



HAL
open science

Full Model Estimation for Non-Parametric Multivariate Finite Mixture Models

Marie Du Roy de Chaumaray, Matthieu Marbac

► **To cite this version:**

Marie Du Roy de Chaumaray, Matthieu Marbac. Full Model Estimation for Non-Parametric Multivariate Finite Mixture Models. *Journal of the Royal Statistical Society: Series B*, 2024, 86 (4), pp.896-921. 10.1093/jrssb/qkae002 . hal-04919421

HAL Id: hal-04919421

<https://hal.science/hal-04919421v1>

Submitted on 29 Jan 2025

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.

Copyright

Full Model Estimation for Non-Parametric Multivariate Finite Mixture Models

Marie Du Roy de Chaumaray¹ and Matthieu Marbac²

¹Univ Rennes, CNRS, IRMAR - UMR 6625, F-35000 Rennes, France

²Univ. Rennes, Ensai, CNRS, CREST - UMR 9194, F-35000 Rennes, France

January 8, 2024

Abstract

This paper addresses the problem of full model estimation for non-parametric finite mixture models. It presents an approach for selecting the number of components and the subset of discriminative variables (*i.e.*, the subset of variables having different distributions among the mixture components) by considering an upper bound on the number of components (this number being allowed to increase with the sample size). The proposed approach considers a discretization of each variable into B bins and a penalization of the resulting log-likelihood. Considering that the number of bins tends to infinity as the sample size tends to infinity, we prove that our estimator of the model (number of components and subset of relevant variables for clustering) is consistent under a suitable choice of the penalty term. The relevance of our proposal is illustrated on simulated and benchmark data.

Keywords: Empirical process, Latent class model, Locally conic model, Model selection, Non-parametric mixture model.

1 Introduction

Finite mixture models permits to achieve clustering by estimating the distribution of the observed variables (McLachlan and Peel, 2000; Fruhwirth-Schnatter et al., 2019). This paper focuses on a full model selection (*i.e.*, estimation of the number of components and detection of the subset of the relevant variables for clustering) for non-parametric mixture models where no assumptions are made on the component distribution except to be defined as a product of univariate densities (see Chauveau et al. (2015) for a review). Thus, we consider a sample composed of n independent observations $\mathbf{X}_1, \dots, \mathbf{X}_n$ where $\mathbf{X}_i = (X_{i1}, \dots, X_{iJ})^\top \in \mathcal{X}$ is the vector composed of the J variables collected on subject i defined over the space $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_J$ where each \mathcal{X}_j is compact¹. Each \mathbf{X}_i is identically distributed according to the non-parametric mixture of K components defined by the density

$$g(\mathbf{x}_i) = \sum_{k=1}^K \pi_k \prod_{j=1}^J \eta_{kj}(x_{ij}), \quad (1)$$

where $\boldsymbol{\pi} = (\pi_1, \dots, \pi_K)^\top$ is a finite dimensional parameter belonging to the simplex of size K , $\mathcal{S}_K = \{\mathbf{u} \in [0, 1]^K : \sum_{k=1}^K u_k = 1\}$ and where the univariate densities η_{kj} constitute infinite dimensional parameters and are supposed to be strictly positive except on a set of Lebesgue measure zero. Model (1) has been used in different fields like in behavioral science (Clogg, 1995), econometry (Hu et al., 2013; Compiani and Kitamura, 2016) or sociology (Hagenaars and McCutcheon, 2002). One standard situation where the conditional independence assumption implied by (1) holds true, is the framework of standard repeated measure random effect model, where the subject-level effect is replaced by a component-level

¹This assumption is made for ease of writing. It could seem restrictive, but real data can always be considered as generated from a compactly supported density. Indication on how to relax this assumption is given in the following sections.

effect. Among the recent developments related to (1), one can cite the papers of Hall and Zhou (2003), Hall et al. (2005) and Allman et al. (2009) who studied the parameter identifiability, while Benaglia et al. (2009), Levine et al. (2011) and Zheng and Wu (2020) proposed different algorithms for estimating the parameters when the number of components K is known.

Even if the estimation of the groups is a key-point for clustering, estimating the number of components in (1) was still an open problem until recently. Thus, before the works of Kasahara and Shimotsu (2014), Bonhomme et al. (2016), Bonhomme et al. (2016) and Kwon and Mbakop (2021), no tools were available for selecting the number of components K . This was the main drawback of non-parametric mixture models compared to parametric mixture models which allow model selection through information criteria. The two recent papers of Kasahara and Shimotsu (2014) and Kwon and Mbakop (2021) introduced two approaches for determining the smallest value of K such that (1) holds true. Kasahara and Shimotsu (2014) provide an estimation of the lower bound of the number of components by considering a partition of the support of each variable (*e.g.*, using a decomposition into bins) and by using the identifiability of the latent class model (*i.e.*, mixture models where each component is a product of multinomial distributions). This discretization allows to consider the tensor defining the probability of each event while the rank of this tensor permits to derive a lower bound on K . Note that previous works on non-parametric mixture models considered a bin decomposition (*i.e.*, a specific discretization method) to estimate non-parametric mixture models or study their identifiability but not for model selection (Hettmansperger and Thomas, 2000; Cruz-Medina et al., 2004; Elmore et al., 2004). However, Kasahara and Shimotsu (2014) do not provide a method for selecting the discretization (*i.e.*, number of elements, locations of those elements). Thus, their method is only consistent to a lower bound of K (see Section 2.3 in Kwon and Mbakop (2021)). The works of Bonhomme et al. (2016), Bonhomme et al. (2016) and Bonhomme et al. (2017) consider transformed moments for density estimation. Moreover, in the Section 5 of Supplementary material of Bonhomme et al. (2016) a rank-based procedure using the approach of Kleibergen and Paap (2006) is proposed to estimate the number of components. Thus, this result generalizes the approach of Kasahara and Shimotsu (2014) by allowing other types of approximating functions than the discretization (*i.e.*, orthogonal serie approximation, *e.g.* Jacobi or Hermite polynomials) to apply rank-based approaches for estimating the number of components. Alternatively, Kwon and Mbakop (2021) consider an integral operator, identified from the distribution of X , that has a rank equal to K . Noting that the singular values of operators are stable under perturbations (to handle the fact that this operator is estimated from the observed sample), a thresholding rule allowing to count the number of non-zero singular values provides a consistent estimator of K . One advantage of the approach of Kwon and Mbakop (2021) is to avoid the use of discretization, even if some connections can be established with the approach of Kasahara and Shimotsu (2014) (see Section 2.3 in Kwon and Mbakop (2021)). One elegant property of the methods of Kasahara and Shimotsu (2014) and Kwon and Mbakop (2021) is that both methods determine an estimator of K without performing the density estimation for different numbers of clusters. Thus, those methods start with a step of model selection followed by the estimation of the parameters for the selected model. This is quite unusual because, when model selection is conducted for parametric mixture models via an information criterion, parameter estimation needs to be first performed for each competing model in order to compute the information criterion. The approaches of Kasahara and Shimotsu (2014) and Kwon and Mbakop (2021) are mainly based on the distribution of a couple of variables. Thus, if the number of variables J is large, computational issues can arise while considering all possible pairs of variables. It restricts the use of their methods to data sets composed of few variables. This is also the case for the other rank-based approaches as they rely on transformed moments of multivariate variables (see Section 3.1 in Bonhomme et al. (2016)). Moreover, the nature of all those approaches makes impossible a variable selection.

Selecting variables is challenging in clustering because the role of a variable (relevant or irrelevant for clustering) is defined from the partition which is unobserved. Thus, the selection of the variables and the clustering need to be performed simultaneously. Note that selecting the variables in clustering has two strong benefits: it improves the accuracy of the estimators (Azizyan et al., 2013) because it reduces the number of estimators to be considered and it facilitates the interpretation of the different components as it only has to be made on the subset of discriminative variables. In a non-model based framework, regularization methods can be used to achieve variable selection in clustering (Friedman and Meulman, 2004; Pan and Shen, 2007; Witten and Tibshirani, 2010). Among these approaches, the sparse K-means (Witten and Tibshirani, 2010) is the most popular because it requires small computational overhead and is able to manage very high-dimensional datasets. The approach uses a lasso-type penalty to select

the set of variables which are relevant for the clustering. The selection of the number of components is a difficult issue since probabilistic tools are not available and its results are sensitive to the structure of the penalty term. The authors proposed to perform model selection with an extension of the gap statistics (Tibshirani et al., 2001). In model-based clustering, Tadesse et al. (2005) define a variable as irrelevant for clustering if its distribution is equal among all the mixture components. A third type of variable, the redundant ones, has been introduced by Raftery and Dean (2006). A redundant variable is a variable that has the same conditional distribution given the relevant variables for each component. However, this type of variable cannot be considered for model (1), as it requires to model intra-component dependencies. The model-based framework implies that the selection of the variables falls in the scope of model selection. Thus, in the paper, we define the objective of full model selection by the double objective of estimating the number of components K and the subset of relevant variables Ω as well. In a parametric framework, variable selection can be performed via information criterion (see Tadesse et al. (2005), Dean and Raftery (2010) and Marbac and Sedki (2017)) but it leads to computational issues because the number of competing models is of order 2^J . To circumvent this issue, Marbac et al. (2019) present a modified EM algorithm (Dempster, A. P. and Laird, N. M. and Rubin, D. B., 1977; Green, 1990) that permits to simultaneously perform the variable selection via the Bayesian Information Criterion (BIC; Schwarz (1978)) and the maximum likelihood estimation, for a fixed number of components. Thus, the algorithm only needs to be run for all the possible numbers of components. Note that all existing methods of variables selection in model-based clustering are restricted to parametric distributions. Thus, if the parametric assumptions are violated, biased estimates of the number of components or on the subset of discriminative variables can be provided.

This paper addresses the issue of full model selection for non-parametric mixture models defined by (1). To the best of our knowledge, this paper presents the first method that permits a full-model selection (*i.e.*, estimation of K and Ω) for non-parametric multivariate mixture models. This double objective is not covered by the rank-based methods (Kasahara and Shimotsu (2014), Bonhomme et al. (2016) or Kwon and Mbakop (2021)) that consider, by nature, that all the variables are relevant for clustering and hence that do not consider the objective of variable selection. Moreover, the proposed method allows many variables to be managed, which makes it a complementary work to the rank-based methods, even in the case where all the variables are considered to be relevant and only the number of components needs to be estimated. As proposed by Tadesse et al. (2005), we consider two types of variables: the relevant variables and the irrelevant variables for clustering. Thus, variable j is said to be irrelevant for clustering if $\eta_{1j} = \dots = \eta_{Kj}$ and the variable j is said to be relevant for clustering otherwise. A model $M = \{K, \Omega\}$ is defined by the number of components K and the indices of the relevant variables $\Omega \subseteq \{1, \dots, J\}$. Therefore, considering the task of full model selection in (1) implies that each \mathbf{X}_i is identically distributed according to a non-parametric mixture of K components defined by the density

$$g_{M, \boldsymbol{\psi}}(\mathbf{x}_i) = \left(\prod_{j \in \bar{\Omega}} \eta_{1j}(x_{ij}) \right) \left(\sum_{k=1}^K \pi_k \prod_{j \in \Omega} \eta_{kj}(x_{ij}) \right), \quad (2)$$

where $\bar{\Omega} = \{1, \dots, J\} \setminus \Omega$ contains the indices of the irrelevant variables for clustering and $\boldsymbol{\psi} \in \Psi_M$ groups the finite dimensional parameters $\boldsymbol{\pi} = (\pi_1, \dots, \pi_K)^\top \in \mathcal{S}_K$ and the infinite dimensional parameters composed of the univariate densities η_{kj} . The space of competing models is defined by an upper bound K_{\max} on the number of components while each variable can be relevant or not. The upper bound on the number of components is supposed to be known but is allowed to tend to infinity as the sample size tends to infinity. To achieve the full model selection, we use a discretization of each continuous variable into B bins. The number of bins tends to infinity with the sample size to ensure the consistency of the approach. Indeed, if B would be fixed, the estimated model could be a sub-model of the true model (*i.e.*, the number of components and the subset of the discriminative variables could be underestimated). The distribution of the resulting discretized variables follows a latent class model where each component is a product of multinomial distributions (Goodman, 1974). Thus, despite the fact that other decompositions (*e.g.*, Jacobi or Hermite polynomials) are used for rank-based methods, we choose to use the discretization to obtain "categorical" variables. This discretization is convenient, because model selection can then be achieved, for the resulting latent class model, by using the penalized likelihood (*e.g.*, BIC) whose consistency has been proven for a large class of finite mixture models in Keribin (2000) and in particular for multivariate multinomial mixtures in Toussile and Gassiat (2009). Moreover, thanks

to the discretization, a specific EM algorithm optimizing the penalized likelihood can be used for simultaneously detecting the subset of the relevant variables and estimating the model parameters (Marbac et al., 2019), for a known number of components. Thus, by considering an upper-bound of the number of components, the full-model selection can be achieved. Unlike in Kasahara and Shimotsu (2014), the procedure provides a consistent estimation on the univariate densities of the components η_{kj} from the discretized data. Therefore, we prove the consistency of the procedure for a wide range of numbers of bins B , at an appropriate rate that we detail in the paper. In this paper, for ease of reading, we consider that the number of bins used for the discretization is the same for each variable, but the method can be extended to different numbers of bins per variables. Note that the proof of the consistency is written by considering an upper-bound of the number of components (see Theorem 1). This assumption was made to ensure the consistency of penalized log-likelihood for parametric mixture (Keribin, 2000). However, this assumption is also implicitly made by the rank-based method of Kasahara and Shimotsu (2014) and Bonhomme et al. (2016), since the number of components cannot exceed the number of elements used in the decomposition (otherwise, Assumption 1 in Bonhomme et al. (2016) as well as Assumption 3 in Bonhomme et al. (2017) cannot hold true for instance). Note that, by the use of the integral operators, Kwon and Mbakop (2021) do not need such an assumption. In this paper, we also provide sufficient conditions on the growth of both B and K_{\max} that lead to the consistency of the procedure by allowing the upper bound on the number of components to tend to infinity as the sample size tends to infinity (see Corollary 1). This case is not covered by the rank-based methods, as in Kasahara and Shimotsu (2014) and Bonhomme et al. (2016), since the number of elements used in their decomposition is fixed. Obviously, their procedures could be developed with an increasing number of elements, but the rate of increasing should be investigated since the results of Kleibergen and Paap (2006) cannot be applied on matrices with increasing dimensions (and fixed rank). Note that the consistency of the proposed procedure cannot be stated directly by using the consistency of information criterion for parametric mixture models (Keribin, 2000; Toussile and Gassiat, 2009) because the parameters space depends on B and thus increases with the sample size. The growth rate of B is mainly driven to avoid underestimation of the model while the range of the penalty is mainly driven to avoid overestimation of the model. Note that for categorical data, a non-asymptotic oracle inequality is obtained in Bontemps and Toussile (2013), which permits a data-driven choice of the penalty through slope heuristics. The case of model underestimation is analyzed by extending the proof of Keribin (2000) in order to deal with the increasing dimension of the parameter space. In the case of model overestimation, the asymptotic distribution of the likelihood ratio is investigated by performing a locally conic parametrization (Dacunha-Castelle and Gassiat, 1997, 1999) of the model obtained on the discretized data. An upper bound of the likelihood ratio is obtained by controlling, on the one hand, the deviation of the likelihood ratio from its asymptotic distribution by using results on empirical processes stated in Chernozhukov et al. (2014) and, on the other hand, the supremum of the asymptotic distribution by applying deviation results on Gaussian processes (Dudley, 2014).

The proposed method uses a discretization that provides an estimator of the densities of the components. However, we advise to use the proposed approach only for model estimation. When the model has been selected, we suggest to use a kernel-based method for density estimation. Indeed, the bin-density estimates are known to be outperformed by kernel-based estimators. Thus, for a real data analysis, we advise to use the proposed approach for model selection then, for the selected model, to perform density estimation with an EM-like algorithm (Benaglia et al., 2009) or by maximizing the smoothed log-likelihood (Levine et al., 2011). The final partition is thus computed from the model selected by the proposed method and the densities estimated via a kernel method.

The paper is organized as follows. Section 2 details the discretization step. Section 3 states the consistency of the procedure. Section 4 presents the algorithm used for the full model selection. Section 5 starts by numerical experiments that compare the proposed approach to standard approaches of variable selection in clustering, then it presents the analysis of benchmark data which illustrates the relevance of the procedure and introduces some extensions of the approach. Section 6 gives a conclusion. Mathematical details and supplementary numerical experiments are presented in Supplementary Materials.

2 Model selection by bin estimation and penalized log-likelihood

This section considers the estimation of the number of components K as well as the subset Ω of relevant variables for clustering for model (2), from a n -sample $\mathbf{X}_1, \dots, \mathbf{X}_n$ with $\mathbf{X}_i \in \mathcal{X}$ with $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_J$, J being fixed and \mathcal{X}_j being compact. The method is based on the discretization of each variable into B non-overlapping bins² I_{Bj1}, \dots, I_{BjB} such that $\cup_{b=1}^B I_{Bjb} = \mathcal{X}_j$ and for any (b, b') with $b \neq b'$, $I_{Bjb} \cap I_{Bjb'} = \emptyset$. Thus, we consider the function σ_{Bjb} with $b \in \{1, \dots, B\}$, such that $\sigma_{Bjb}(x_{ij}) = 1$ if $x_{ij} \in I_{Bjb}$ and $\sigma_{Bjb}(x_{ij}) = 0$ if $x_{ij} \notin I_{Bjb}$, and we denote by l_{Bjb} the size of the bin I_{Bjb} . The discretized variables follow a latent class model where each component is a product of J multinomial distributions having B levels each. Therefore, the pdf of the discretized subject i is

$$f_{M,B,\theta}(\mathbf{x}_i) = \prod_{j \in \bar{\Omega}} \prod_{b=1}^B \left(\frac{\alpha_{B1jb}}{l_{Bjb}} \right)^{\sigma_{Bjb}(x_{ij})} \left(\sum_{k=1}^K \pi_k \prod_{j \in \Omega} \prod_{b=1}^B \left(\frac{\alpha_{Bkjb}}{l_{Bjb}} \right)^{\sigma_{Bjb}(x_{ij})} \right), \quad (3)$$

where θ groups the component proportions π_k and the probabilities α_{Bkjb} that one subject arisen from component k takes level b for the variable j when this variable is discretized into B bins. The parameter space is given by the product of simplexes $S_K \times S_B^{K|\Omega|+(J-|\Omega|)}$, where $|\Omega|$ denotes the cardinal of the set of discriminative variables Ω . Note that $f_{M,B,\theta}$ is an approximation of $g_{M,\psi}$ and that this approximation becomes more accurate when B tends to infinity. The discretized version f_{M_0,B,θ_0} of the true density g_0 will be denoted by $f_{0,B}$. The limitation to densities with compact support may seem to be restrictive, but real data can always be considered as generated from a compactly supported density, *i.e.*, over the range of the data, since we cannot get any information about what happens outside this range without additional assumptions. Thus, the proposed method can be directly extended to distributions defined on non-compact spaces by focusing on density estimation made on the convex hull of the data and assuming that the density is zero outside. Indeed, as discussed at the end of Section 3.1, this will not impact the results of likelihood-based methods, as they will never be evaluated outside this range. Note that, in such a case, assumptions should be made on the tails of the distribution in order to ensure that the growth of the convex hull still permits to satisfy the assumptions of the theorem and in particular Assumptions 4.

The probabilities α_{Bkjb} are unknown and must be estimated from the observed sample. This estimation can be achieved by maximizing the log-likelihood defined by

$$\ell_n(f_{M,B,\theta}) = \sum_{i=1}^n \ln f_{M,B,\theta}(\mathbf{x}_i).$$

The maximum likelihood statistics for a model M and B bins per variables is

$$T_{n,M,B} = \sup_{\theta \in \Theta_{M,B,\varepsilon}} \ell_n(f_{M,B,\theta}),$$

where, in order to avoid numerical issues, we introduced a threshold ε such that the parameter space becomes $\Theta_{M,B,\varepsilon} = S_{K,\varepsilon} \times S_{B,\varepsilon}^{K|\Omega|+(J-|\Omega|)}$, with $\varepsilon > 0$ being the minimal value of all the elements defined in the simplexes, *i.e.*, $S_{B,\varepsilon} = \{\mathbf{u} \in \mathbb{R}^B : u_b > \varepsilon, \sum_{b=1}^B u_b = 1\}$. Note that the use of such a threshold is quite usual in this framework (see Toussile and Gassiat (2009); Bontemps and Toussile (2013)). Under the condition that $B\varepsilon$ tends to zero as B goes to infinity and ε to zero, the parameter space $\Theta_{M,B,\varepsilon}$ converges to the whole parameter space. Note that, due to the growth rate of B which will be stated by Assumption 4(i) in the next section, it is sufficient to set $\varepsilon^{-1} = O(n^{\alpha+1})$ for some $\alpha > 0$. This maximization can be achieved via an Expectation-Maximization algorithm (EM algorithm; Dempster, A. P. and Laird, N. M. and Rubin, D. B. (1977)).

The penalized likelihood is defined by subtracting from the maximum likelihood statistics a penalty term $a_{n,M,B}$ which takes into account the sample size and the complexity of model M . Thus, we obtain the following information criterion

$$W_{n,M,B} = T_{n,M,B} - a_{n,M,B}. \quad (4)$$

²Alternatively, one could use different numbers of bins per variables as long as all those numbers respect Assumptions 3 and 4 given in the next section. Here, for ease of reading, we focus on the case where all variables are discretized with the same number of bins.

Depending on the choice of $a_{n,M,B}$ in (4), different well-known criteria can be considered. Among them one can cite the Akaike criterion (AIC; Akaike (1970)) or the Bayesian Information Criterion (BIC; Schwarz (1978)) which are obtained with $a_{n,M,B} = \nu$ and $a_{n,M,B} = \nu \log(n)/2$ respectively, where ν denotes the model complexity that is defined for (3) by $\nu = (K - 1) + |\Omega|K(B - 1) + (J - |\Omega|)(B - 1)$.

We consider the set of competing models \mathcal{M} defined by all the mixture models with at most K_{\max} components and at least three relevant variables (for identifiability reasons), so that

$$\mathcal{M} = \{M = \{K, \Omega\} : K \leq K_{\max}, \Omega \subseteq \{1, \dots, J\} \text{ and } |\Omega| \geq 3\}.$$

The definition of the set of the competing models \mathcal{M} implies that a knowledge of an upper bound of the number of components K_{\max} is known. Note that this assumption was also made to prove the consistency of the estimator provided by penalized log-likelihood for parametric mixture models (Keribin, 2000). Moreover, for the approach of Kasahara and Shimotsu (2014), the number of bins used in the discretization of each variable is an upper bound on the number of components (see Proposition 1 in Kasahara and Shimotsu (2014)). Similarly, the approaches based on transformed moments implicitly consider an upper bound of the number of components that is defined as number of moments considered in the procedure (otherwise Assumption 3 in Bonhomme et al. (2017) cannot hold true). To the best of our knowledge, only the approach of Kwon and Mbakop (2021) does not require an upper bound on the number of components thanks to the use of integral operators. In the next section, we give sufficient conditions to consistently estimate the model when K_{\max} is fixed. However, we also provide a corollary that gives sufficient conditions to maintain the consistency when the upper bound of the number of components is allowed to tend to infinity with the sample size (note that the true model, including the number of components, does not change with the sample size).

The estimator $\widehat{M}_{n,B}$ of the model maximizes the penalized likelihood as follows

$$\widehat{M}_{n,B} = \arg \max_{M \in \mathcal{M}} W_{n,M,B}.$$

The study of the asymptotic properties of the estimator $\widehat{M}_{n,B}$ is covered by the approach of Keribin (2000) only if the number of intervals B does not increase with the sample size n . However, in such a case, due to the discretization, the approach would provide an estimator that converges to a model included into the true model. Indeed, we only obtain a lower bound on the number of components and a subset of the discriminative variables. By increasing the number of intervals with n , we avoid the issues due to the loss of identifiability. However, we need to investigate the behavior of the statistics $T_{n,M,B}$ and to study the convergence of $T_{n,M,B}/n$ to the minimum Kullback divergence, which requires controlling empirical processes defined on space having increasing dimension. The next section presents statistical guarantees of the proposed approach.

3 Convergence in probability of the estimator

This section investigates the convergence in probability of $\widehat{M}_{n,B}$. It starts by presenting the assumptions required to obtain this convergence, which is then stated.

3.1 Assumptions

The consistency of the estimator is established under four sets of assumptions described below. Assumption 1 and Assumption 2 state the constraints on the model and on the distribution of the components respectively. Assumption 3 gives some conditions on the penalty term. Finally, Assumption 4 gives some conditions on the discretization.

Assumption 1. *The number of variables is at least three (i.e., $3 \leq J$) and each proportion $\pi_k > 0$ is not zero. Moreover, there exists $\Upsilon \subseteq \{1, \dots, J\}$ such that $|\Upsilon| = 3$ and for any $j \in \Upsilon$ the univariate densities η_{k_j} are linearly independent.*

Assumption 2. *(i) There exists a function τ in $L_1(g_0\nu)$ such that: $\forall M \in \mathcal{M}$ and $\forall \psi \in \Psi_M$, $|\ln g_{M,\psi}| < \tau$ ν -a.e.*

- (ii) There exists a positive constant $L < \infty$ such that $\forall j \in \{1, \dots, J\}$ and $\forall x_j \in \mathcal{X}_j$, $|\eta'_{kj}(x_j)| \leq L$.
- (iii) Each variable j is defined on a compact space \mathcal{X}_j and its densities for each component k , denoted by η_{kj} , are strictly positive except on a set of Lebesgue measure zero.

Assumption 3. (i) $a_{n,M,B}$ is an increasing function of K , $|\Omega|$ and B .

- (ii) For any model M , $a_{n,M,B}/n$ tends to 0 as n tends to infinity.
- (iii) For any model M , $B/a_{n,M,B}$ tends to 0 as n tends to infinity.
- (iv) For any models M and \widetilde{M} with $M \subset \widetilde{M}$, the limit inf of $a_{n,\widetilde{M},B}/a_{n,M,B}$ is strictly larger than one as n tends to infinity.

Assumption 4. (i) The number of bins B tends to infinity with n in the following way $\lim_{n \rightarrow \infty} B = \infty$ and $\lim_{n \rightarrow \infty} B(\ln^2 n)/n^{1/2} = 0$.

- (ii) The length of each bin is not zero and satisfies, for any $j \in \{1, \dots, J\}$ and any $b \in \{1, \dots, B\}$, $l_{Bjb}^{-1} = O(B)$.
- (iii) Let \mathcal{I}_{jB} be the set of the upper bounds of the B intervals, then, for any value $x_j \in \mathcal{X}_j$, $d(x_j, \mathcal{I}_{jB})$ tends to zero as B tends to infinity.

Assumption 1 is derived from the conditions of identifiability for finite mixtures of non-parametric product measures (see Theorems 8 and 9 in Allman et al. (2009)). Because Theorems 8 and 9 in Allman et al. (2009) consider all the variables as relevant for clustering, we need to extend their assumptions such that there are at least three relevant variables to obtain the identifiability of the model (2). Note that this assumption implies the identifiability of the parameters that is required for the proposed method since it relies on the penalized likelihood. Remark that this assumption is not required by the rank-based methods.

Assumption 2 gives sufficient conditions on the component distributions to ensure that the results of Dacunha-Castelle and Gassiat (1999) can be applied to the mixture model obtained after discretization. Note that Assumption 2(iii) was used to obtain a lower bound on $\min_{j,b,k} \int_{I_{Bjb}} \eta_{kj}(x_j) dx_j$ (see Section A.3.4). However, unlike Kasahara and Shimotsu (2014) and Kwon and Mbakop (2021), Assumption 2(iii) does not allow for the support of relevant variables to vary across different mixture components and does not allow for some of the variables in X to be discrete. In econometrics, support variations are often used to justify the full-rank condition needed for the identification of finite mixture models (see An et al. (2010)). Since we do not have guarantees on the proposed procedure when Assumption 2(iii) does not hold true, we investigate its performances on numerical experiments by considering mixture of uniform distributions with disjoint supports in Section B.4 of the supplementary material.

Assumption 3 presents standard conditions for penalized likelihood model selection in the case of embedding models. It generalizes the usual conditions for selecting the number of components (Keribin, 2000; Gassiat, 2002; Chambaz et al., 2006) to the case of feature selection for mixture models. Conditions (i) and (iii) permit avoiding the overestimation of the model (*i.e.*, overestimation of the number of components or of the support of the relevant variables), while condition (ii) permits avoiding the underestimation of the model by making the penalty term negligible with respect to the model bias. Note that Assumption 3 allows the BIC penalty to be considered.

Even if Assumption 1 provides sufficient conditions to obtain the identifiability of model (2), after the discretization, model (3) could be not identifiable if the number of intervals B were fixed. As an example, one can consider a bi-component mixture model with equal proportions defined with a first component following a product of $J \geq 3$ beta distributions $\mathcal{B}e(\alpha, \alpha)$ and a second component following a product of $J \geq 3$ beta distributions $\mathcal{B}e(2\alpha, 2\alpha)$, with $\alpha \geq 1$. This model is identifiable but the model (3) defined after the discretization of each variable into two bins of equal size (*e.g.*, for any j , $\sigma_{j1}(u) = 1$ if $0 \leq u \leq 1/2$, $\sigma_{j1}(u) = 0$ if $1/2 < u \leq 1$ and $\sigma_{j2}(u) = 1 - \sigma_{j1}(u)$) is not identifiable (*i.e.*, the two mixture components follow the same distribution for the discretized data). However, if the number of bins is strictly larger than two and if each interval has a length not equal to zero, then the model (3) becomes identifiable.

The model identifiability is obtained by Assumption 4 that states conditions on the discretization. In particular, the number of levels has to tend to infinity when the sample size increases in such a way that

the size of the largest interval tends to zero when the sample size tends to infinity. However, its growth rate has to be upper bounded, which is a key point to control the convergence of the estimators. Note that Assumption 4(iii) uses the same ideas as Lemma 17 in Allman et al. (2009) and that this condition is not stringent. For instance the bounds of the intervals can be determined by the quantiles $1/B, \dots, B/B$. In addition, the sizes l_{jb} can vary from one bin to another. This is for instance the case when we consider the quantiles. However, we cannot allow a bin to be exponentially small with n , in order to keep the asymptotic behavior of our estimator which is stated in the next subsection. Note that Assumption 4 allows to consider the rate $B = n^{1/3}$ that is usual for bin-density estimation. Finally, since we have an upper bound K_{\max} on the number of components, it could be natural to consider discretizations with at most K_{\max} bins. Indeed, from Kasahara and Shimotsu (2014), it is known that there exists at least one discretization ensuring the identification of K . Therefore, considering K_{\max} bins for each discretization could be reasonable. However, because the discretization ensuring the identification of K is unknown (see the counter-example described in the previous paragraph), this approach is not feasible in practice and so we make B tends to infinity (at an appropriate rate to control the convergence of the maximum likelihood statistics). If a procedure would be developed to obtain the discretization ensuring the identification of K , then it would permit to consider a fixed-size discretization. Note that consistency of the procedure of Kasahara and Shimotsu (2014) (and not consistency to a lower bound of K) could be obtained by allowing the number of bins to increase with the sample size, if a control of the statistic of Kleibergen and Paap (2006) could be obtained with an increasing dimension matrix having fixed rank. Similar reasoning could be extended to methods based on transformed moments of the univariate distributions (Bonhomme et al. (2016), Bonhomme et al. (2016) and Bonhomme et al. (2017)) to avoid the assumptions made on rank of some matrices (see for instance Assumption 3 in Bonhomme et al. (2017)). Indeed, such an assumption could not be satisfied for a fixed number of transformed moments. Thus, considering an increasing number of transformed moments would permit avoiding this assumption but would imply to control the accuracy of the estimator and also to adapt the procedure described in Section 5 of the supplementary material of Bonhomme et al. (2016) to the case of matrices with increasing dimensions. Finally, note that increasing B ensures that the model is not underestimated because of the discretization. Increasing B can thus be seen as a reduction of the bias. However, by giving an upper bound on the increase of B , a control of the supremum of the maximum likelihood statistics can be obtained. Upper bounding B can thus be seen as a sufficient condition to control the variance of the estimators.

Finally, the assumption on the compactness of \mathcal{X}_j can be relaxed if some densities defined on \mathbb{R} are wanted to be considered. In such case, the estimates of the densities are considered on the compact $[\min_i x_{ij}, \max_i x_{ij}]$ defined from the observed sample, and the estimates of the densities are zero outside this interval. It could be seen as adding two bins outside the range of the data, with associated parameters α_{Bkjb} being equal to zero. However, note that if the densities are defined on \mathbb{R} then the convex hull of the data will grow with the sample size. Thus, assumptions on the tails of the distributions should be made to ensure that Assumptions 4 still hold. Note that those assumptions are similar to those ensuring the consistency of density estimation per histogram.

3.2 Convergence in probability of the estimator

We state the consistency of the estimator $\widehat{M}_{n,B}$ then we give its (sketch of) proof. Note that the proof of all the numbered equations are given in Section A of the Supplementary Materials. Finally, we explain the key points of the proof which are different from the proof of the consistency of information criteria for parametric mixture models stated in Keribin (2000).

Theorem 1. *Assume that independent data arise from (2) with the true model $M_0 = \{K_0, \Omega_0\}$, that Assumptions 1, 2, 3 and 4 hold true, and that the set of competing models \mathcal{M} is defined with a known upper bound for the number of clusters K_{\max} . Then, $\widehat{M}_{n,B}$ converges in probability to M_0 .*

Proof of Theorem 1. The proof is divided into three parts: the case where M_0 is underestimated (*i.e.*, $K < K_0$ or $\Omega_0 \not\subseteq \Omega$), the case where the subset of the relevant variables is overestimated with K_0 (*i.e.*, $K = K_0$ and $\Omega_0 \subsetneq \Omega$) and the case where the number of components and the subset of relevant variables are overestimated (*i.e.*, $K > K_0$ and $\Omega_0 \subseteq \Omega$).

- **Part 1:** We consider the case where M_0 is underestimated. Thus, we consider the set of models

$$\mathcal{N}_1 = \{M = \{K, \Omega\} \in \mathcal{M} : K < K_0 \text{ or } \Omega_0 \not\subseteq \Omega\}.$$

The probability to underestimate the true model can be upper-bounded as follows

$$\mathbb{P}(\widehat{M}_{n,B} \in \mathcal{N}_1) \leq \sum_{M \in \mathcal{N}_1} \mathbb{P}(W_{n,M,B} - W_{n,M_0,B} \geq 0).$$

For any $g_{M,\psi}$ given by model (2), Assumption 2(i) implies that $\mathbb{E}_{g_0}[\ln g_{M,\psi}]$ is defined. The Kullback-Leibler divergence from model M to the true distribution g_0 is defined by

$$\text{KL}(g_0, \mathcal{G}_M) := \inf_{\psi \in \Psi_M} \mathbb{E}_{g_0} \left[\ln \frac{g_0}{g_{M,\psi}} \right].$$

Noting that by definition $\mathcal{N}_1 = \{M \in \mathcal{M} : K < K_0\} \cup \{M \in \mathcal{M} : \exists j_0 \in \Omega_0 \text{ such that } j_0 \notin \Omega\}$, the model M_0 is not included in any model of \mathcal{N}_1 . Therefore, using the definition of \mathcal{N}_1 and the identifiability of g_0 (ensured by Assumption 1), for each $M \in \mathcal{N}_1$, there exists some $\delta_M > 0$ such that $\text{KL}(g_0, \mathcal{G}_M) \geq \delta_M$. In Section A.1 of the Supplementary Materials, we prove the following asymptotic bound in probability:

$$\frac{1}{n} (T_{n,M,B} - \ell_n(g_0)) \leq -\delta_M + o_{\mathbb{P}}(1). \quad (5)$$

The proof of (5) holds true for the true model M_0 but, in this case, $\delta_{M_0} = 0$, which implies that $\frac{1}{n}|T_{n,M_0,B} - \ell_n(g_0)| = o_{\mathbb{P}}(1)$. Noting that $\frac{1}{n}(W_{n,M,B} - W_{n,M_0,B}) = \frac{1}{n}(T_{n,M,B} - \ell_n(g_0) + \ell_n(g_0) - T_{n,M_0,B} - a_{n,M,B} + a_{n,M_0,B})$ and using the properties of the penalty (see Assumption 3), we deduce that for any $M \in \mathcal{N}_1$

$$\frac{1}{n} (W_{n,M,B} - W_{n,M_0,B}) \leq -\delta_M + o_{\mathbb{P}}(1).$$

Therefore, noting that $\delta_M > 0$ and that the cardinal of \mathcal{N}_1 is fixed and finite, we have

$$\lim_{n \rightarrow \infty} \mathbb{P}(\widehat{M}_{n,B} \in \mathcal{N}_1) = 0. \quad (6)$$

Thus, the probability of underestimating the model tends to zero as n tends to infinity.

• **Part 2:** We consider the case where the number of components is correct but the subset of the relevant variables is overestimated. Thus, we consider a model $M \in \mathcal{N}_2$ where

$$\mathcal{N}_2 = \{M = \{K, \Omega\} \in \mathcal{M} : K = K_0 \text{ and } \Omega_0 \subsetneq \Omega\}.$$

We have the following upper-bound

$$\mathbb{P}(\widehat{M} \in \mathcal{N}_2) \leq \sum_{M \in \mathcal{N}_2} \mathbb{P}(W_{n,M,B} \geq W_{n,M_0,B}) = \sum_{M \in \mathcal{N}_2} \mathbb{P}\left(\frac{T_{n,M,B} - T_{n,M_0,B}}{a_{n,M_0,B}} \geq \frac{a_{n,M,B}}{a_{n,M_0,B}} - 1\right).$$

Noting that $\mathcal{N}_2 = \{M = \{K, \Omega\} \in \mathcal{M} : K = K_0; \forall j_0 \in \Omega_0, j_0 \in \Omega \text{ and } \exists j_1 \in \Omega, j_1 \notin \Omega_0\}$, we have that true model M_0 is included in any model of \mathcal{N}_2 leading that the Kullback-Leibler divergence from any model M to g_0 is null. Since for any $M \in \mathcal{N}_2$, we have $\delta_M = 0$, the reasoning used to demonstrate that M_0 is not underestimated cannot be used here. However, the models on \mathcal{N}_2 are identifiable because $K = K_0$. Thus, using usual results on likelihood ratio, for a fixed value of B , $2(T_{n,M,B} - T_{n,M_0,B})$ is asymptotically distributed like a $\chi^2(\Delta)$ where Δ is given by the difference of the dimensions of both parameter spaces $\Delta = (B - 1)(K - 1)(|\Omega| - |\Omega_0|)$. However, since B tends to infinity as n tends to infinity, the control of $(T_{n,M,B} - T_{n,M_0,B})$ has to be made carefully, see Section A.2, in order to state that under Assumption 3 (iii)

$$\frac{T_{n,M,B} - T_{n,M_0,B}}{a_{n,M_0,B}} = o_{\mathbb{P}}(1). \quad (7)$$

For any $M \in \mathcal{N}_2$, Assumption 3 (iv) implies that $a_{n,M,B}/a_{n,M_0,B} - 1 > 0$, thus, as the cardinal of \mathcal{N}_2 is finite and does not depend on B , we can conclude that

$$\lim_{n \rightarrow \infty} \mathbb{P}(\widehat{M} \in \mathcal{N}_2) = 0. \quad (8)$$

• **Part 3:** We consider the case where the number of components and the subset of the relevant variables are overestimated. Thus, we consider a model $M \in \mathcal{N}_3$ where

$$\mathcal{N}_3 = \{M = \{K, \Omega\} \in \mathcal{M} : K_0 < K \leq K_{\max} \text{ and } \Omega_0 \subseteq \Omega\}.$$

Note that $\mathcal{N}_3 = \mathcal{M} \setminus \{\mathcal{N}_1 \cup \mathcal{N}_2 \cup M_0\}$. The probability of overestimating the model (*i.e.*, $\widehat{M} \in \mathcal{N}_3$) can be upper-bounded by

$$\mathbb{P}(\widehat{M} \in \mathcal{N}_3) \leq \sum_{M \in \mathcal{N}_3} \mathbb{P}(W_{n,M,B} \geq W_{n,M_0,B}).$$

Note that for any $M \in \mathcal{N}_3$, we have $\delta_M = 0$ and thus the reasoning used to demonstrate that M_0 is not underestimated cannot be used. Moreover, because for $M \in \mathcal{N}_3$, $K > K_0$, the model suffers from a loss of identifiability implying that the likelihood ratio must be carefully investigated. We have, for any $M \in \mathcal{N}_3$

$$\mathbb{P}(W_{n,M,B} \geq W_{n,M_0,B}) = \mathbb{P}\left(\frac{T_{n,M,B} - \ell_n(f_{0,B})}{a_{n,M_0,B}} \geq \frac{a_{n,M,B}}{a_{n,M_0,B}} - 1\right).$$

First, we establish an upper-bound of the probability of $\{W_{n,M,B} \geq W_{n,M_0,B}\}$ which holds true for any value of B (see (11)). We then show that, under Assumptions 1-4, the terms of this upper-bound tend to zero with the sample size even if B grows with the sample size.

The distribution of $T_{n,M,B} - \ell_n(f_{0,B})$ can be established for any value of B by applying the locally-conic parametrization proposed by Dacunha-Castelle and Gassiat (1997, 1999) on model (3), and noting that Assumption 2 holds true, we can rewrite the log-likelihood ratio as in the proof of Lemma 3.3 in Keribin (2000)

$$\begin{aligned} & T_{n,M,B} - \ell_n(f_{0,B}) \\ &= \sup \left\{ \sup_{d \in D_B} \frac{1}{2} \mathcal{G}_n^2(d) \mathbb{1}_{\mathcal{G}_n(d) \geq 0}; \sup_{d_1 \in D_{1B}, d_2 \in D_{2B}} \frac{1}{2} (\mathcal{G}_n^2(d_1) + \mathcal{G}_n^2(d_2) \mathbb{1}_{\mathcal{G}_n(d_2) \geq 0}) \right\} (1 + o_{\mathbb{P}}(1)) \end{aligned} \quad (9)$$

where, for each function d , $\mathcal{G}_n(d) = n^{-1/2} \sum_{i=1}^n d(X_i)$; the considered spaces of functions as well as the definition of $f_{0,B}$ are detailed in Section A.3 of the Supplementary Material. Note that

$$\sup_{d \in D_B} \frac{1}{2} \mathcal{G}_n^2(d) \mathbb{1}_{\mathcal{G}_n(d) \geq 0} \leq \frac{1}{2} \left(\sup_{d \in D_B} \mathcal{G}_n(d) \right)^2.$$

In addition, as D_{1B} and D_{2B} are subspaces of D_B , we have

$$\sup_{d_1 \in D_{1B}, d_2 \in D_{2B}} \frac{1}{2} (\mathcal{G}_n^2(d_1) + \mathcal{G}_n^2(d_2) \mathbb{1}_{\mathcal{G}_n(d_2) \geq 0}) \leq \left(\sup_{d \in D_{B,s}} \mathcal{G}_n(d) \right)^2,$$

where $D_{B,s}$ is the symmetrized space $D_B \cup (-D_B)$. Therefore, we deduce that

$$T_{n,M,B} - \ell_n(f_{0,B}) \leq \left(\sup_{d \in D_{B,s}} \mathcal{G}_n(d) \right)^2 (1 + o_{\mathbb{P}}(1)).$$

Thus, using the fact that $D_{B,s}$ is a symmetric space, we obtain that, for any $\varepsilon > 0$, for n sufficiently large,

$$\left\{ \frac{T_{n,M,B} - \ell_n(f_{0,B})}{a_{n,M_0,B}} > 4\varepsilon \right\} \subset \left\{ \left| \sup_{d \in D_{B,s}} \mathcal{G}_n(d) \right| > 2\sqrt{\varepsilon a_{n,M_0,B}} \right\}. \quad (10)$$

Note that under our Assumptions, for a fixed value B^* of B , $\sup_{d \in D_{B^*,s}} \mathcal{G}_n(d)$ converges in distribution to $\sup_{d \in D_{B^*,s}} \xi_d$ (see Section A.3 in the Supplementary Materials) as n goes to infinity. However, as in our context, B goes to infinity with n , leading that $D_{B,s}$ grows with n . We cannot directly conclude that the probability of the event $\left\{ \left| \sup_{d \in D_{B,s}} \mathcal{G}_n(d) \right| > 2\sqrt{\varepsilon a_{n,M_0,B}} \right\}$ tends to zero as n tends to infinity. The previous equation implies that for any value of B ,

$$\begin{aligned} & \mathbb{P}\left(\frac{T_{n,M,B} - \ell_n(f_{0,B})}{a_{n,M_0,B}} > \varepsilon\right) \leq \\ & \mathbb{P}\left(\left| \sup_{d \in D_{B,s}} \xi_d \right| > \sqrt{\varepsilon a_{n,M_0,B}}\right) + \mathbb{P}\left(\left| \sup_{d \in D_{B,s}} \mathcal{G}_n(d) - \sup_{d \in D_{B,s}} \xi_d \right| > \sqrt{\varepsilon a_{n,M_0,B}}\right), \end{aligned} \quad (11)$$

where $(\xi_d)_d$ is a Gaussian process indexed by $D_{B,s}$, with covariance the usual Hilbertian product on $L^2_{f_{0,B}}$.

In Section A.4, we show that, under Assumption 2, we can control the deviation of the supremum of a Gaussian process indexed by a space with increasing dimension. Thus, we can control the first term on the right-hand side of (11) by using existing deviation bounds for the supremum of Gaussian processes, leading that

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\sup_{d \in D_{B,s}} \xi_d > \sqrt{\varepsilon a_{n,M_0,B}} \right) = 0. \quad (12)$$

In Section A.5, using the results on the approximation of suprema of general empirical processes by a sequence of suprema of Gaussian processes (Chernozhukov et al., 2014), we control the approximation of suprema of general empirical processes by a sequence of suprema of Gaussian processes indexed on a space with increasing dimension. Thus, we can control the second term on the right-hand side of (11) by using existing deviation bounds for the supremum of Gaussian processes, which leads to

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\left| \sup_{d \in D_{B,s}} \mathcal{G}_n(d) - \sup_{d \in D_{B,s}} \xi_d \right| > \sqrt{\varepsilon a_{n,M_0,B}} \right) = 0. \quad (13)$$

Combining (12) and (13), we have for any $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\frac{T_{n,M,B} - \ell_n(f_{0,B})}{a_{n,M_0,B}} > \varepsilon \right) = 0. \quad (14)$$

Noting that for any $M \in \mathcal{N}_3$, Assumption 3 implies that $a_{n,M,B}/a_{n,M_0,B} - 1 > 0$ and noting that the cardinal of \mathcal{N}_3 is finite and fixed (it does not depend on B), we obtain that

$$\lim_{n \rightarrow \infty} \mathbb{P}(\widehat{M} \in \mathcal{N}_3) = 0, \quad (15)$$

which concludes the third part.

Combining equations (6), (8) and (15) leads to the convergence in probability of \widehat{M} to M_0 . \square

Some comments: Note that the arguments used in Keribin (2000) to prove that underestimation is avoided, cannot be used here. Indeed, the proof of Theorem 2.1 in Keribin (2000) considers parameters that are defined on fixed dimensional space and thus cannot be used to obtain (5). In our context, we require that B tends to infinity with n (see Assumption 4) to ensure the identifiability and thus the convergence of $\inf_{\theta} \mathbb{E}_{g_0} [\ln f_{M_0,B,\theta^*_{M_0,B}} - \ln f_{M,B,\theta}]$ to a quantity lower-bounded by δ_M where $\theta^*_{M_0,B} = \arg \max_{\theta \in \Theta_{M_0,B}} \mathbb{E}_{g_0} [\ln f_{M_0,B,\theta}]$ (this convergence is ensured by Assumption 2(ii), as discussed in the proof of (5)). Note that the existence of $\theta^*_{M_0,B}$ is ensured by the fact that $\Theta_{M_0,B}$ is compact and that the Kullback-Leibler divergence is continuous.

Note also that the arguments used in Keribin (2000) to prove that overestimation is avoided cannot be used here either. Indeed, despite the fact that $T_{n,M,B^*} - \ell_n(f_{0,B^*})$ converges in distribution for a fixed B^* and that $1/a_{n,M,B^*}$ tends to 0, we cannot directly conclude that $\mathbb{P}([T_{n,M,B} - \ell_n(f_{0,B})]/a_{n,M,B} > \varepsilon)$ tends to zero as n tends to infinity, for any $\varepsilon > 0$ and any $M \in \mathcal{N}_3$. Indeed, in our context, the event $\{[T_{n,M,B} - \ell_n(f_{0,B})]/a_{n,M,B} > \varepsilon\}$ is included in an event which involves a supremum over a functional space with increasing dimension (see (10)) since we require that B tends to infinity with n to avoid the underestimation (see Assumption 4).

Note that Part 2 is proven by making use of some results of Part 3, because any model in \mathcal{N}_2 can be included in a model of \mathcal{N}_3 , see Section A.2. However, an alternative proof of Part 2 could be obtained by following the same reasoning as in the proof of Part 3 by noting that an equation similar to (9) should be computed for a model belonging to \mathcal{N}_2 . Such an equation would be slightly different from (9) since there is no loss of identifiability for models in \mathcal{N}_2 while such a loss occurs for models in \mathcal{N}_3 .

3.3 Relaxing the knowledge on the upper bound of the number of components

Theorem 1 considers an upper bound on the number of components K_{\max} that holds fixed with the sample size. However, in practical applications, the range of the possible values of K_0 can be unknown leading that K_{\max} is unknown too. Therefore, in this section, we extend the proposed method by allowing the upper bound on the number of components K_{\max} to increase with the sample size. Thus, the consistency of the procedure is ensured without any knowledge on the number of components, since there exists a sample size that ensures that $K_{\max} > K_0$. Since considering an increasing value of K_{\max} implies that the space of the competing models as well as the space of the parameters have increasing dimensions, the increase of both K_{\max} and B must be done at rates that are related to each other. The following corollary gives sufficient condition to state the consistency of $\widehat{M}_{n,B,K}$ with an increasing upper bound of the number of components. Note that despite the fact that the upper bound of components is allowed to increase with the sample size, the true number of components K_0 as well as the true subset of relevant variables Ω_0 do not change with the sample size. Thus, M_0 stays fixed as n grows.

Corollary 1. *Assume that independent data arise from (2) with true model $M_0 = \{K_0, \Omega_0\}$ and that Assumptions 1, 2, 3 and 4 hold true except Assumption 3.(iii) and Assumption 4.(i). If K_{\max} and B are two integers which tend to infinity with n , $\lim_{n \rightarrow \infty} K_{\max} = \infty$ and $\lim_{n \rightarrow \infty} B = \infty$, and which satisfy $\lim_{n \rightarrow \infty} (BK_{\max} + \ln K_{\max})/a_{n,M,B} = 0$ and $\lim_{n \rightarrow \infty} K_{\max}^2 B (\ln^2 n)/n^{1/2} = 0$ then $\widehat{M}_{n,B}$ converges in probability to M_0 .*

To prove Corollary 1, first note that the spaces \mathcal{N}_1 and \mathcal{N}_2 hold fixed as K_{\max} tends to infinity and thus (6) and (8) still hold true. Note that \mathcal{N}_3 increases as K_{\max} tends to infinity and that for any model $M \in \mathcal{N}_3$, we have from Section A.4 and Section A.5 in the Supplementary Materials that for some positive constant C

$$\mathbb{P}\left(\frac{T_{n,M,B} - \ell_n f_{0,B}}{a_{n,M_0,B}} > \varepsilon\right) \leq 2 \exp\left(-\frac{\sqrt{\varepsilon a_{n,M_0,B}} - \tilde{\kappa} \sqrt{(K_{\max} - 1) + K_{\max} B J}}{2\zeta_\xi^2}\right) + C\left(\gamma_n + \frac{\ln n}{n}\right),$$

with $\gamma_n = \frac{4(bcK_{\max}BJ \ln n)^2}{\varepsilon a_{n,M_0,B} n^{1/2}}$. Noting that $\text{card}(\mathcal{N}_3) < \text{card}(\mathcal{M}) < K_{\max} 2^J$, we have

$$\mathbb{P}(\widehat{M} \in \mathcal{N}_3) \leq K_{\max} 2^J \left(2 \exp\left(-\frac{\sqrt{\varepsilon a_{n,M_0,B}} - \tilde{\kappa} \sqrt{(K_{\max} - 1) + K_{\max} B J}}{2\zeta_\xi^2}\right) + C\left(\gamma_n + \frac{\ln n}{n}\right)\right).$$

The proof is concluded by noting that the conditions $\lim_{n \rightarrow \infty} (BK_{\max} + \ln K_{\max})/a_{n,M,B} = 0$ and $\lim_{n \rightarrow \infty} K_{\max}^2 B (\ln^2 n)/n^{1/2} = 0$ imply that the term in the right-hand side of the previous equation tends to zero as n tends to infinity.

Note that the previous corollary only allows for a slow increase of K_{\max} . For instance, if a BIC-like penalty is used, since K_{\max} does not appear in the penalty term, then K_{\max} should satisfy $\lim_{n \rightarrow \infty} (K_{\max} + \ln K_{\max})/\ln n = 0$. However, even if only a slow increase of K_{\max} can be considered, the previous Corollary avoids the use of a fixed upper bound for the number of components (as n tends to infinity) and thus a consistent approach can be obtained even if there is no information about the range of the possible values of K_0 .

4 Estimation of the best model

The estimation of $\widehat{M}_{n,B}$ requires an optimization over a discrete space whose cardinal is of order $2^J K_{\max}$. Thus, an exhaustive approach computing $W_{n,M,B}$ for each M in \mathcal{M} is not doable in practice. As the combinatorial issue is mainly due to the feature selection, we follow the approach of Marbac et al. (2019) that consists of simultaneously performing feature selection and parameter estimation, with a fixed number of components, via a specific EM algorithm optimizing the penalized likelihood. Thus, for a fixed value of K , the goal of the algorithm is to estimate

$$\widehat{M}_{n,B,K} = \arg \max_{\{M=\{K,\Omega\} \text{ with } \Omega \subseteq \{1,\dots,J\} \text{ and } |\Omega| \geq 3\}} W_{n,M,B}.$$

The following EM algorithm permits the estimation of the model parameters and the detection of the subset of relevant variables, for a fixed number of components K . This algorithm ensures that the penalized log-likelihood increases at each iteration. Thus, parameter estimation is achieved by maximum likelihood and model selection is done with an information criterion with penalty $a_{n,K,\Omega,M} = \nu_{K,\Omega,M} c_n$ where $\nu_{K,\Omega,M} = (K-1) + |\Omega|K(B-1) + (J-|\Omega|)(B-1)$ is the number of model parameters. The algorithm considers a fixed number of components K and starts at an initial point $\{\Omega^{[0]}, \theta^{[0]}\}$. Its iteration $[r]$ is composed of two steps:

E-step Computation of the fuzzy partition

$$t_{ik}^{[r]} := \frac{\pi_k^{[r-1]} \prod_{j \in \Omega^{[r-1]}} \prod_{b=1}^B \left(\alpha_{Bkj b}^{[r-1]} \right)^{\sigma_{Bjb}(x_{ij})}}{\sum_{\ell=1}^K \pi_\ell^{[r-1]} \prod_{j \in \Omega^{[r-1]}} \prod_{b=1}^B \left(\alpha_{B\ell j b}^{[r-1]} \right)^{\sigma_{Bjb}(x_{ij})}},$$

M-step Maximization of the expectation of the penalized complete-data log-likelihood over Ω and θ such

$$\Omega^{[r]} = \{j : \Delta_j^{[r]} > 0\}, \quad \pi_k^{[r]} = \frac{n_k^{[r]}}{n} \quad \text{and} \quad \alpha_{Bkj b}^{[r]} = \begin{cases} \tilde{\alpha}_{Bkj b}^{[r]} & \text{if } j \in \Omega^{[r]} \\ \bar{\alpha}_{Bkj b} & \text{otherwise} \end{cases},$$

where

$$\Delta_j^{[r]} = \sum_{i=1}^n \sum_{b=1}^B \sigma_{Bjb}(x_{ij}) \sum_{k=1}^K t_{ik}^{[r]} \ln \left(\frac{\tilde{\alpha}_{Bkj b}^{[r]}}{\bar{\alpha}_{Bkj b}} \right) - (K-1)(B-1)c_n$$

is the difference between the maximum of the expected value of the penalized complete-data log-likelihood obtained when variable j is relevant and when it is irrelevant, with

$$\tilde{\alpha}_{Bkj b}^{[r]} = \frac{1}{n_k^{[r]}} \sum_{i=1}^n t_{ik}^{[r]} \sigma_{jb}(x_{ij}), \quad \bar{\alpha}_{Bkj b} = \frac{1}{n} \sum_{i=1}^n \sigma_{jb}(x_{ij}) \quad \text{and} \quad n_k^{[r]} = \sum_{i=1}^n t_{ik}^{[r]}.$$

Note that, when less than three variables get a positive value for $\Delta_j^{[r]}$, the M-step selects in $\Omega^{[r]}$ the three variables having the largest values of $\Delta_j^{[r]}$. To obtain the pair Ω and θ maximizing the penalized observed-data log-likelihood, for a fixed number of components, many random initializations of this algorithm should be done. Hence, the model (*i.e.*, K and Ω) and the parameters maximizing the penalized observed-data log-likelihood are obtained by performing this algorithm for every values of K between 1 and K_{\max} . By considering $c_n = (\ln n)/2$, this algorithm carries out the model selection according to the BIC.

From the previous algorithm, we obtain an estimator of the model and of its parameters. Indeed, $\hat{\alpha}_{Bkj b}/l_{Bjb}$ estimates the density $\eta_{kj}(u)$ for any u such that $\sigma_{jb}(u) = 1$. However, the bin-based density estimators are generally outperformed by kernel-based estimators. Thus, we advise to use the proposed approach only for model estimation. Then, for the selected model, kernel-based density estimates provided by the EM-like algorithm (Benaglia et al., 2009) or by maximizing the smoothed log-likelihood (Levine et al., 2011) should be considered. However, note that establishing asymptotic properties of those kernel-based density estimators is still an open question.

5 Numerical experiments

5.1 Investigating the impact of the discretization

The proposed procedure implies to define a number of bins B as well as the location of the bins. Despite the fact that the discretization must respect Assumptions 3 and 4, there are many ways to perform this discretization. Two natural procedures can be considered for the discretization: one can consider a discretization of each $[\min_i x_{ij}, \max_i x_{ij}]$ into B bins of equal sizes (as it is done for histograms) or one can consider a discretization based on the empirical quantiles $1/B, \dots, B/B$ implying that $\sigma_{jb}(u) = 1$ only if $u \in [q_{jb-1}, q_{jb}]$ and $\sigma_{jb}(u) = 0$ otherwise, with $q_{j0} = \min_i x_{ij}$, $q_{jB} = \max_i x_{ij}$ and q_{jb} is the empirical quantile of order b/B for $b = 1, \dots, B-1$. Moreover, only a growth rate for the number of bins B is stated by Assumption 4. Thus, we could consider different rates for the growth of B with n : $[n^{1/7}]$,

$[n^{1/6}]$, $[n^{1/5}]$, $[n^{1/4}]$ and $[n^{1/3}]$. This section presents a numerical experiment that aims at providing some advises on the procedure of discretization.

Data are generated from a mixture with three components considering well-balanced clusters (*i.e.*, equal proportions $\pi_k = 1/3$) and unbalanced clusters (*i.e.*, $\pi_1 = 7/12$, $\pi_2 = 3/12$ and $\pi_3 = 2/12$). The density of X_i given Z_i is a product of J univariate densities such that $X_{ij} = \sum_{k=1}^K z_{ik} \delta_{kj} + \xi_{ij}$ where all the ξ_{ij} are independent and where $\delta_{ku} = \tau$ if u is a multiple of k and $\delta_{kj} = 0$ otherwise, which implies all the variables are relevant for clustering. Three distributions are considered for the ξ_{ij} (standard Gaussian, Student with three degrees of freedom and Laplace) and different values of J are considered while the value of τ is defined to obtain a theoretical misclassification rate of 5%. For each setup, 100 replicates are generated. Table 1 presents the percentage of replicates where the method selects the true number of components obtained for well-balanced clusters, by considering $K_{\max} = 6$ and different procedures of discretization and $J = 6$ relevant variables. Table 2 presents the percentage of replicates where the method selects the true number of components obtained for the unbalanced clusters, by considering $K_{\max} = 6$ and different procedures of discretization $J = 6$ relevant variables. Section B.1 presents similar results obtained with $J = 12$ and $J = 24$ variables.

Table 1: Percentage where the method selects the true number of components from data generated from the well-balanced clusters when a discretization is done with quantiles or equal-size bins and different number of bins $[n^{1/7}]$, $[n^{1/6}]$, $[n^{1/5}]$, $[n^{1/4}]$ and $[n^{1/3}]$, for different families of components and $J = 6$ variables (all the variables being relevant).

Method	Band.	Gaussian sample size				Student sample size				Laplace sample size			
		100	250	500	1000	100	250	500	1000	100	250	500	1000
quantile	$n^{1/7}$	72	100	100	100	91	100	100	100	86	100	100	100
	$n^{1/6}$	72	100	100	100	91	100	100	100	86	100	100	100
	$n^{1/5}$	66	100	100	100	79	100	100	100	77	100	100	100
	$n^{1/4}$	66	100	100	100	79	100	100	100	77	100	100	100
	$n^{1/3}$	1	92	100	100	6	96	100	100	3	95	100	100
equal-size	$n^{1/7}$	9	81	100	100	0	3	78	98	0	0	47	98
	$n^{1/6}$	9	69	100	100	0	0	11	98	0	0	15	98
	$n^{1/5}$	0	69	100	100	0	0	11	79	0	0	15	98
	$n^{1/4}$	0	69	100	100	0	0	1	47	0	1	78	100
	$n^{1/3}$	0	4	100	100	0	1	30	91	0	10	91	100

Table 2: Percentage where the method selects the true number of components from data generated from the unbalanced clusters when a discretization is done with quantiles or equal-size bins and different number of bins $[n^{1/7}]$, $[n^{1/6}]$, $[n^{1/5}]$, $[n^{1/4}]$ and $[n^{1/3}]$, for different families of components and $J = 6$ variables (all the variables being relevant).

Method	Band.	Gaussian sample size				Student sample size				Laplace sample size			
		100	250	500	1000	100	250	500	1000	100	250	500	1000
quantile	$n^{1/7}$	3	72	98	100	19	79	100	100	16	72	100	100
	$n^{1/6}$	3	73	100	100	19	92	100	100	16	88	100	100
	$n^{1/5}$	3	73	100	100	19	92	100	100	11	88	100	100
	$n^{1/4}$	3	51	97	100	19	67	100	100	11	74	100	100
	$n^{1/3}$	0	5	22	99	0	8	49	100	0	15	57	100
equal-size	$n^{1/7}$	1	21	68	100	0	0	7	15	0	0	3	7
	$n^{1/6}$	1	4	60	100	0	0	1	15	0	0	0	7
	$n^{1/5}$	0	4	60	100	0	0	1	0	0	0	0	8
	$n^{1/4}$	0	3	37	100	0	0	0	0	0	0	6	93
	$n^{1/3}$	0	0	8	97	0	0	0	10	0	0	4	68

Results illustrate the consistency of the method and show that the choice of the number of bins impacts the results for small samples. Moreover, the case of unbalanced data is more complex and thus the proposed method requires a larger sample size than for the well-balance case. Based on this experiment, we advise to use B levels defined by the B empirical quantiles. Moreover, considering a large increasing rate for B seems to deteriorate the results, so we advise to use $B = \lceil n^{1/6} \rceil$ or $B = \lceil n^{1/7} \rceil$. Note that those advises are only based on this numerical experiment. For small sample sizes, other procedures should be more efficient for some specific distributions of the components. However, the discretization procedure that we advise, respects Assumption 3 and 4 and seems to have a good behavior in all the situations considered during the experiment.

5.2 Comparing the methods for selecting the number of components

In this section, we assess the performance of our estimator of the number of components without considering the task of variable selection (*i.e.*, $\Omega = \{1, \dots, J\}$ is supposed to be known). We consider the estimator \hat{K} obtained with a BIC penalty and B levels defined by the empirical quantiles $1/B, \dots, B/B$ implying that $\sigma_{jb}(u) = 1$ only if $u \in [q_{jb-1}, q_{jb}]$ and $\sigma_{jb}(u) = 0$ otherwise, with $q_{j0} = \min_i x_{ij}$, $q_{jB} = \max_i x_{ij}$ and q_{jb} the empirical quantile of order b/B for $b = 1, \dots, B - 1$. Three values of B are investigated: $\lceil n^{1/5} \rceil$, $\lceil n^{1/6} \rceil$ and $\lceil n^{1/7} \rceil$, where $\lceil \cdot \rceil$ denotes the closest integer. Our estimator is compared to the procedure of Kwon and Mbakop (2021) SVT and to the procedures of Kasahara and Shimotsu (2014) max ave-rk⁺ statistics with $M_0 = 4$. For the comparison, we consider the simulation setup of Kasahara and Shimotsu (2014). This setup originally considers four designs with two variables and one design with eight variables (Design 5). Our approach cannot be used when only two variables are observed for identifiability issues. This is a drawback of the proposed method that requires the parameter identifiability (ensured by Assumption 1) since this method relies on the penalized log-likelihood. Note that the case of bivariate data can be handled by Kasahara and Shimotsu (2014), despite the fact that the model parameters cannot be non-parametrically identified for such mixture models (Hall and Zhou, 2003). Indeed, Kasahara and Shimotsu (2014) show that the number of components K can be identified for bivariate data under weaker conditions than Assumption 1 which do not ensure the parameter identifiability. Note that the approaches of Kasahara and Shimotsu (2014) and Kwon and Mbakop (2021) are based on estimators of bivariate density and we advise to use these methods for bivariate data (in this case variable selection does not make sense). However, these methods suffer from computational issues if the number of variables is large (explaining that we do not run these methods on the next section) while the proposed method does not suffer from this problem. This illustrates the complementarity of the three approaches. We simulate 1000 samples of size $n = 500$ and of size $n = 2000$ from a mixture of three Gaussian distributions (see numerical experiments in Kasahara and Shimotsu (2014) and Kwon and Mbakop (2021)) defined with equal proportions, centers $\mu_1 = (0, 0, 0, 0, 0, 0, 0, 0)^\top$, $\mu_2 = (1, 2, 0.5, 1, 0.75, 1.25, 0.25, 0.5)^\top$ and $\mu_3 = (2, 1, 1, 0.5, 1.25, 0.75, 0.5, 0.25)^\top$ and covariance matrices equal to the identity matrix of size eight.

Table 3: Empirical probabilities of selecting the different numbers of clusters.

	$n = 500$				$n = 2000$			
	$K = 1$	$K = 2$	$K = 3$	$K \geq 4$	$K = 1$	$K = 2$	$K = 3$	$K \geq 4$
BIC with $B = n^{1/7}$	0.000	0.996	0.004	0.000	0.000	0.517	0.483	0.000
BIC with $B = n^{1/6}$	0.001	0.999	0.000	0.000	0.000	0.876	0.124	0.000
BIC with $B = n^{1/5}$	0.001	0.999	0.000	0.000	0.000	0.992	0.008	0.000
SVT	0.000	0.992	0.008	0.000	0.000	0.493	0.507	0.000
ave-rk	0.142	0.810	0.047	0.001	0.005	0.776	0.214	0.005
AIC by ave-rk	0.012	0.867	0.119	0.003	0.000	0.587	0.399	0.013
BIC by ave-rk	0.284	0.715	0.001	0.000	0.035	0.942	0.023	0.000
HQ by ave-rk	0.078	0.909	0.013	0.000	0.004	0.878	0.117	0.001

Table 3 presents the outcome of the simulation. Note that the results of the approaches of Kasahara and Shimotsu (2014) and Kwon and Mbakop (2021) arise from Section 4 in Kwon and Mbakop (2021) and thus are not obtained on the same samples. For $n = 500$, all the procedures performs poorly by mainly selecting two components. The best approach is the AIC (ave-risk) procedure. The proposed

approach obtains similar results than the SVT approach which are worse than those of the AIC procedure. However, when the sample size increases ($n = 2000$), the SVT approach obtains the best results (the true number of components is detected 50.7% of time). The proposed method obtains also good results (the true number of components is detected 48.3% of time) when $B = \lceil n^{1/7} \rceil$ but its results deteriorate when the number of levels grows faster (*i.e.*, $B = \lceil n^{1/6} \rceil$ and $B = \lceil n^{1/5} \rceil$). This illustrates that the results of the proposed method are sensitive to the choice of the number of levels used for the discretization, for a fix sample size. However, all the growth rates of the number of levels respecting Assumption 4.(i) provide consistent estimator as illustrated by the additional experiments presented in Section B of the Supplementary Materials.

5.3 Comparing the methods for a full model selection

This section compares approaches for a full model selection (*i.e.*, estimation of the subset of the relevant variables and on the number of components) on simulated data. We compare the proposed approach with a BIC applied on a Gaussian mixture model, with the sparse K -means approach and with the non-parametric approach considering all the variables as relevant. The results of the proposed clustering method are obtained by performing full model selection with B levels defined by the empirical quantiles $1/B, \dots, B/B$ where $B = \lceil n^{1/6} \rceil$ and a BIC like penalty and then by estimating the mixture components for the selected model by maximizing the smoothed log-likelihood with a bandwidth, for variable j , equal to $\hat{\sigma}_j n^{-1/5}$ where $\hat{\sigma}_j$ is the empirical standard deviation of variable j . Thus, when the discretization is performed, the model selection can be achieved via the R package VarSelLCM (Marbac and Sedki, 2019) then, when the best model is selected, the maximization of the smoothed log-likelihood is achieved via the R package mixtools (Benaglia et al., 2009). The parametric mixture model considers that all the components are Gaussian (this approach is also implemented in the R package VarSelLCM) and uses the BIC to perform model selection. The sparse K -means approach is implemented in the R package sparcl (Witten and Tibshirani, 2010) and consists in the sparse K -means algorithm initialized with the partition provided by the sparse hierarchical ascendant classification with the ‘‘average’’ method. The sparse K -means estimates weights for the variables and thus we consider that a variable is claimed to be relevant if its weight is more than ι and that a variable is irrelevant if its weight is less than ι , where the small threshold $\iota = 10^{-6}$ is introduced to avoid numerical issues. Finally, the non-parametric mixture model is implemented the R package mixtools (Benaglia et al., 2009) and considers the estimator maximizing the smoothed log-likelihood with a bandwidth, for variable j , equal to $\hat{\sigma}_j n^{-1/5}$ where $\hat{\sigma}_j$ is the empirical standard deviation of variable j . To compare the different methods of clustering, we generate data from a mixture with three components and equal proportions ($\pi_k = 1/3$). The density of X_i given Z_i is a product of univariate densities such that $X_{ij} = \sum_{k=1}^K z_{ik} \delta_{kj} + \xi_{ij}$ where all the ξ_{ij} are independent and where $\delta_{11} = \delta_{12} = \delta_{23} = \delta_{24} = \delta_{35} = \delta_{36} = \tau$, while all remaining $\delta_{kj} = 0$, which implies that only the first six variables are relevant for clustering. Three distributions are considered for the ξ_{ij} (standard Gaussian, Student with three degrees of freedom and Laplace) and the value of τ is defined to obtain a theoretical misclassification rate of 5% (τ is equal to 1.94, 2.60 and 2.52 for the Gaussian, Student and Laplace distributions respectively). In the Section B of the Supplementary Materials, all the experiments are also run with theoretical misclassification rates equal to 10% and 15%.

Selection of the discriminative features To investigate the performances of the competing methods for feature selection, we first consider the situation with a known number of components. Thus, the model selection consists in performing the feature selection. We consider the methods which automatically provide an estimator of the relevant variables (*i.e.*, the proposed method, sparse K -means and VarSelLCM). Accuracy of this selection is measured by sensitivity (probability to detect as relevant a true discriminative variable) and specificity (probability to detect as irrelevant a true non discriminative variable). Table 4 and 5 present the sensibility and the specificity obtained by the proposed approach and the parametric approach. They exhibit an advantage of the parametric method when the distribution is well-specified, but only for small samples ($n = 100$). The reason is that, for such samples, the proposed method only finds a part of the relevant variables. However, both methods perform well for larger samples. Moreover, the proposed method obtains similar results for the two other distributions of the components while the results of the parametric approach are strongly deteriorated for both sensibility and specificity, especially for heavy tailed distributions (*e.g.*, Student distribution).

Table 4: Mean of the sensitivity (Sen.: $\text{card}(\widehat{\Omega} \cap \Omega)/6$) for the feature selection obtained by the proposed method (Proposed method), the parametric method (VarSelLCM) and the sparse K-means (Sparcl) on 100 replicates for each scenario with theoretical misclassification rate of 5%, when the number of components is known.

Component	J	Proposed method			VarSelLCM			Sparcl		
		n			n			n		
		100	250	500	100	250	500	100	250	500
Gaussian	20	0.81	1.00	1.00	1.00	1.00	1.00	0.86	0.91	0.95
	50	0.66	0.99	1.00	0.96	1.00	1.00	0.91	0.95	0.97
	100	0.35	0.70	1.00	0.69	1.00	1.00	0.78	0.97	0.98
Student	20	0.82	1.00	1.00	0.35	0.42	0.50	0.74	0.80	0.81
	50	0.69	1.00	1.00	0.10	0.13	0.21	0.72	0.74	0.79
	100	0.53	0.90	1.00	0.08	0.15	0.15	0.56	0.74	0.79
Laplace	20	0.86	1.00	1.00	0.93	1.00	1.00	0.81	0.83	0.81
	50	0.72	1.00	1.00	0.55	1.00	1.00	0.89	0.91	0.91
	100	0.52	0.89	1.00	0.19	0.89	1.00	0.82	0.95	0.94

Table 5: Mean of the specificity (Spe.: $\text{card}(\widehat{\Omega}^c \cap \Omega^c)/(J - 6)$) for the feature selection obtained by the proposed method (Proposed method), the parametric method (VarSelLCM) and the sparse K-means (Sparcl) on 100 replicates for each scenario with theoretical misclassification rate of 5%, when the number of components is known.

Component	J	Proposed method			VarSelLCM			Sparcl		
		n			n			n		
		100	250	500	100	250	500	100	250	500
Gaussian	20	0.98	1.00	1.00	1.00	1.00	1.00	0.79	0.66	0.50
	50	0.98	1.00	1.00	1.00	1.00	1.00	0.71	0.45	0.30
	100	0.98	1.00	1.00	1.00	1.00	1.00	0.75	0.39	0.22
Student	20	0.97	1.00	1.00	0.70	0.57	0.47	0.77	0.58	0.61
	50	0.98	1.00	1.00	0.74	0.66	0.58	0.79	0.83	0.62
	100	0.98	1.00	1.00	0.76	0.69	0.63	0.85	0.85	0.76
Laplace	20	0.98	1.00	1.00	0.94	0.96	0.97	0.76	0.85	0.83
	50	0.98	1.00	1.00	0.92	0.97	0.97	0.75	0.78	0.74
	100	0.99	1.00	1.00	0.92	0.97	0.98	0.75	0.62	0.64

Full model selection We now compare both non-parametric and parametric approaches on their performances for full model selection. Table 6 presents the statistics on the number of components selected by both approaches. Again, when the distribution of the components is well-specified, the parametric approach obtains better results on small samples because the proposed approach tends to underestimate the number of components. However, when the sample size increases, both methods perform similarly. When the distribution of the components is not Gaussian, the parametric method performs poorly and asymptotically overestimates the number of components with probability one. The proposed method is consistent for any number of variables, however, it tends to underestimate the number of components for small samples.

Table 7 presents the sensitivity and the specificity for feature selection obtained by both approaches when the number of components is also estimated. Again, results show the benefits of the proposed approach when the parametric assumptions are violated. In such a case, the parametric approach overestimates the number of components and, for heavy tail distributions (*e.g.*, Student distribution), this approach tends to overestimate the subset of relevant variables. Moreover, for the small samples, the sensitivity is quite low explaining the tendency of overestimating the number of components.

Accuracy of the partition We are now interested in investigating the accuracy of the estimated partition. Thus, we compute the Adjusted Rand index (Hubert and Arabie, 1985) between the true

Table 6: Probability to select the true number of components (Tr.) and to overestimate it (Ov.) obtained by the proposed method (Proposed method) and the parametric method (VarSelLCM) on 100 replicates for each scenario with theoretical misclassification rate of 5%, by performing a selection of the variables.

Component	J	Proposed method						VarSelLCM					
		$n = 100$		$n = 250$		$n = 500$		$n = 100$		$n = 250$		$n = 500$	
		Tr.	Ov.	Tr.	Ov.	Tr.	Ov.	Tr.	Ov.	Tr.	Ov.	Tr.	Ov.
Gaussian	20	0.46	0.00	1.00	0.00	1.00	0.00	0.94	0.00	1.00	0.00	1.00	0.00
	50	0.27	0.00	0.98	0.00	1.00	0.00	0.86	0.01	1.00	0.00	1.00	0.00
	100	0.05	0.00	0.62	0.00	1.00	0.00	0.54	0.00	1.00	0.00	1.00	0.00
Student	20	0.56	0.00	1.00	0.00	1.00	0.00	0.62	0.15	0.18	0.80	0.00	1.00
	50	0.37	0.00	1.00	0.00	1.00	0.00	0.81	0.13	0.36	0.64	0.02	0.98
	100	0.13	0.00	0.82	0.01	1.00	0.00	0.70	0.29	0.22	0.78	0.01	0.99
Laplace	20	0.57	0.00	1.00	0.00	1.00	0.00	0.77	0.08	0.34	0.66	0.00	1.00
	50	0.33	0.00	1.00	0.00	1.00	0.00	0.40	0.00	0.51	0.49	0.01	0.99
	100	0.15	0.00	0.88	0.01	1.00	0.00	0.09	0.00	0.61	0.21	0.01	0.99

Table 7: Mean of the sensitivity (Sen.: $\text{card}(\widehat{\Omega} \cap \Omega)/6$) and the specificity (Spe.: $\text{card}(\widehat{\Omega}^c \cap \Omega^c)/(J-6)$) for the feature selection obtained by the proposed method (Proposed method) and the parametric method (VarSelLCM) on 100 replicates for each scenario with theoretical misclassification rate of 5%, when the number of components also is estimated.

Component	J	Proposed method						VarSelLCM					
		$n = 100$		$n = 250$		$n = 500$		$n = 100$		$n = 250$		$n = 500$	
		Sen.	Spe.	Sen.	Spe.	Sen.	Spe.	Sen.	Spe.	Sen.	Spe.	Sen.	Spe.
Gaussian	20	0.97	0.82	1.00	1.00	1.00	1.00	1.00	0.98	1.00	1.00	1.00	1.00
	50	0.95	0.74	1.00	0.99	1.00	1.00	1.00	0.96	1.00	1.00	1.00	1.00
	100	0.94	0.59	1.00	0.94	1.00	1.00	0.99	0.87	1.00	1.00	1.00	1.00
Student	20	0.96	0.86	1.00	1.00	1.00	1.00	0.67	0.37	0.66	0.72	0.56	0.97
	50	0.96	0.79	1.00	1.00	1.00	1.00	0.73	0.09	0.68	0.15	0.59	0.24
	100	0.95	0.70	1.00	0.98	1.00	1.00	0.76	0.09	0.72	0.13	0.67	0.13
Laplace	20	0.97	0.87	1.00	1.00	1.00	1.00	0.92	0.90	0.96	1.00	0.91	1.00
	50	0.97	0.80	1.00	1.00	1.00	1.00	0.86	0.44	0.97	1.00	0.96	1.00
	100	0.95	0.72	1.00	0.99	1.00	1.00	0.84	0.13	0.96	0.83	0.98	1.00

partition and the estimators of the partition given by the non-parametric and the parametric methods when K is known and then when it is estimated. Moreover, to illustrate the benefit of feature selection, we also estimate the partition by considering the full variables as relevant and the true number of components. Results are presented in Table 8. Thus, when the parametric assumptions are satisfied, the parametric approach outperforms the proposed approach only on small samples (few observations with respect to the number of variables), whenever the number of components is known or not. However, when the parametric assumptions are violated, the proposed approach strongly outperforms the parametric approach. Note that, when the number of irrelevant variables increases, the approach considering all the variables for clustering performs poorly (see results with $J = 100$), illustrating the benefit of feature selection for clustering.

5.4 Benchmark data

This section illustrates our procedure on four real data sets. The first data set illustrates the advantage of the procedure for selecting the number of components while the second data set sheds light on the importance of variable selection. The third data set shows that the procedure can be easily extended to the case of mixed-type data sets (a data set composed of continuous and categorical data). The last data set shows that the procedure can also be used to cluster data with non-ignorable missingness by considering the semi-parametric mixture models defined in Du Roy de Chaumaray and Marbac (2023).

Table 8: Mean values of the Adjusted Rand Index (ARI) obtained on the resulting partition when feature selection is performed with the true number of components by the proposed method (proposed.K-known) and by the parametric method (VarSelLCM.K-known), by the sparse K-means (Sparcl.K-known) and by the model considering all the variables as relevant components (mixtools.K-known) and when the full model selection (feature selection and estimation of the number of components) is achieved by the proposed approach (proposed.K-unknown) and the parametric approach (VarSelLCM.K-unknown). Data are generated with theoretical misclassification rate of 5%

J	Method	K known	Gaussian			Student			Laplace		
			sample size			sample size			sample size		
			100	250	500	100	250	500	100	250	500
20	Proposed	yes	0.67	0.84	0.85	0.66	0.81	0.83	0.70	0.83	0.84
	VarSelCLM	yes	0.81	0.85	0.86	0.15	0.12	0.10	0.65	0.77	0.79
	sparcl	yes	0.70	0.72	0.76	0.54	0.61	0.60	0.64	0.62	0.61
	mixtools	yes	0.71	0.81	0.83	0.59	0.72	0.78	0.68	0.80	0.82
	Proposed	no	0.61	0.85	0.85	0.62	0.82	0.84	0.66	0.83	0.84
	VarSelCLM	no	0.79	0.85	0.86	0.15	0.38	0.53	0.62	0.70	0.62
50	Proposed	yes	0.56	0.84	0.85	0.56	0.81	0.84	0.62	0.84	0.84
	VarSelCLM	yes	0.79	0.85	0.85	0.01	0.00	0.00	0.34	0.75	0.77
	sparcl	yes	0.66	0.78	0.80	0.45	0.51	0.58	0.61	0.67	0.68
	mixtools	yes	0.33	0.71	0.79	0.22	0.51	0.66	0.30	0.66	0.77
	Proposed	no	0.51	0.84	0.85	0.56	0.81	0.84	0.56	0.84	0.84
	VarSelCLM	no	0.77	0.85	0.85	0.00	0.00	0.03	0.24	0.69	0.64
100	Proposed	yes	0.30	0.61	0.85	0.43	0.73	0.83	0.42	0.75	0.84
	VarSelCLM	yes	0.55	0.84	0.86	0.00	0.00	0.00	0.10	0.65	0.77
	sparcl	yes	0.41	0.78	0.82	0.27	0.47	0.52	0.42	0.70	0.71
	mixtools	yes	0.12	0.36	0.74	0.09	0.18	0.41	0.11	0.29	0.64
	Proposed	no	0.34	0.69	0.85	0.42	0.75	0.83	0.44	0.79	0.84
	VarSelCLM	no	0.62	0.84	0.86	0.00	0.00	0.00	0.04	0.58	0.64

Swiss banknotes data We consider the Swiss banknotes data set (Flury, 1988) containing six measurements (length of bill, width of left edge, width of right edge, bottom margin width, top margin width and length of diagonal) made on 100 genuine and 100 counterfeit old-Swiss 1000-franc bank notes. This data set is available in the R package `mclust` (Scrucca et al., 2016). The status of the banknote (genuine or counterfeit) is also known. We perform the clustering of the bills based on the six morphological measurements and we evaluate the resulting partition with the status of the bills. Considering all the six morphological measurements as relevant for clustering, the proposed method detects two clusters which are strongly similar to the status of the bill (the ARI is 0.98 and only one genuine bill is assigned to the cluster grouping all the counterfeit bills). Clustering with Gaussian mixture models provides more components (`mclust` selects three components and `VarSelLCM` selects four components) and a partition related but different to the status of the bill (the ARI is 0.84 and 0.48 for `mclust` and `VarSelLCM` respectively). When a full model selection (feature selection and estimation of the number of components) is performed, the proposed method still selects two components and detects all the variables as relevant. Thus, a full model selection approach provides the same results as a method used for selecting the clusters by considering all the variables as relevant. Moreover, the Gaussian mixture models obtains less relevant results because `VarSelLCM` considers that all the measurements are relevant and thus obtains the same results as without performing feature selection.

Chemical properties of coffees We consider the data set collected by Streuli (1973) that reports on the chemical composition of coffee samples collected from around the world. A total of 43 samples were collected from 29 countries, with beans from both Arabica and Robusta species, which is often considered as a pertinent partition. This data is available in the R package `pgmm` (McNicholas et al., 2015). We cluster the different coffees based on twelve chemical constituents. A full Gaussian mixture clustering implemented in `mclust` estimates three clusters and provided an ARI of 0.38. The same partition is

obtained when the clustering is performed by VarSelLCM with a full model selection conducted according to the BIC (all the variables are detected as relevant for clustering). Again, similar results are obtained by the non-parametric mixture if the proposed method is used to select the number of components. However, if we perform a full model selection, only five of the twelve variables are detected as relevant for clustering and only two components are estimated. Moreover, this simpler model provides a perfect recovery of the species (ARI=1.00). This illustrates the importance of variable selection for clustering. Note that McNicholas and Murphy (2008) proposed a parsimonious Gaussian mixture model where parsimonious constraints are made on the covariance matrices of the Gaussian distributions, that also provides a perfect recovery of the partition.

Cleveland data set We consider the Cleveland dataset (available at <https://www.kaggle.com/ronitf/heart-disease-uci/version/1>). This data set is composed of $n = 303$ subjects. Each subject is described by eight categorical variables having between two and six levels and five continuous variables. The "goal" field refers to the presence of heart disease in the patient (no presence vs presence). Model (2) can be easily extended to the case of mixed-type data (data set composed of continuous and categorical variables). Indeed, if variable j is categorical then η_{kj} is the probability mass function of a multinomial distribution. Thus, the discretization procedure used for model selection is applied only on the continuous variables while the number of levels for the categorical variables is fixed (*i.e.*, it is not defined from the sample size). When the model is selected, the estimation of the extension of model (2) can be easily achieved by maximizing the smoothed log-likelihood via an MM algorithm. The proposed approach detects the true number of clusters (*i.e.*, two) while the approach implemented in VarSelLCM overestimates it since it selects six components. Moreover, by considering ten variables as relevant for clustering, our procedure returns a more relevant partition with respect to the occurrence of heart disease because it obtains an ARI equals to 0.37 while the procedure implemented in VarSelLCM obtains an ARI equal to 0.12.

Echocardiogram data set We consider the Echocardiogram Data Set (Salzberg, 1988) freely available in the R package MNARclust. This data set is composed by $n = 132$ subjects who suffered from heart attack at some point in the past. The task is generally to determine from the other variables whether or not the patient will survive at least one year. The data set is composed by 5 continuous variables: *age at heart attack* (missing rate 4.5%), *fractional shortening* (a measure of contractility around the heart, lower numbers are increasingly abnormal, missing rate 6.0%), *epss* (E-point septal separation, another measure of contractility, larger numbers are increasingly abnormal, missing rate 11.4%), *lvdd* (left ventricular end-diastolic dimension; this is a measure of the size of the heart at end-diastole; large hearts tend to be sick hearts, missing rate 8.3%) and *wall-motion-score* (a measure of how the segments of the left ventricle are moving, missing rate 3.0%); one binary variable *pericardial effusion* (pericardial effusion is fluid around the heart, 0=no fluid, 1=fluid, missing rate 0.7%). We also have one binary variable which can be used as a partition among the subjects: *still alive* (0=dead at end of survival period, 1=still alive). This binary variable is not used for clustering but permits to evaluate the accuracy of the estimated partition. Among the variables used for clustering there is 5.7% of missing values and 19.1% of the subjects have at least one missing value. Moreover, the variable *still alive* has only one missing value. Du Roy de Chaumaray and Marbac (2023) perform the cluster analysis of this data set by considering an extension of (2) to the case of mixed-type data (data set composed of one binary and five continuous variables) with non-ignorable missingness process. Again, to use the proposed approach, we only discretize the continuous data. Moreover, to deal with the non-ignorable missingness, the absence of response is defined as a level of each variable. Performing a full model selection, the proposed approach detects three clusters. This results was suggested in Du Roy de Chaumaray and Marbac (2023) with empirical analysis of the evolution of the smoothed log-likelihood with respect to the number of clusters. Moreover, the proposed approach selects three relevant variables (fractional.shortening, epss and lvdd). Note that these three variables are detected as the most discriminative one for the missingness process and for their conditional distribution within component given the fact that the variable is observed (see Figures D.7 and D.8 in Du Roy de Chaumaray and Marbac (2023)). Finally, note that the procedure returns a partition which is partially related to the indicator of surviving after the study since the ARI is 0.19.

6 Conclusion

In this paper, we introduced a novel approach for a full model selection (number of components and subset of relevant variables for clustering) in multivariate finite mixture models. Thus, this approach is the first that allows variable selection for non-parametric mixture models. Under mild assumptions on the distributions of the observed variables, we showed that the model can be identified using a discretization of the data into bins and the penalized log-likelihood of the resulting distribution. Ranges are given for the number of bins and the penalty term to ensure the consistency of the procedure. With a careful reading of our proof, a finite-sample size control of the probability of overestimating the model can be obtained.

Since different discretizations satisfy the sufficient conditions leading the consistency of the estimators, numerical experiments have been conducted to provide advice about the choice of this tuning parameters. Based on the results, we recommend to use a discretization of the data based on the empirical quantiles where the number of bins is defined by $\lceil n^{1/6} \rceil$. Note that due to the discretization, the proposed method implies an approximation of the density via histograms. Thus, we recommend to use the method only for model selection, then to use an usual kernel-based method for estimating the parameters of the selected model.

In this paper, we consider that each X_{ij} is a scalar variable. However, the conditional independence structure of the mixture model may be more plausible for groups of components of \mathbf{X}_i . As suggested by a reviewer, the proposed method can be extended to this framework. Hence, the model defined by (2) holds true when each x_{ij} is a vector. The proposed method can be used for selecting the number of components and the groups of components of \mathbf{X}_i that are relevant for clustering, as long as the groups of components of \mathbf{X}_i which are conditionally independent are known. In such case, a discretization into B cells is performed on each vector x_{ij} . Thus, model (3) still holds true where l_{Bjb} corresponds to the volume of the b -th cell in the partition of the support of the vector X_{ij} . Despite the fact that the procedure remains consistent, the performances of this method could be deteriorated if the dimension of X_{ij} is large, because it involves non-parametric estimators for multivariate densities.

7 Data availability statement

The data underlying the illustration of this article were derived from sources in the public domain: the Swiss banknotes data set (Flury, 1988) is available in the R package `mclust` (Scrucca et al., 2016), the data set collected by Streuli (1973) that reports on the chemical composition of coffee samples is available in the R package `pgmm` (McNicholas et al., 2015), the Cleveland data set is available at <https://www.kaggle.com/ronitf/heart-disease-uci/version/1> and the Echocardiogram Data Set (Salzberg, 1988) is available in the R package `MNARclust`.

8 Fundings and conflict of interests

The authors have no funding to report. They have no conflict of interest to disclose.

References

- Akaike, H. (1970). Statistical predictor identification. *Annals of the Institute of Statistical Mathematics* 22(1), 203–217.
- Allman, E. S., C. Matias, J. A. Rhodes, et al. (2009). Identifiability of parameters in latent structure models with many observed variables. *The Annals of Statistics* 37(6A), 3099–3132.
- An, Y., Y. Hu, and M. Shum (2010). Estimating first-price auctions with an unknown number of bidders: A misclassification approach. *Journal of Econometrics* 157(2), 328–341.
- Azizyan, M., A. Singh, and L. Wasserman (2013). Minimax theory for high-dimensional gaussian mixtures with sparse mean separation. *Neural Information Processing Systems, NIPS*.

- Benaglia, T., D. Chauveau, D. Hunter, and D. Young (2009). mixtools: An r package for analyzing finite mixture models. *Journal of Statistical Software* 32(6), 1–29.
- Benaglia, T., D. Chauveau, and D. R. Hunter (2009). An EM-like algorithm for semi-and nonparametric estimation in multivariate mixtures. *Journal of Computational and Graphical Statistics* 18(2), 505–526.
- Bonhomme, S., K. Jochmans, and J.-M. Robin (2016). Non-parametric estimation of finite mixtures from repeated measurements. *Journal of the Royal Statistical Society: Series B: Statistical Methodology*, 211–229.
- Bonhomme, S., K. Jochmans, and J.-M. Robin (2017). Nonparametric estimation of non-exchangeable latent-variable models. *Journal of Econometrics* 201(2), 237–248.
- Bonhomme, S., K. Jochmans, J.-M. Robin, et al. (2016). Estimating multivariate latent-structure models. *The Annals of Statistics* 44(2), 540–563.
- Bontemps, D. and W. Toussile (2013). Clustering and variable selection for categorical multivariate data. *Electronic Journal of Statistics* 7(none), 2344 – 2371.
- Chambaz, A. et al. (2006). Testing the order of a model. *Annals of statistics* 34(3), 1166–1203.
- Chauveau, D., D. R. Hunter, M. Levine, et al. (2015). Semi-parametric estimation for conditional independence multivariate finite mixture models. *Statistics Surveys* 9, 1–31.
- Chernozhukov, V., D. Chetverikov, and K. Kato (2014). Gaussian approximation of suprema of empirical processes. *The Annals of Statistics* 42(4), 1564–1597.
- Clogg, C. C. (1995). Latent class models. In *Handbook of statistical modeling for the social and behavioral sciences*, pp. 311–359. Springer.
- Compiani, G. and Y. Kitamura (2016). Using mixtures in econometric models: a brief review and some new results. *The Econometrics Journal* 19(3), C95–C127.
- Cruz-Medina, I., T. Hettmansperger, and H. Thomas (2004). Semiparametric mixture models and repeated measures: the multinomial cut point model. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 53(3), 463–474.
- Dacunha-Castelle, D. and E. Gassiat (1997). Testing in locally conic models, and application to mixture models. *ESAIM: Probability and Statistics* 1, 285–317.
- Dacunha-Castelle, D. and E. Gassiat (1999). Testing the order of a model using locally conic parametrization: population mixtures and stationary arma processes. *The Annals of Statistics* 27(4), 1178–1209.
- Dean, N. and A. E. Raftery (2010). Latent class analysis variable selection. *Annals of the Institute of Statistical Mathematics* 62(1), 11–35.
- Dempster, A. P. and Laird, N. M. and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* 39(1), 1–38.
- Du Roy de Chaumaray, M. and M. Marbac (2023). Clustering data with non-ignorable missingness using semi-parametric mixture models assuming independence within components. *Advances in Data Analysis and Classification*, 1–42.
- Dudley, R. M. (2014). *Uniform central limit theorems*, Volume 142. Cambridge university press.
- Elmore, R. T., T. P. Hettmansperger, and H. Thomas (2004). Estimating component cumulative distribution functions in finite mixture models. *Communications in Statistics-Theory and Methods* 33(9), 2075–2086.
- Flury, B. (1988). *Multivariate statistics: a practical approach*. Chapman & Hall, Ltd.

- Friedman, J. and J. Meulman (2004). Clustering objects on subsets of attributes (with discussion). *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 66(4), 815–849.
- Fruhwirth-Schnatter, S., G. Celeux, and C. P. Robert (2019). *Handbook of mixture analysis*. CRC press.
- Gassiat, E. (2002). Likelihood ratio inequalities with applications to various mixtures. In *Annales de l’Institut Henri Poincaré (B) Probability and Statistics*, Volume 38, pp. 897–906. Elsevier.
- Goodman, L. A. (1974). Exploratory latent structure analysis using both identifiable and unidentifiable models. *Biometrika* 61(2), 215–231.
- Green, P. J. (1990). On use of the em for penalized likelihood estimation. *Journal of the Royal Statistical Society. Series B (Methodological)*, 443–452.
- Hagenaars, J. A. and A. L. McCutcheon (2002). *Applied latent class analysis*. Cambridge University Press.
- Hall, P., A. Neeman, R. Pakyari, and R. Elmore (2005). Nonparametric inference in multivariate mixtures. *Biometrika* 92(3), 667–678.
- Hall, P. and X.-H. Zhou (2003). Nonparametric estimation of component distributions in a multivariate mixture. *The Annals of Statistics* 31(1), 201–224.
- Hettmansperger, T. and H. Thomas (2000). Almost nonparametric inference for repeated measures in mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 62(4), 811–825.
- Hu, Y., D. McAdams, and M. Shum (2013). Identification of first-price auctions with non-separable unobserved heterogeneity. *Journal of Econometrics* 174(2), 186–193.
- Hubert, L. and P. Arabie (1985). Comparing partitions. *Journal of classification* 2(1), 193–218.
- Kasahara, H. and K. Shimotsu (2014). Non-parametric identification and estimation of the number of components in multivariate mixtures. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 76(1), 97–111.
- Keribin, C. (2000). Consistent estimation of the order of mixture models. *Sankhyā: The Indian Journal of Statistics, Series A*, 49–66.
- Kleibergen, F. and R. Paap (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of econometrics* 133(1), 97–126.
- Kwon, C. and E. Mbakop (2021). Estimation of the number of components of nonparametric multivariate finite mixture models. *The Annals of Statistics* 49(4), 2178 – 2205.
- Levine, M., D. R. Hunter, and D. Chauveau (2011). Maximum smoothed likelihood for multivariate mixtures. *Biometrika*, 403–416.
- Marbac, M. and M. Sedki (2017). Variable selection for model-based clustering using the integrated complete-data likelihood. *Statistics and Computing* 27(4), 1049–1063.
- Marbac, M. and M. Sedki (2019). Varsellcm: an r/c++ package for variable selection in model-based clustering of mixed-data with missing values. *Bioinformatics* 35(7), 1255–1257.
- Marbac, M., M. Sedki, and E. Patin (2019). Variable selection for mixed data clustering: application in human population genomics. *Journal of Classification*, 1–19.
- McLachlan, G. and D. Peel (2000). *Finite mixture models*. Wiley-Interscience, New York: Wiley Series in Probability and Statistics: Applied Probability and Statistics.
- McNicholas, P. D., A. ElSherbiny, K. R. Jampani, A. F. McDaid, T. B. Murphy, L. Banks, and M. P. D. McNicholas (2015). Package ‘pgmm’.

- McNicholas, P. D. and T. B. Murphy (2008). Parsimonious gaussian mixture models. *Statistics and Computing* 18(3), 285–296.
- Pan, W. and X. Shen (2007). Penalized Model-Based Clustering with Application to Variable Selection. *Journal of Machine Learning Research* 8, 1145–1164.
- Raftery, A. E. and N. Dean (2006). Variable selection for model-based clustering. *Journal of the American Statistical Association* 101(473), 168–178.
- Salzberg, S. L. (1988). *Exemplar-based learning: Theory and implementation*. Harvard University, Center for Research in Computing Technology, Aiken.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics* 6(2), 461–464.
- Scrucca, L., M. Fop, T. B. Murphy, and A. E. Raftery (2016). mclust 5: clustering, classification and density estimation using gaussian finite mixture models. *The R journal* 8(1), 289.
- Streuli, H. (1973). Der heutige stand der kaffeechemie. In *ASSIC, 6e. Colloque, Bogota*, Volume 61.
- Tadesse, M. G., N. Sha, and M. Vannucci (2005). Bayesian variable selection in clustering high-dimensional data. *Journal of the American Statistical Association* 100(470), 602–617.
- Tibshirani, R., G. Walther, and T. Hastie (2001). Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63(2), 411–423.
- Toussile, W. and E. Gassiat (2009, August). Variable selection in model-based clustering using multilocus genotype data. *Advances in Data Analysis and Classification* 3(2), 109–134.
- Witten, D. and R. Tibshirani (2010). A Framework for Feature Selection in Clustering. *Journal of the American Statistical Association* 105(490), 713–726.
- Zheng, C. and Y. Wu (2020). Nonparametric estimation of multivariate mixtures. *Journal of the American Statistical Association* 115(531), 1456–1471.