

# Sparse Bayesian Unsupervised Learning

Stéphane Gaïffas \*

Bertrand Michel †

February 3, 2014

## Abstract

This paper is about variable selection, clustering and estimation in an unsupervised high-dimensional setting. Our approach is based on fitting constrained Gaussian mixture models, where we learn the number of clusters  $K$  and the set of relevant variables  $S$  using a generalized Bayesian posterior with a sparsity inducing prior. We prove a sparsity oracle inequality which shows that this procedure selects the optimal parameters  $K$  and  $S$ . This procedure is implemented using a Metropolis-Hastings algorithm, based on a clustering-oriented greedy proposal, which makes the convergence to the posterior very fast.

## 1 Introduction

This paper is about variable selection, clustering and estimation for case where we observe unlabelled i.i.d data  $X_1, \dots, X_n$  in  $\mathbb{R}^d$ , denoted  $X_i = (X_i^1, \dots, X_i^d)$ , with  $d$  being eventually much larger than  $n$ . Clustering is now an important tool for the analysis of high-dimensional data. An example of application is gene function discovery and cancer subtype discovery, where one wants to construct groups of genes with their expression levels across different conditions or across several patients tissue samples, see [?] and [?] for instance.

In the high-dimensional setting, clustering becomes challenging because of the presence of a large number of noise variables, that can hide the cluster structure. So, one must come up with an algorithm that, at the same time, selects relevant variables and constructs a clustering based only on these variables. This interconnection between variable selection and clustering makes the problem challenging, and contrasts with supervised problems.

Among many approaches, model-based clustering becomes increasingly popular. It benefits from a well-understood probabilistic framework, see [?], [?], [?], [?]. Because of its flexibility and interpretability, the most popular is the Gaussian Mixture Model (GMM). Indeed, GMM can be classified into 28 models [? ?], where various constraints on the covariance matrices allow to control the shapes and orientations of the clusters. Variable selection in this context is based on the decomposition  $\mu_k = \bar{\mu} + m_k$  of the mean of the  $k$ -th component of the GMM, where  $\bar{\mu}$  is a global mean and  $m_k$  contains specific information about the  $k$ -th cluster. A sparsity assumption on the  $m_k$ 's is natural in this setting: it means that only a few variables contribute to the clustering in terms of location. Even further, since the clustering is totally determined by the estimation of the GMM on the relevant variables, it is natural to assume that the  $m_k$ 's share the *same* support in order to get an accurate description of these variables.

---

\*CMAP – Ecole Polytechnique. Email: [stephane.gaiffas@cmap.polytechnique.fr](mailto:stephane.gaiffas@cmap.polytechnique.fr)

†Université Pierre et Marie Curie, Paris 6. Email: [bertrand.michel@upmc.fr](mailto:bertrand.michel@upmc.fr)

There are roughly two approaches for variable selection in the context of high-dimensional model-based clustering: the Bayesian approach, see [? ], [? ], [? ], [? ], [? ], [? ] and the penalization approach, see [? ], [? ]. The Bayesian approach is very flexible and allows complex modelings of the covariance matrices for the GMM. However, this approach is computationally demanding since it requires heavy MCMC stochastic search on continuous parameter sets. The penalization approach is based on penalizing the log-likelihood by the  $\ell_1$ -norm of the mean vectors and inverse covariance matrices of the GMM. It leads to a soft-thresholding operation in the M-step of the Expectation-Maximization (EM) algorithm, see [? ], [? ]. By doing so, many coordinates of the mean vectors are shrunk towards zero, which helps, hopefully, to remove noise variables. A problem with this thresholding (or equivalently  $\ell_1$ -penalization) approach is that the resulting estimated mean vectors have no reason to share the same supports. As mentioned above, this property is strongly suitable since it models precisely the fact that only a few variables makes a distinction across the clusters. Moreover, another problem is that, as far as we know, there exists no mathematical result about the statistical properties of these penalized GMM methods for high-dimensional unsupervised learning, such as upper bounds on the estimation error, or sparsity oracle inequalities (for mixture models in the regression setting, one can see [? ] and [? ]).

The motivations for this work are two-fold. First, we propose a new approach for model-based unsupervised learning, by combining the GMM with a learning procedure based on the PAC-Bayesian approach. The PAC-Bayesian approach was originally developed for classification by [? ], [? ] and [? ? ], see also [? ? ], [? ] and [? ? ] for further developments. This approach has proved successful for sparse regression problems, see [? ? ? ], [? ], [? ], [? ]. In this work, we use prior distributions that suggest a small support  $S \subset \{1, \dots, d\}$  for significant variables and a small number  $K$  of clusters. Then, we learn  $S$  and  $K$  using a randomized aggregation rule based on the Gibbs posterior distribution, see Section 2.2. Our methodology is based on a Metropolis-Hastings (MH) exploration algorithm that explores a discrete (but large) set for the “meta” parameter  $\eta = (K, S)$ . The exploration can be done in a very efficient way, thanks to the use of a proposal which is particularly relevant for the clustering problem, see Section 4. As shown in our empirical study (see Sections 4 and 5), an order of 300 steps in the MH algorithm is sufficient for convergence on a large scale problem.

Second, our methodology is supported by strong theoretical guarantees: using PAC-Bayesian tools [? ? ? ? ], we prove a sparsity oracle inequality for our procedure. This oracle inequality shows that our procedure automatically selects the parameter  $\eta = (K, S)$  which is optimal in terms of estimation error. Note that this is the first result of this kind for sparse model-based unsupervised learning.

## 2 Our procedure

We assume from now on that we have  $n$  observations. First, we split at random the whole sample  $\mathbb{X} = (X_1, \dots, X_n)$  into a *learning* sample  $\mathbb{X}_1$  and an *estimation* sample  $\mathbb{X}_2$ , of sizes  $n_1$  and  $n_2$  such that  $n_1 + n_2 = n$ . Then, the two main ingredients of our procedure are fitting constrained Gaussian Mixture Models (GMM) and learning the number of clusters and relevant variables using a generalized Bayesian Posterior with a sparsity inducing prior, which are described respectively in Sections 2.1 and 2.2. The main steps of our procedure are finally summarized in Section 2.3.

## 2.1 Constrained Gaussian mixtures models

Let us denote by  $\phi_{(d)}(\cdot|\mu_k, \Sigma_k)$  the density of the multivariate Gaussian  $N_d(\mu_k, \Sigma_k)$  distribution. The density of a Gaussian mixture model (GMM) with  $K$  components is given by

$$f_\theta = \sum_{k=1}^K p_k \phi_{(d)}(\cdot|\mu_k, \Sigma_k),$$

where  $p_1, \dots, p_K$  are the mixture proportions,  $\mu_1, \dots, \mu_K$  are the mean vectors and  $\Sigma_1, \dots, \Sigma_K$  are the covariance matrices. Let  $\theta = (p_1, \dots, p_K, \mu_1, \dots, \mu_K, \Sigma_1, \dots, \Sigma_K)$  be the parameter vector of the GMM. An interest of GMM is that once the model is fitted, one can easily obtain a clustering using the Maximum A Posteriori (MAP) rule, which is recalled in Section 4.1 below. As explained in Introduction, the following structure on the  $\mu_k$ 's is natural:

$$\mu_k = \bar{\mu} + m_k,$$

where  $\bar{\mu}$  is the global mean and where  $m_1, \dots, m_K$  are sparse vectors that share the *same* support. The idea behind this structure is that we want only a few variables to have an impact on the clustering, and even further, we want an informative variable to be informative for *all* clusters. In the following it is assumed without loss of generality that the data have been centered and normalized so that we can take  $\bar{\mu} = 0$ . Note that this is a classical assumption in the model-based clustering context.

Let  $\mathcal{P}(A)$  denote the set of all the subsets of a finite set  $A$ , and  $|A|$  denote the cardinality of  $A$ . Let us introduce the set of GMM configurations

$$\Upsilon = \mathbb{N}^* \times \mathcal{P}(\{1, \dots, d\}).$$

A parameter  $\eta = (K, S) \in \Upsilon$  is a “meta-parameter” for  $\theta$ . It fixes the number of clusters  $K$  and the common support  $S$  of the vectors  $m_k$ . The set  $S$  is thus the set of indexes of the active variables. If  $\mu \in \mathbb{R}^d$  and  $S \subset \{1, \dots, d\}$  we define  $\mu_S = (\mu_j)_{j \in S} \in \mathbb{R}^{|S|}$ . If  $\Sigma$  is a  $d \times d$  matrix, we also define the  $|S| \times |S|$  matrix  $\Sigma_S = (\Sigma_{j,j'})_{j,j' \in S}$ .

Let  $\text{Shape}(\eta)$  denotes the shape of the GMM restricted to the active variables, namely  $\text{Shape}(\eta)$  is a particular set of constrained  $K$ -vector of  $|S| \times |S|$  covariance matrices. An example is, using the notation introduced in [? ], the shape

$$\text{Shape}_{LB}(\eta) = \left\{ ((\Sigma_1)_S, \dots, (\Sigma_K)_S) : (\Sigma_1)_S = \dots = (\Sigma_K)_S = \text{diag}(\sigma_1^2, \dots, \sigma_{|S|}^2), \right. \\ \left. (\sigma_1, \dots, \sigma_{|S|}) \in (\mathbb{R}^+)^{|S|} \right\}, \quad (1)$$

which corresponds to identical and diagonal matrices with normalized noise variables. Note that in the high-dimensional setting, a simple structure on noise variables is suitable, since fitting a GMM with large covariance matrices is prohibitive. The theoretical result given in Section 3 is valid for all possible shapes. Even more than that, since our approach is based on an aggregation algorithm, one could perfectly mix several GMM fits with different shapes (leading to an increased computational cost). From now on, we fix a family of shapes  $(\text{Shape}(\eta))_{\eta \in \Upsilon}$ . We also assume that active variables and non-active variables are independent. This is only for the sake of simplicity,

since we could use more elaborated models, such as the ones from [? ]. All these assumptions lead to the following set of parameters associated to a configuration  $\eta$  and a shape  $\text{Shape}(\eta)$ :

$$\Theta(\eta, \text{Shape}(\eta)) = \left\{ (p_1, \dots, p_K, \mu_1, \dots, \mu_K, \Sigma_1, \dots, \Sigma_K) : (\mu_1)_{S^c} = \dots = (\mu_K)_{S^c} = 0, \right. \\ \left. ((\Sigma_1)_S, \dots, (\Sigma_K)_S) \in \text{Shape}(\eta), \quad (\Sigma_1)_{S^c} = \dots = (\Sigma_K)_{S^c} = I_{d-|S|} \right\},$$

where  $(\mu_k)_A$  is the projection of  $\mu_k$  on the set of coordinates  $A$ , where  $(\Sigma_k)_A$  is the restriction of the matrix  $\Sigma_k$  on  $A \times A$  and where  $I_q$  stands for the identity matrix on  $\mathbb{R}^q$ . Then, for  $\theta \in \Theta(\eta, \text{Shape}(\eta))$ , the density  $f_\theta$  can be decomposed as follows :

$$f_\theta = \phi_{(d-|S|)}(\cdot | 0_{d-|S|}, I_{d-|S|}) \sum_{k=1}^K p_k \phi_{(S)}(\cdot | (\mu_k)_S, (\Sigma_k)_S). \quad (2)$$

For instance, if  $\eta = (K, S)$  and  $\theta \in \Theta(\eta, \text{Shape}_{LB})$ , then  $f_\theta$  is a Gaussian mixture with  $K$  components, with variables outside of  $S$  that are uncorrelated and standard  $N(0, 1)$  and a shape  $\text{Shape}_{LB}$ .

We consider the maximum likelihood estimator of  $\theta$  over the set of constraints  $\Theta(\eta)$  for the observations in the estimation sample  $\mathbb{X}_2$ :

$$\hat{\theta}(\eta, \text{Shape}(\eta)) \in \underset{\theta \in \Theta(\eta, \text{Shape}(\eta))}{\text{argmin}} L_\theta(\mathbb{X}_2) \quad \text{where} \quad L_\theta(\mathbb{X}_2) = \sum_{i: X_i \in \mathbb{X}_2} -\ln f_\theta(X_i). \quad (3)$$

An approximation of this estimator can be computed using the Expectation-Maximization (EM) algorithm, see Section 4.1 for more details. In the following we assume that a shape has been fixed and then use the notation  $\Theta(\eta)$  and  $\hat{\theta}(\eta)$ .

## 2.2 Learning $K$ and $S$

Now, we want to learn the number of clusters  $K$  and the support  $S$  based on the learning data  $\mathbb{X}_1$ . We use a randomized aggregation procedure, which corresponds in this setting to a generalized Bayesian posterior. This method relies on the choice of prior distributions. For  $K$ , we have in mind to force the number of clusters to remain small, so we simply consider the Poisson prior

$$\pi_{\text{clust}}(K) = \frac{e^{-1}}{K!} \mathbf{1}_{K \geq 0}.$$

This prior gives nice results and always recovers smoothly the correct number of clusters in most settings. The use of another intensity parameter (we simply take 1 here) does not make significant differences. This comes from the fact that the parameter  $\lambda$  of the generalized Bayesian posterior (see (4) below) already tunes the degree of sparsity over the whole learning process.

As explained above, we want only the most significant variables to have an impact on the final clustering. So, we consider a prior that downweights supports with a large cardinality exponentially fast. Namely, we consider

$$\pi_{\text{supp}}(S) = \frac{1}{\binom{d}{|S|} e^{|S|} C_d},$$

where  $C_d = \sum_{k=0}^d e^{-k}$ . This prior is used in [? ] for sparse regression learning. This choice is far from being the only one, since we observed empirically that any other prior inducing a small cardinality gives similar results. Finally, the prior for a meta-parameter  $\eta = (K, S) \in \Upsilon$  is

$$\pi(\eta) = \pi_{\text{clust}}(K) \times \pi_{\text{supp}}(S).$$

From a Bayesian point of view, this means that we assume  $K$  and  $S$  to be independent, which is reasonable in this setting. This choice has an impact on the risk bound we obtain, this is discussed further. The Gibbs posterior distribution is now defined by

$$\hat{\pi}_\lambda(d\eta) = \frac{\exp\left(-\lambda L_{\hat{\theta}(\eta)}(\mathbb{X}_1)\right)}{\mathbb{E}_{\eta \sim \pi} \exp\left(-\lambda L_{\hat{\theta}(\eta)}(\mathbb{X}_1)\right)} \pi(d\eta), \quad (4)$$

where  $\lambda > 0$  is the *temperature* parameter and where we recall that  $\hat{\theta}(\eta)$  is given by (3) and that  $L_\theta(\mathbb{X}_1) = \sum_{i: X_i \in \mathbb{X}_1} \ln f_\theta(X_i)$ . Note that written in the following way:

$$\hat{\pi}_\lambda(d\eta) = \frac{\prod_{i: X_i \in \mathbb{X}_1} f_{\hat{\theta}(\eta)}^\lambda(X_i)}{\mathbb{E}_{\eta \sim \pi} \prod_{i: X_i \in \mathbb{X}_1} f_{\hat{\theta}(\eta)}^\lambda(X_i)} \pi(d\eta), \quad (5)$$

the Gibbs posterior is the Bayesian posterior when  $\lambda = 1$ , and it is a so-called generalized Bayesian posterior otherwise. We use a Metropolis-Hastings algorithm to pick at random  $\eta$  with distribution  $\hat{\pi}_\lambda$ , see Section 4.2. The final randomized aggregated estimator is given by

$$f_{\hat{\theta}(\eta)} \quad \text{with} \quad \eta \sim \hat{\pi}_\lambda. \quad (6)$$

The parameter  $\lambda$  is a smoothing parameter that makes the balance between goodness-of-fit on  $\mathbb{X}_1$  and the Kullback-Leibler divergence to the prior (see Equation (13) below). Hence, a careful data-driven choice for  $\lambda$  is important, we observed that AIC or BIC criteria gives satisfying results, see Section 4.3.

### 2.3 Main steps

The main steps of our procedure can be summarized as follow. We fix a number  $B$  of splits (say 20).

1. Repeat the following  $B$  times:
  - (a) Split the whole sample  $\mathbb{X} = (X_1, \dots, X_n)$  at random into a learning sample  $\mathbb{X}_1$  and an estimation sample  $\mathbb{X}_2$ , of size  $n_1$  and  $n_2$ . We call  $b$  this split in the following.
  - (b) For each  $\lambda$  in a grid  $\Lambda$ , use the Metropolis-Hastings (MH) algorithm described in Section 4.2 below to pick at random  $\eta = (K, S)$  distributed according to  $\hat{\pi}_\lambda$ , see (5). Along the MH exploration, the likelihoods used in the generalized Bayes posterior are computed using  $\mathbb{X}_1$  whereas, to compute approximations of the  $\hat{\theta}(\eta)$ 's, the EM algorithm uses  $\mathbb{X}_2$ .
  - (c) Select a temperature  $\lambda(b)$  from the grid, see Section 4.3.
  - (d) Choose a configuration  $\eta(b)$  for temperature  $\lambda(b)$  and select a final meta-parameter  $\hat{\eta}$  using the  $\eta(b)$  chosen by each split  $b = 1, \dots, B$ , see Section 4.3. Fit on  $\mathbb{X}$  a final GMM with configuration  $\hat{\eta}$ .

2. Use the MAP rule to obtain a final clustering.

Several alternatives to these steps are possible since the  $B$  splits bring a lot of information that can be analyzed in different ways. For instance an ‘‘aggregated clustering’’ based on a proximity matrix summarizing all the links pointed out by the  $B$  clusterings can be easily computed (see Section 4.3).

### 3 Main results

If  $f$  and  $g$  are probability density functions, we denote respectively by  $\mathcal{H}(f, g)$  and  $\mathcal{K}(f, g)$  the Hellinger and Kullback divergences between  $f$  and  $g$ . We denote by  $\mathbb{E}_{\mathbb{X}_1}$  and  $\mathbb{E}_{\mathbb{X}_2}$  the expectations with respect to the learning sample  $\mathbb{X}_1$  and  $\mathbb{X}_2$ . If the samples  $\mathbb{X}_1$  and  $\mathbb{X}_2$  are i.i.d with a common density  $f^*$ , a measure of statistical risk for (6) is simply given by  $\mathbb{E}_{\mathbb{X}_1} \mathbb{E}_{\eta \sim \hat{\pi}_\lambda} \mathcal{H}^2(f^*, f_{\hat{\theta}(\eta)})$ . The next Theorem is a sparsity oracle inequality for the aggregated estimator (6).

**Theorem 3.1.** *Let  $\lambda \in (0, 1)$  and consider the randomized aggregated estimator (6), where we recall that  $\hat{\theta}(K, S)$  is a constrained maximum likelihood estimator (3) for a fixed shape. Conditionally to  $\mathbb{X}_2$  we have*

$$\begin{aligned} \mathbb{E}_{\mathbb{X}_1} \mathcal{H}^2(f^*, \mathbb{E}_{\eta \sim \hat{\pi}_\lambda} f_{\hat{\theta}(\eta)}) &\leq \mathbb{E}_{\mathbb{X}_1} \mathbb{E}_{\eta \sim \hat{\pi}_\lambda} \mathcal{H}^2(f^*, f_{\hat{\theta}(\eta)}) \\ &\leq c_\lambda \inf_{\substack{K \in \mathbb{N}^* \\ S \subset \{1, \dots, d\}}} \left\{ \lambda \mathcal{K}(f^*, f_{\hat{\theta}(K, S)}) + \frac{\ln K! + 1 + 2|S| \ln(ed/|S|)}{n} \right\} \end{aligned}$$

and furthermore

$$\begin{aligned} \mathbb{E}_{\mathbb{X}_1 \mathbb{X}_2} \mathcal{H}^2(f^*, \mathbb{E}_{\eta \sim \hat{\pi}_\lambda} f_{\hat{\theta}(\eta)}) &\leq \mathbb{E}_{\mathbb{X}_1 \mathbb{X}_2} \mathbb{E}_{\eta \sim \hat{\pi}_\lambda} \mathcal{H}^2(f^*, f_{\hat{\theta}(\eta)}) \\ &\leq c_\lambda \inf_{\substack{K \in \mathbb{N}^* \\ S \subset \{1, \dots, d\}}} \left\{ \lambda \mathbb{E}_{\mathbb{X}_2} \mathcal{K}(f^*, f_{\hat{\theta}(K, S)}) + \frac{\ln K! + 1 + 2|S| \ln(ed/|S|)}{n} \right\} \end{aligned} \quad (7)$$

where  $c_\lambda = 2/\min(\lambda, 1 - \lambda)$ .

The proof of Theorem 3.1 is given in Section 6 below. Theorem 3.1 entails that the aggregated estimator (6) has a estimation error close to the one of the GMM with the best number of clusters  $K$  and the best support  $S$ . Hence, it proves that the generalized Bayesian posterior (5) automatically selects the correct  $K$  and  $S$  in terms of estimation error. Note that the residual term is of order  $|S| \ln(d/|S|)/n$ , which coincides with the optimal residual term for the model-selection aggregation problem in supervised settings, see [? ], [? ], [? ]. Broadly, this computational cost is  $|S| \ln |S| + K \ln d$ , this additive structure additive in  $|S|$  and  $K$  is a direct consequence of the assumption of independence between  $S$  and  $K$  we make in the prior. The use of the squared Hellinger distance on the left hand side and of the Kullback divergence on the right hand side is a standard technical problem for the proof of oracle inequalities for density estimation, see for instance [? ], [? ], and [? ]. The restriction on  $\lambda$  is for the sake of simplicity, since more general (but sub-optimal) bounds for  $\lambda > 1$  can be established, see [? ], and requires extra technicalities that are beyond the scope of this paper.

The upper bound given in Theorem 3.1 can be made more explicit by computing the Kullback Leibler risk of the MLE in a family of GMMs of fixed shape, we investigate this question in following of this section. It is known since the works are [? ] (see also Section 2.4 in [? ]) that for well specified parametric models in  $\mathbb{R}^p$ , the  $\mathcal{K}$  risk of the MLE is of the order of  $\frac{p}{n}$ . However, the model has no reason to be well specified in our context and moreover we need a more precise bound than these asymptotic results. It is also well known (see for instance [? ]) that rates of convergence for estimators should be related to the metric structure with  $\mathcal{H}$ . Indeed, rates of convergence of the MLE are usually given in term of Hellinger risk. In the context of GMMs, rates of convergence of MLE for the Hellinger risk have first been investigated in [? ] and [? ]. To compute such rates of

convergence for our models, we need to bound the parameters of the sets  $\Theta(\eta)$ . We assume as before that the GMM shape is fixed. For a configuration  $\eta$  and some positive constants  $\bar{\mu}$ ,  $\underline{\sigma} < 1 < \bar{\sigma}$ ,  $\underline{L} < \bar{L}$  we consider restricted sets of parameters with the following constraints:

$$\Theta_r(\eta) = \left\{ \theta \in \Theta(\eta) : \forall k \in \{1 \dots K\}, |\mu_k| \leq \bar{\mu}, \text{sp}(\Sigma_k) \subset [\underline{\sigma}^2, \bar{\sigma}^2]^d, \underline{L} \leq (2\pi)^{d/2} |\Sigma_k| \leq \bar{L} \right\}$$

where  $\text{sp}(\Sigma)$  is the spectrum of  $\Sigma$ . Let  $\mathcal{F}_\eta$  be the set of GMM densities  $f_\theta$  for  $\theta \in \Theta_r(\eta)$ . Let  $\hat{\theta}_r(\eta)$  denotes the MLE of  $\theta$  in the restricted set  $\Theta_c(\eta)$ . We then consider the randomized restricted estimator

$$f_{\hat{\theta}_r(\eta)} \quad \text{with} \quad \eta \sim \hat{\pi}_\lambda.$$

We also define

$$\mathcal{K}(f^*, \mathcal{F}_\eta) := \inf_{\theta \in \Theta_r(\eta)} \mathcal{K}(f^*, f_\theta)$$

and the number of free parameters  $D(\eta)$  of the Gaussian mixture model parametrized by  $\Theta_r(\eta)$ . Note that this number of free parameters depends of the chosen shape. For instance, if the shape is  $\text{Shape}_{LB}$  then  $D(\eta) = (K + 1)|S| - 1$ .

In order to upper bound the  $\mathcal{K}$ -risk, we use the following standard result: for any measure  $P$  and  $Q$  defined on the same probability space (see Section 7.6 in [?] for instance),

$$2\mathcal{H}^2(P, Q) \leq \mathcal{K}(P, Q) \leq 2 \left( 1 + \ln \left\| \frac{dP}{dQ} \right\|_\infty \right) \mathcal{H}^2(P, Q). \quad (8)$$

To avoid "unbounded problem" we will assume here that  $f^*$  is bounded and has a compact support. Under these hypotheses it is then possible to upper bound the  $\mathcal{K}$ -risk of the MLE on the spaces defined  $\Theta_r(\eta)$ .

**Theorem 3.2.** *Assume that  $\mathbb{X}_1$  and  $\mathbb{X}_2$  are both i.i.d. with a common density  $f^*$  such that  $\|f^*\|_\infty < \infty$  and such that the support of  $f^* < \infty$  is included in  $B(0, \bar{\mu})$ . There exists an absolute constant  $\kappa$  such that for any  $\lambda \in (0, 1)$ ,*

$$\mathbb{E}_{\mathbb{X}_1 \mathbb{X}_2} \mathcal{H}^2(f^*, \mathbb{E}_{\eta \sim \hat{\pi}_\lambda} f_{\hat{\theta}(\eta)}) \leq c_\lambda \inf_{\substack{K \in \mathbb{N}^* \\ S \subset \{1, \dots, d\}}} \left\{ \lambda C \left[ \mathcal{K}(f^*, \mathcal{F}_\eta) + \kappa \frac{D(\eta)}{n_2} \left\{ \mathcal{A}^2 + \ln^+ \left( \frac{n_2}{\mathcal{A}^2 D(\eta)} \right) \right\} + \frac{\kappa}{n_2} \right] + \frac{\ln K! + 1 + 2|S| \ln(ed/|S|)}{n_1} \right\}$$

where  $C = \frac{8}{2 \ln 2 - 1} \left( 1 + \ln(\|f^*\|_\infty L^+) + \frac{2\bar{\mu}^2}{\bar{\sigma}^2} \right)$  and where the constant  $\mathcal{A}$  only depends on the GMM shape and the bounding parameters.

Of course, Theorem 3.2 is only meaningful if the true distribution  $f^*$  can be arbitrarily well-approximated in the  $\mathcal{K}$  divergence sense by the Gaussian mixtures of the model collection. Roughly, the model dimension  $D(\eta)$  is of the order of  $K|S|$  or  $K|S|^2$  according to the chose shape. The upper bound given by this result is thus of the order of

$$\inf_{\substack{K \in \mathbb{N}^* \\ S \subset \{1, \dots, d\}}} \left\{ \mathcal{K}(f^*, \mathcal{F}_\eta) + \frac{D(\eta)}{n} \ln \frac{n}{D(\eta)} \right\}$$

if we take for instance  $n_1 = n_2$ . A similar bound has been found in the same framework by [?] for an  $l_0$ -penalization procedure in the spirit of [?]. Moreover, using recent results of [?] about the

approximation of log-Holder densities using univariate Gaussian mixtures models, [?] has shown the optimality of such risk bound, in the minimax sense.

Note that the assumptions on the true density  $f^*$  could be probably relaxed. However these assumptions are not too strong for the clustering framework of this paper. Finally, it is not possible to give one simple expression for the constant  $\mathcal{A}$  since it depends on the shape chosen. The interested reader is referred to Lemma 6.3 in the Appendix and references therein.

## 4 Implementation

This section details the whole implementation of our method. We illustrate the main steps on the a simulated example presented further.

### 4.1 The EM algorithm and the MAP rule

An approximation of the maximum likelihood estimator of GMMs can be computed thanks to the EM algorithm [?]. In this section we briefly present the principle of the algorithm for our context. Assuming that a split  $(\mathbb{X}_1, \mathbb{X}_2)$  has been chosen, only the estimation sample  $\mathbb{X}_2$  can be used to compute the estimators. Let  $\hat{\theta}(\eta)$  be the maximum likelihood estimator of the GMM with configuration  $\eta$ , see Equation (3). Let  $Z_i$  be the (unknown) random vector giving the cluster of the observation  $i$ :

$$Z_{i,k} = \begin{cases} 1 & \text{if the observation } i \text{ is in cluster } k, \\ 0 & \text{otherwise.} \end{cases}$$

The complete log-likelihood of the observations  $(\mathbb{X}, \mathbb{Z})$  is defined by

$$L_{\theta}^{(c)}(\mathbb{X}_2, \mathbb{Z}_2) = \sum_{i: X_i \in \mathbb{X}_2} \sum_{k=1}^K Z_{i,k} (\ln p_k + \ln \phi_{\mu_k, \Sigma_k}(X_i)).$$

The algorithm consists of maximizing the expected value of the log-likelihood with respect to the conditional distribution of  $Z$  given  $X$  under a current estimate of the parameters  $\theta^{(r)}$ :

$$Q(\theta|\theta^{(r)}) = \mathbb{E}_{Z|X, \theta^{(r)}} [L_{\theta}^{(c)}(X, Z)].$$

More precisely, the EM algorithm iterates the two following steps:

- Expectation step: compute the conditional probabilities

$$t_{i,k}^{(r)} := \mathbb{P}(Z_{ik} = 1 | X, \theta^{(r)}) = \frac{p_k^{(r)} \phi_{\mu_k^{(r)}, \Sigma_k^{(r)}}(X_i)}{\sum_{s=1}^K p_s^{(r)} \phi_{\mu_s^{(r)}, \Sigma_s^{(r)}}(X_i)}, \quad (9)$$

and then compute  $Q(\theta|\theta^{(r)})$ .

- Maximization step : find  $\theta^{(r+1)} \in \operatorname{argmax}_{\theta \in \Theta(\eta)} Q(\theta|\theta^{(r)})$ .

After some iterations,  $\theta^{(r)}$  is a good approximation of the maximum likelihood estimator  $\hat{\theta}(\eta)$ .

At several steps of our method, we use the *Maximum A Posteriori* (MAP) rule to construct a clustering based on a GMM fit  $\hat{\theta}$ . The MAP rule consists of attributing each observation  $i$  to the



class which maximizes the conditional probability  $\mathbb{P}(Z_{i,k} = 1 | X_i, \hat{\theta})$ . According to (9), it reduces to the choice of the class  $\hat{k}(i)$  such that

$$\hat{k}(i) = \operatorname{argmax}_{k=1,\dots,K} \left\{ \hat{p}_k \phi_{\hat{\mu}_k, \hat{\Sigma}_k}(X_i) \right\}.$$

## 4.2 Metropolis-Hastings with a particular proposal

To draw  $\eta$  at random according to the generalized Bayes posterior (5), we use the Metropolis-Hastings (MH) algorithm which is of standard use for Bayesian statistics, see [?] and for PAC-Bayesian algorithms, see [?]. The MH algorithm is typically slow, so the careful choice of a proposal is often suitable to obtain fast but statistically pertinent exploration.

In a supervised setting, such as regression, a natural idea for variable selection is to add iteratively the variables that are the most correlated with the residuals coming from a previous fit. This idea leads to the so-called greedy algorithms, which are known to be computationally efficient in a high-dimensional setting, see for instance [?].

In the unsupervised setting, however, no residuals are available, so one must come up with another idea. What is available though in our case is the clustering coming from a previous GMM fit. So, when exploring the variables space, it seems natural to give a stronger importance to the variables that best explain the previous clustering. This importance can be measured using the between-variance.

### The between-variance

Let  $\mathcal{C} = \{\mathcal{C}_1, \dots, \mathcal{C}_K\}$  be a clustering into  $K$  groups of the observations  $\mathbb{X}$ . The between-variance of this clustering is defined by

$$\operatorname{Var}_B(\mathcal{C}) = \frac{1}{n} \sum_{k=1}^K n_k \|G_k - G\|^2,$$

where  $n_k$  is the size and  $G_k = (G_{k,1}, \dots, G_{k,d})^\top$  is the barycenter of cluster  $\mathcal{C}_k$ , and where  $G = (G_1, \dots, G_d)^\top$  is the barycenter of the whole sample and  $\|\cdot\|$  is the Euclidean norm on  $\mathbb{R}^d$ . The between-variance of the variable  $j$  is

$$\operatorname{Var}_B(\mathcal{C}, j) = \frac{1}{n} \sum_{k=1}^K n_k (G_{k,j} - G_j)^2.$$

### A particular proposal

Assume that we are at step  $u$  of the MH algorithm (see Algorithm 1 below), the GMM estimated at this step has configuration  $\eta = (K, S)$  and is denoted by  $g_u$ . The clustering of  $\mathbf{X}$  deduced from this configuration with the MAP rule is denoted by  $\mathcal{C}(\eta)$ . To propose a new configuration  $\tilde{\eta} = (\tilde{K}, \tilde{S})$  for the step  $u + 1$ , we use a transition kernel  $W$  defined as follows:

$$\forall \tilde{\eta} \in T, \quad W(\eta, \tilde{\eta}) = H(K, \tilde{K}) M_{K, \tilde{K}}(S, \tilde{S}).$$

The kernel  $H$  determines how the number of clusters can change along the trajectory, it is defined by

$$H(K, \cdot) = H^+(1, \cdot) \mathbf{1}_{K=1} + \frac{H^-(K, \cdot) + H^+(K, \cdot)}{2} \mathbf{1}_{1 < K < K_{\max}} + H^-(K_{\max}, \cdot) \mathbf{1}_{K=K_{\max}}$$

with  $H^+(K, K) = H^+(K, K+1) = 1/2$  and  $H^-(K, K) = H^-(K, K-1) = 1/2$ .

The kernel  $M_{K, \tilde{K}}$  determines how the number of active variables can change conditionally to the moves  $K \rightarrow \tilde{K}$ , it is defined by as follows:

- if  $\tilde{K} \neq K$  :

$$M_{K, \tilde{K}}(S, \cdot) = \mathbf{1}_{S=\tilde{S}}$$

- if  $\tilde{K} = K$  :

$$M_{K, K}(S, \cdot) = M_K^+(\emptyset, \cdot) \mathbf{1}_{|S|=0} + \frac{M_K^-(S, \cdot) + M_K^+(S, \cdot)}{2} \mathbf{1}_{0 < |S| < d} + W_{2, K}^-(\{1, \dots, d\}, \cdot) \mathbf{1}_{|S|=d}$$

with

$$\begin{aligned} M_K^+(S, S \cup \{j\}) &= \frac{\text{Var}_B(\mathcal{C}(K, S), j)}{\sum_{j \notin S} \text{Var}_B(\mathcal{C}(K, S), j)} && \text{if } j \in S \text{ and 0 otherwise} \\ M_K^-(S, S \cup \{j\}) &= \frac{\text{Var}_B^{-1}(\mathcal{C}(K, S), j)}{\sum_{j \in S} \text{Var}_B^{-1}(\mathcal{C}(K, S), j)} && \text{if } j \notin S \text{ and 0 otherwise.} \end{aligned}$$

According to the transition kernel  $W$ , the set of variables can only change when the number of clusters is unchanged. When  $\tilde{K} = K$ , we decide to add or remove one variable with probability 1/2. When adding a variable, we pick a variable at random outside of  $S$  according to the distribution proportional to the between-variances  $\text{Var}_B(\mathcal{C}_{u-1}, j)$ . When removing a variable, we choose a variable at random inside of  $S$  according to the distribution proportional to the inverse between-variances  $\text{Var}_B^{-1}(\mathcal{C}_{u-1}, j)$ . We observed empirically that this proposal helps the MH exploration to focus quickly on interesting variables.

Algorithm 1 below details our MH algorithm for one split and one given temperature  $\lambda$ . This algorithm is similar to a stochastic greedy algorithm, which decides to add or to remove a variable according to a criterion based on an improvement of the likelihood, and constrained by the sparsity of  $K$  and  $S$ .

### Preliminary pruning

Algorithm 1 requires to compute an EM algorithm at each step of the Markov chain. EM algorithms can be time consuming, in particular for GMM based on a large family of active variables. To quickly focus on a reasonably large set of active variables, we speed up the procedure using a rough ‘‘preliminary pruning’’. We replace  $W_2^-$  by a proposal that allows to remove several variable in one step, until the number of active variables is smaller than a ‘‘target number’’ fixed by the user, we used  $d/3$  in our experiments. Moreover, the number of variables to be removed is drawn at random, and cannot be larger than half of the remaining active variables. Once the target number of active variables is reached, we use the proposal  $W_2^-$  described above. Note that this rough preliminary pruning may remove sometimes true variables, but, fortunately, they can be recovered later.

---

**Algorithm 1** Metropolis-Hastings algorithm to draw  $\eta \sim \hat{\pi}_\lambda$

---

**Require:**  $\mathbb{X}_1, \mathbb{X}_2, K_0, \lambda$

Initialize  $K \leftarrow K_0$

**if**  $n \leq d$  **then**

Take  $S_0 \leftarrow \{1 \dots d\}$

**else**

Find a clustering  $\mathcal{C}$  with the  $K$ -means algorithm

Take  $S_0$  as the set of  $n$  variables  $j$  maximizing  $\text{Var}_B(\mathcal{C}, j)$

**end if**

Fit  $g_0$  on  $\mathbb{X}_2$  for the configuration  $(K_0, S_0)$

Find a clustering  $\mathcal{C}_0$  of  $\mathbb{X}$  with  $g_0$

**for**  $u = 1$  to convergence **do**

Draw  $\eta_{\text{new}} = (K_{\text{new}}, S_{\text{new}})$  distributed as  $W(\eta_{u-1}, \cdot)$

Put  $r \leftarrow \exp \left[ \lambda (L_{\hat{\theta}(\eta_{\text{new}})}(\mathbb{X}_1) - L_{\hat{\theta}(\eta_{u-1})}(\mathbb{X}_1)) \right] \times \frac{\pi(\eta_{\text{new}})}{\pi(\eta_{u-1})} \times \frac{W(\eta_{\text{new}}, \eta_{u-1})}{W(\eta_{u-1}, \eta_{\text{new}})}$

Draw  $U$  distributed uniformly on  $[0, 1]$

**if**  $U > r$  **then**

Put  $\eta_u \leftarrow \eta_{\text{new}}$

Fit  $g_u$  on  $\mathbb{X}_2$  for the configuration  $\eta_u$

Find a clustering  $\mathcal{C}_u$  of  $\mathbb{X}$  with  $g_u$

**else**

Keep  $\eta_u \leftarrow \eta_{u-1}$ ,  $g_u \leftarrow g_{u-1}$  and  $\mathcal{C}_u \leftarrow \mathcal{C}_{u-1}$

**end if**

**end for**

**return**  $\eta_u$

---

## An illustrative example

We illustrate the complete procedure on a simulated GMM based on two components in dimension 100. Only the 15 first variables are active and the others are i.i.d. with distribution  $N(0, 1)$ . The mean of the first 15 variables for each component are respectively  $\mu_1 = 1$  for cluster 1 and  $\mu_2 = 0$  for cluster 2. Both clusters have 100 observations and the whole sample has been centered and standardized. We use only here GMM with shapes  $\text{Shape}_{LB}(\eta)$  defined by (1).

Figure 1 shows nine trajectories of active variables, each trajectory being relative to a temperature taken from the grid  $\Lambda = \{1, 2, 5, 10, 20, 30, 50, 75, 100\}$ . All the trajectories are based on the same split. The chains have been initialized on the configuration  $K_0 = 2$  and  $S_0 = \{1, \dots, d\}$ . We used the preliminary pruning here to quickly reduce the number of variables. After a few dozen of steps, the number of active variables has been dramatically reduced. Note that we use chains of length 300 only to reach convergence. There is no need to wait any longer since the chains are already stabilized at this stage around the correct configuration.

### 4.3 Post-treatment of the trajectories

The MH algorithm presented in the previous section is proceeded for several splits and several temperatures. Leaving aside for the moment the temperature choice issue, we then have  $B$  available

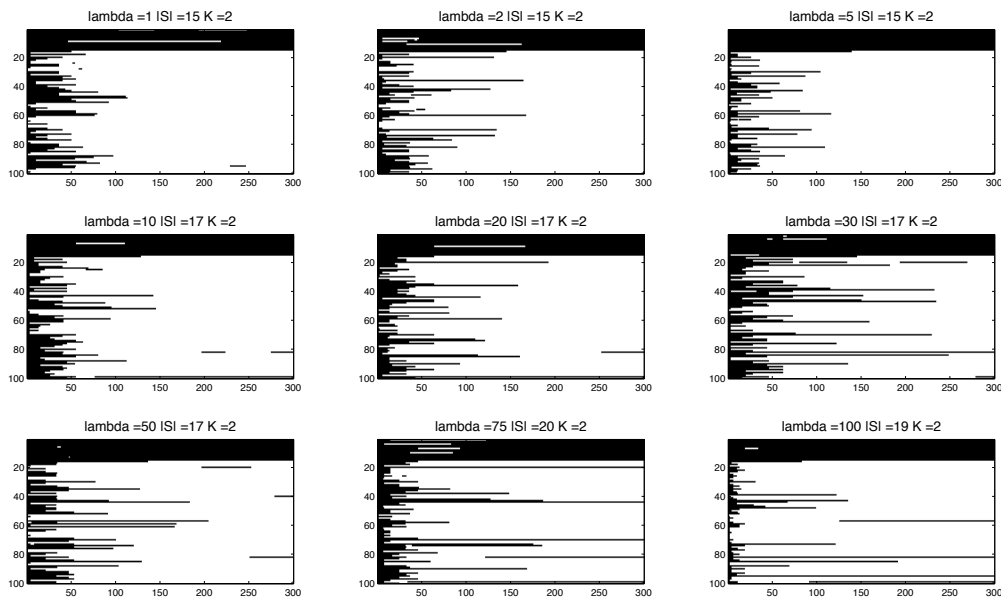


Figure 1: MH trajectories with temperatures in  $\Lambda$  and for a same split. These graphs show the actives variables in black. The number of components at the end of the chain is given in the title of each graph.

GMM estimators of the density. This generates a large amount of information that can be used to cluster the data  $\mathbb{X}$ . One first idea is to aggregate all these estimators to provide one final estimator of the density. Nevertheless, this estimator can not be easily used to produce a relevant clustering of the data since this last is a GMM with all the composants of each GMM. We thus need to aggregate the information provided by the splits in another way. We propose here two alternative methods : the first one consists in selecting one final GMM and the second method consists in aggregating the clustering provided by the  $B$  clusterings.

### Selection of one configuration for each chain

We associate one configuration  $\eta(b, \lambda)$  to each split  $b$  and each temperature  $\lambda$  by choosing the most visited configuration at the end of the chain. This choice is indeed reasonable because the chain state is very stable when  $u$  is large enough as shown in Figure 1 (in practice we consider the last 100 visited configurations). Note that most of the time this configuration is also the last configuration visited by the chain.

### Temperature choice

We choose a temperature  $\lambda$  for each split  $b$  in the following way. We fit a GMM for the configuration  $\eta(b, \lambda)$  using  $\mathbb{X}$ . Let us denote by  $\hat{\theta}(b, \lambda)$  the parameters of this fit. For a given split  $b$ , we choose a temperature according to BIC or AIC criteria:

$$\lambda_{\text{aic}}(b) = \underset{\lambda \in \Lambda}{\operatorname{argmin}} -2L_{\hat{\theta}(b, \lambda)}(\mathbb{X}) + |\eta(b, \lambda)|$$

and

$$\lambda_{\text{bic}}(b) = \underset{\lambda \in \Lambda}{\operatorname{argmin}} -2L_{\hat{\theta}(b,\lambda)}(\mathbb{X}) + \ln n|\eta(b, \lambda)|,$$

where  $|\eta|$  is the number of free parameters of the GMM associated to the configuration  $\eta$ . This gives for each split  $b$  one configuration  $\eta(b)$ . Figure 2 shows the supports chosen for each split using AIC and BIC criteria.

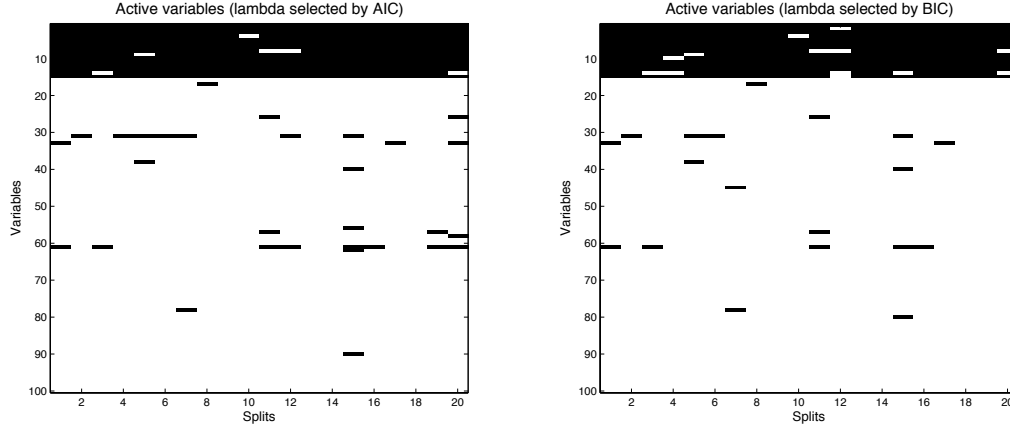


Figure 2: Chosen supports for each split using AIC and BIC criteria.

### Final configuration selection

The information brought by the family of splits can be used to measure the “importance” of each variable by considering the proportion of splits for which this variable is active, see the Figure 3 (right). Note that this measure could also be used to produce a ranking of the variables. One final configuration  $\hat{\eta}$  is finally chosen using a majority vote: for each variable, we vote over the  $B$  splits to decide if this variable is active or not. We choose  $\hat{K}$  as the integer the closest to  $\frac{1}{B} \sum_{b=1 \dots B} K(b)$ .

The method has been applied on the illustrative example using 20 splits. Figure 3 (left) shows the number of components  $K$  chosen over the 20 splits:  $K = 2$  is majoritarian if the temperature is chosen whether by AIC or BIC. For this experiment, the exact family of active variable is also correctly recovered by the voting method with a temperature selected by AIC or BIC.

### Final Clustering

One can think of two strategies to define a final clustering.

- *Direct clustering from  $\hat{\eta}$* : we fit a GMM model on  $\mathbb{X}$  for the configuration  $\hat{\eta}$  chosen by the method detailed above. We then propose a clustering for  $\mathbb{X}$  using the MAP rule.
- *Aggregated clustering*: each split provides a configuration  $\eta(b)$  and an associated clustering thanks to the MAP rule. Many methods exists to produce an aggregated clustering using these  $B$  clusterings available, we propose two versions :
  - *Aggregated clustering by CAH*: define  $A$  as the similarity matrix with entries  $a_{i,j}$  equal to the number of times  $i$  and  $j$  are in the same cluster across the splits. Then, a

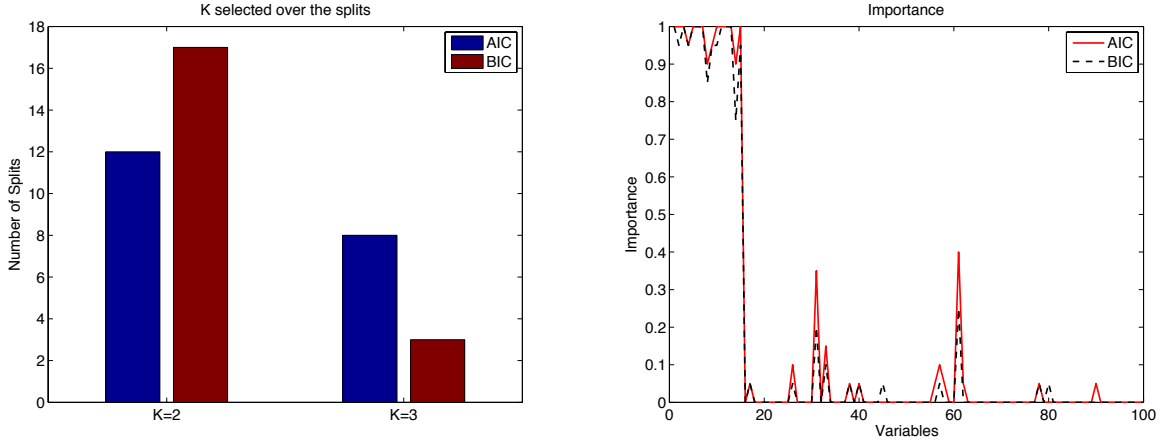


Figure 3: Left: number of components chosen over the 20 splits using AIC or BIC to set the temperature for each split. Right: variable importance measured over the 20 splits.

Method	ARI
Direct clustering (temperature set by AIC)	0.9020
Direct clustering (temperature set by BIC)	0.9020
Aggregated clustering (temperature set by AIC)	0.8090
Aggregated clustering (temperature set by BIC)	0.7557

Table 1: Adjusted Random Indexes (ARI) for the four possible clustering methods applied to the illustrative example.

hierarchical clustering method gives us a final clustering by using for instance  $\exp(-A)$  as a dissimilarity matrix.

To assess a clustering method, we use the Adjusted Random Index (ARI) from [? ], which is an established standard to measure the correspondences between a given clustering and the true one. The ARI's of our methods on the illustrative example are reported in Table 1. In this case, a direct clustering deduced from the configuration  $\hat{\eta}$  gives the best results. Indeed, we observed on many examples that, quite surprisingly, the aggregated clustering generally does not provide the best ARI's compared to direct clustering. This phenomenon certainly deserves further investigations, to be considered in another work.

## 5 Numerical experiments

We compare variable selection and clustering results for the Lasso method for GMM from [? ] (denoted Lasso-GMM) and our method (called MH-GMM). On each simulation, the observations are centered and standardized before being used by both methods. We use the BIC criterion to tune the temperature for our method and the smoothing parameter for Lasso-GMM. In the following experiments, we only deal with diagonal and shared covariance matrices (the shape (1)), which is also the setting considered in [? ].

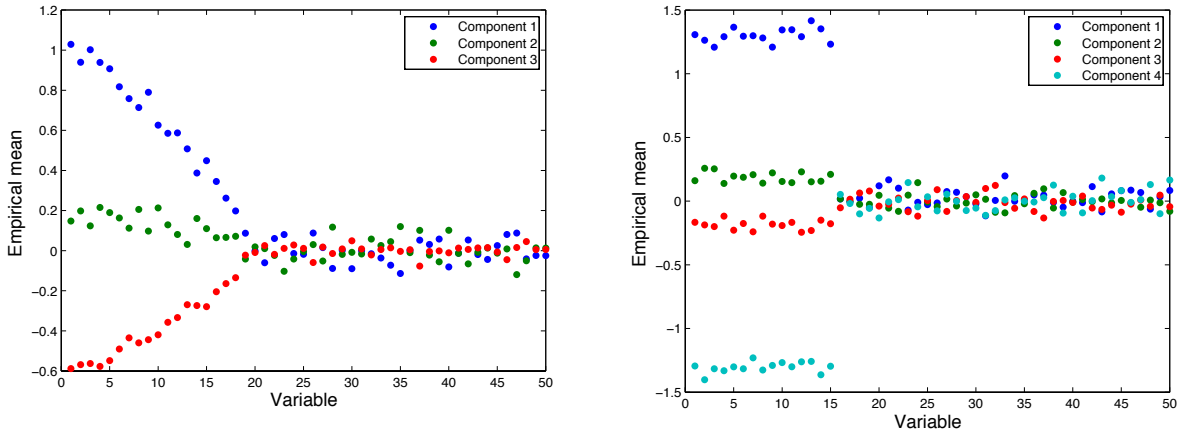


Figure 4: Empirical means of the variables in each component for Experiment 1 (left) and Experiment 2 (right).

### Experiment 1

We simulate 100 replications of a GMM based on three components in dimension 100. All the variables are independent, only the 20 first variables are active and the others are i.i.d.  $N(0, 1)$ . Let  $a = (1, 0.95, 0.9, 0.85, \dots, 0.1, 0.05)$ , the distribution of the vector of the 20 active variables of the first component is  $N_{20}(a, I_{20})$ . The distribution of the vector of the second component is  $N_{20}(0_{20}, I_{20})$  and the third is  $N_{20}(-a, I_{20})$ . There are 200 observations in the two first clusters and 400 in the last one.

On this example, the discriminant power of the clustering variables decreases with respect to the variable index: the three sub-populations of the mixture are progressively gathered together into a unique Gaussian distribution after the 20th variable, as illustrated in the left-hand side of Figure 4. Note that the means of the second component are close to zero.

### Experiment 2

We simulate 100 replications of a GMM based on four components in dimension 100. These four clusters are based on 15 active variables (from 1 to 15), the variances of the active variables are equal to one for the four components. The other variables are i.i.d.  $N(0, 1)$ . All the variables are independent and there are 100 observations for components 1 and 4 and 200 for components 2 and 3. The means of the four components are given by

$$\begin{aligned} \mu_1 &= (2, \dots, 2, 0, \dots, 2), & \mu_2 &= (0.3, \dots, 0.3, 0, \dots, 0), \\ \mu_3 &= (-0.3, \dots, -0.3, 0, \dots, 0), & \mu_4 &= (-2, \dots, -2, 0, \dots, 0), \end{aligned}$$

with 15 non-zero coordinates for each component. For this experiment, selecting  $K$  is more difficult than in the previous experiment. The mean of each cluster is illustrated in the right-hand side of Figure 4.

<b>Experiment 1</b>	Variable selection				Choice of $K$	
	true active	false active	true non-active	false non-active	$K = 3$	$k = 4$
<i>True model</i>	20	0	80	0	100	0
Lasso-GMM	18.56	2.17	77	1.44	99	1
MH-GMM	16.82	0	80	3.18	100	0
<b>Experiment 2</b>	Variable selection				Choice of $K$	
	true active	false active	true non-active	false non-active	$K = 3$	$k = 4$
<i>True model</i>	15	0	85	0	100	0
Lasso-GMM	15.03	0.81	84.19	0	72	28
MH-GMM	15	0	85	0	0	100

Table 2: Variables and  $K$  selected by Lasso-GMM and MH-GMM for Experiments 1 and 2.




Figure 5: ARIs for Lasso-GMM and MH-GMM for Experiments 1 and 2

## Results

We compare variable selection results and clustering results using the ARI for Lasso-GMM and MH-GMM on Experiments 1 and 2. A variable selected by one method is called active. Concerning Lasso-GMM, a variable is active if at least one  $\mu_{j,k}$ ,  $k = 1, \dots, K$  is non-zero while for MH-GMM all the coordinates associated to an active variable variables are non-zero. The terms true and false refers to the true configuration. The results are summarized in Table 2 and Figure 5.

On Experiment 1, Lasso-GMM finds the correct active variables but most of the times it does not activate the coordinates corresponding to the second component. This leads to a non-optimal estimation of the GMM even if it has the correct configuration. As a consequence, a large proportion of observations are misclassified. The MH-GMM method does not shrink coordinates, so it shows better ARI rates.

On Experiment 2, Lasso-GMM fails to select  $K = 4$  correctly 72 times out of 100 replications while MH-GMM chooses  $K = 4$  every time. Indeed, Lasso-GMM hardly separates Cluster 2 from Cluster 3. Consequently, MH-GMM has a much better ARI clustering rate than Lasso-GMM, see the right-hand side of Figure 5.



## 6 Proof of Theorem 3.1

In this section, we assume that a shape has been fixed.

### 6.1 Some preliminary tools from information theory

A measure of statistical complexity of a randomized estimator  $\hat{\pi}$  with respect to a prior  $\pi$  is given by the Kullback-Leibler divergence, defined by

$$\mathcal{K}(\hat{\pi}, \pi) = \int_{\Upsilon} \ln \frac{d\hat{\pi}}{d\pi}(\eta) \hat{\pi}(d\eta),$$

assuming that it exists. The statistical risk, or generalization error of a procedure  $\hat{\pi}$  associated to a loss function  $L_\eta(\cdot)$  on  $\Upsilon \times \mathcal{X}$  is given by

$$\mathbb{E}_{\eta \sim \hat{\pi}} \mathbb{E}_X L_\eta(X) = \int \int L_\eta(x) \hat{\pi}(d\eta) P_X(dx).$$

Let  $\pi'$  and  $\pi$  be probability measures on  $\Upsilon$ , and let  $g : \Upsilon \rightarrow \mathbb{R}$  be any measurable function such that  $\mathbb{E}_{\eta \sim \pi} \exp(g(\eta)) < +\infty$ . An easy computation gives, for  $\pi_g(d\eta) = \frac{\exp(g(\eta))}{\mathbb{E}_{\eta \sim \pi} \exp(g(\eta))} \pi(d\eta)$ , that

$$\mathcal{K}(\pi', \pi_g) = \mathcal{K}(\pi', \pi) + \ln \mathbb{E}_{\eta \sim \pi} \exp(g(\eta)) - \mathbb{E}_{\eta \sim \pi'} g(\eta). \quad (10)$$

Since  $\mathcal{K}(\pi', \pi_g) \geq 0$  (using Jensen's inequality) (10) entails

$$\mathbb{E}_{\eta \sim \pi'} g(\eta) \leq \mathcal{K}(\pi', \pi) + \ln \mathbb{E}_{\eta \sim \pi} \exp(g(\eta)). \quad (11)$$

Inequality (11) is a well-known convex duality inequality, see for instance [? ], that can be used as follows. We denote  $\mathbb{X}_1 = (X_1, \dots, X_n)$  and let  $L_\eta(\mathbb{X}_1)$  be some loss function. Using (11) with a randomized estimator  $\hat{\pi}$  and with the application

$$g(\eta) = -L_\eta(\mathbb{X}_1) - \ln \mathbb{E}_{\mathbb{X}_1} e^{-L_\eta(\mathbb{X}_1)}$$

leads to

$$\mathbb{E}_{\mathbb{X}_1} \exp \left( \mathbb{E}_{\eta \sim \hat{\pi}} \left[ -L_\eta(\mathbb{X}_1) - \ln \mathbb{E}_{\mathbb{X}_1} e^{-L_\eta(\mathbb{X}_1)} - \mathcal{K}(\hat{\pi}, \pi) \right] \right) \leq 1,$$

where we apply  $x \mapsto e^x$ , take the expectation  $\mathbb{E}_{\mathbb{X}_1}$  on both sides of (11), and use Fubini's Theorem. Now, if  $L_\eta(\mathbb{X}_1) = \sum_{i=1}^n \ell_\eta(X_i)$ , we have using Jensen's inequality on the left-hand side, and rearranging the terms, that

$$- \mathbb{E}_{\mathbb{X}_1} \mathbb{E}_{\eta \sim \hat{\pi}} \ln \mathbb{E}_{\mathbb{X}_1} e^{-\ell_\eta(X_1)} \leq \frac{\mathbb{E}_{\mathbb{X}_1} [\mathbb{E}_{\eta \sim \hat{\pi}} L_\eta(\mathbb{X}_1) + \mathcal{K}(\hat{\pi}, \pi)]}{n}. \quad (12)$$

This leads to the so-called *information complexity minimization*, see [? ]: the statistical procedure obtained by minimizing

$$\mathbb{E}_{\eta \sim \hat{\pi}} L_\eta(\mathbb{X}_1) + \mathcal{K}(\hat{\pi}, \pi) \quad (13)$$

over every probability measures  $\hat{\pi}$  on  $\Upsilon$  is explicitly given by

$$\hat{\pi}_{\text{sol}}(d\eta) = \frac{\exp(-L_\eta(\mathbb{X}_1))}{\mathbb{E}_{\eta \sim \pi} \exp(-L_\eta(\mathbb{X}_1))} \pi(d\eta),$$

since (10) entails

$$\mathcal{K}(\hat{\pi}, \hat{\pi}_{\text{sol}}) = \mathbb{E}_{\eta \sim \hat{\pi}} L_{\eta}(\mathbb{X}_1) + \mathcal{K}(\hat{\pi}, \pi) + \ln \mathbb{E}_{\eta \sim \pi} \exp(-L_{\eta}(\mathbb{X}_1)).$$

Note that if  $L_{\eta}(\mathbb{X}_1) = \sum_{i=1}^n \ell_{\eta}(X_i)$  with  $\ell_{\eta}(x) = -\lambda \ln f_{\hat{\theta}(\eta)}(x)$  with  $\lambda > 0$ , the solution is given by the Gibbs posterior (4), which is in this case the generalized Bayesian posterior distribution (5).

For any  $\rho \in (0, 1)$  we can define the  $\rho$ -divergence between probability measures  $P, Q$  as

$$\mathcal{D}_{\rho}(P, Q) = \frac{1}{\rho(1-\rho)} \mathbb{E}_P \left[ 1 - \left( \frac{dQ}{dP} \right)^{\rho} \right],$$

the case  $\rho = 1/2$  giving  $\mathcal{D}_{\rho}(P, Q) = 4\mathcal{H}(P, Q)$ , where  $\mathcal{H}(P, Q) = \mathbb{E}_P(1 - \sqrt{dQ/dP})$  is the Hellinger distance. Note that

$$\frac{1}{2} \mathcal{H}^2(P, Q) \leq \max(\rho, 1-\rho) \mathcal{D}_{\rho}(P, Q) \quad (14)$$

for any  $\rho \in (0, 1)$ , see [? ].

## 6.2 Proof of Theorem 3.1

Since  $1-x \leq -\ln x$  for any  $x \geq 0$ , we have

$$\lambda(1-\lambda) \mathcal{D}_{\lambda}(f^*, f_{\hat{\theta}(\eta)}) \leq -\ln \mathbb{E}_{X_1} \exp(-\ell_{\eta}(X_1)),$$

where we consider  $\ell_{\eta}(x) = \lambda \ln \frac{f^*(X)}{f_{\hat{\theta}(\eta)}(X)}$ . So, using (12) and the definition of  $\hat{\pi}_{\lambda}$ , we have

$$\begin{aligned} \lambda(1-\lambda) \mathbb{E}_{\mathbb{X}_1} \mathbb{E}_{\eta \sim \hat{\pi}_{\lambda}} \mathcal{D}_{\lambda}(f^*, f_{\hat{\theta}(\eta)}) &\leq \mathbb{E}_{\mathbb{X}_1} \left[ \frac{\lambda}{n} \mathbb{E}_{\eta \sim \hat{\pi}_{\lambda}} \sum_{i=1}^n \ln \frac{f^*(X_i)}{f_{\hat{\theta}(\eta)}(X_i)} + \frac{1}{n} \mathcal{K}(\hat{\pi}_{\lambda}, \pi) \right] \\ &= \mathbb{E}_{\mathbb{X}_1} \left[ \inf_{\hat{\pi}} \left( \frac{\lambda}{n} \mathbb{E}_{\eta \sim \hat{\pi}} \sum_{i=1}^n \ln \frac{f^*(X_i)}{f_{\hat{\theta}(\eta)}(X_i)} + \frac{1}{n} \mathcal{K}(\hat{\pi}, \pi) \right) \right] \\ &\leq \inf_{\hat{\pi}} \left( \lambda \mathbb{E}_{\eta \sim \hat{\pi}} \mathcal{K}(f^*, f_{\hat{\theta}(\eta)}) + \frac{1}{n} \mathcal{K}(\hat{\pi}, \pi) \right), \end{aligned}$$

where the infimum is taken among any probability measure on  $\Upsilon$ . Using (14), this leads to

$$\begin{aligned} \frac{1}{2} \mathbb{E}_{\mathbb{X}_1} \mathbb{E}_{\eta \sim \hat{\pi}_{\lambda}} \mathcal{H}(f^*, f_{\hat{\theta}(\eta)}) &\leq \max(\lambda, 1-\lambda) \mathbb{E}_{\mathbb{X}_1} \mathbb{E}_{\eta \sim \hat{\pi}_{\lambda}} \mathcal{D}_{\lambda}(f^*, f_{\hat{\theta}(\eta)}) \\ &\leq \frac{\max(\lambda, 1-\lambda)}{\lambda(1-\lambda)} \inf_{\hat{\pi}} \left( \lambda \mathbb{E}_{\eta \sim \hat{\pi}} \mathcal{K}(f^*, f_{\hat{\theta}(\eta)}) + \frac{1}{n} \mathcal{K}(\hat{\pi}, \pi) \right). \end{aligned}$$

By considering only the subset of Dirac distributions over  $\Upsilon$ , we obtain

$$\mathbb{E}_{\mathbb{X}_1} \mathbb{E}_{\eta \sim \hat{\pi}_{\lambda}} \mathcal{H}(f^*, f_{\hat{\theta}(\eta)}) \leq c_{\lambda} \inf_{\substack{K \in \mathbb{N}^* \\ S \subset \{1, \dots, d\}}} \left( \lambda \mathcal{K}(f^*, f_{\hat{\theta}(K, S)}) + \frac{1}{n} \ln \left( \frac{1}{\pi(K, S)} \right) \right),$$

where  $c_{\lambda}$  is defined in the statement of Theorem 3.1. Since

$$\ln \left( \frac{1}{\pi(K, S)} \right) = \ln \left( \frac{1}{\pi_{\text{clust}}(K)} \right) + \ln \left( \frac{1}{\pi_{\text{supp}}(S)} \right) \leq \ln K! + 1 + 2|S| \ln \left( \frac{ed}{|S|} \right),$$

it gives that

$$\mathbb{E}_{\mathbb{X}_1} \mathbb{E}_{\eta \sim \hat{\pi}_\lambda} \mathcal{H}^2(f^*, f_{\hat{\theta}(\eta)}) \leq c\lambda \inf_{\substack{K \in \mathbb{N}^* \\ S \subset \{1, \dots, d\}}} \left\{ \lambda \mathcal{K}(f^*, f_{\hat{\theta}(K,S)}) + \frac{\ln K! + 1 + 2|S| \ln(ed/|S|)}{n} \right\}.$$

The other inequalities given in Theorem 3.1 are straightforward using the convexity properties of the Hellinger distance (see for instance Lemma 7.25 in [? ]).

### 6.3 Proof of Theorem 3.2

We start with an elementary lemma for bounding sup-norm of density ratios.

**Lemma 6.1.** *Let  $f^*$  be density in  $\mathbb{R}^d$  such that  $\|f^*\|_\infty < \infty$  and such that the support of  $f^*$  is included in  $B(0, \bar{\mu})$ . Then, for any GMM shape, for any  $\eta \in \Upsilon$  and any  $\theta \in \Theta_r$ :*

$$1 \leq \left\| \frac{f^*}{f_\theta} \right\|_\infty \leq \|f^*\|_\infty L^+ \exp\left(\frac{2\bar{\mu}^2}{\bar{\sigma}^2}\right)$$

*Proof.* Let  $\eta = (K, S) \in \Upsilon$  and  $\theta \in \Theta_c$ . We have  $\left\| \frac{f^*}{f_\theta} \right\|_\infty \geq 1$  because  $f^* \ll f_\theta$  since  $f^*$  has a bounded support. For the other inequality, first note that

$$\left\| \frac{f^*}{f_\theta} \right\|_\infty \leq \frac{\|f\|_\infty}{\inf_{x \in B(0, \mu)} f_\theta} \quad (15)$$

According to the constraints on the determinant of the covariance matrices, for any  $x \in B(0, \mu)$ :

$$\begin{aligned} f_\theta(x) &\geq \frac{1}{L^+} \sum_{k=1}^K p_k \exp\left[-\frac{1}{2}(x - \mu_k)' \Sigma_k^{-1} (x - \mu_k)\right] \\ &\geq \frac{1}{L^+} \exp\left[-\frac{1}{2} \sum_{k=1}^K p_k \|x - \mu_k\|^2 \max(\text{sp}(\Sigma^{-1}))\right] \\ &\geq \frac{1}{L^+} \exp\left[-\frac{2\bar{\mu}^2}{\bar{\sigma}^2}\right]. \end{aligned}$$

and the Lemma is proved using (15). □

It is well known that rates of convergence of ML-estimators can be stated by computing bracketing entropies of the statistical models involved, (see [? ? ] among others). Let  $\mathcal{F}$  be a set of densities with respect of the Lebesgue measure. An  $\varepsilon$ -bracketing for  $\mathcal{F}$  with respect to  $\mathcal{H}$  is a set of integrable function pairs  $(l_1, m_1), \dots, (l_N, m_N)$  such that for each  $f \in \mathcal{F}$ , there exists  $j \in \{1, \dots, N\}$  such that  $l_j \leq f \leq m_j$  and  $\mathcal{H}(l_j, m_j) \leq \varepsilon$ . The bracketing number  $N_{[\cdot]}(\varepsilon, \mathcal{F}, \mathcal{H})$  is the smallest number of  $\varepsilon$ -brackets necessary to cover  $\mathcal{F}$  and the bracketing entropy is defined by  $H_{[\cdot]}(\varepsilon, \mathcal{F}, \mathcal{H}) = \ln \{N_{[\cdot]}(\varepsilon, \mathcal{F}, \mathcal{H})\}$ .

Rates of convergences of MLE in GMM was first studied by [? ] and [? ], following a method introduced by [? ]. Here we use the following result which can be, easily rewritten from the proof of Theorem 7.11 in [? ]. This result gives an exponential deviation bound for the Hellinger risk of a maximum likelihood estimator. Note that for our problem we do not need an uniform control of the risk of the estimators over the model collection.

Assume that there exists a nondecreasing function  $\Psi$  such that  $x \rightarrow \Psi(x)/x$  is nonincreasing on  $]0, +\infty[$  and such that for any  $\xi \in \mathbb{R}_+$  and any  $u \in \mathcal{F}$ :

$$\int_0^\xi \sqrt{\mathcal{H}_{[\cdot]}(x, \mathcal{F}(g, \xi), \mathcal{H})} dx \leq \Psi(\xi) \quad (16)$$

where  $\mathcal{F}(g, \xi) := \{t \in \mathcal{F}; \mathcal{H}(t, g) \leq \xi\}$ .

**Theorem 6.2** (Adapted from the proof of Theorem 7.11 in [? ]). *Under the previous assumptions, let  $\hat{f}$  be a MLE on  $\mathcal{F}$  defined using a sample  $X_1, \dots, X_n$  of i.i.d. random variables with density  $f^*$ . Let  $\bar{f} \in \mathcal{F}$  such that  $\mathcal{H}^2(f^*, \bar{f}) \leq 2 \inf_{g \in \mathcal{F}} \mathcal{H}^2(f^*, g)$  and let  $\xi_n$  denotes the unique positive solution of the equation*

$$\Psi(\xi_n) = \sqrt{n} \xi_n^2. \quad (17)$$

Then, there exists an absolute constant  $\kappa'$  such that, except on a set of probability  $\exp(-x)$ ,

$$\mathcal{H}^2(f^*, \hat{f}) \leq \frac{4}{2 \ln 2 - 1} \mathcal{K}(f^*, \mathcal{F}) + \kappa' \left( \xi_n^2 + \frac{x}{n} \right) + (P_n - P) \left( \frac{1}{2} \ln \frac{\bar{f}}{f^*} \right)$$

where  $P$  is the probability measure of density  $f^*$  and  $P_n$  is the empirical measure for the observations  $X_1, \dots, X_n$ .

Since  $(P_n - P) \left( \frac{1}{2} \ln \frac{\bar{f}}{f^*} \right)$  is centered at expectation, by integrating this tail bound we find that

$$\mathbb{E}_{f^*} \mathcal{H}^2(f^*, \hat{f}) \leq \frac{4}{2 \ln 2 - 1} \mathcal{K}(f^*, \mathcal{F}) + \kappa' \left( \xi_n^2 + \frac{1}{n} \right). \quad (18)$$

For a fixed shape, a configuration  $\eta$  and the bounding parameters  $\bar{\mu}, \underline{\sigma} < 1 < \bar{\sigma}, \underline{L} < \bar{L}$ , remember that  $\mathcal{F}_\eta$  is the set of Gaussian mixture densities parametrized by  $\Theta_r(\eta)$ . The following control of the bracketing entropy can be found in [? ].

**Lemma 6.3.** *For a fixed GMM shape and for all  $u \in (0, 1)$ ,*

$$H_{[\cdot]} \left( \frac{u}{9}, \mathcal{F}_\eta, \mathcal{H} \right) \leq \mathcal{I}(\eta) + D(\eta) \ln \frac{1}{u}$$

where  $\mathcal{I}$  is a constant depending on the GMM shape and the bounding parameters. Moreover, for all  $\xi > 0$ ,

$$\int_0^\xi \sqrt{\mathcal{H}_{[\cdot]}(u, \mathcal{F}_\eta, \mathcal{H})} du \leq \Psi_\eta(\xi) := \xi \sqrt{D(\eta)} \left\{ \mathcal{A} + \sqrt{\ln \left( \frac{1}{1 \wedge \xi} \right)} \right\} \quad (19)$$

where the constant  $\mathcal{A}$  depends on the GMM shape and the bounding parameters.

We are now in position to finish the proof of Theorem 3.2. Remember that the sample  $\mathbb{X}_2$  is used for computing the maximum likelihood estimators. For a fixed shape, and a given  $\eta \in \Upsilon$ , let  $\xi_{n_2}$  satisfying (17) :  $\Psi_\eta(\xi_{n_2}) = \sqrt{n_2} \xi_{n_2}^2$ . Note that  $\sqrt{\frac{D(\eta)}{n_2}} \mathcal{A} \leq \xi_{n_2}$ , thus we have

$$\xi_{n_2}^2 \leq \frac{D(\eta)}{n_2} \left\{ 2\mathcal{A}^2 + 2 \ln^+ \left( \frac{n_2}{\mathcal{A}^2 D(\eta)} \right) \right\}.$$

Finally, using Theorem 3.1, Lemma 6.1 and Inequalities (18) and (8), we find that

$$\mathbb{E}_{\mathbb{X}_1 \times \mathbb{X}_2} \mathcal{H}^2(f^*, \mathbb{E}_{\eta \sim \hat{\pi}_\lambda} f_{\hat{\theta}(\eta)}) \leq c_\lambda \inf_{\substack{K \in \mathbb{N}^* \\ S \subset \{1, \dots, d\}}} \left\{ \frac{\ln K! + 1 + 2|S| \ln(ed/|S|)}{n_1} \right. \\ \left. + \lambda C \left[ \mathcal{K}(f^*, \mathcal{F}_\eta) + \kappa \left( \frac{D(\eta)}{n_2} \left\{ \mathcal{A}^2 + \ln^+ \left( \frac{n_2}{\mathcal{A}^2 D(\eta)} \right) \right\} + \frac{1}{n_2} \right) \right] \right\}$$

where  $C = \frac{8}{2 \ln 2 - 1} \left( 1 + \ln(\|f^*\|_\infty L^+) + \frac{2\bar{\mu}^2}{\sigma^2} \right)$  and  $\kappa = \kappa' \frac{2 \ln 2 - 1}{2}$ .

## Acknowledgements

The authors wish to thank , C. Meynet and C. Maugis for helpful discussions.