\delta = \frac{\epsilon}{\sqrt{2}(p^2 q R + 3/4q)}$$

we get the wanted bound.

Now, explain how to construct $\nu_{J,R}$ to get (13).

$$\begin{aligned} \|\beta - \nu_{J,R}\|_2^2 &= \sum_{j=1}^p \sum_{z=1}^q \sum_{l=1}^R |\sigma_l u_{l,j} v_{l,z} - \tilde{\sigma}_l \tilde{u}_{l,j} \tilde{v}_{l,z}|^2 \\ &= \sum_{j=1}^p \sum_{z=1}^q \sum_{l=1}^R \left| \sigma_l - \tilde{\sigma}_l |u_{j,l} v_{z,l}| - \tilde{\sigma}_l |\tilde{u}_{l,j} - u_{l,j}| |\tilde{v}_{z,l}| - \tilde{\sigma}_l u_{l,j} |v_{l,z} - \tilde{v}_{l,z}| \right|^2 \\ &\leq \sum_{j=1}^p \sum_{z=1}^q \sum_{l=1}^R \left(|\sigma_l - \tilde{\sigma}_l| + A_\sigma |\tilde{u}_{l,j} - u_{l,j}| + A_\sigma |v_{l,z} - \tilde{v}_{l,z}| \right)^2 \\ &\leq 2pqr \left(\max_l |\sigma_l - \tilde{\sigma}_l|^2 + A_\sigma \max_{l,j} |\tilde{u}_{l,j} - u_{l,j}|^2 + A_\sigma \max_{l,z} |\tilde{v}_{l,z} - v_{l,z}|^2 \right) \end{aligned}$$

We need $\|\beta - \nu_{J,R}\|_2^2 \leq pqr \frac{\delta^2}{2} a_\Sigma^2 (1 - 2^{-1/4})$.

If we choose $\tilde{\sigma}_l$ such that

$$|\sigma_l - \tilde{\sigma}_l| \leq \frac{\delta}{\sqrt{12}} a_\Sigma \sqrt{1 - 2^{-1/4}},$$

and $\tilde{u}_{l,j}$ such that

$$|u_{l,j} - \tilde{u}_{l,j}| \leq \frac{\delta}{\sqrt{12} A_\sigma} a_\Sigma \sqrt{1 - 2^{-1/4}},$$

and $\tilde{v}_{l,z}$ such that

$$|v_{l,z} - \tilde{v}_{l,z}| \leq \frac{\delta}{\sqrt{12} A_\sigma} a_\Sigma \sqrt{1 - 2^{-1/4}},$$

then it works.

To get this, we let, for $\lfloor \cdot \rfloor$ begin the floor function,

$$S = \mathbb{Z} \cap \left[0, \left\lfloor \frac{A_\sigma}{\frac{\delta}{\sqrt{12}} a_\Sigma \sqrt{1 - 2^{-1/4}}} \right\rfloor \right]$$

and

$$\tilde{\sigma}_l = \operatorname{argmin}_{\varsigma \in S} \left| \sigma_l - \frac{\delta}{\sqrt{12}} a_\Sigma \sqrt{1 - 2^{-1/4} \varsigma} \right|,$$

and

$$U = \mathbb{Z} \cap \left[0, \left\lfloor \frac{A_\sigma}{\frac{\delta}{\sqrt{12} A_\sigma} a_\Sigma \sqrt{1 - 2^{-1/4}}} \right\rfloor \right]$$

and

$$\tilde{u}_{l,j} = \operatorname{argmin}_{\mu \in U} \left| u_{l,j} - \frac{\delta}{\sqrt{12} A_\sigma} a_\Sigma \sqrt{1 - 2^{-1/4} \mu} \right|,$$

and

$$V = \mathbb{Z} \cap \left[0, \left\lfloor \frac{A_\sigma}{\frac{\delta}{\sqrt{12} A_\sigma} a_\Sigma \sqrt{1 - 2^{-1/4}}} \right\rfloor \right]$$

and

$$\tilde{v}_{l,j} = \operatorname{argmin}_{\nu \in V} \left| v_{l,j} - \frac{\delta}{\sqrt{12} A_\sigma} a_\Sigma \sqrt{1 - 2^{-1/4} \nu} \right|.$$

Remark that we just need to determine vectors $((\tilde{u}_{l,j})_{1 \leq j \leq J-l})_{1 \leq l \leq R}$ and $((\tilde{v}_{l,z})_{1 \leq z \leq q-l})_{1 \leq l \leq R}$ because those vectors are unit.

Then, we let

$$\begin{aligned} \forall j \in J^c, \forall z \in [1, p], (\nu_{J,R})_{j,z} &= 0 \\ \forall j \in J, \forall z \in [1, p], (\nu_{J,R})_{j,z} &= \sum_{l=1}^R \tilde{\sigma}_l \tilde{u}_{l,j} \tilde{v}_{l,z} \end{aligned}$$

We have defined our bracket.

We want to control the entropy.

$$\begin{aligned} & |\mathcal{B}_\epsilon(\mathcal{F}_{(K,J,R)})| \\ & \leq N \prod_{k=1}^K \left(\frac{A_\sigma}{\frac{\delta}{\sqrt{12} A_\sigma} a_\Sigma \sqrt{1 - 2^{-1/4}}} \right)^{R(k)} \left(\frac{A_\sigma}{\frac{\delta}{\sqrt{12} A_\sigma} a_\Sigma \sqrt{1 - 2^{-1/4}}} \right)^{R(k)(|J|+q-R(k)-1)} \\ & \leq N \left(\frac{A_\sigma \sqrt{12}}{\delta a_\Sigma \sqrt{1 - 2^{-1/4}}} \right)^{\sum_{k=1}^K R(k)(|J|+q-R(k))} \\ & \leq 4 \left(\frac{A_\Sigma}{a_\Sigma} + \frac{1}{2} \right) \left(\frac{A_\sigma \sqrt{12}}{a_\Sigma \sqrt{1 - 2^{-1/4}}} \right)^{\sum_{k=1}^K R(k)(|J|+q-R(k))} \delta^{-\sum_{k=1}^K R(k)(|J|+q-R(k))+1} \end{aligned}$$

Then,

$$\mathcal{N}_{[\cdot]}(\epsilon, \mathcal{S}_{(K,J,R)}^{\mathcal{B}}, d_H) \leq C \left(\frac{1}{\epsilon} \right)^{D_{(K,J,R)}}$$

where

$$(14) \quad \begin{aligned} C &= 4 \left(\frac{A_\Sigma}{a_\Sigma} + \frac{1}{2} \right) \left(\frac{A_\sigma \sqrt{12}}{a_\Sigma \sqrt{1 - 2^{-1/4}}} \right)^{\sum_{k=1}^K R(k)(|J|+q-R(k))} \\ & \quad \times \left(\sqrt{2}(p^2 q R + 3/4 q) \right)^{\sum_{k=1}^K R(k)(|J|+q-R(k))} \end{aligned}$$

and $D_{(K,J,R)} = \sum_{k=1}^K R(k)(|J| + q - R(k))$.

5.2.3. *For the mixture.* We have to determine $\phi_{(K,J,R)}$ such that

$$(15) \quad \int_0^\sigma \sqrt{H_{[\cdot]}(\epsilon, S_{(K,J,R)}(s_{(K,J,R)}), \sigma), d_H^{\otimes n}} d\epsilon \leq \phi_{(K,J,R)}(\sigma).$$

Let compute the integral.

$$\begin{aligned} & \int_0^\sigma \sqrt{H_{[\cdot]}(\epsilon, S_{(K,J,R)}(s_{(K,J,R)}), \sigma), d_H^{\otimes n}} d\epsilon \\ & \leq \sigma \sqrt{\log(C_2)} + \sqrt{D_{(K,J,R)}} \int_0^\sigma \sqrt{\log\left(\frac{1}{\epsilon}\right)} d\epsilon \\ & \leq \sqrt{D_{(K,J,R)}} \sigma \left[\sqrt{\pi} + \sqrt{\frac{\log(C_2)}{D_{(K,J,R)}}} + \sqrt{\log\left(\frac{1}{\sigma \wedge 1}\right)} \right] \end{aligned}$$

with, according to (14),

$$\begin{aligned} \log(C_2) &= \log(4) + \frac{k}{2} \log(2\pi e) + (K-1) \log(3) + \log\left(\frac{A_\Sigma}{a_\Sigma} + \frac{1}{2}\right) \\ &+ D_{(K,J,R)} \log\left(\frac{A_\sigma \sqrt{12}}{a_\Sigma \sqrt{1-2^{-1/4}}}\right) \\ &+ D_{(K,J,R)} \log\left(\sqrt{2}(p^2 q R + \frac{3}{4}q)\right) + \log(K) \\ &\leq D_{(K,J,R)} \left(\log(4) + \frac{\log(2\pi e)}{2} + \log 3 \right) \\ &+ D_{(K,J,R)} \left(\log\left(\frac{A_\Sigma}{a_\Sigma} + \frac{1}{2}\right) + \log\left((p^2 q R + \frac{3}{4}q)\sqrt{2}\right) + 1 \right) \\ &\leq D_{(K,J,R)} \left(\log\left(\frac{24\sqrt{6}\pi e}{\sqrt{1-2^{-1/4}}}\right) + \log\left(\frac{A_\Sigma}{a_\Sigma} + \frac{1}{2}\right) \left(\frac{A_\sigma}{a_\Sigma}\right) \right. \\ &\quad \left. + \log\left(p^2 q R + \frac{3}{4}q\right) \right) \end{aligned}$$

Then,

$$\begin{aligned} & \int_0^\sigma \sqrt{H_{[\cdot]}(\epsilon, S_{(K,J,R)}(s_{(K,J,R)}), \sigma), d_H^{\otimes n}} d\epsilon \\ & \leq \sqrt{D_{(K,J,R)}} \left(3 + \sqrt{\log\left(\frac{A_\Sigma}{a_\Sigma} + \frac{1}{2}\right) \left(\frac{A_\sigma}{a_\Sigma}\right)} \right. \\ & \quad \left. + \sqrt{\log\left(p^2 q R + \frac{3}{4}q\right)} + \sqrt{\log\left(\frac{1}{\sigma \wedge 1}\right)} \right) \end{aligned}$$

Consequently, by putting

$$B = 3 + \sqrt{\log\left(\frac{A_\Sigma}{a_\Sigma} + \frac{1}{2}\right) \left(\frac{A_\sigma}{a_\Sigma}\right)} + \sqrt{\log\left(p^2 q R + \frac{3}{4}q\right)}$$

we get that the function $\Psi_{(K,J,R)}$ defined on \mathbb{R}_+^* by

$$\psi_{(K,J,R)}(\sigma) = \sqrt{D_{(K,J,R)}} \sigma \left(B + \sqrt{\log\left(\frac{1}{\sigma \wedge 1}\right)} \right)$$

satisfies (15). Besides, $\Psi_{(K,J,R)}$ is nondecreasing and $\sigma \mapsto \frac{\Psi_{(K,J,R)}(\sigma)}{\sigma}$ is nonincreasing, then $\Psi_{(K,J,R)}$ is convenient.

Finally, we need to find an upper bound of σ_* satisfying

$$\Psi_{(K,J,R)}(\sigma_*) = \sqrt{n} \sigma_*^2.$$

Consider σ_* such that $\Psi_{(K,J,R)}(\sigma_*) = \sqrt{\sigma_*^2}$.

This is equivalent to solve

$$\sigma_* = \sqrt{\frac{D}{n}} \left(B + \sqrt{\log \left(\frac{1}{\sigma_* \wedge 1} \right)} \right)$$

and then we could choose

$$\sigma_*^2 \leq \frac{D}{n} \left(2B^2 + \log \left(\frac{1}{1 \wedge \frac{D}{n} B^2} \right) \right).$$

5.3. Assumption (K). Let $x_{(K,J,R)} = D_{(K,J)} \log \left(\frac{4epq}{(D_{(K,J)} - q^2) \wedge pq} \right) + \max_{k \in \{1, \dots, K\}} R(k)$. Then, we could compute the sum

$$\sum_{(K,J,R)} e^{-x_{(K,J,R)}} = \underbrace{\left(\sum_{K>0} \sum_{1 \leq |J| \leq pq} e^{-D_{(K,J)} \log \left(\frac{4epq}{(D_{(K,J)} - q^2) \wedge pq} \right)} \right)}_A \sum_R e^{-R}$$

The term A is controlled in proposition 4.5 in [7] by 2. Then,

$$\sum_{(K,J,R)} e^{-x_{(K,J,R)}} \leq 2.$$

6. ACKNOWLEDGMENT

I am grateful to Pascal Massart and Christophe Giraud for stimulating discussions.

REFERENCES

- [1] M. Aldrin. Moderate projection pursuit regression for multivariate response data. *Computational Statistics & Data Analysis*, 21(5):501–531, May 1996.
- [2] J-P Baudry, C. Maugis, and B. Michel. Slope heuristics: overview and implementation. *Stat. Comput.*, 22(2):455–470, 2012.
- [3] L. Birgé and P. Massart. Minimal penalties for gaussian model selection. 138:33773, 2006.
- [4] F. Bunea. Consistent selection via the Lasso for high dimensional approximating regression models. In *Pushing the limits of contemporary statistics: contributions in honor of Jayanta K. Ghosh*, volume 3 of *Inst. Math. Stat. Collect.*, pages 122–137. Inst. Math. Statist., Beachwood, OH, 2008.
- [5] S. Cohen and E. Le Pennec. Conditional density estimation by penalized likelihood model selection. Submitted, 2011.
- [6] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. Discussion. *J. R. Stat. Soc., Ser. B*, 39:1–38, 1977.
- [7] E. Devijver. Finite mixture regression: A sparse variable selection by model selection for clustering.
- [8] E. Devijver. Model-based clustering for high-dimensional data. application to functional data.
- [9] C. Giraud. Low rank Multivariate regression. 23 pages, September 2010.
- [10] J. Guo, E. Levina, G. Michailidis, and J. Zhu. Pairwise variable selection for high-dimensional model-based clustering. *Biometrics*, 66(3):793–804, 2010.
- [11] A. Izenman. Reduced-rank regression for the multivariate linear model. *Journal of Multivariate Analysis*, 5(2):248–264, 1975.
- [12] Z. Ma and T. Sun. Adaptive sparse reduced-rank regression. 2014.
- [13] P. Massart. *Concentration inequalities and model selection*. Lecture Notes in Mathematics. Springer, 33, 2003, Saint-Flour, Cantal, 2007.
- [14] C. Maugis and B. Michel. A non asymptotic penalized criterion for Gaussian mixture model selection. *ESAIM Probab. Stat.*, 15:41–68, 2011.
- [15] C. Meynet and C. Maugis-Rabusseau. A sparse variable selection procedure in model-based clustering. Research report, September 2012.
- [16] W. Pan and X. Shen. Penalized model-based clustering with application to variable selection. *J. Mach. Learn. Res.*, 8:1145–1164, December 2007.
- [17] N. Städler, P. Bühlmann, Sara Van de Geer, and Rejoinder. Comments on ℓ_1 -penalization for mixture regression models. *Test*, 19(2):209–256, 2010.
- [18] W. Sun, J. Wang, and Y. Fang. Regularized k-means clustering of high-dimensional data and its asymptotic consistency. *Electronic Journal of Statistics*, 6:148–167, 2012.
- [19] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. R. Stat. Soc., Ser. B*, 58(1):267–288, 1996.
- [20] H. Zhou, W. Pan, and X. Shen. Penalized model-based clustering with unconstrained covariance matrices. *Electronic Journal of Statistics*, 3:1473–1496, 2009.
- [21] N. Zhou and J. Zhu. Group variable selection via a hierarchical lasso and its oracle property. 3:557–574+, 2010.

INRIA SELECT, UNIVERSITÉ PARIS SUD, BÂT. 425, 91405 ORSAY CEDEX, FRANCE
 E-mail address: emilie.devijver@math.u-psud.fr