

Variable selection methods for model-based clustering

Michael Fop* and Thomas Brendan Murphy*

University College Dublin
e-mail: michael.fop@ucd.ie; brendan.murphy@ucd.ie

Abstract: Model-based clustering is a popular approach for clustering multivariate data which has seen applications in numerous fields. Nowadays, high-dimensional data are more and more common and the model-based clustering approach has adapted to deal with the increasing dimensionality. In particular, the development of variable selection techniques has received a lot of attention and research effort in recent years. Even for small size problems, variable selection has been advocated to facilitate the interpretation of the clustering results. This review provides a summary of the methods developed for variable selection in model-based clustering. Existing R packages implementing the different methods are indicated and illustrated in application to two data analysis examples.

Keywords and phrases: Gaussian mixture model, latent class analysis, model-based clustering, R packages, variable selection.

Received July 2017.

Contents

1	Introduction	19
2	Methods of variable selection	20
3	Model-based clustering	22
3.1	Gaussian mixture model	24
3.2	Latent class analysis model	24
4	Variable selection methods for Gaussian mixture models	25
4.1	Bayesian approaches	25
4.2	Penalization approaches	27
4.3	Model selection approaches	32
4.4	Other approaches	36
4.5	R packages and data example	39
5	Variable selection methods for latent class analysis	44
5.1	Bayesian approaches	44
5.2	Penalization approaches	45
5.3	Model selection approaches	46
5.4	Other approaches	49
5.5	R packages and data example	50

*This work was supported by the Science Foundation Ireland funded Insight Research Centre (SFI/12/RC/2289)

6 Discussion	53
Acknowledgments	57
References	57

1. Introduction

Model-based clustering is a well established and popular tool for clustering multivariate data. In this approach, clustering is formulated in a modeling framework and the data generating process is represented through a finite mixture of probability distributions. Often, all the variables at a user's disposal are employed in the modeling. Nonetheless, in many situations considering all the variables unnecessarily increases the model complexity. Moreover, some variables may not possess any clustering information and are of no use in the detection of the group structure. Rather, they could be detrimental to the clustering. Likewise, the case were all of the variables contain clustering information can also be problematic. Along with the increasing number of dimensions comes the *curse of dimensionality* (Bellman, 1957) and including superfluous variables in the model leads to identifiability problems and over-parameterization (Bouveyron and Brunet-Saumard, 2014a; Bartholomew, Knott and Moustaki, 2011). Therefore, resorting to variable selection techniques can facilitate model fitting, ease the interpretation of the results and lead to data classification of better quality. Even in situations of moderate or low dimensionality, reducing the set of variables employed in the clustering process can be beneficial (Fowlkes, Gnanadesikan and Kettenring, 1988).

This article gives a review of the available methods for variable selection in model-based clustering. Starting from early works on the topic, we will give a summary of the various approaches up to the most recent developments. The focus will be on variable selection when clustering multivariate continuous and categorical data, the two most common data types in practice. References to the existing R packages that implement the different methods are provided. As illustrative examples, we will apply the variable selection methods to two illustrative datasets.

The exposition is structured as follows. Section 2 introduces the main features of the variable selection methods for mixture model clustering. Section 3 recalls the model-based clustering framework, with a brief description of the two principal models for multivariate continuous and categorical data: Gaussian mixture models and latent class analysis. In Section 4, classical and recent variable selection methods for Gaussian mixture models are reviewed. The section terminates with the list of available R packages and a data analysis example concerning historical mortality rates. Section 5 is dedicated to variable selection methods for latent class analysis. R packages and an application to voting data are presented in conclusion. The paper ends with a discussion in Section 6.

2. Methods of variable selection

When performing variable selection for clustering, the goal is to retain only the relevant variables. With this aim, it is crucial to define what it means for a variable to be, or not to be, “relevant”. In supervised learning the matter has been the object of rigorous discussion for a long time; we mention the works of Yu and Liu (2004); Blum and Langley (1997); Kohavi and John (1997); Koller and Sahami (1996) and John, Kohavi and Pfleger (1994). Within the model-based clustering approach, the question has been addressed recently in Ritter (2014). Model-based clustering places the clustering task into a formal modeling framework and the group structure is embedded in the group membership variable (McLachlan and Basford, 1988). In this case, the definition of relevance can be expressed in terms of probabilistic dependence (or independence) statements with respect to this variable (Ritter, 2014). *Relevant variables* contain the essential clustering information. In a model-based clustering context, their distribution directly depends on the group membership variable. On the other hand, *irrelevant variables* do not convey any beneficial information. These can be further divided into redundant and uninformative variables. *Redundant variables* provide information of similar quality to that already available in the relevant ones, therefore are not needed for a parsimonious modeling. In many situations, they contain similar information because they are correlated with the relevant ones. In terms of distributional representation, we may think of redundant variables as conditionally independent of the grouping variable given the relevant variables; they could be useful for clustering, but only if the relevant ones are not present. On the contrary, *uninformative variables* possess no discriminative information whatsoever. They correspond to noise and their distribution is completely independent of the group structure.

Different model-based clustering and variable selection strategies can be delineated by distributional assumptions on relevant and irrelevant variables. Two main assumptions are peculiar to the task of variable selection and mixture model clustering; these are:

- *Local independence assumption.* The relevant variables are conditionally independent within the groups.
- *Global independence assumption.* The irrelevant variables are independent of relevant clustering variables.

Figure 1 gives a simple sketch of the two independence statements. In specifying a general strategy for model-based clustering and selection of variables, these assumptions can be combined together or employed separately. They have different implications on the clustering model and the model for the relations among the variables. In particular, the local independence assumption helps to simplify the modeling of the joint distribution of relevant variables and is especially useful in high-dimensional data settings. It is a standard assumption of the latent class analysis model (Clogg, 1988). For Gaussian mixture models, the statement corresponds to assuming components with diagonal covariance matrices. The global independence assumption implies that the joint distribution

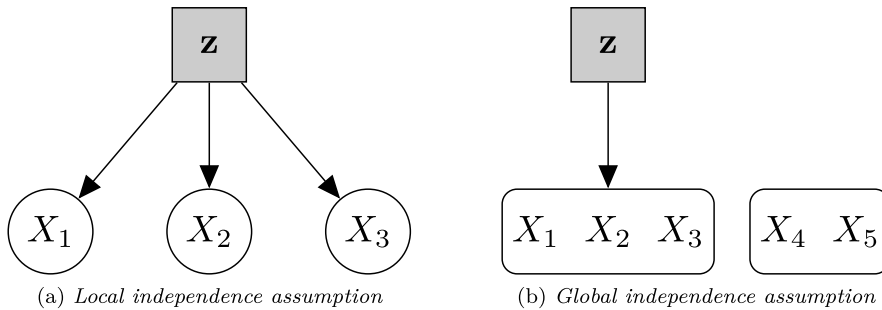


FIG 1. Local and global independence assumptions. In the example, \mathbf{z} is the group membership variable, X_1, X_2 and X_3 are relevant clustering variables, while X_4 and X_5 are irrelevant and not related to \mathbf{z} . Under the local independence assumption there are no edges among the relevant variables. Under the global independence assumption there is no edge between the set of relevant variables and the set of irrelevant ones.

of the variables factors into the product of the mixture distribution of the relevant variables and the distribution of the irrelevant variables. The term “global” is used because the independence statement affects the distribution of *all* the variables, not only the clustering ones. This assumption simplifies the modeling of the relation between relevant and irrelevant variables. However, it limits the capability of taking into account the presence of redundant variables (Law, Figueiredo and Jain, 2004; Raftery and Dean, 2006). Many of the methods that are presented in this review make use of one of the two assumptions above, either implicitly or explicitly.

Likewise, how the variable selection algorithm interacts with the model fitting process defines the overall approach to the problem. For a general learning task, the principal distinction is in whether the selection is carried out separately or jointly to the learning procedure (John, Kohavi and Pfleger, 1994; Dash and Liu, 1997; Dy and Brodley, 2004; Saeys, Inza and Larrañaga, 2007; Liu and Motoda, 2007). The first case corresponds to *filter methods*, where the selection is performed as a pre (or sometimes post) processing step. The second corresponds to *wrapper methods*, that combine learning and variable selection at the same time. In this case the selection procedure is “wrapped” around the learning algorithm. Filter approaches are easy to implement and computationally efficient. However, wrapper methods often provide superior results, despite being more involved (Blum and Langley, 1997; Kohavi and John, 1997; Guyon and Elisseeff, 2003).

In model-based clustering, filter methods perform the variable selection before (or after) the model has been estimated. The inferred classification is then used to evaluate the quality of the variables. In contrast, with a wrapper method model estimation and selection are conducted simultaneously. Variables selected via filter methods can miss important information as the selection is external to the model estimation. Indeed, wrapper methods have recently attracted most of the attention and are the most widespread. Their advantages lie in the fact

that they can be naturally included in the model fitting, can lead to a better classification and can provide a better representation of the data generating process (Dy and Brodley, 2004; Law, Figueiredo and Jain, 2004).

Within the class of wrappers, the various methods for variable selection in mixture model clustering can be further distinguished according to the type of statistical approach used. Three major approaches can be found in the literature:

- *Bayesian approaches.* In this class of methods, the problem of variable selection is addressed assuming the existence of a latent variable indicating if an observed variable characterizes a mixture distribution or not. Variable selection is conducted by making inference about the posterior distribution of such a latent variable.
- *Penalization approaches.* Within this category, variable selection is performed by using a penalized log-likelihood approach. The penalization term is a function of the mixture parameters and acts to shrink the estimates towards an overall common value. Variables whose estimates take this common value across the different mixture components are considered irrelevant and are discarded.
- *Model selection approaches.* Here the task of variable selection is re-formulated as a model selection problem. Different models are specified according to the role of the variables towards the clustering structure. Consequently, relevant variables are selected by comparing different models using some predefined criterion.

The above categorization is not exhaustive nor the definitions are intended to be mutually exclusive. Indeed, many of the methods that will be presented in the rest of the paper have some degree of overlap and a method belonging to one type of approach could be easily rephrased in terms of the other ones. These three general approaches are the predominant in the literature and the classification will be used for ease of exposition and a systematic presentation.

In summary, the existing variable selection methods for model-based clustering are differentiated in relation to the distributional assumptions for relevant and irrelevant variables, the interaction between variable selection and model fitting, and the general statistical approach. We will give an overview of these methods after a short description of the model-based clustering framework.

3. Model-based clustering

Let \mathbf{X} be the $N \times J$ data matrix, where each row $\mathbf{x}_i = (x_{i1}, \dots, x_{ij}, \dots, x_{iJ})$ is the realization of a J -dimensional vector of random variables $\mathbb{X} = (X_1, \dots, X_j, \dots, X_J)$. Model-based clustering assumes that each observation arises from a finite mixture of G probability distributions, each representing a different cluster or group (Fraley and Raftery, 2002; Melnykov and Maitra, 2010; McNicholas, 2016). The general form of a finite mixture distribution is specified as follows:

$$p(\mathbf{x}_i; \Theta) = \sum_{g=1}^G \tau_g p(\mathbf{x}_i; \Theta_g), \quad (1)$$

where the τ_g are the mixing probabilities and Θ_g is the parameter set corresponding to component g ; Θ denotes the set of all parameters of the mixture. The component densities fully characterize the group structure of the data and each observation belongs to the corresponding cluster according to a latent cluster membership indicator variable $\mathbf{z}_i = (z_{i1}, \dots, z_{ig}, \dots, z_{iG})$, such that $z_{ig} = 1$ if \mathbf{x}_i arises from the g th subpopulation (McLachlan and Peel, 2000; McLachlan and Basford, 1988).

For a fixed number of components, parameters are usually estimated using the EM algorithm (Dempster, Laird and Rubin, 1977; McLachlan and Krishnan, 2008; Bartholomew, Knott and Moustaki, 2011; O'Hagan, Murphy and Gormley, 2012). Moreover, generally model selection corresponds to the selection of the number of components G and to accomplish the task a plethora of methods have been suggested in the literature, the *Bayesian Information Criterion* (BIC, Schwarz, 1978) being the most popular one. Another popular approach for mixture model selection is the *Integrated Complete-data Likelihood* criterion (ICL, Biernacki, Celeux and Govaert, 2000), which gives more importance to models with well separated clusters. See McLachlan and Rathnayake (2014) for a detailed review of the various methods.

After parameters have been estimated, each observation is assigned to the corresponding cluster using the *maximum a posteriori* (MAP) rule (McLachlan and Peel, 2000; McNicholas, 2016). The posterior probabilities $u_{ig} = \Pr(z_{ig} = 1 | \mathbf{x}_i)$ of observing cluster g given the data point i are estimated as follows:

$$\hat{u}_{ig} = \frac{\hat{\tau}_g p(\mathbf{x}_i; \hat{\Theta}_g)}{\sum_{h=1}^G \hat{\tau}_h p(\mathbf{x}_i; \hat{\Theta}_h)},$$

Then observation \mathbf{x}_i is assigned to cluster g if

$$g = \text{MAP}(\hat{\mathbf{u}}_i) = \underset{h}{\operatorname{argmax}} \{ \hat{u}_{i1}, \dots, \hat{u}_{ih}, \dots, \hat{u}_{iG} \}.$$

According to the nature of the data, different specifications for the component densities in (1) have been proposed. In the following sections we focus on the cases of continuous and categorical data, taking in consideration the two most popular distributions: Gaussian and Multinomial. However, various and more flexible distributional assumptions can be specified, enabling to take into account for skewness, heavy tails and different data types; for example, see McNicholas (2016) for a review on non-Gaussian components, McLachlan and Peel (1998) for mixtures of multivariate t -Student distributions, Lee and McLachlan (2013) and Lee and McLachlan (2016) for mixtures of skew- t and skew normal distributions, Karlis and Meligkotsidou (2007) and Rau et al. (2015) for clustering multivariate count data, McParland and Gormley (2016) for mixed data model-based clustering, Kosmidis and Karlis (2016) for the use of copulas in model-based clustering, DeSantis et al. (2008) for latent class analysis of ordinal data.

3.1. Gaussian mixture model

When clustering multivariate continuous data, a common approach is to model each component density by a multivariate Gaussian distribution. For a *Gaussian mixture model* (GMM) the mixture density in (1) becomes:

$$p(\mathbf{x}_i; \Theta) = \sum_{g=1}^G \tau_g \phi(\mathbf{x}_i; \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g), \quad (2)$$

where ϕ is the multivariate Gaussian density and $\boldsymbol{\mu}_g$ and $\boldsymbol{\Sigma}_g$ are the mean and covariance parameters respectively; see Fraley and Raftery (2002), Melnykov and Maitra (2010) and McNicholas (2016) for reviews. To attain parsimony, several approaches involving re-parameterizations of the covariance matrix $\boldsymbol{\Sigma}_g$ have been presented; for example Banfield and Raftery (1993); Celeux and Govaert (1995); Bouveyron, Girard and Schmid (2007); McNicholas and Murphy (2008); Biernacki and Lourme (2014). We refer to Bouveyron and Brunet-Saumard (2014a) for a review.

Model based-clustering via GMMs can be performed in R (R Core Team, 2017) using the packages `mclust` (Scrucca et al., 2016), `Rmixmod` (Lebre et al., 2015), `EMCluster` (Chen and Maitra, 2015), `mixtools` (Benaglia et al., 2009) and `flexmix` (Leisch, 2004); also the `EMMIX` software (McLachlan et al., 1999) can be used for fitting mixtures of Gaussian distributions. Lastly, we note that in Python (Python Software Foundation, 2017), GMMs estimation can be executed via the packages `scikit-learn`¹ and `PyMixmod`².

3.2. Latent class analysis model

For clustering multivariate categorical data, the common approach is to use the *latent class analysis model* (LCA, Bartholomew, Knott and Moustaki, 2011). Under this model, the mixture density in (1) is a mixture of Multinomial distributions as follows:

$$p(\mathbf{x}_i; \Theta) = \sum_{g=1}^G \tau_g \mathcal{C}(\mathbf{x}_i; \boldsymbol{\pi}_g), \quad (3)$$

where

$$\mathcal{C}(\mathbf{x}_i; \boldsymbol{\pi}_g) = \prod_{j=1}^J \prod_{c=1}^{C_j} \pi_{gjc}^{\mathbb{1}\{x_{ij}=c\}},$$

with π_{gjc} representing the probability of occurrence of category c for variable X_j in class g , and C_j the number of categories of variable j . The factorization in $\mathcal{C}(\mathbf{x}_i; \boldsymbol{\pi}_g)$ is due to the local independence assumption, stating that the variables are independent within each latent class (Clogg, 1988). More details about the model can be found in Vermunt and Magdison (2002); Agresti (2002); Collins

¹<http://scikit-learn.org/stable/>

²<http://www.mixmod.org/spip.php?article62>

and Lanza (2010). For different values of G , not all the models can be fitted and constraints on the parameters need to be placed in order to ensure identifiability (Goodman, 1974); for examples see Formann (1985).

R packages implementing the LCA model for clustering categorical data are `BayesLCA` (White and Murphy, 2014), `poLCA` (Linzer and Lewis, 2011) and `flexmix` (Leisch, 2004); package `e1071` (Meyer et al., 2017) contains function `lca` for fitting the LCA model on binary data.

4. Variable selection methods for Gaussian mixture models

In this section we provide an overview of the available methods for clustering and variable selection using Gaussian mixture models. Steinley and Brusco (2008) and Celeux et al. (2014) compare and evaluate the performances of some of the methods described in the subsequent sections. In particular, Steinley and Brusco (2008) perform an empirical comparison of different procedures for clustering and variable selection, also those not based on mixture models. While Celeux et al. (2014) compare the model selection approach of Maugis, Celeux and Martin-Magniette (2009a) (see Section 4.3) and the regularization approach of Witten and Tibshirani (2010) (see Section 4.2); the authors found that, in the case of correlated variables, the model selection approach was substantially more accurate in terms of both classification and variable selection than the regularization approach, and that both gave more accurate classifications than K-means without variable selection.

4.1. Bayesian approaches

Various methods have been developed within the Bayesian paradigm for simultaneous model-based clustering and variable selection. A common feature among them is the introduction of a variable φ , usually following a distribution $p(\varphi)$. The variable splits \mathbf{X} into two sets: \mathbf{X}^C , the set of variables defining a Gaussian mixture distribution, and \mathbf{X}^{NC} , the set of variables indicating a single multivariate Gaussian distribution. In its most general form, the distribution of the data, conditional on \mathbf{z} , can be expressed as:

$$p(\mathbf{X} | \mathbf{z}, G, \varphi, \Omega) = p(\mathbf{X}^C | \mathbf{z}, G, \varphi, \Theta) p(\mathbf{X}^{NC} | \varphi, \Gamma),$$

where $p(\mathbf{X}^C | \mathbf{z}, G, \varphi, \Theta)$ is a mixture distribution whose parameters are denoted by Θ , $p(\mathbf{X}^{NC} | \varphi, \Gamma)$ is a single multivariate Gaussian distribution with parameters Γ , and Ω denotes the collection $\{\Theta, \Gamma\}$ of all parameters. Then, with the aim of variable selection and clustering, the focus is in drawing inference from the posterior distribution:

$$p(\mathbf{z}, G, \varphi, \Omega | \mathbf{X}) \propto p(\mathbf{X} | \mathbf{z}, G, \varphi, \Omega) p(\mathbf{z} | \Omega) p(G) p(\varphi) p(\Omega).$$

In this context, Liu et al. (2003) propose the *anchor mode model*, where variable selection is performed by selecting the most informative principal components (or factors) of the data. In this approach, a preliminary dimension reduction through principal component analysis is performed and the first k_0 factors

are retained. Then it is assumed that only a subset of these factors is informative for clustering and distributed according to a mixture of Gaussians, while the remaining components follow a simple Gaussian distribution. The subset consists of the first φ principal components, where φ is a random variable distributed according to the prior distribution $p(\varphi)$. Inference on this number of relevant factors is conducted employing a *Markov Chain Monte-Carlo* (MCMC) scheme where the prior $p(\varphi)$ is taken to be the Uniform distribution. The method is shown to perform well on high-dimensional gene expression data, however, selection is performed on a set of features derived from the original variables and in general the principal components with the larger eigenvalues will not necessarily contain the most useful information about the clustering (see Chang, 1983).

As an alternative to the “hard selection” approach (a variable is either selected or not), Law, Figueiredo and Jain (2004) suggest a Bayesian framework where the concept of *feature saliency* is introduced. Let $\varphi = (\varphi_1, \dots, \varphi_j, \dots, \varphi_J)$ be a binary variable such that $\varphi_j = 1$ if X_j is relevant and $\varphi_j = 0$ otherwise. Then the saliency of variable X_j is the quantity $\rho_j = \Pr(\varphi_j = 1)$ and can be interpreted as the importance of the variable in characterizing the cluster structure of the data. Assuming conditional independence of the variables, it follows that the likelihood of a data-point can be expressed as:

$$p(\mathbf{x}_i | \Omega) = \sum_{g=1}^G \tau_g \prod_{j=1}^J [\rho_j \phi(x_{ij} | \mu_g, \sigma_g^2) + (1 - \rho_j) \phi(x_{ij} | \Gamma)],$$

with clear use of the notation. Rather than φ , the interest here is in recovering the probabilities ρ_j and a Dirichlet prior is placed on the corresponding vector. To encourage the saliences of some variables to converge to zero, the authors adopt a *minimum message length* criterion (Wallace and Freeman, 1987) and utilize an EM algorithm for maximum a posteriori estimation. Within the same framework, Constantinopoulos, Titsias and Likas (2006) consider a variational learning method for estimation of the saliences.

The same idea of a binary clustering-relevance indicator variable φ is employed in Tadesse, Sha and Vannucci (2005). The authors assume a prior on φ of the form:

$$p(\varphi | \eta) = \prod_{j=1}^J \eta^{\varphi_j} (1 - \eta)^{1 - \varphi_j},$$

with η the hyper-parameter interpreted as the proportion of variables expected to discriminate the groups. A MCMC scheme is used for inference, and the vector φ is updated using a Metropolis search where a new candidate is generated from the previous state by adding, removing and swapping at random its entries. Posterior inference on φ is drawn after integrating out the parameters and considering the marginal posterior $p(\varphi_j = 1 | \mathbf{X})$. Then the best clustering variables can be identified as those with largest marginal posterior, $p(\varphi_j = 1 | \mathbf{X}) > t$, with a specified t . Alternatively, the selection can be performed by taking into account the complete vector φ with largest posterior probability among all visited vectors throughout the chain, thus considering the joint density of φ . The

method is shown to perform well in clustering high-dimensional microarray data. Furthermore, in subsequent work, Kim, Tadesse and Vannucci (2006) extend the approach by formulating the clustering in terms of an infinite mixture of Gaussian distributions via Dirichlet process mixtures, while Swartz et al. (2008) expand it to the modeling of data with a known structure imposed from an experimental design.

4.2. Penalization approaches

In this context, a penalization term is introduced on the model parameters and variable selection is performed by inducing sparsity in the estimates. The aim is to maximize a penalized version of the log-likelihood under a Gaussian mixture model and discard those variables whose parameter estimates are shrunk to zero or to a common value across the mixture components. In its general form, this penalized log-likelihood is as follows:

$$\ell_Q = \sum_{i=1}^N \log \left\{ \sum_{g=1}^G \tau_g \phi(\mathbf{x}_i; \Theta_g) \right\} - Q_\lambda(\Theta), \quad (4)$$

where the penalization term $Q_\lambda(\Theta)$ is a function of the Gaussian densities parameters Θ and λ , a generic penalty parameter (here in the notation we denoted with Θ the collection of all Gaussian density parameters and with Θ_g the subset corresponding to component g). Generally, the various methods are differentiated by the form of the function $Q_\lambda(\cdot)$, having different implications on the selection of variables.

Seminal work in this class of approaches is the method introduced by Pan and Shen (2007). The authors use a L_1 penalty function of the form:

$$\lambda \sum_{g=1}^G \sum_{j=1}^J |\mu_{gj}|.$$

After centering the data, the method realizes variable selection by automatically shrinking several small estimates of μ_{gj} towards zero. Indeed, if $\mu_{gj} = 0$ for all g , the component means for variable j are equal to the overall data mean and variable j does not contribute to the clustering. The authors show an application to the Golub et al. (1999) leukemia gene expression dataset, demonstrating the usefulness of the approach in “high dimensions – small sample size” settings.

A closely related approach is the one suggested by Bhattacharya and McNicholas (2014), where the penalizing function accounts for the size of the clusters and is given by:

$$\lambda \sum_{g=1}^G N \tau_g \sum_{j=1}^J |\mu_{gj}|.$$

The parameter λ depends on the sample size and the authors derive a BIC-type model selection criterion for high-dimensional settings.

The L_1 penalty function treats each μ_{gj} individually, not using the information that, across the mixture components, the parameters $(\mu_{1j}, \dots, \mu_{gj}, \dots, \mu_{Gj})$ corresponding to the same variable X_j are grouped. This results in the fact that, if for a fixed variable j and some component g , we have $\mu_{gj} \neq 0$ while $\mu_{kj} = 0$ for all the remaining components, then the variable would not be excluded. Wang and Zhu (2008) suggest a solution to the problem by replacing the L_1 norm with the L_∞ norm. Thus the penalty function is given by:

$$\lambda \sum_{j=1}^J \max_g \{ |\mu_{1j}|, \dots, |\mu_{gj}|, \dots, |\mu_{Gj}| \}.$$

After re-parameterizing $\mu_{gj} = \alpha_j \beta_{gj}$ ($\alpha_j \geq 0$), the authors consider also a hierarchical penalization function of the form:

$$\lambda_1 \sum_{j=1}^J \alpha_j + \lambda_2 \sum_{g=1}^G \sum_{j=1}^J |\beta_{gj}|,$$

with λ_1 controlling the amount of shrinkage on the μ_{gj} as a group for $g = 1, \dots, G$, and λ_2 controlling the shrinkage effect within variable j . This function has the advantage of being more flexible and inducing a less ‘‘hard’’ penalization than the L_∞ norm. Both penalty functions take into account the fact that component means corresponding to the same variable can be treated as grouped and tend to conduct a more effective variable selection.

The idea of grouped parameters is also accounted in Xie, Pan and Shen (2008a). Here the authors suggest the use of two planes of grouping: *vertical* and *horizontal mean grouping*. For the first, mean parameters afferent to the same variable are treated as a whole and the penalty function is:

$$\lambda \sqrt{G} \sum_{j=1}^J \|\boldsymbol{\mu}_j\|_2$$

where $\boldsymbol{\mu}_j = (\mu_{1j}, \dots, \mu_{gj}, \dots, \mu_{Gj})$ and $\|\cdot\|_2$ denotes the L_2 norm. The same idea is exploited in the horizontal grouping, where prior knowledge that some variables work in groups is introduced in the penalization. Here the variables can be grouped in M groups indexed by m and each one of size H_m . Then the function is given by:

$$\lambda \sum_{g=1}^G \sum_{m=1}^M \sqrt{H_m} \|\boldsymbol{\mu}_{gm}\|_2,$$

with $\boldsymbol{\mu}_{gm}$ the vector of means of component g for variables in group m . The two planes of penalization can be combined together and Xie, Pan and Shen (2008a) show their superior performance in comparison to the standard L_1 penalty.

An approach allowing to identify which variables are discriminative for which specific pairs of clusters is the one proposed by Guo et al. (2010). They introduce

a *pairwise fusion* penalty of the form:

$$\lambda \sum_{j=1}^J \left(\sum_{g=1}^G \sum_{h < g} |\mu_{gj} - \mu_{gh}| \right).$$

The penalty function shrinks the mean parameters toward each other and if $\mu_{gj} = \mu_{gh}$, variable X_j is not useful in discriminating components g and h , but may be useful in discriminating the other components. A variable will be considered as non-informative and discarded only in the case where all the cluster means are shrunken to the same value.

A shared characteristic of the penalization methods presented to this point is the assumption that the data are generated by a Gaussian mixture distribution with a common diagonal covariance matrix. Consequently, the mixture components are differentiated only by their mean parameters and variables with cluster-specific means all equal can be discarded since non-informative to clustering. Nonetheless, in some situations, the assumption of a common isotropic covariance matrix might be too restrictive, as it implies that all the clusters are of spherical shape and have the same volume. For example, the Golub et al. (1999) data contain a sample of 38 patients known to have three types of leukemia, with different variability in the gene expressions. In this case, assuming a common isotropic covariance matrix would likely lead to the selection of a number of clusters larger than the actual number of types of leukemia in order to accommodate for the extra within-cluster variability.

Xie, Pan and Shen (2008b) move away from this assumption and present an approach for Gaussian mixture distributions with cluster-specific diagonal covariance matrices. They introduce two penalization terms of the form:

$$\begin{aligned} \lambda_1 \sum_{g=1}^G \sum_{j=1}^J |\mu_{gj}| + \lambda_2 \sum_{g=1}^G \sum_{j=1}^J |\sigma_{gj}^2 - 1|, \\ \lambda_1 \sum_{g=1}^G \sum_{j=1}^J |\mu_{gj}| + \lambda_2 \sum_{g=1}^G \sum_{j=1}^J |\log \sigma_{gj}^2|. \end{aligned}$$

After the data have been standardized to have mean 0 and variance 1 for all the variables, variable selection is performed by removing those variables having common mean 0 and common variance 1 across the clusters. The authors also expand the penalization to account for grouped variables in a similar fashion as in Xie, Pan and Shen (2008a). The method is applied to the Golub et al. (1999) dataset and it is shown to perform a more parsimonious selection than the approach of Pan and Shen (2007) and to detect clusters where gene expressions have different variances.

Zhou, Pan and Shen (2009) further extend the previous method to the case of unconstrained covariance matrices. The authors suggest the following penalization:

$$\lambda_1 \sum_{g=1}^G \sum_{j=1}^J |\mu_{gj}| + \lambda_2 \sum_{g=1}^G \|\mathbf{W}_g\|_1,$$

where \mathbf{W}_g is the precision matrix, $\mathbf{W}_g = \boldsymbol{\Sigma}_g^{-1}$, and $\|\mathbf{W}\|_1 = \sum_{j=1}^J \sum_{h=1}^J |w_{jh}|$. The term involving the mean parameters is used for variable selection, while that for the precision is a regularization for dealing with high-dimensional data without imposing a diagonal covariance matrix. In this framework, covariances between the variables are taken into account and clusters are also allowed to be elliptical and of different shapes. The authors apply their method on the Golub et al. (1999) data, improving the classification performance with respect to the previous approaches.

By introducing a penalization term in the mixture of factor analyzers model (see Bouveyron and Brunet-Saumard, 2014a, for example) dimension reduction and variable selection can be performed simultaneously. Xie, Pan and Shen (2010) introduce a penalization approach in the mixture of factor analyzers model where the penalty function is given by:

$$\lambda_1 \sum_{g=1}^G \sum_{j=1}^J |\mu_{gj}| + \lambda_2 \sum_{g=1}^G \sum_{j=1}^J \sqrt{\sum_{r=1}^R \gamma_{gjr}^2},$$

with γ_{gjr} the cluster-specific loading corresponding to variable j and latent dimension r . The penalization on the loadings serves as a grouped variable penalty similarly to Xie, Pan and Shen (2008a). The author apply the method for clustering gene expression data of lung cancer patients, uncovering clusters of subjects with distinct risks of mortality.

Galimberti, Montanari and Viroli (2009) introduce a mixture of factor analyzers model with common factor loading matrix that projects the cluster means in a low dimensional space. They consider a penalized log-likelihood of the form:

$$\ell_Q = \sum_{i=1}^N \log \left\{ \sum_{g=1}^G \tau_g \phi(\mathbf{x}_i; \mathbf{V}\boldsymbol{\mu}_g, \mathbf{V}\boldsymbol{\Sigma}_g\mathbf{V}' + \mathbf{B}) \right\} - \lambda N \sum_{j=1}^J \sum_{r=1}^R |\gamma_{jr}|,$$

where \mathbf{V} is the (orthogonal) $J \times R$ matrix of loadings ($R < J$) with entries γ_{jr} and \mathbf{B} is a diagonal covariance matrix. Therefore, variables whose estimated loadings are $\gamma_{jr} = 0 \forall r$ will be non influential to the clustering and discarded.

A related method is suggested by Bouveyron and Brunet-Saumard (2014b), who propose to perform selection of the relevant variables by inducing sparsity in the loading matrix of the Fisher-EM algorithm (see Bouveyron and Brunet, 2012; Bouveyron and Brunet-Saumard, 2014a, for details). The method assumes that the observed data lie in a discriminative latent space defined by latent features which are linear combinations of the original variables. As before, the density of each data point is expressed by:

$$p(\mathbf{x}_i; \boldsymbol{\Theta}) = \sum_{g=1}^G \tau_g \phi(\mathbf{x}_i; \mathbf{V}\boldsymbol{\mu}_g, \mathbf{V}\boldsymbol{\Sigma}_g\mathbf{V}' + \mathbf{B}).$$

The main idea is to obtain a sparse estimate of \mathbf{V} such that variables whose loadings are shrunk to zero will not be relevant in defining the discriminative

latent subspace and hence not useful for clustering. The authors propose three different approaches. The simplest, although easy to implement and competitive with the others, aims at obtaining a sparse approximation $\tilde{\mathbf{V}}$ of the estimate $\hat{\mathbf{V}}$ computed during an iteration of the Fisher-EM. To this purpose, the following minimization problem is posed:

$$\tilde{\mathbf{V}} = \underset{\mathbf{V}}{\operatorname{argmin}} \left\| \mathbf{X}'\hat{\mathbf{V}} - \mathbf{X}'\mathbf{V} \right\|_F + \lambda \sum_{r=1}^R \|\mathbf{v}_r\|_1$$

where \mathbf{v}_r is a column of \mathbf{V} and $\|\cdot\|_F$ denotes the Frobenius norm. The other two are more involved and estimate a matrix \mathbf{V} that is directly sparse and defines a discriminative subspace; we point to the referenced work for details.

For all the methods described above, maximization of the penalized likelihood for different specifications of $Q_\lambda(\cdot)$ is usually carried out resorting to a form of penalized EM algorithm (Green, 1990). Moreover, model selection consists in the selection of the number of mixture components and of the optimal shrinkage parameter λ , and to accomplish the task a modified BIC is employed; see the referenced works for details.

A method that differs in some way from the general form of (4) is the *sparse k-means* proposed by Witten and Tibshirani (2010). Let $\mathbb{C} = \{C_1, \dots, C_g, \dots, C_G\}$ be a partition of the observations into disjoint subsets, and let $\mathbf{w} = (w_1, \dots, w_j, \dots, w_J)$ be a vector of weights for each variable. Clustering and variable selection is performed by solving the optimization problem:

$$\underset{\mathbb{C}, \mathbf{w}, \boldsymbol{\mu}}{\operatorname{argmax}} \left\{ \sum_{j=1}^J w_j \left[\sum_{i=1}^N (x_{ij} - \bar{x}_j)^2 - \sum_{g=1}^G \sum_{i \in C_g} (x_{ij} - \mu_{gj})^2 \right] \right\}$$

subject to $\|\mathbf{w}\|^2 \leq 1, \quad \|\mathbf{w}\|_1 \leq s,$
with $w_j \geq 0 \quad \forall j,$

where \bar{x}_j is the sample mean of variable j and s is a tuning parameter. The L_2 penalty will force the weights in the interval $[0, 1]$, while the L_1 norm is used to shrink the values to 0. The weight w_j is interpreted as the contribution of X_j to the resulting clustering: a large value of w_j indicates a variable useful for clustering, while $w_j = 0$ means that variable j does not contribute to the clustering. Note that the above objective function is related to the log-likelihood of a Gaussian mixture model with a common isotropic covariance matrix of the form $\boldsymbol{\Sigma}_g = \sigma^2 \mathbb{I}$, an implicit assumption of the method. Indeed, in this exposition we only considered the case associated to Gaussian mixture modeling and we adapted equation (10) of Witten and Tibshirani (2010) to the purpose, using the connection of k -means to the fitting of mixture of Gaussians with common spherical covariance matrix (see the cited work for additional details). Furthermore, the method needs the number of clusters G to be known in advance. However, the approach is particularly suitable for high-dimensional data and the authors describe a more general framework for sparse clustering. Related work on sparse k -means is in Sun, Wang and Fang (2012), where consistency properties are investigated.




FIG 2. The two competing models proposed in Raftery and Dean (2006).

4.3. Model selection approaches

Within this class of approaches, the problem of selecting the relevant clustering variables is recast as a model selection problem. Different models are specified by the role attributed to the variables in connection to their relation with the clustering variable \mathbf{z} . Then these models are compared by means of a model selection criterion and relevant clustering variables are chosen accordingly to the best model.

The framework was pioneered by the work of Raftery and Dean (2006), where three main roles are specified for the variables. The authors propose a procedure where \mathbf{X} is partitioned into the following subsets of variables:

- \mathbf{X}^C , the set of current clustering variables;
- \mathbf{X}^P , the variable(s) proposed to be added or removed from the set of clustering variables;
- \mathbf{X}^{NC} , the set of other variables not relevant for clustering.

Then the decision for inclusion or exclusion of \mathbf{X}^P is taken by comparing the models (Fig 2):

$$\begin{aligned} \mathcal{M}_A: \quad & p(\mathbf{X} | \mathbf{z}) = p(\mathbf{X}^C, \mathbf{X}^P | \mathbf{z}) p(\mathbf{X}^{NC} | \mathbf{X}^C, \mathbf{X}^P), \\ \mathcal{M}_B: \quad & p(\mathbf{X} | \mathbf{z}) = p(\mathbf{X}^C | \mathbf{z}) p(\mathbf{X}^P | \mathbf{X}^C) p(\mathbf{X}^{NC} | \mathbf{X}^C, \mathbf{X}^P). \end{aligned}$$

In model \mathcal{M}_A , \mathbf{X}^P is useful for clustering and the joint distribution $p(\mathbf{X}^C, \mathbf{X}^P | \mathbf{z})$ corresponds to a Gaussian mixture distribution; on the other hand, \mathcal{M}_B states that \mathbf{X}^P does not depend on the clustering \mathbf{z} and the conditional distribution $p(\mathbf{X}^P | \mathbf{X}^C)$ corresponds to a linear regression. An important feature of the framework formulation is that in \mathcal{M}_B the irrelevant variables are not required to be independent of the clustering variables. This allows to discard redundant variables related to the clustering ones but not to the clustering itself. Another important characteristic of the model construction lies in the avoidance of the unrealistic assumption of the independence between clustering and irrelevant




FIG 3. The two models compared in Maugis, Celeux and Martin-Magniette (2009a).

variables. The competing models are compared using the BIC approximation to their marginal likelihoods:

$$\begin{aligned} \text{BIC}_A &= \text{BIC}_{\text{clust}}(\mathbf{X}^C, \mathbf{X}^P), \\ \text{BIC}_B &= \text{BIC}_{\text{no clust}}(\mathbf{X}^C) + \text{BIC}_{\text{reg}}(\mathbf{X}^P | \mathbf{X}^C), \end{aligned}$$

where $\text{BIC}_{\text{clust}}(\mathbf{X}^C, \mathbf{X}^P)$ is the BIC of a GMM in which \mathbf{X}^P adds useful information to the clustering, $\text{BIC}_{\text{no clust}}(\mathbf{X}^C)$ is the BIC of the GMM on the current set of clustering variables and $\text{BIC}_{\text{reg}}(\mathbf{X}^P | \mathbf{X}^C)$ is the BIC of the regression of X^P on X^C . Then if the difference $\text{BIC}_A - \text{BIC}_B$ is greater than zero, \mathbf{X}^P is added to the set of clustering variables. To perform the selection, variables are added/removed and different models are compared using a stepwise algorithm.

The method has been further extended by Maugis, Celeux and Martin-Magniette (2009a) and Maugis, Celeux and Martin-Magniette (2009b). Maugis, Celeux and Martin-Magniette (2009a) build on the modeling framework of Raftery and Dean (2006), however, in the conditional distribution $p(\mathbf{X}^P | \mathbf{X}^C)$, \mathbf{X}^P can be related only to a subset $\mathbf{X}^R \subseteq \mathbf{X}^C$ of the clustering variables. Therefore, in the regression only a subset \mathbf{X}^R of predictors are used to describe the dependency between \mathbf{X}^P and X^C . In this way, the authors avoid the inclusion of unneeded parameters that would over-penalize the log-likelihood with the consequence of favoring models that declare some irrelevant variables as relevant when model comparison is performed using the BIC. The models compared in Maugis, Celeux and Martin-Magniette (2009a) are depicted in Fig. 3, with \mathcal{M}_A the same as previously stated and \mathcal{M}_C defined as follows:

$$\mathcal{M}_C: \quad p(\mathbf{X} | \mathbf{z}) = p(\mathbf{X}^C | \mathbf{z}) p(\mathbf{X}^P | \mathbf{X}^R \subseteq \mathbf{X}^C) p(\mathbf{X}^{NC} | \mathbf{X}^C, \mathbf{X}^P),$$

with $p(\mathbf{X}^P | \mathbf{X}^R \subseteq \mathbf{X}^C)$ the regression term of \mathbf{X}^P on the relevant predictors \mathbf{X}^R . For this model the BIC is given by:

$$\text{BIC}_C = \text{BIC}_{\text{no clust}}(\mathbf{X}^C) + \text{BIC}_{\text{reg}}(\mathbf{X}^P | \mathbf{X}^R \subseteq \mathbf{X}^C),$$

with $\text{BIC}_{\text{reg}}(\mathbf{X}^P | \mathbf{X}^R \subseteq \mathbf{X}^C)$ the BIC of the regression of \mathbf{X}^P on \mathbf{X}^C after selection of the optimal set of predictors \mathbf{X}^R . Again, to decide if \mathbf{X}^P is relevant




FIG 4. The three different models presented in Maugis, Celeux and Martin-Magniette (2009b).

for clustering, the difference $\text{BIC}_A - \text{BIC}_C$ is computed and \mathbf{X}^P is added to the set of clustering ones if this quantity is greater than zero. Model search is performed using a backward search and for selecting the relevant predictors in the regression the authors propose a standard stepwise procedure.

The framework is subsequently expanded in Maugis, Celeux and Martin-Magniette (2009b), where the authors consider an additional role for the variables proposed to be added or removed. They explicitly account for the case where \mathbf{X}^P can be independent of the clustering variables \mathbf{X}^C . The assumption leads to further circumvent the over-penalization problem of the log-likelihood when parsimonious Gaussian mixture models are involved in the comparison and penalized likelihood criteria are used for model selection. The authors suggest a more flexible framework where three different models are compared, all represented in Fig. 4. Models \mathcal{M}_A and \mathcal{M}_C are as before, while model \mathcal{M}_D is specified as follows:

$$\mathcal{M}_D: \quad p(\mathbf{X} | \mathbf{z}) = p(\mathbf{X}^C | \mathbf{z}) p(\mathbf{X}^P) p(\mathbf{X}^{NC} | \mathbf{X}^C, \mathbf{X}^P),$$

with $p(\mathbf{X}^P)$ the density of a multivariate Gaussian distribution. Then, the corresponding BIC is given by:

$$\text{BIC}_D = \text{BIC}_{\text{no clust}}(\mathbf{X}^C) + \text{BIC}_{\text{indep}}(\mathbf{X}^P),$$

where $\text{BIC}_{\text{indep}}(\mathbf{X}^P)$ is the BIC of a multivariate Gaussian model. For an efficient model search, the authors suggest to rewrite BIC_D as:

$$\text{BIC}_D^* = \text{BIC}_{\text{no clust}}(\mathbf{X}^C) + \text{BIC}_{\text{reg}}(\mathbf{X}^P \mid \tilde{\mathbf{X}}^R \subseteq \mathbf{X}^C),$$

where $\tilde{\mathbf{X}}^R$ is allowed to be the empty set \emptyset and thus $\text{BIC}_{\text{reg}}(\mathbf{X}^P \mid \tilde{\mathbf{X}}^R \subseteq \mathbf{X}^C) = \text{BIC}_{\text{indep}}(\mathbf{X}^P)$ when $\tilde{\mathbf{X}}^R = \emptyset$. As in Maugis, Celeux and Martin-Magniette (2009a), the difference $\text{BIC}_A - \text{BIC}_D^*$ is computed at each step of the algorithm and \mathbf{X}^P is added to \mathbf{X}^C if this quantity is positive. In the algorithm the variables are added and removed using a stepwise searching method and relevant predictors are selected by a backward stepwise algorithm. Maugis-Rabusseau, Martin-Magniette and Pelletier (2012) extend the variable selection framework to handle data with missing values without resorting to any imputation procedure.

Note that the mentioned methods not only perform the selection of the relevant clustering variables, but at the same time they select the optimal number of clusters and the best parsimonious Gaussian clustering model from the family of models of Celeux and Govaert (1995); we refer to the cited works for details. Lastly, it is worth mentioning two further extensions of the above modeling frameworks: Scrucca (2016) proposes the use of genetic algorithms for searching over the whole model space to overcome the sub-optimality of a stepwise search, while Galimberti, Manisi and Soffritti (2017) suggest a framework where the relevant variables can provide information about more than one clustering structure of interest.

In the above approaches, at each step of the variable selection algorithm, the likelihood of a GMM has to be optimized several times. To deal with this computational issue, Marbac and Sedki (2017a) propose a procedure that relies on the ICL criterion and does not require multiple calls of the EM algorithm. In the framework, a variable is declared as irrelevant to the clustering if its one-dimensional marginal distributions are equal between the mixture components. The authors introduce a binary indicator variable $\boldsymbol{\omega} = (\omega_1, \dots, \omega_j, \dots, \omega_J)$ such that $\omega_j = 1$ if variable X_j is relevant, 0 otherwise. In this context we would denote $\mathbf{X}^C = \{X_j \in \mathbf{X} : \omega_j = 1, \forall j\}$ and $\mathbf{X}^{NC} = \{X_j \in \mathbf{X} : \omega_j = 0, \forall j\}$, and different models are specified by different configurations of the binary vector $\boldsymbol{\omega}$. These models are compared by means of a criterion based on the following integrated complete-data likelihood:

$$p(\mathbf{X}, \mathbf{z} \mid \boldsymbol{\omega}) = \int p(\mathbf{X}, \mathbf{z} \mid \boldsymbol{\Omega}, \boldsymbol{\omega}) p(\boldsymbol{\Omega} \mid \boldsymbol{\omega}) d\boldsymbol{\Omega},$$

where $\boldsymbol{\Omega}$ denotes the collection composed of the parameters related to the mixture distribution for \mathbf{X}^C and the parameters related to the distribution for \mathbf{X}^{NC} . After assuming local and global independence among the variables and placing conjugate priors on the parameters, the above integral reduces to:

$$p(\mathbf{X}, \mathbf{z} \mid \boldsymbol{\omega}) = p(\mathbf{z}) \prod_{j=1}^J p(X_j \mid \omega_j, \mathbf{z}),$$

where all the quantities have closed analytical expressions (see authors' paper for more details). For a non relevant variable, $p(X_j | \omega_j, \mathbf{z}) = p(X_j | \omega_j)$ since this quantity does not depend on the data classification, while for a relevant one this quantity will assume different values over the mixture components. Consequently, variable selection is performed by finding the vector ω that maximizes the Maximum Integrated Complete-data Likelihood (MICL) criterion, given by:

$$\text{MICL}(\omega) = \log p(\mathbf{X}, \mathbf{z}_\omega^* | \omega),$$

with $\mathbf{z}_\omega^* = \text{argmax}_{\mathbf{z}} \log p(\mathbf{X}, \mathbf{z} | \omega)$. Maximization of $\text{MICL}(\omega)$ is carried out iteratively using an algorithm that alternates between two optimizations steps of the ICL: optimization on the classification \mathbf{z} given the data and ω , and maximization on ω given the data and the classification \mathbf{z} . As in Raftery and Dean (2006), Maugis, Celeux and Martin-Magniette (2009a) and Maugis, Celeux and Martin-Magniette (2009b), the algorithm also returns the optimal number of clusters and we refer to the paper for details. The approach is fast and scales well for problems with a large number of variables. However, the optimization on \mathbf{z} can be computationally demanding for large sample sizes; the authors suggest that for sample sizes smaller than 10^4 this optimization is still doable. Furthermore, in a situation where variables are correlated and many redundant ones are present, the local and global independence assumptions could be too restrictive.

4.4. Other approaches

Other variable selection methods that do not univocally fit in the previous sections have been developed in the literature. Within a wrapper approach, Dy and Brodley (2004) propose to embed in the EM a forward selection algorithm for the maximization of two simple alternative criteria for variable selection. The first is based on a measure of separability and unimodality of the clusters, and consists in the quantity $\text{tr}(S_W^{-1} S_B)$, where S_W is the within-cluster scatter matrix and S_B is the between class scatter matrix. The second is based on the quantification of how well the model fits the data and the natural choice is the likelihood of a GMM itself.

A related method is the one of Andrews and McNicholas (2014). The authors suggest an hybrid filter-wrapper approach based on a variable within-group variance, given by:

$$W_j = \frac{\sum_{g=1}^G \sum_{i=1}^N z_{ig}(x_{ij} - \mu_j)^2}{N}.$$

First, an initial estimate of W_j is found by running a preliminary clustering step and the variables are listed in ascending order according to their values. The top variable in the list is automatically selected (as the one with the minimum W_j) and placed in the set of selected variables \mathcal{S} . Subsequently, a variable j is selected if

$$|\rho_{rj}| < (1 - W_j)^c \quad \forall X_r \in \mathcal{S},$$

with ρ_{rj} the correlation between X_j and X_r , and c a coefficient used to weight the within group variation. The authors suggest to use the integer values from

1 to 5 and the procedure will tend to include more variables in the model as c increases. Lastly, clustering is performed by fitting a GMM on the final set S . Note that before running the procedure the data need to be standardized to have mean 0 and variance 1.

Another method using the concept of separability and unimodality of the clusters as in Dy and Brodley (2004) is the one found in Lee and Li (2012). Here GMMs are used to estimate the density of the data and the variable selection is based on a cluster separability measure defined using ridgelines, one-dimensional parametric curves linking the modes of two clusters as function of their densities. Let $y_{gk}(a_t)$, be the set of points individuating the ridgeline between cluster g and k ($a_0 = 0 < a_1 < \dots < a_T = 1, t = 0, \dots, T$); the pairwise separability between these clusters is measured by:

$$1 - \frac{\min_{t=1}^T \left\{ \sum_{h=1}^G \phi(y_{gk}(a_t); \boldsymbol{\mu}_h, \boldsymbol{\Sigma}_h) \right\}}{\min \left\{ \sum_{h=1}^G \phi(y_{gk}(0); \boldsymbol{\mu}_h, \boldsymbol{\Sigma}_h), \sum_{h=1}^G \phi(y_{gk}(1); \boldsymbol{\mu}_h, \boldsymbol{\Sigma}_h) \right\}}.$$

The aim is to find the set of variables that achieves the maximum aggregated separability across all the pairs of clusters and the authors suggest two alternative forward selection algorithms for the task. The method tends to select variables indicating well-separated clusters.

A simple filter approach used to pre-select the variables and reduce the dimensionality of the data is the one suggested by McLachlan, Bean and Peel (2002), within the context of clustering microarray expression data. In this approach, first a univariate mixture model is fitted to each variable in turn. Then the likelihood ratio statistic for the test of a single normal distribution versus a mixture of Gaussians is computed and those variables for which this statistic is significant are retained. The selected variables are subsequently clustered using a k-means procedure in order to find those centroids that represent subsets of variables. Model-based clustering is consequently performed on these representative centroids, which represent the data in a lower dimensional subspace.

The stepwise algorithm presented in Maugis, Celeux and Martin-Magniette (2009b) could be very slow in high-dimensional settings. To overcome the computational burden, Celeux et al. (2018) suggest an hybrid approach where a LASSO-like procedure and the model selection algorithm are employed together. The approach consists of two subsequent steps. First, the method of Zhou, Pan and Shen (2009) is applied to the data (see Section 4.2) and a ranking of the variables is obtained. For fixed G and a set of different combinations $L_{\lambda_1} \times L_{\lambda_2}$ of values (λ_1, λ_2) , the *clustering score* of each variable X_j is defined as:

$$\mathcal{O}_G(j) = \sum_{L_{\lambda_1} \times L_{\lambda_2}} \mathcal{B}_{G, \lambda_1, \lambda_2}(j),$$

where

$$\mathcal{B}_{G, \lambda_1, \lambda_2}(j) = \begin{cases} 0 & \text{if } \hat{\mu}_{1j}(\lambda_1, \lambda_2) = \dots = \hat{\mu}_{gj}(\lambda_1, \lambda_2) = \dots = \hat{\mu}_{Gj}(\lambda_1, \lambda_2) = 0, \\ 1 & \text{if } \hat{\mu}_{gj}(\lambda_1, \lambda_2) \neq 0, g \in \{1, \dots, G\}. \end{cases}$$

The larger the value of $\mathcal{O}_G(j)$, the more a variable is likely to be relevant for the clustering. Then the variables are ranked in decreasing order according to this quantity; denote this rank as \mathcal{I}_G . Ideally, in the rank relevant variables are in the top positions, then redundant variables and non-informative ones at the end. In the second step, Maugis, Celeux and Martin-Magniette (2009b) method is applied using this rank. The variable set is scanned according to the \mathcal{I}_G order and variables are declared as relevant until for c consecutive variables there is no evidence of being clustering variables. Then irrelevant variables are determined scanning the set in the reverse order of \mathcal{I}_G . The process stops as soon as c consecutive variables are not declared as non-informative. The redundant variables are thus determined by the remaining variables. The tuning parameter c is a buffer to reduce the effect of mistakenly ranked variables in the LASSO-like step.

Another hybrid contribution is the work of Malsiner-Walli, Frühwirth-Schnatter and Grün (2016). The authors define a full Bayesian framework for estimation of sparse Gaussian mixtures where the assumption of a local shrinkage prior on the component means allows the identification of cluster-relevant variables. They suggest a Normal-Gamma prior of the form:

$$\mu_{gj} \mid \lambda_j, b_{0j} \sim \mathcal{N}(b_{0j}, \lambda_j s_0),$$

where

$$\lambda_j \sim \mathcal{G}(\nu_1, \nu_2), \quad b_{0j} \sim \mathcal{N}(m_0, M_0),$$

with ν_1, ν_2, m_0, M_0 and s_0 some hyperparameters. For values of $\nu_1 < 1$, the prior shrinks all the component means of variable j towards the common value b_{0j} , with the effect that variables uninformative for the clustering are effectively fit by a single mixture component. Note that this approach could be interpreted as an hybrid form of Bayesian and penalized variable selection procedures.

Although not directly related to the Gaussian mixture modeling framework, it is worth to mention two other approaches that can be thought of a cross-over between Bayesian and penalization methods. The first is the work from Hoff (2006), where it is proposed a Bayesian framework for finding subsets of variables that distinguish the clusters from each other. Thus, every cluster is characterized by its own set of discriminating variables that differentiates it from the rest of the clusters. In the method, each data point is modeled as:

$$x_{ij} = \bar{\mu}_j + \varphi_{gj} \mu_{gj} \sigma_{gj} + \sigma_{gj}^{\varphi_{gj}} \epsilon_{ij},$$

where $\bar{\mu}_j$ is the overall data mean of variable j , φ_{gj} is a binary variable indicating if X_j is “active” on cluster g and ϵ_{ij} is an error component such that $\epsilon_{ij} \sim \mathcal{N}(0, \bar{\sigma}_j^2)$, with $\bar{\sigma}_j^2$ the overall data variance of X_j . The model results in a multivariate Dirichlet process mixture model and the use of conjugate priors allows the implementation of a Gibbs sampling scheme for inference. Rather than performing a genuine variable selection, the method identifies the cluster-specific relevant variables by detecting the shifts in means and variances from the corresponding overall data quantities. The approach belongs to the class of *subspace clustering* methods and is related to the work of Friedman and Meulman (2004).

The second is the one proposed by Nia and Davison (2015) and is closely related. In this case, the authors consider the following linear model for the data points:

$$x_{gij} = \bar{\mu}_j + \gamma_j \varphi_{gj} \mu_{gj} + \epsilon_{gij},$$

where x_{gij} is the data of clustering observation i measured on variable j in cluster g ; γ_j and φ_{gj} are binary variables such that $\gamma_j \sim \text{Bernoulli}(\nu)$ and $\varphi_{gj} \sim \text{Bernoulli}(\eta)$ respectively. Probability ν may be interpreted as the prior proportion of relevant variables, while η as the prior proportion of non-overlapping component means for a clustering variable. Two families of prior distributions are suggested for μ_{gj} : Gaussian and asymmetric Laplace. In both cases, the marginal distribution of the observations is given in closed form and results in a *spike and slab* density in which the data is modeled by a mixture of two densities. When $(\gamma_j = 1, \varphi_{gj} = 1)$, the variable is active and the data is modeled by the slab density which drives the clustering procedure when a variable is relevant for clustering. When $\gamma_j = 0$ or $\varphi_{gj} = 0$, the density of the data corresponds to the spike density, which reduces the effect of non-informative variables. Variables not important for any of the clusters can be discarded. The setting allows the use of Bayes factors for computing the importance of each variable and perform the selection, and clustering of the data is achieved via an agglomerative hierarchical procedure.

4.5. R packages and data example

The available R packages for variable selection for Gaussian mixture models are: `sparcl` (Witten and Tibshirani, 2013), `clustvarsel` (Scrucca and Raftery, 2018), `VarSelLCM` (Marbac and Sedki, 2017b), `vsc` (Andrews and McNicholas, 2013), `SelvarMix` (Sedki, Celeux and Maugis-Rabusseau, 2017), and `bclust` (Nia and Davison, 2012). Table 1 lists them, with information about the method and the type of approach. It is also worth to mention the package `ClustOfVar` (Chavent et al., 2012). Rather than performing variable selection and obtain a classification of the data points, this package aims to find clusters of variables linked to a central synthetic variable obtained from a principal components decomposition of the data; for this reason, it will not be considered in the subsequent analysis. Lastly, we also point to the C++ softwares `SelvarClust`³ and `SelvarClustIndep`⁴, which implement the modeling framework of Raftery and Dean (2006) and Maugis, Celeux and Martin-Magniette (2009a,b). We will not consider these software in the following analysis, since the aim of this paper is in reviewing R packages.

The adjusted Rand index (ARI, Hubert and Arabie, 1985) will be employed to assess and compare the clustering performance of the packages. Let us consider two different partitions of the N data points, one into G clusters and the other into K clusters; this index is defined as:

³<http://www.math.univ-toulouse.fr/~maugis/SelvarClustHomepage.html>

⁴<http://www.math.univ-toulouse.fr/~maugis/SelvarClustIndepHomepage.html>

TABLE 1
R packages for GMMs variable selection.

Package	Type	Method
sparcl	penalized	Witten and Tibshirani (2010)
clustvarsel	model selection	Maugis, Celeux and Martin-Magniette (2009a) Maugis, Celeux and Martin-Magniette (2009b) Scrucca and Raftery (2018)
VarSelLCM	model selection	Marbac and Sedki (2017a)
vsc	filter/wrapper	Andrews and McNicholas (2014)
SelvarMix	model selection/penalized	Celeux et al. (2018) Sedki, Celeux and Maugis (2014)
bclust	other/Bayesian	Nia and Davison (2015, 2012)

$$\text{ARI} = \frac{\sum_g^G \sum_k^K \binom{N_{gk}}{2} - \left[\sum_g^G \binom{N_{g\cdot}}{2} \sum_k^K \binom{N_{\cdot k}}{2} \right] / \binom{N}{2}}{\frac{1}{2} \left[\sum_g^G \binom{N_{g\cdot}}{2} + \sum_k^K \binom{N_{\cdot k}}{2} \right] - \left[\sum_g^G \binom{N_{g\cdot}}{2} \sum_k^K \binom{N_{\cdot k}}{2} \right] / \binom{N}{2}},$$

where N_{gk} is the number of observations falling in cluster g and cluster k , while $N_{g\cdot} = \sum_k N_{gk}$ and $N_{\cdot k} = \sum_g N_{gk}$. The ARI corrects the comparison for the fact that two partitions could match merely by chance (Hubert and Arabie, 1985). The index has a maximum value of 1 and attains this value when there is perfect matching between the two partitions; on the other hand, it has expected value of 0 when the two partitions are completely independent. Therefore, the index is used to measure the agreement between two different classifications of the data.

We apply the different variable selection methods to data from the Human Mortality Database (Human Mortality Database, 2017). The database contains life tables from several countries, spanning a period of time from the middle of the 18th century to 2015. We consider central mortality rates for both genders over 10-year time intervals and 5-year age intervals. The central mortality rate ${}_n m_x$ is interpreted as the average 10-year interval rate at which people die during the period between age x and $x+n$, normalized by the number of those living. The data contain 389 age patterns of mortality, each over a 10-year period time for 24 age intervals from 0 to 110+ (Figure 5). The aim is to cluster the mortality schedules and individuate those age groups relevant for the clustering structure. There are 27 non overlapping 10-year periods, some present multiple times (especially recent ones), and the patterns of mortality tend to cluster according to the historical period (Clark and Sharro, 2011). To assess the quality of the clustering, we will compare the estimated partition with the information regarding the 10-year period each schedule belongs to using the ARI.

First we fit a Gaussian mixture model using the `mclust` package. The model selected is a 3-component mixture, with an ARI of 0.13 between the estimated classification and the historical period. Then we apply the variable selection




FIG 5. Central mortality rates (values on logarithmic scale).

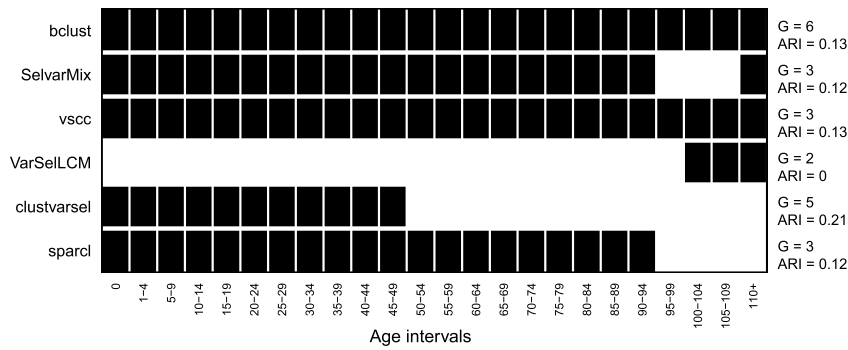


FIG 6. Variables selected via the R packages for GMMs variable selection listed in Table 1. A dark square indicates the variable has been selected. For each method, also the selected number of clusters and the ARI between the estimated classification and the 10-year historical period are reported.

methods listed in Table 1. Figure 6 displays the selected age groups for each method, alongside the chosen number of clusters and the ARI between the estimated classification and the historical period. Note that for `sparcl` the number of components needs to be set in advance and we fixed it to the number selected by `mclust`. Package `bclust` did not discard any variables, most likely due to the fact that adjacent age groups are highly correlated and the independence assumption of this method is unrealistic in this case. Also package `vsccl` did not remove any of the variables. Package `VarSelLCM` retained the smallest subset of

TABLE 2
Cross-tabulation between the historical period classification and the `clustvare` partition on the selected age groups.

	Group				
	1	2	3	4	5
1750-1759	1				
1760-1769	1				
1770-1779	1				
1780-1789	1				
1790-1799	1				
1800-1809	1				
1810-1819	2				
1820-1829	2				
1830-1839	4				
1840-1849	6				
1850-1859	7				
1860-1869	7				
1870-1879	10				
1880-1889	10				
1890-1899	8	2			
1900-1909	7	4			
1910-1919	5	6			
1920-1929	3	11			
1930-1939		15			
1940-1949	4	15	1		
1950-1959		5	21		1
1960-1969		3	25		6
1970-1979			26	2	7
1980-1989		1	8	19	10
1990-1999			2	26	12
2000-2009				34	7
2010-2015				30	9

variables. However, the package obtained an ARI of zero and surprisingly selected only the oldest age intervals, those expected to have the least information, thus it seems to have detected a spurious solution. We ran the package multiple times and the other best solutions found were a model with 3 clusters and the same subset of variables, and a model with only one cluster and no variables discarded. The package seems to be sensitive to the initialization and, also in this case, this is an indication that the strong independence assumptions at the basis of the method implemented could be too restrictive. The other packages, `sparcl`, `SelvarMix`, and `clustvare` discarded age groups representing the old population. Among these, `clustvare` performed the most parsimonious selection, retaining only 11 age groups in the range from 0 to 49 years old, and with the highest ARI. The clusters are discriminated along the young and middle ages, as it would be expected since for elder ages the differences among death rates level off. Table 2 reports the cross-tabulation between the historical periods and the `clustvare` partition on the selected age intervals. The clusters have a meaningful interpretation in terms of the time dimension. Figure 7 shows the estimated cluster means on the age groups selected by `clustvare`.




FIG 7. Average cluster mortality rate estimated by `clustvarsel` on the selected age groups. The full lines denote the `mclust` estimated cluster means on the whole data. Values are on logarithmic scale.

The `mclust` estimated cluster means are added for comparison. In particular, the mortality patterns of clusters 1, 3 and 4 estimated on the selected variables are close to the average patterns estimated on all the data. The extra cluster 2 seems to capture the mortality pattern related to the period corresponding to the two World wars. Moreover, cluster 5 has a mean rate similar to cluster 3, but higher death rate for age intervals above 20 years old and captures the extra-heterogeneity present in the more recent years.

Table 3 reports the computing time of the packages considered in Table 1 on a Dell machine with Intel Core i7-3770 CPU @3.40GHz×8. The packages are listed from the fastest to the slowest. `vsc` is the fastest among the packages considered, but it did not discard any of the age intervals. Package `sparcl` is the second fastest. In the case of highly correlated variables, `clustvarsel` performs the most parsimonious selection and obtains the best classification, with an acceptable computing time. Note, however, that in using `clustvarsel` we considered only 4 out of the 14 covariance models available, as some of them may require a long computational time to be fitted and to consider a set of models consistent with `SelvarMix`. Therefore, aside from `vsc`, `sparcl` is the fastest and as such it is particularly suitable to perform variable selection in high-dimensional settings where it is reasonable to assume local independence among the variables. Nevertheless, the package requires the number of clusters to be specified in advance, while `clustvarsel` and `SelvarMix` automatically select such number. `bclust` and `VarSelLCM` are the slowest in this example, and care may need to be taken when using these packages on highly correlated data, since

TABLE 3
Computing time (in seconds) and relative computing time of the R packages for variable selection for GMMs.

Package	Time (sec.)	Relative
<code>vsc</code>	10	1.00
<code>sparcl</code>	42	4.09
<code>clustvarsel</code>	51	5.01
<code>SelvarMix</code>	61	5.90
<code>VarSelLCM</code>	438	42.58
<code>bclust</code>	916	89.14

they make use of independence assumptions that could be too restrictive. Lastly, we point that `clustvarsel`, `SelvarMix` and `VarSelLCM` allow for parallelization of the computations.

5. Variable selection methods for latent class analysis

After the description of the main variable selection methods for GMMs, in this section we present the different approaches for variable selection in latent class analysis.

5.1. Bayesian approaches

Only recently, attention has been posed on the use of the Bayesian framework for variable selection in latent class analysis. The overall setting is similar to the one provided in Section 4.1 and we refer there for its description. Moreover, generally the methods borrow ideas from Bayesian variable selection for Gaussian mixture models.

Indeed, Silvestre, Cardoso and Figueiredo (2015) propose an adaptation of Law, Figueiredo and Jain (2004) method to categorical data clustering and variable selection. The mixture of Gaussian densities is replaced by the mixture in (3) and the authors use the same concept of feature saliency. Analogously, a Dirichlet distribution is assumed on the saliences and an EM algorithm for MAP estimation is employed.

White, Wyse and Murphy (2016) present an approach inspired by the work of Tadesse, Sha and Vannucci (2005). The authors propose the use of a collapsed Gibbs sampler algorithm in which all the mixture parameters are integrated out. A prior of the form $p(\varphi | \eta) = \prod_j \eta^{\varphi_j} (1 - \eta)^{1 - \varphi_j}$ is assumed for the indicator variable φ . Then, assuming conjugate priors on all the parameters, it is possible to marginalize them out analytically, obtaining the posterior:

$$p(G, \mathbf{z}, \varphi | \mathbf{X}) \propto p(G) \int p(\mathbf{z}, \varphi, \boldsymbol{\Omega} | \mathbf{X}, G) d\boldsymbol{\Omega},$$

with $\boldsymbol{\Omega}$ denoting the collection of parameters of the latent class analysis model and the parameters of the distribution for the irrelevant variables. The above

quantity has a closed form expression and allows to implement a simple MCMC scheme. Subsequently, inference is performed using a post-hoc procedure. Variable selection and selection of G is carried out jointly, by calculating the proportion of time the sampler spent in a certain number of groups and including a certain variable. The authors employ the method on different datasets and it is shown to outperform a more involved reversible jump MCMC algorithm for variable selection.

5.2. Penalization approaches

Despite of the large amount of work developed for penalized model-based clustering with GMMs (confront Section 4.2), little work has been done in the context of LCA. Additionally, more attention has been given to the aim of model identifiability and regularization, rather than to the purpose of variable selection per se.

Within this framework enters the work of Houseman, Coull and Betensky (2006). The authors propose a penalization approach for latent class regression analysis of high-dimensional genomic data where a set of categorical response variables is related to a set of numerical covariates. Here the class-conditional probabilities π_g in (3) are expressed as function of a vector of covariates and a set of regression coefficients. Then a ridge or LASSO penalty is used in order to obtain sparse estimates for the vectors of regression coefficients. Consequently, by discarding those predictors with coefficients shrunken to zero, the clusters would be characterized by different sets of selected covariates used to model the categorical responses. However, variable selection is not performed on the variables directly involved in the clustering.

Wu (2013) presents a related approach, but this time focusing on the selection of variables involved in the clustering. In the method, the class conditional probabilities π_{gjc} of observing category c for variable j within class g are reparameterized in terms of the logit transform as:

$$\pi_{gjc} = \frac{\exp(\alpha_{jc} + \beta_{gjc})}{\sum_{l=1}^{C_j} \exp(\alpha_{jl} + \beta_{gjl})}, \quad \text{with} \quad \alpha_{j1} = \beta_{gj1} = 1.$$

Then, similarly to (4) in Section 4.2, the following penalized log-likelihood is considered:

$$\ell_Q = \sum_{i=1}^N \log \left\{ \sum_{g=1}^G \tau_g \mathcal{C}(\mathbf{x}_i; \boldsymbol{\alpha}, \boldsymbol{\beta}_g) \right\} - \lambda \sum_{g=1}^G \sum_{j=1}^J \sum_{c=1}^{C_j} |\beta_{gjc}|,$$

where $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}_g$ are the collections of parameters α_{jc} and β_{gjc} respectively, and we made explicit the dependence on them of the Multinomial probability mass function. For a variable X_j , the quantities β_{gjc} measure the difference from the overall probabilities $\exp(\alpha_{jc}) / \left[\sum_{l=1}^{C_j} \exp(\alpha_{jl}) \right]$ and its contribution to the clustering. Consequently, if $\sum_{g=1}^G \sum_{c=1}^{C_j} |\beta_{gjc}| = 0$, variable X_j is considered




FIG 8. The two models presented in Dean and Raftery (2010).

not influent to the classification and discarded. As for GMMs, the estimation in this setting is carried out by means of a penalized EM algorithm and the optimal λ and the number of clusters G are selected by BIC.

5.3. Model selection approaches

Model selection approaches for variable selection in LCA draw from the work developed for GMMs. In Dean and Raftery (2010) the authors suggest a framework similar to the one presented in Raftery and Dean (2006). As in section 4.3, the following partition of \mathbf{X} is considered:

- \mathbf{X}^C , the set of current clustering variables;
- X^P , the variable proposed to be added or removed from the set of clustering variables;
- \mathbf{X}^{NC} , the set of variables not relevant for clustering.

Then the authors propose to compare the two models represented in Fig. 8. Model \mathcal{M}_A states that X^P is a clustering variable and there is no edge between \mathbf{X}^C and X^P due to the local independence assumption of LCA that implies $p(\mathbf{X}^C, X^P | \mathbf{z}) = p(\mathbf{X}^C | \mathbf{z})p(X^P | \mathbf{z})$. In model \mathcal{M}_B , X^P is not useful for clustering (the missing edge between \mathbf{z} and X^P) and is assumed to be independent from the set \mathbf{X}^C . Under these models, the joint distribution of \mathbf{X} is expressed as:

$$\begin{aligned} \mathcal{M}_A: \quad p(\mathbf{X} | \mathbf{z}) &= p(\mathbf{X}^C, X^P | \mathbf{z}) p(\mathbf{X}^{NC} | \mathbf{X}^C, X^P), \\ \mathcal{M}_B: \quad p(\mathbf{X} | \mathbf{z}) &= p(\mathbf{X}^C | \mathbf{z}) p(X^P) p(\mathbf{X}^{NC} | \mathbf{X}^C, X^P), \end{aligned}$$

with $p(\mathbf{X}^C, X^P | \mathbf{z})$ the LCA model on the clustering variables and X^P , $p(X^P)$ the Multinomial distribution and $p(\mathbf{X}^{NC} | \mathbf{X}^C, X^P)$ the distribution for the non-informative variables. The two models are compared via the BIC approximations:

$$\text{BIC}_A = \text{BIC}_{\text{clust}}(\mathbf{X}^C, X^P),$$




FIG 9. The models proposed in Fop, Smart and Murphy (2017).

$$\text{BIC}_B = \text{BIC}_{\text{no clust}}(\mathbf{X}^C) + \text{BIC}(X^P),$$

where $\text{BIC}_{\text{clust}}(\mathbf{X}^C, X^P)$ is the BIC of a LCA model in which X^P is a relevant clustering variable, $\text{BIC}_{\text{no clust}}(\mathbf{X}^C)$ is the BIC of the LCA model on the current clustering variables and $\text{BIC}(X^P)$ is the BIC for a Multinomial distribution fit. Then X^P is considered useful for clustering if the difference $\text{BIC}_A - \text{BIC}_B$ is greater than zero. To search through the model space, the authors propose a forward headlong search algorithm (Badsberg, 1992) consisting of inclusion and removal steps. The algorithm has the advantage of being more computationally efficient than a backward search, but the disadvantage of being sensitive to the initialization of the set of clustering variables.

To deal with the problems of multimodality of the LCA log-likelihood and the sensitivity to the initialization, Bartolucci, Montanari and Pandolfi (2016) propose to add an extra step in the variable selection procedure of Dean and Raftery (2010). They consider a *random check* step aimed at initializing the estimation algorithm with a large number of random starting values, so as to prevent the problem of incurring in local optima. In an application to nursing home evaluation, the authors also perform a sensitivity analysis of the selected variables with respect to the initialization of the set of the clustering variables. The extra step is shown to be beneficial in increasing the chances of finding the global optimum.

A major drawback of Dean and Raftery (2010) framework is the independence assumption between the proposed variable and the set of clustering ones in model \mathcal{M}_B . In fact, because of this assumption, the model does not take into account that the proposed variable could be redundant for clustering given the set of already selected relevant variables. This results in the method not being capable of discarding variables related to the clustering ones, but not directly related to the group structure itself. To overcome the problem, Fop, Smart and Murphy (2017) propose the modeling framework depicted in Fig. 9. Following Maugis, Celeux and Martin-Magniette (2009a,b) (consult Section 4.3), in model \mathcal{M}_C the proposed variable is not relevant for clustering and it is assumed to be

dependent on the clustering variables \mathbf{X}^C . Model \mathcal{M}_A is as before, while under model \mathcal{M}_C the density of \mathbf{X} is given by:

$$\mathcal{M}_C: \quad p(\mathbf{X} | \mathbf{z}) = p(\mathbf{X}^C | \mathbf{z}) p(X^P | \mathbf{X}^R \subseteq \mathbf{X}^C) p(X^{NC} | \mathbf{X}^C, \mathbf{X}^P),$$

where the conditional distribution $p(X^P | \mathbf{X}^R \subseteq \mathbf{X}^C)$ is modeled using a multinomial logistic regression. Only a subset $\mathbf{X}^R \subseteq \mathbf{X}^C$ of relevant predictors enter in the regression, avoiding the inclusion of extra parameters without necessarily increasing the likelihood. Moreover the subset \mathbf{X}^R is allowed to be the empty set, and in the case of $\mathbf{X}^R = \emptyset$ the Dean and Raftery (2010) method is automatically recovered. For \mathcal{M}_C the BIC is given by:

$$\text{BIC}_C = \text{BIC}_{\text{no clust}}(\mathbf{X}^C) + \text{BIC}_{\text{reg}}(X^P | \mathbf{X}^R \subseteq \mathbf{X}^C),$$

with $\text{BIC}_{\text{reg}}(X^P | \mathbf{X}^R \subseteq \mathbf{X}^C)$ the BIC of the multinomial logistic regression after the selection of the set \mathbf{X}^R , accomplished using a simple backward stepwise search, similarly to Maugis, Celeux and Martin-Magniette (2009a,b). Models are compared by means of the difference $\text{BIC}_A - \text{BIC}_C$ and X^P is added to \mathbf{X}^C if this quantity is greater than zero. To perform the model search, the authors suggest a backward *swap-stepwise* selection algorithm where standard removal and inclusion steps are alternated to a *swap step*. In this step two different configurations of model \mathcal{M}_C are compared and they differ in the fact that one clustering variable is swapped with one of the non-clustering variables. The extra step allows the algorithm to not incur in local optima in the case of highly correlated variables. An application to the clustering of patients suffering low back pain shows that the method is capable of performing a parsimonious variable selection with a good classification performance and interpretable results. As in the case of GMMs, both Dean and Raftery (2010) and Fop, Smart and Murphy (2017) methods return the best clustering variables and the optimal number of latent classes; see the cited works for details.

Toussile and Gassiat (2009) present an approach where the modeling problem of simultaneously selecting the number of latent classes and the relevant variables is considered in a density estimation framework. Let \mathcal{S} be the set of clustering variables and let $\mathcal{M}_{(G,\mathcal{S})}$ denote a model of probability distributions such that:

$$\mathcal{M}_{(G,\mathcal{S})}: \quad p(\mathbf{X}) = p(\mathbf{X}^C) p(\mathbf{X}^{NC}),$$

with $p(\mathbf{X}^C)$ denoting a LCA model on the relevant variables and $p(\mathbf{X}^{NC})$ the Multinomial distribution model for the non-clustering variables. The aim is to select the couple (G, \mathcal{S}) , which defines the model that generated the data. The selection procedure is implemented using an algorithm that first finds the optimal subset \mathcal{S} for a range of possible values of G and then selects the optimal G for a fixed set of clustering variables. The authors employ penalized log-likelihood criteria for model selection and consistency results for BIC type criteria are derived. A further generalization and discussion is found in Bontemps and Toussile (2013), where they derive a criterion that minimizes a risk function based on the Kullback-Leibler divergence of the estimated density with respect to the

true density. The general form of this penalized log-likelihood model selection criterion is given by:

$$\ell_{(G,S)}^* - \kappa \frac{\mathcal{D}_{(G,S)}}{N},$$

with $\ell_{(G,S)}^*$ the maximized log-likelihood under model $\mathcal{M}(G, \mathcal{S})$, $\mathcal{D}_{(G,S)}$ indicating the dimension of the model and κ is a parameter depending on the sample size and the collection of the models. The authors suggest the use of *slope heuristics* (see Baudry, Maugis and Michel, 2012, for an overview) to calibrate the value of κ .

Marbac and Sedki (2017c) extend the approach of Marbac and Sedki (2017a) (see Section 4.3) to clustering and variable selection of data of mixed type, with data on only categorical variables as a particular case. The method is based again on the MICL criterion and does not require multiple calls of the EM algorithm. In addition, the approach can manage situations where data have missing values. However, global and local independence need to be assumed to compute the MICL, and optimization is carried out over the variable indicator ω and the cluster membership indicator \mathbf{z} .

5.4. Other approaches

As for GMMs, other methods have been suggested for variable selection in LCA. Zhang and Ip (2014) present a filter method where the absolute and relative importance of a variable towards the clustering structure is assessed by means of two measures. One is the *expected posterior gradient*, which is a measure of the relative change between the entropy of the prior distribution of the latent classes and the entropy of the posterior distribution of the latent classes after observing the data. This quantity is bounded between 0 and 2 and higher values indicate higher discriminative power of a variable. The other measure is the *Kolmogorov variation of posterior distribution*, which is based on the Kolmogorov distance between the class-conditional distributions. The measure is linked to the classification accuracy and the total variation of the posterior distribution, and, if there is a substantial reduction in this quantity without variable X_j in the model, then X_j can be considered a relevant clustering variable. Both measures are presented in the context of LCA for mixed type data, and the observation of only categorical data is a particular case.

Another filter method where the selection is carried out by evaluating the impact of a variable on the clustering solution is proposed in Bartolucci, Montanari and Pandolfi (2017). Here, the initial set of variables is assumed to provide an optimal clustering of the data and the objective is to select a minimum subset of variables that is sufficient to reproduce the inferred classification. To achieve this, the authors suggest a selection algorithm that removes the variables whose presence does not significantly impact the clustering. Suppose a LCA model has been fitted to the data and the estimated posterior probabilities \hat{u}_{ig} (consult Section 3) have been obtained. In the procedure, for each variable j , the proportion \mathcal{F}_{-j} of observations whose class assignment is modified when X_j is

removed with respect to the initial clustering is computed. By (3), the updated estimate $\hat{u}_{ig}^{(-j)}$ after removing X_j is simply given by:

$$\hat{u}_{ig}^{(-j)} = \frac{\hat{\tau}_g \prod_{h=1, h \neq j}^J \prod_{c=1}^{C_h} \hat{\pi}_{ghc}^{\mathbb{1}\{x_{ih}=c\}}}{\sum_{g=1}^G \hat{\tau}_g \prod_{l=1, l \neq j}^J \prod_{c=1}^{C_l} \hat{\pi}_{glc}^{\mathbb{1}\{x_{il}=c\}}}.$$

Therefore:

$$\mathcal{F}_{-j} = \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{\text{MAP}(\hat{\mathbf{u}}_i^{(-j)}) \neq \text{MAP}(\hat{\mathbf{u}}_i)\}$$

Then the variable with the minimum \mathcal{F}_{-j} is removed. If some variables have the same value of \mathcal{F}_{-j} , the variable to be removed is the one with the minimum Kullback-Leibler distance:

$$\sum_{i=1}^N \sum_{g=1}^G \hat{u}_{ig} \log \frac{\hat{u}_{ig}}{\hat{u}_{ig}^{(-j)}}.$$

The method is developed in connection to item selection for questionnaires and is fast and simple to implement. However, it requires the somewhat strong assumption that the inferred classification and the number of latent classes does not change when removing the variables.

It is worth to mention an approach that lies outside the modeling framework of LCA and is closely related to the work of Hoff (2006). Hoff (2005) presents a method for subset clustering of binary sequences where each latent class is characterized by its own set of discriminating variables that differentiates it from the rest of the classes. Let φ_{gj} be a binary variable indicating the relevance of variable j to cluster g . The author suggests to parameterize the class conditional probability of occurrence of X_j as:

$$\pi_{gj} = \varphi_{gj} \tilde{\pi}_{gj} + (1 - \varphi_{gj}) \pi_j,$$

where $\tilde{\pi}_{gj}$ is the probability that $X_j = 1$ within class g and π_j is the probability of observing variable X_j in the data. Thus, if φ_{gj} is active, the corresponding variable is relevant and the π_{gj} differs from the overall data value. The framework involves a Pólya urn scheme for the parameters and the indicator variables and results in a Dirichlet process mixture model. Inference is performed using MCMC and allows to recover the subset of informative variables of each latent class.

5.5. R packages and data example

The R packages for variable selection for latent class analysis are `ClustMMDD` (Toussile, 2016), `LCAvareSel` (Fop and Murphy, 2017) and `VarSelLCM` (Marbac and Sedki, 2017b). In particular, we note that `VarSelLCM` implements a more general framework for clustering and variable selection of data of mixed type

TABLE 4
R packages for LCA variable selection.

Package	Type	Method
ClustMMDD	model selection	Toussile and Gassiat (2009) Bontemps and Toussile (2013)
LCAvare1	model selection	Dean and Raftery (2010) Fop, Smart and Murphy (2017)
VarSelLCM	model selection	Marbac and Sedki (2017a) Marbac and Sedki (2017c)

TABLE 5
Variables in the `uscongress` dataset.

1. Handicapped infants	9. MX missile
2. Water project cost-sharing	10. Immigration
3. Adoption of the budget resolution	11. Synfuels corporation cutback
4. Physician fee freeze	12. Education spending
5. El Salvador aid	13. Superfund right to sue
6. Religious groups in schools	14. Crime
7. Anti-satellite test ban	15. Duty-free exports
8. Aid to Nicaraguan contras	16. Export administration act/South Africa

(Marbac and Sedki, 2017c). Table 4 lists the packages, with information regarding the type and the method; all implement a model selection approach.

In this section we consider the `uscongress` data available in the UCI repository at the web page <https://archive.ics.uci.edu/ml/datasets/congressional+voting+records>. The data contain votes of $N = 435$ members of the U.S. House of Representatives, 267 Democrats and 168 Republicans. Each record expresses the position of a member on 16 key votes regarding major issues selected by the Congressional Quarterly Almanac (1984); the votes are presented in Table 5. A vote can take three possible outcomes: `y` (yea), if in favor of a specific matter, `n` (nay), if against, or `u` (unknown), if the position is unknown. A graphical representation of the data is in Figure 10. The observations are classified into two groups corresponding to the main parties, but it is reasonable to expect the presence of more than two classes in the data because of internal subdivisions. Indeed, at the time the Democratic party was split into Northern and Southern Democrats, and a large portion of the second often voted in agreement with the Republican party (Congressional Quarterly Almanac, 1984).

First, we fit a LCA model on all the variables, determining the number of components using BIC. The selected model is a 4-class model and Table 6 contains a cross-tabulation of the estimated classification and the party membership. Class 1 and Class 2 are fairly polarized into Democrats and Republicans. Class 3 is predominantly characterized by Democrats, while Class 4 contains an equal amount of representatives from both sides.

Then we apply the variable selection procedures `ClustMMDD-bic` (Toussile and Gassiat, 2009), `ClustMMDD` (Bontemps and Toussile, 2013), `LCAvare1-ind`




FIG 10. Key votes in the `uscongress` data. For each vote, a paired stacked-barplot represents the observed relative frequency of a response (n), (y) or (u) within the Democrats (D) and the Republicans (R) parties.

TABLE 6
Cross-tabulation between the classification estimated by the LCA model on all the variables and the political affiliation.

	Class			
	1	2	3	4
Dem.	170	23	65	9
Rep.	1	132	25	10

TABLE 7
Cross-tabulation between the classification inferred on the variables selected by `LCAvarsel` procedure and the political affiliation.

	Class			
	1	2	3	4
Dem.	159	29	74	5
Rep.	3	129	30	6

(Dean and Raftery, 2010), `LCAvarsel` (Fop, Smart and Murphy, 2017) and `VarSelLCM` (Marbac and Sedki, 2017a,c). The variables selected by each method are displayed in Figure 11. Also the chosen number of latent classes and the ARI between the estimated classification and the party affiliation are reported. `VarSelLCM` does not discard any variable in this example. The method underlying the package makes use of both local and global independence assumptions, and they are too restrictive for this type of data. `ClustMMDD` and `LCAvarsel-ind` select the same model and discard 2 variables. `ClustMMDD` discards 3 variables and selects a 3-component model. Note that the common discarded variables 2 and 10 are votes with close call results, and thus likely uninformative with respect to the two parties or any clustering structure. `LCAvarsel` performs a more parsimonious selection, choosing a model with 4 latent classes and retaining 10 variables, but attaining a lower ARI. Figure 12 displays the class-conditional probabilities of the outcomes for the votes selected by `ClustMMDD-bic`, `ClustMMDD`, `LCAvarsel-ind` and `LCAvarsel`; we did not include `VarSelLCM`, as it discarded none of the variables. We focus the attention on the `LCAvarsel` result, as it is the package performing the most parsimonious selection. Table 7 reports a tabulation of the estimated classification and the party affiliation, with an interpretation similar to Table 6. Figure 12 (d) suggests that Class 1 and 2 are denoted by very polarized outcomes and opposite voting positions. Class 3 includes members who likely expressed a vote not loyal to their party line. Class

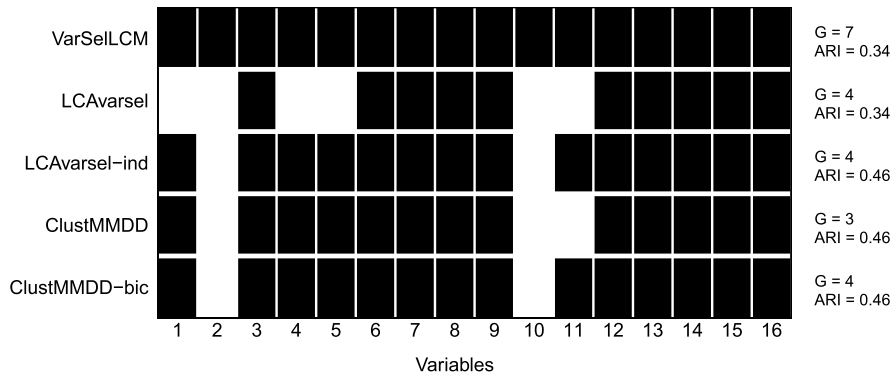


FIG 11. Variables selected via the `ClustMMDD-bic`, `ClustMMDD`, `LCAvarsel-ind`, `LCAvarsel` and `VarSelLCM` methods. A dark square indicates the variable has been selected. For each method, also the selected number of latent classes and the ARI between the estimated classification and the party affiliation are reported.

TABLE 8

Computing time (in seconds) and relative computing time of the R packages for variable selection for LCA.

Package	Time (sec.)	Relative
<code>LCAvarsel-ind</code>	114.94	1.00
<code>VarSelLCM</code>	211.20	1.84
<code>LCAvarsel</code>	591.62	5.15
<code>ClustMMDD</code>	11585.65	100.80
<code>ClustMMDD-bic</code>	11585.65	100.80

4 is characterized by an higher probability of an unknown position regarding the selected key votes.

Table 8 reports the computing time of the packages listed in Table 4 on a Dell machine with Intel Core i7-3770 CPU @3.40GHz×8. The methods are listed from the fastest to the slowest. `LCAvarsel-ind` is the fastest, followed by `VarSelLCM`, although the last did not discard any of the variables. `LCAvarsel` is slower, but still with an acceptable computing time compared to the two. Methods `ClustMMDD` and `ClustMMDD-bic` are the slowest in this example and have the same computing time, since they implement the same selection procedure. The algorithm at the basis of package `ClustMMDD` performs a more extensive search than the greedy ones implemented in `LCAvarsel` and `VarSelLCM`, however at a larger computational cost. Lastly, it is worth to notice that packages `LCAvarsel` and `VarSelLCM` can implement parallel computations.

6. Discussion

Stimulated by the wide popularity of the approach and the diffusion of high-dimensional data, the topic of variable selection for model-based clustering has seen increasing attention and rapid development. In this paper we gave an

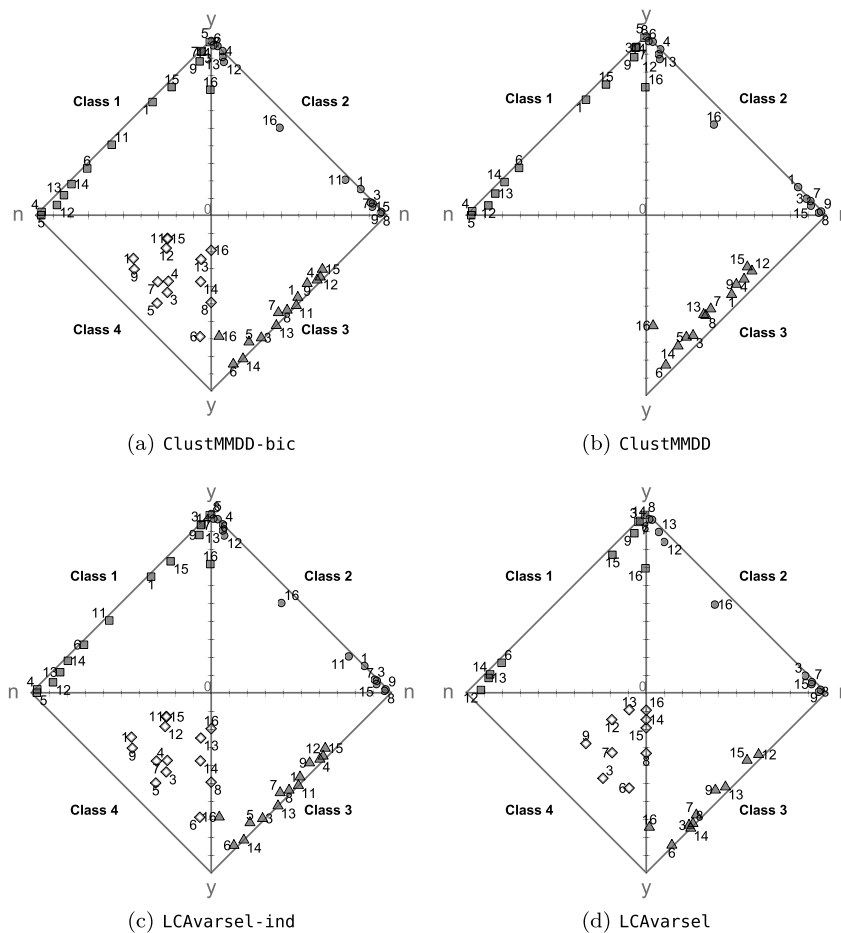


FIG 12. *Class-conditional probabilities of the outcomes for the votes selected by ClustMMDD-bic, ClustMMDD, LCAvarsel-ind and LCAvarsel. Each quadrant contains the 2-dimensional simplex representing, for every selected variable, the class-conditional probabilities of outcome n , y and u by difference. All the quadrants have values ranging from 0 to 1. Points close to the axis vertices denote high probability of outcome n or y , while points close to the center indicate high probability of outcome u .*

overview of the available variable selection methods, starting from early works to the most recent state of the art. We suggested a general and systematic picture of the features characterizing the different methods. The exposition focused on variable selection methods for Gaussian mixture models and latent class analysis, the most common model-based clustering approaches. Illustrative data examples have been used to show some of the methods in action and we provided references to the available R packages implementing them. The datasets and the R code used for the analyses are available at the web page <https://michaelfop.github.io/>. We conclude with some final comments.

Distributional assumptions

Assumptions concerning the dependence among relevant and irrelevant variables play a crucial role in the variable selection procedure. The global independence assumption simplifies the model for the joint distribution of irrelevant and relevant variables. Directly or indirectly, Bayesian and penalization approaches make use of this assumption, resulting in frameworks with a simpler association structure and that allow to discard uninformative variables. However, in the case where the variables are highly correlated, the global independence assumption can prevent to accomplish a parsimonious selection, in which also redundant variables are discarded (Law, Figueiredo and Jain, 2004; Tadesse, Sha and Vannucci, 2005; Raftery and Dean, 2006). On the other hand, model selection methods allow to depict a flexible and realistic structure for the relations among relevant, redundant and uninformative variables, but at the cost of a more complex model for the joint distribution of relevant and irrelevant variables. We remark that, when using a variable selection method, the trade-off between model complexity and selection performance needs to be taken into account.

The local independence assumption is commonly used in latent variable models, especially in the case of multivariate categorical data clustering. This assumption notably simplifies the model for the joint distribution of the relevant clustering variables (Bartholomew, Knott and Moustaki, 2011). However, it can hardly hold in practice and several approaches in the literature have been proposed to overcome it. For example, Gollini and Murphy (2014) propose a setting with continuous latent variables that allow to model the dependences among the observed categorical variables; Marbac, Biernacki and Vandewalle (2015) develop a framework where the categorical variables are grouped into blocks, each one following a specific distribution that takes into account the dependency between variables. Extending these frameworks to allow for variable selection would be interest of future research.

In this review, we mainly focused on Gaussian and Multinomial mixture models for clustering of continuous and categorical data. Variable selection for model-based clustering with other distributions and/or of data of different nature is still an open research area. The topic has only recently started to attract attention: Wallace et al. (2017) propose a framework for clustering and variable selection with the multivariate skew-Normal distribution, while Marbac and Sedki (2017c) define a method for data of mixed types. There is certainly a wide scope for further investigations and developments in this direction.

Computational aspects

Solving a variable selection problem is a task that requires a noticeable computational cost. For fixed number of mixture components, in general there are 2^J possible combinations of variables that could be considered as relevant clustering variables (Miller, 2002). The problem is worsened when the variable selection is concomitant to the problem of mixture components selection and model estimation. As already noted, filter approaches keep the two tasks separated, since the

variable selection procedure is executed before or after estimation of the clustering model. Wrapper approaches combine model estimation, variable selection and number of components determination; hence, usually multiple models need to be estimated for different combinations of clustering variables and number of components. For these reasons, the firsts are usually faster, while the seconds are computationally expensive, but often give superior results (Guyon and Elisseeff, 2003). To speed up the computations, a filter method could be employed as a preliminary step to reduce the number of variables, then the wrapper method could be used on the reduced set of variables.

Regarding the statistical approach used to perform the variable selection, the various methods present different computational characteristics and issues. Bayesian methods provide a solid ground for uncertainty evaluation of the variable selection process. However, the MCMC schemes employed within this class of methods often require runs with a large number of iterations to explore the enormous space and ensure convergence (Tadesse, Sha and Vannucci, 2005, indeed, all the works examined in this review considered a number of iterations in the order of $10^5 - 10^6$). Unfortunately, the MCMC algorithms used do not allow parallelization of the computations in order to speed up the process. Model selection approaches usually implement greedy stepwise algorithms that have a general computational complexity of $\mathcal{O}(KJ^2)$, hence they become rapidly impractical as the number of variables and mixture components increases. If the number of variable is large, and is expected that only a small number of variables $J_0 \ll J$ are relevant, forward-type algorithms can be applied, thus reducing the complexity to $\mathcal{O}(KJJ_0)$. However, in this case the difficult issue of how to initialize the algorithm and the set of clustering variables arises. Backward-type procedures do not require initialization of the clustering set, but can be very computationally demanding if the actual set of clustering variables is small. Headlong strategies could be employed to mitigate the problem (Badsberg, 1992). Nonetheless, the advantage of these stepwise methods is that they can be implemented in parallel, saving computational time. Penalization approaches are usually faster than Bayesian and model selection methods. In fact, the computational complexity of these approaches mainly depends on the form of the penalty function adopted and, in most cases, analytical solution to the optimization problem or efficient numerical procedures are already available. For this reason, they tend to scale particularly well to high-dimensional settings and have been proven to perform well also in the case of data with thousands of variables. Nevertheless, these methods do not allow for parallel computations.

Given the constant increase of the data dimensionality, the development of efficient algorithms allowing for scalable model-based clustering and variable selection is a relevant research topic that is likely to attract considerable attention in the future.

Missing data

Missing observations are a common issue in many data analysis problems. Among the works examined in this review, only few present methods for model-

based clustering and variable selection in the presence of missing data. Maugis-Rabusseau, Martin-Magniette and Pelletier (2012) extend the model selection method of Maugis, Celeux and Martin-Magniette (2009b) with the aim of avoiding a preliminary imputation procedure. The authors consider the case of missing values generated under the *missing at random* mechanism (Little and Rubin, 2002, MAR). Under this assumption, the missing responses are ignorable for likelihood-based inference. Hence, the framework and the model selection criterion can be stated in terms of only the observed data and the imputation process is avoided. Bartolucci, Montanari and Pandolfi (2016) use similar arguments for clustering categorical data via the latent class analysis model. The same MAR assumption is also considered in Marbac and Sedki (2017c), where the MICL criterion used for variable selection is computed considering only the observed entries. In application to data related to the quality-of-life of elderly patients hosted in nursing homes, Bartolucci, Montanari and Pandolfi (2017) move away from the MAR assumption and suggest to add an extra category corresponding to the missing responses, circumventing the assumption of ignorability of the missing data.

In general, in the literature multiple works exist for mixture model estimation with missing data, see for example: McLachlan and Peel (2000); Little and Rubin (2002); Hunt and Jorgensen (2003); Formann (2007); Chen, Prentice and Wang (2014). These approaches could be incorporated in a general framework for model-based clustering and variable selection for data with missing entries and may be interest of future developments.

Software

The last remark regards software availability. Despite the quantity of theory and methods developed, not many R packages for variable selection in model-based clustering are available. In particular, there is lack of packages implementing the various penalization methods for GMMs, especially useful for clustering high-dimensional data. This is somewhat surprising given the practical importance of the topic.

Acknowledgments

We would like to thank the Editor, Associate Editor and Referees whose suggestions and comments helped to improve the quality of the work.

References

- AGRESTI, A. (2002). *Categorical Data Analysis*. Wiley. [MR1914507](#)
- ANDREWS, J. L. and MCNICHOLAS, P. D. (2013). *vsc*: Variable selection for clustering and classification R package version 0.2, <https://cran.r-project.org/package=vsc>.

- ANDREWS, J. L. and MCNICHOLAS, P. D. (2014). Variable selection for clustering and classification. *Journal of Classification* **31** 136–153. [MR3232592](#)
- BADSBERG, J. H. (1992). Model search in contingency tables by CoCo. In *Computational Statistics* (Y. Dodge and J. Whittaker, eds.) Vol. 1, 251–256. Heidelberg: Physica Verlag.
- BANFIELD, J. D. and RAFTERY, A. E. (1993). Model-based Gaussian and non-Gaussian clustering. *Biometrics* **49** 803–821. [MR1243494](#)
- BARTHOLOMEW, D., KNOTT, M. and MOUSTAKI, I. (2011). *Latent Variable Models and Factor Analysis*. Wiley. [MR2849614](#)
- BARTOLUCCI, F., MONTANARI, G. E. and PANDOLFI, S. (2016). Item selection by latent class-based methods: an application to nursing home evaluation. *Advances in Data Analysis and Classification* **10** 245–262. [MR3505059](#)
- BARTOLUCCI, F., MONTANARI, G. E. and PANDOLFI, S. (2017). Latent ignorability and item selection for nursing home case-mix evaluation. *Journal of Classification*.
- BAUDRY, J.-P., MAUGIS, C. and MICHEL, B. (2012). Slope heuristics: overview and implementation. *Statistics and Computing* **22** 455–470. [MR2865029](#)
- BELLMAN, R. (1957). *Dynamic Programming*. Princeton University Press. [MR0090477](#)
- BENAGLIA, T., CHAUVEAU, D., HUNTER, D. R. and YOUNG, D. (2009). mixtools: An R package for analyzing finite mixture models. *Journal of Statistical Software* **32** 1–29.
- BHATTACHARYA, S. and MCNICHOLAS, P. D. (2014). A LASSO-penalized BIC for mixture model selection. *Advances in Data Analysis and Classification* **8** 45–61. [MR3168679](#)
- BIERNACKI, C., CELEUX, G. and GOVAERT, G. (2000). Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE transactions on pattern analysis and machine intelligence* **22** 719–725.
- BIERNACKI, C. and LOURME, A. (2014). Stable and visualizable Gaussian parsimonious clustering models. *Statistics and Computing* **24** 953–969. [MR3253847](#)
- BLUM, A. L. and LANGLEY, P. (1997). Selection of relevant features and examples in machine learning. *Artificial Intelligence* **97** 245–271. [MR1605908](#)
- BONTEMPS, D. and TOUSSILE, W. (2013). Clustering and variable selection for categorical multivariate data. *Electronic Journal of Statistics* **7** 2344–2371. [MR3108816](#)
- BOUYEYRON, C. and BRUNET, C. (2012). Simultaneous model-based clustering and visualization in the Fisher discriminative subspace. *Statistics and Computing* **22** 301–324. [MR2865072](#)
- BOUYEYRON, C. and BRUNET-SAUMARD, C. (2014a). Model-based clustering of high-dimensional data: A review. *Computational Statistics & Data Analysis* **71** 52–78. [MR3131954](#)
- BOUYEYRON, C. and BRUNET-SAUMARD, C. (2014b). Discriminative variable selection for clustering with the sparse Fisher-EM algorithm. *Computational Statistics* **29** 489–513. [MR3261825](#)

- BOUYEYRON, C., GIRARD, S. and SCHMID, C. (2007). High-dimensional data clustering. *Computational Statistics & Data Analysis* **52** 502–519. [MR2409998](#)
- CELEUX, G. and GOVAERT, G. (1995). Gaussian parsimonious clustering models. *Pattern Recognition* **28** 781–793.
- CELEUX, G., MARTIN-MAGNIETTE, M. L., MAUGIS-RABUSSEAU, C. and RAFTERY, A. E. (2014). Comparing model selection and regularization approaches to variable selection in model-based clustering. *Journal de la Société Française de Statistique* **155** 57–71. [MR3211754](#)
- CELEUX, G., MAUGIS-RABUSSEAU, C. and SEDKI, M. (2018). Variable selection in model-based clustering and discriminant analysis with a regularization approach. *Advances in Data Analysis and Classification*.
- CHANG, W.-C. (1983). On Using Principal Components Before Separating a Mixture of Two Multivariate Normal Distributions. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **32** 267–275. [MR0770316](#)
- CHAVENT, M., KUENTZ-SIMONET, V., LIQUET, B. and SARACCO, J. (2012). ClustOfVar: An R package for the clustering of variables. *Journal of Statistical Software, Articles* **50** 1–16.
- CHEN, W.-C. and MAITRA, R. (2015). EMCluster: EM Algorithm for model-based clustering of finite mixture Gaussian distribution. R Package, URL <http://cran.r-project.org/package=EMCluster>.
- CHEN, L. S., PRENTICE, R. L. and WANG, P. (2014). A penalized EM algorithm incorporating missing data mechanism for Gaussian parameter estimation. *Biometrics* **70** 312–322. [MR3258036](#)
- CLARK, S. J. and SHARROW, D. J. (2011). Contemporary model life tables for developed countries: An application of model-based clustering. *Working Paper, Center for Statistics and the Social Sciences, University of Washington*.
- CLOGG, C. C. (1988). Latent class models for measuring. In *Latent Trait and Latent Class Models* (R. Langeheine and J. Rost, eds.) 8, 173–205. Plenum Press.
- COLLINS, L. M. and LANZA, S. T. (2010). *Latent Class and Latent Transition Analysis*. Wiley. [MR2420348](#)
- CONGRESSIONAL QUARTERLY ALMANAC (1984). 98th Congress, 2nd session.
- CONSTANTINOPOULOS, C., TITSIAS, M. K. and LIKAS, A. (2006). Bayesian feature and model selection for Gaussian mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **28** 1013–1018.
- DASH, M. and LIU, H. (1997). Feature selection for classification. *Intelligent data analysis* **1** 131–156.
- DEAN, N. and RAFTERY, A. E. (2010). Latent class analysis variable selection. *Annals of the Institute of Statistical Mathematics* **62** 11–35. [MR2577437](#)
- DEMPSTER, A. P., 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* **39** 1–38. [MR0501537](#)
- DESANTIS, S. M., HOUSEMAN, E. A., COULL, B. A., STEMMER-RACHAMIMOV, A. and BETENSKY, R. A. (2008). A penalized latent class model for ordinal data. *Biostatistics* **9** 249.
- DY, J. G. and BRODLEY, C. E. (2004). Feature selection for unsupervised learning. *Journal of Machine Learning Research* **5** 845–889. [MR2248002](#)

- FOP, M. and MURPHY, T. B. (2017). LCAvarsel: Variable selection for latent class analysis R package version 1.1, <https://cran.r-project.org/package=LCAvarsel>.
- FOP, M., SMART, K. M. and MURPHY, T. B. (2017). Variable selection for latent class analysis with application to low back pain diagnosis. *Annals of Applied Statistics* **11** 2085–2115. [MR3743289](#)
- FORMANN, A. K. (1985). Constrained latent class models: Theory and applications. *British Journal of Mathematical and Statistical Psychology* **38** 87–111. [MR0793251](#)
- FORMANN, A. K. (2007). Mixture analysis of multivariate categorical data with covariates and missing entries. *Computational Statistics & Data Analysis* **51** 5236–5246. [MR2370868](#)
- FOWLKES, E. B., GNANADESIKAN, R. and KETTENRING, J. R. (1988). Variable selection in clustering. *Journal of Classification* **5** 205–228. [MR0971153](#)
- FRALEY, C. and RAFTERY, A. E. (2002). Model-based clustering, discriminant analysis and density estimation. *Journal of the American Statistical Association* **97** 611–631. [MR1951635](#)
- FRIEDMAN, J. H. and MEULMAN, J. J. (2004). Clustering objects on subsets of attributes (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **66** 815–849. [MR2102467](#)
- GALIMBERTI, G., MANISI, A. and SOFFRITTI, G. (2017). Modelling the role of variables in model-based cluster analysis. *Statistics and Computing* 1–25. [MR3741643](#)
- GALIMBERTI, G., MONTANARI, A. and VIROLI, C. (2009). Penalized factor mixture analysis for variable selection in clustered data. *Computational Statistics & Data Analysis* **53** 4301–4310. [MR2744326](#)
- GOLLINI, I. and MURPHY, T. B. (2014). Mixture of latent trait analyzers for model-based clustering of categorical data. *Statistics and Computing* **24** 569–588. [MR3223542](#)
- GOLUB, T. R., SLONIM, D. K., TAMAYO, P., HUARD, C., GAASENBEEK, M., MESIROV, J. P., COLLER, H., LOH, M. L., DOWNING, J. R., CALIGIURI, M. A. and BLOOMFIELD, C. D. (1999). Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science* **286** 531–537.
- GOODMAN, L. A. (1974). Exploratory latent structure analysis using both identifiable and unidentifiable models. *Biometrika* **61** 215–231. [MR0370936](#)
- GREEN, P. J. (1990). On use of the EM for penalized likelihood estimation. *Journal of the Royal Statistical Society. Series B (Methodological)* **52** 443–452. [MR1086796](#)
- GUO, J., LEVINA, E., MICHAILIDIS, G. and ZHU, J. (2010). Pairwise variable selection for high-dimensional model-based clustering. *Biometrics* **66** 793–804. [MR2758215](#)
- GUYON, I. and ELISSEEFF, A. (2003). An introduction to variable and feature selection. *Journal of Machine Learning Research* **3** 1157–1182.
- HOFF, P. D. (2005). Subset clustering of binary sequences, with an application to genomic abnormality data. *Biometrics* **61** 1027–1036. [MR2216196](#)

- HOFF, P. D. (2006). Model-based subspace clustering. *Bayesian Analysis* **1** 321–344. [MR2221267](#)
- HOUSEMAN, E. A., COULL, B. A. and BETENSKY, R. A. (2006). Feature-specific penalized latent class analysis for genomic data. *Biometrics* **62** 1062–1070. [MR2297677](#)
- HUBERT, L. and ARABIE, P. (1985). Comparing partitions. *Journal of Classification* **2** 193–218.
- HUMAN MORTALITY DATABASE (2017). University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany). <http://www.mortality.org/>.
- HUNT, L. and JORGENSEN, M. (2003). Mixture model clustering for mixed data with missing information. *Computational Statistics & Data Analysis* **41** 429–440. [MR1973722](#)
- JOHN, G. H., KOHAVI, R. and PFLEGER, K. (1994). Irrelevant features and the subset selection problem. In *Proceedings of the Eleventh International Conference on Machine Learning* 121–129.
- KARLIS, D. and MELIGKOTSIDOU, L. (2007). Finite mixtures of multivariate Poisson distributions with application. *Journal of Statistical Planning and Inference* **137** 1942–1960. [MR2323875](#)
- KIM, S., TADESSE, M. G. and VANNUCCI, M. (2006). Variable selection in clustering via Dirichlet process mixture models. *Biometrika* **93** 877–893. [MR2285077](#)
- KOHAVI, R. and JOHN, G. H. (1997). Wrappers for feature subset selection. *Artificial Intelligence* **97** 273–324.
- KOLLER, D. and SAHAMI, M. (1996). Toward optimal feature selection. In *Proceedings of the Thirteenth International Conference on Machine Learning (ICML)* (L. SAITTA, ed.) 284–292. Morgan Kaufmann Publishers.
- KOSMIDIS, I. and KARLIS, D. (2016). Model-based clustering using copulas with applications. *Statistics and Computing* **26** 1079–1099. [MR3535046](#)
- LAW, M. H. C., FIGUEIREDO, M. A. T. and JAIN, A. K. (2004). Simultaneous feature selection and clustering using mixture models. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **26** 1154–1166.
- LEBRET, R., IOVLEFF, S., LANGROGNET, F., BIERNACKI, C., CELEUX, G. and GOVAERT, G. (2015). Rmixmod: The R package of the model-based unsupervised, supervised, and semi-supervised classification mixmod library. *Journal of Statistical Software, Articles* **67** 1–29.
- LEE, H. and LI, J. (2012). Variable selection for clustering by separability based on ridgelines. *Journal of Computational and Graphical Statistics* **21** 315–336. [MR2945469](#)
- LEE, S. X. and MCLACHLAN, G. J. (2013). On mixtures of skew normal and skew t -distributions. *Advances in Data Analysis and Classification* **7** 241–266. [MR3103965](#)
- LEE, S. X. and MCLACHLAN, G. J. (2016). Finite mixtures of canonical fundamental skew t -distributions: The unification of the restricted and unrestricted skew t -mixture models. *Statistics and Computing* **26** 573–589. [MR3489858](#)

- LEISCH, F. (2004). FlexMix: A general framework for finite mixture models and latent class regression in R. *Journal of Statistical Software, Articles* **11** 1–18.
- LINZER, D. A. and LEWIS, J. B. (2011). poLCA: An R package for polytomous variable latent class analysis. *Journal of Statistical Software* **42** 1–29.
- LITTLE, R. J. A. and RUBIN, D. B. (2002). *Statistical Analysis with Missing Data*. Wiley. [MR1925014](#)
- LIU, H. and MOTODA, H. (2007). *Computational Methods of Feature Selection*. CRC Press.
- LIU, J. S., ZHANG, J. L., PALUMBO, M. J. and LAWRENCE, C. E. (2003). Bayesian clustering with variable and transformation selections. In *Bayesian Statistics 7: Proceedings of the Seventh Valencia International Meeting* (J. M. BERNARDO, M. J. BAYARRI, J. O. BERGER, A. P. DAWID, D. HECKERMAN, A. F. M. SMITH and M. WEST, eds.) 249–275. Oxford University Press.
- MALSINER-WALLI, G., FRÜHWIRTH-SCHNATTER, S. and GRÜN, B. (2016). Model-based clustering based on sparse finite Gaussian mixtures. *Statistics and Computing* **26** 303–324. [MR3439375](#)
- MARBAC, M., BIERNACKI, C. and VANDEWALLE, V. (2015). Model-based clustering for conditionally correlated categorical data. *Journal of Classification* **32** 145–175. [MR3369412](#)
- MARBAC, M. and SEDKI, M. (2017a). Variable selection for model-based clustering using the integrated complete-data likelihood. *Statistics and Computing* **27** 1049–1063. [MR3627562](#)
- MARBAC, M. and SEDKI, M. (2017b). VarSelLCM: Variable selection for model-based clustering of continuous, count, categorical or mixed-type data set with missing values R package version 2.0.1, <https://CRAN.R-project.org/package=VarSelLCM>.
- MARBAC, M. and SEDKI, M. (2017c). Variable selection for mixed data clustering: A model-based approach. arXiv:1703.02293.
- MAUGIS, C., CELEUX, G. and MARTIN-MAGNIETTE, M. L. (2009a). Variable selection for clustering with Gaussian mixture models. *Biometrics* **65** 701–709. [MR2649842](#)
- MAUGIS, C., CELEUX, G. and MARTIN-MAGNIETTE, M. L. (2009b). Variable selection in model-based clustering: A general variable role modeling. *Computational Statistics and Data Analysis* **53** 3872–3882. [MR2749931](#)
- MAUGIS-RABUSSEAU, C., MARTIN-MAGNIETTE, M.-L. and PELLETIER, S. (2012). SelvarClustMV: Variable selection approach in model-based clustering allowing for missing values. *Journal de la Société Française de Statistique* **153** 21–36. [MR3008597](#)
- McLACHLAN, G. J. and BASFORD, K. E. (1988). *Mixture models: Inference and applications to clustering*. Marcel Dekker. [MR0926484](#)
- McLACHLAN, G. J., BEAN, R. W. and PEEL, D. (2002). A mixture model-based approach to the clustering of microarray expression data. *Bioinformatics* **18** 413–422.
- McLACHLAN, G. and KRISHNAN, T. (2008). *The EM Algorithm and Extensions*. Wiley. [MR2392878](#)

- MCLACHLAN, G. J. and PEEL, D. (1998). *Robust cluster analysis via mixtures of multivariate t -distributions* In *Advances in Pattern Recognition: Joint IAPR International Workshops SSPR'98 and SPR'98 Sydney, Australia, August 11–13, 1998 Proceedings* 658–666. Springer Berlin Heidelberg. [MR1682742](#)
- MCLACHLAN, G. and PEEL, D. (2000). *Finite Mixture Models*. Wiley. [MR1789474](#)
- MCLACHLAN, G. J. and RATHNAYAKE, S. (2014). On the number of components in a Gaussian mixture model. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **4** 341–355.
- MCLACHLAN, G. J., PEEL, D., BASFORD, K. E. and P., A. (1999). The EMMIX software for the fitting of mixtures of normal and t -components. *Journal of Statistical Software* **4** 1–14.
- MCNICHOLAS, P. D. (2016). Model-based clustering. *Journal of Classification* **33** 331–373. [MR3575621](#)
- MCNICHOLAS, D. P. and MURPHY, T. B. (2008). Parsimonious Gaussian mixture models. *Statistics and Computing* **18** 285–296. [MR2413385](#)
- MCPARLAND, D. and GORMLEY, I. C. (2016). Model based clustering for mixed data: clustMD. *Advances in Data Analysis and Classification* **10** 155–169. [MR3505054](#)
- MELNYKOV, V. and MAITRA, R. (2010). Finite mixture models and model-based clustering. **4** 80–116. [MR2644009](#)
- MEYER, D., DIMITRIADOU, E., HORNIK, K., WEINGESSEL, A. and LEISCH, F. (2017). e1071: Misc functions of the Department of Statistics, Probability Theory Group (formerly: E1071), TU Wien R package version 1.6-8, <https://CRAN.R-project.org/package=e1071>.
- MILLER, A. (2002). *Subset selection in regression*. CRC Press. [MR2001193](#)
- NIA, V. P. and DAVISON, A. C. (2012). High-dimensional Bayesian clustering with variable selection: The R package bclust. *Journal of Statistical Software* **47** 1–22.
- NIA, V. P. and DAVISON, A. C. (2015). A simple model-based approach to variable selection in classification and clustering. *Canadian Journal of Statistics* **43** 157–175. [MR3353378](#)
- O'HAGAN, A., MURPHY, T. B. and GORMLEY, I. C. (2012). Computational aspects of fitting mixture models via the expectation-maximization algorithm. *Computational Statistics & Data Analysis* **56** 3843–3864. [MR2957835](#)
- PAN, W. and SHEN, X. (2007). Penalized model-based clustering with application to variable selection. *Journal of Machine Learning Research* **8** 1145–1164.
- PYTHON SOFTWARE FOUNDATION (2017). Python: A dynamic, open source programming language <https://www.python.org/>.
- R CORE TEAM (2017). R: A Language and Environment for Statistical Computing R Foundation for Statistical Computing, Vienna, Austria <https://www.R-project.org/>.
- RAFTERY, A. E. and DEAN, N. (2006). Variable selection for model-based clustering. *Journal of the American Statistical Association* **101** 168–178. [MR2268036](#)

- RAU, A., MAUGIS-RABUSSEAU, C., MARTIN-MAGNIETTE, M.-L. and CELEUX, G. (2015). Co-expression analysis of high-throughput transcriptome sequencing data with Poisson mixture models. *Bioinformatics* **31** 1420–1427.
- RITTER, G. (2014). *Robust cluster analysis and variable selection*. CRC Press. [MR3241693](#)
- SAEYS, Y., INZA, I. N. and LARRAÑAGA, P. (2007). A review of feature selection techniques in bioinformatics. *Bioinformatics* **23** 2507–2517.
- SCHWARZ, G. (1978). Estimating the dimension of a model. *Annals of Statistics* **6** 461–464. [MR0468014](#)
- SCRUCCA, L. (2016). Genetic algorithms for subset selection in model-based clustering. In *Unsupervised Learning Algorithms* (M. E. Celebi and K. Aydin, eds.) 55–70. Springer.
- SCRUCCA, L. and RAFTERY, A. E. (2018). clustvarsel: A package implementing variable selection for Gaussian model-based clustering in R. *Journal of Statistical Software* **84** 1–28. [MR2268036](#)
- SCRUCCA, L., FOP, M., MURPHY, T. B. and RAFTERY, A. E. (2016). mclust 5: Clustering, classification and density estimation using Gaussian finite mixture models. *The R Journal* **8** 289–317.
- SEDKI, M., CELEUX, G. and MAUGIS, C. (2014). SelvarMix: A R package for variable selection in model-based clustering and discriminant analysis with a regularization approach. *INRIA Technical report*.
- SEDKI, M., CELEUX, G. and MAUGIS-RABUSSEAU, C. (2017). SelvarMix: Regularization for variable selection in model-based clustering and discriminant analysis R package version 1.2.1, <https://CRAN.R-project.org/package=SelvarMix>.
- SILVESTRE, C., CARDOSO, M. G. M. S. and FIGUEIREDO, M. (2015). Feature selection for clustering categorical data with an embedded modelling approach. *Expert Systems* **32** 444–453.
- STEINLEY, D. and BRUSCO, M. J. (2008). Selection of variables in cluster analysis: An empirical comparison of eight procedures. *Psychometrika* **73** 125–144. [MR2395296](#)
- SUN, W., WANG, J. and FANG, Y. (2012). Regularized k-means clustering of high-dimensional data and its asymptotic consistency. *Electronic Journal of Statistics* **6** 148–167. [MR2879675](#)
- SWARTZ, M. D., MO, Q., MURPHY, M. E., LUPTON, J. R., TURNER, N. D., HONG, M. Y. and VANNUCCI, M. (2008). Bayesian variable selection in clustering high-dimensional data with substructure. *Journal of Agricultural, Biological, and Environmental Statistics* **13** 407. [MR2590937](#)
- TADESSE, M. G., SHA, N. and VANNUCCI, M. (2005). Bayesian variable selection in clustering high-dimensional data. *Journal of the American Statistical Association* **100** 602–617. [MR2160563](#)
- TOUSSILE, W. (2016). ClustMMDD: Variable selection in clustering by mixture models for discrete data R package version 1.0.4.
- TOUSSILE, W. and GASSIAT, E. (2009). Variable selection in model-based clustering using multilocus genotype data. *Advances in Data Analysis and Classification* **3** 109–134. [MR2551051](#)

- VERMUNT, J. K. and MAGDISON, J. (2002). Latent class cluster analysis. In *Applied Latent Class Analysis* (J. A. Hagenaaers and A. L. McCutcheon, eds.) 3, 89–106. Cambridge University Press. [MR1927666](#)
- WALLACE, C. S. and FREEMAN, P. R. (1987). Estimation and inference by compact coding. *Journal of the Royal Statistical Society. Series B (Methodological)* **49** 240–265. [MR0928937](#)
- WALLACE, M. L., BUYSE, D. J., GERMAIN, A., HALL, M. H. and IYENGAR, S. (2017). Variable selection for skewed model-based clustering: Application to the identification of novel sleep phenotypes. *Journal of the American Statistical Association* **0** 0-0.
- WANG, S. and ZHU, J. (2008). Variable selection for model-based high-dimensional clustering and its application to microarray data. *Biometrics* **64** 440–448. [MR2432414](#)
- WHITE, A. and MURPHY, T. B. (2014). BayesLCA: An R package for Bayesian latent class analysis. *Journal of Statistical Software* **61** 1–28.
- WHITE, A., WYSE, J. and MURPHY, T. B. (2016). Bayesian variable selection for latent class analysis using a collapsed Gibbs sampler. *Statistics and Computing* **26** 511–527. [MR3439388](#)
- WITTEN, D. M. and TIBSHIRANI, R. (2010). A framework for feature selection in clustering. *Journal of the American Statistical Association* **105** 713–726. [MR2724855](#)
- WITTEN, D. M. and TIBSHIRANI, R. (2013). sparcl: Perform sparse hierarchical clustering and sparse k-means clustering R package version 1.0.3, <https://CRAN.R-project.org/package=sparcl>.
- WU, B. (2013). Sparse cluster analysis of large-scale discrete variables with application to single nucleotide polymorphism data. *Journal of Applied Statistics* **40** 358–367. [MR3047154](#)
- XIE, B., PAN, W. and SHEN, X. (2008a). Variable selection in penalized model-based clustering via regularization on grouped parameters. *Biometrics* **64** 921–930. [MR2526644](#)
- XIE, B., PAN, W. and SHEN, X. (2008b). Penalized model-based clustering with cluster-specific diagonal covariance matrices and grouped variables. *Electronic Journal of Statistics* **2** 168–212. [MR2386092](#)
- XIE, B., PAN, W. and SHEN, X. (2010). Penalized mixtures of factor analyzers with application to clustering high-dimensional microarray data. *Bioinformatics* **26** 501.
- YU, L. and LIU, H. (2004). Efficient feature selection via analysis of relevance and redundancy. *Journal of Machine Learning Research* **5** 1205–1224. [MR2248015](#)
- ZHANG, Q. and IP, E. H. (2014). Variable assessment in latent class models. *Computational Statistics & Data Analysis* **77** 146–156. [MR3210054](#)
- ZHOU, H., PAN, W. and SHEN, X. (2009). Penalized model-based clustering with unconstrained covariance matrices. *Electronic Journal of Statistics* **3** 1473–1496. [MR2578834](#)