



Better than the best? Answers via model ensemble in density-based clustering

Alessandro Casa¹  · Luca Scrucca² · Giovanna Menardi¹

Received: 6 March 2020 / Revised: 24 July 2020 / Accepted: 14 September 2020 /
Published online: 2 October 2020
© The Author(s) 2020

Abstract

With the recent growth in data availability and complexity, and the associated outburst of elaborate modelling approaches, model selection tools have become a lifeline, providing objective criteria to deal with this increasingly challenging landscape. In fact, basing predictions and inference on a single model may be limiting if not harmful; ensemble approaches, which combine different models, have been proposed to overcome the selection step, and proven fruitful especially in the supervised learning framework. Conversely, these approaches have been scantily explored in the unsupervised setting. In this work we focus on the model-based clustering formulation, where a plethora of mixture models, with different number of components and parametrizations, is typically estimated. We propose an ensemble clustering approach that circumvents the single best model paradigm, while improving stability and robustness of the partitions. A new density estimator, being a convex linear combination of the density estimates in the ensemble, is introduced and exploited for group assignment. As opposed to the standard case, where clusters are typically associated to the components of the selected mixture model, we define partitions by borrowing the modal, or nonparametric, formulation of the clustering problem, where groups are linked with high-density regions. Staying in the density-based realm we thus show how blending together parametric and nonparametric approaches may be beneficial from a clustering perspective.

Keywords Cluster analysis · Model averaging · Ensemble learning · Density-based clustering · Density estimation

✉ Alessandro Casa
casa@stat.unipd.it

Luca Scrucca
luca.scrucca@unipg.it

Giovanna Menardi
menardi@stat.unipd.it

¹ Department of Statistical Sciences, University of Padova, Via C. Battisti, 241, 35121 Padua, Italy

² Department of Economics, University of Perugia, Via A. Pascoli, 20, 06123 Perugia, Italy

1 Introduction

In virtually any scientific domain we are witnessing an explosion in the availability of the data, coupled with a tremendous growth in their complexity. As a straightforward consequence, the number of choices that we have to make is increasing as well as the number of sophisticated modelling strategies proposed to deal with such newly introduced challenges. These choices are practically involved in any phase of the modelling process, spanning a wide landscape of possible options: from choosing a class of models or an appropriate approach to analyze a set of data, to more specific decisions as the selection of subsets of relevant variables or suitable parametrizations. Therefore, nowadays model selection steps, helping to formally extricate ourselves from the labyrinth of all these possible alternatives, are ubiquitous in any data analysis routine. Some commonly considered ways forward hence consist in estimating a set of different models and then selecting the best one according to some criterion (Claeskens and Hjort 2008), or resorting to penalization schemes aimed at balancing fit and complexity (see Tibshirani et al. 2015, for an introduction).

Nevertheless basing predictions and inference on a single model could turn out to be sub-optimal. Hence, in order to propose viable alternatives to this paradigm, model averaging and ensemble techniques have been thoroughly studied in literature. Even if these two approaches focus on different phases of the modelling process, respectively estimation and prediction, they share the same founding rationale as they aim to improve performances of the base models by combining their strengths and simultaneously relieving their limits. For this reason the two expressions will be used interchangeably in the rest of the paper. With inferential goals in mind, model averaging approaches have been proposed, intended to estimate quantities by computing weighted averages of different estimates. Such strategies may lead to improvements in the estimation process by accounting for model uncertainty. Similarly, from a predictive point of view, ensemble techniques have shown remarkable performances in a lot of different applications by building predictions as combinations of the ones given by a set of different models. Well established methods as *bagging*, *stacking*, *boosting* or the *random forests* (see Friedman et al. 2001, for a review) have become the state of the art in the supervised learning framework.

While extensively studied in the classification context, ensemble techniques have been scarcely pursued in the clustering one. A possible explanation can be found in the unsupervised nature of the problem itself; the absence of a response variable introduces relevant issues in evaluating the quality of a model and of the corresponding partition. As a consequence, weighting models in order to combine them turns out to be an awkward problem. Nonetheless mixing different partitions in a final one could in principle allows combining clustering techniques based on different focuses to give a multiresolution view of the data and possibly improves the stability and the robustness of the solutions. In this direction Fern and Brodley (2003) exploit the concept of *similarity matrix* in order to aggregate partitions obtained on multiple random projections, and a similar approach is followed by Kuncheva and Hadjitodorov (2004)

to study the concept of diversity among partitions. Monti et al. (2003) consider again a similarity matrix in order to evaluate the robustness of a discovered cluster under random resampling. In turn, the work by Strehl and Ghosh (2002) introduces three different solutions to the ensemble problem in the unsupervised setting by exploiting hypergraph representations of the partitions.

In this work we focus mainly on the parametric, or model-based, approach to cluster analysis where a one-to-one correspondence between clusters and components of an appropriate mixture model is drawn. Here the usual working routine is based on the *single best model paradigm*, i.e. a set of models is fitted and only the best one is chosen and considered to obtain a partition. Our aim is to go beyond this paradigm by introducing a model averaging methodology to give partitions resulting from an ensemble of models, thus possibly achieving a greater accuracy and robustness. Averaging is pursued directly on the estimated mixture densities in order to build a new and more accurate estimate. Clustering is then addressed, building on the resulting estimate, within a density-based formulation, yet with a shift to a modal, or nonparametric approach, where clusters are associated to the domains of attraction of the density modes. We turn out with a partition where cluster shapes are not constrained by some distributional assumption, as in the model-based framework, but having arbitrary, possibly non-convex and skewed shapes. Therefore, by combining the strengths of the parametric and nonparametric frameworks, our proposal results in a hybrid method which enjoys the advantages of both.

The rest of the paper is structured as follows. In Sect. 2 we outline the proposed methodology and describe in details the estimation procedure. In Sect. 3 we discuss some specific aspects of our proposal and highlight connections with other models. Lastly in Sect. 4 we show the performances of our method on both simulated and real datasets, and compare it with some competitors while Sect. 5 presents some concluding remarks.

2 Model averaging in model-based clustering

2.1 Framework and model specification

The goal of partitioning a set of data into some groups, diffusely known as clustering, has been pursued by proposing a lot of techniques with different rationales behind. While most of them are based on a vague notion of clusters, associated to some measure of similarity, an attempt to obtain a precise formalization of the problem is given by the so called *density-based* approach. Here the concept of cluster finds a formal definition by linking it to some specific features of the density $f : \mathbb{R}^d \rightarrow \mathbb{R}$ assumed to underlie the data $\mathcal{X} = \{x_1, \dots, x_n\}$ and consequently inducing a partition of the whole sample space. Furthermore this assumption allows framing the clustering problem in a standard inferential context where, having a “ground truth” to aim at, several tools can be used in order to evaluate and compare alternative clustering configurations.

The idea behind density-based clustering has been developed by taking two distinct paths. In the modal, or nonparametric, clustering formulation, clusters are defined as the “domains of attraction” of the modes of the density f (Stuetzle 2003) usually

estimated by means of some nonparametric density estimator (see, for a recent account, Chacón and Duong 2018). The operational identification of the modal regions can be addressed following different routes (for a comprehensive review readers can refer to Menardi 2016) where the most common one consists in finding explicitly the local maxima of the density by exploiting numerical optimization methods. Most of these methods can be seen as refinements or slight modifications of the mean-shift algorithm (Fukunaga and Hostetler 1975) that, starting from a generic point, shifts it recursively along the steepest ascend path of the gradient of the density estimate until converging to a mode; the final partition of the data is then obtained by grouping together those observations ascending to the same mode.

On the other hand the model-based, or parametric, approach (Banfield and Raftery 1993; Fraley and Raftery 2002) represents the other, more widespread, formulation of density-based clustering. In this framework f is assumed to be adequately described by means of finite mixture models. Therefore the density of a generic observation $x_i \in \mathbb{R}^d$ is written as

$$f(x_i|\Psi) = \sum_{k=1}^K \pi_k \varphi_k(x_i|\theta_k),$$

where K is the number of mixture components, $\varphi_k(\cdot)$ the k th component density, while $\Psi = (\pi_1, \dots, \pi_{K-1}, \theta_1, \dots, \theta_K)$ is the vector of parameters where π_k are the mixing proportions with $\pi_k > 0$, $\forall k = 1, \dots, K$ and $\sum_{k=1}^K \pi_k = 1$. When Gaussian densities are employed as mixture components, we may write $\varphi_k(\cdot) = \phi_k(\cdot)$ and $\theta_k = \{\mu_k, \Sigma_k\}$. The concept of cluster here is defined by drawing a one-to-one correspondence between the group itself and a component of the mixture. Operationally, after having estimated the model, usually via maximum likelihood by means of the EM algorithm (Dempster et al. 1977), the allocation is obtained via maximum a posteriori (MAP) classification by assigning the i th observation to cluster k^* if

$$k^* = \arg \max_k \frac{\hat{\pi}_k \varphi_k(x_i|\hat{\theta}_k)}{\sum_{k=1}^K \hat{\pi}_k \varphi_k(x_i|\hat{\theta}_k)}.$$

When practically resorting to model-based clustering in order to obtain a partition, several choices have to be made as, for example, the number of groups K , the parametric specification of the mixture components or a specific parsimony-inducing parametrization of Σ_k in the Gaussian case. Since each combination of these possibilities can be seen as a different model, it is clear how model selection steps play an essential role in this framework. Indeed several different models corresponding to such combinations are usually estimated, the best one is then chosen according to an information criterion such as the BIC (Schwarz 1978) or the ICL (Biernacki et al. 2000) and successively used to obtain the final partition.

The way of proceeding, usually referred to as the *single best model paradigm*, could be sub-optimal especially when differences among values of the information criterion across competing models are close. As an illustrative example we analyze the widely known Iris dataset by employing Gaussian mixture models with $K = 1, \dots, 9$ with

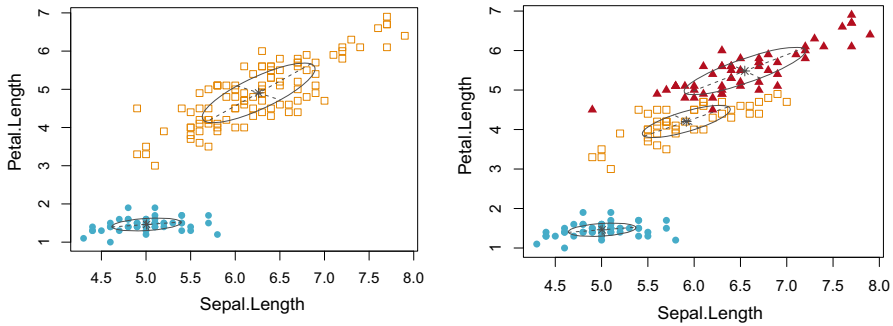


Fig. 1 Example on Iris data: on the left the partition induced by the best model according to the Bayesian information criterion ($\text{BIC} = -561.72$). On the right the partition induced by the second best model ($\text{BIC} = -562.55$)

all the possible parametrizations of the component covariance matrices available in `mclust` package (see Scrucca et al. 2016 for further details). In Fig. 1 the results obtained by the two best models according to the BIC, obtained on the four dimensional space, are shown on the subspace spanned by the variables sepal length and petal length. Even if no formal criterion is available in order to check if the difference between the values of the BIC is significant, they appear quite close. In fact, the true labels indicate the presence of three classes, here adequately captured by the second best model. Even if a proper quantification of the number of groups in these data is still controversial, since the presence of three species does not necessarily imply the presence of three groups, it seems natural to ask if, by discarding completely the second best model, useful information on the data is thrown away.

In this setting combining competitive models together may lead to a gain in robustness, stability and in the quality of the partition, as often witnessed in the supervised framework.

In a parametric clustering framework the idea of combining different models has been developed in order to obtain partitions based on an average of models rather than on a single one. Both the works of Russell et al. (2015) and Wei and McNicholas (2015) propose a Bayesian model averaging approach to postprocess the results of model-based clustering. A key issue pointed out in both the proposals consists in the need of selecting an invariant quantity, i.e. a quantity having the same meaning across all the models in the ensemble, to average on. In parametric clustering this represents a cumbersome problem since the models to mix together could possibly have different number of groups as it happens, for example, if the ensemble is built with the two models in Fig. 1; as a consequence, parameters spaces have different dimensions, thus preventing the chance to average directly parameters estimates. Wei and McNicholas (2015) overcome this issue by introducing a component merging step in the procedure. Alternatively, Russell et al. (2015) consider a similarity matrix as the invariant quantity, built on the agreement of cluster assignment of pairs of observations. They obtain an ensemble similarity matrix by averaging the candidate models ones. Afterwards the resulting matrix, where the (i, j) th entry represents the averaged probability of x_i

and x_j to belong to the same cluster, is considered to obtain partitions adopting a hierarchical clustering approach.

In this work we take a different path with respect to the ones mentioned above. The issue is tackled directly at its roots, by exploiting the essential role assumed by the density in the considered framework. Therefore, recasting the problem to a density estimation one, the density itself is chosen as the invariant quantity to be averaged. Note that some promising results, from a density estimation perspective only, have been obtained by Glodek et al. (2013) whose work shares some similarities with our proposal.

Let $\{f_m(\cdot|\hat{\Psi}_m)\}_{m=1,\dots,M}$ be a set of estimated candidate mixture models. In this section the number M of models to average is considered as given, and we refer the reader to Sect. 3 for a discussion about this aspect. In the rest of the work we focus specifically on mixtures of Gaussian densities. This choice is not binding for subsequent developments since, in principle, the ensemble may be populated by mixture models with different parametric specifications for the component densities. A new estimator, being a convex linear combination of the estimated densities $f_m(\cdot|\hat{\Psi}_m)$, is introduced:

$$\tilde{f}(x; \alpha) = \sum_{m=1}^M \alpha_m f_m(x|\hat{\Psi}_m), \quad (1)$$

with $\alpha_m > 0$, $\sum_m \alpha_m = 1$, representing the weight assigned to the m th model for all $m = 1, \dots, M$. A key aspect, as it will be discussed in Sect. 2.2, consists in properly estimating the model weights in order to guarantee that models describing more adequately the underlying density will count more in the resulting estimator.

The rationale behind our proposal draws strength from some results obtained by Rigollet and Tsybakov (2007). Here the authors show that, under some fairly general regularity assumptions, linearly aggregating density estimators leads asymptotically to an improvement in the resulting density under L_2 -loss perspective. Hence, by possibly improving the quality of the density estimates, we aim at obtaining better characterizations of the relevant patterns in the data, leading to more refined partitions.

Even if the estimator (1) is still a mixture model we cannot obtain a partition as usually carried out in parametric clustering, thus resorting to the one-to-one correspondence among groups and components. As an illustrative example, consider an ensemble formed by two mixture models, with two and three components respectively. In this situation $\tilde{f}(\cdot; \alpha)$ will result in a five component mixture model hence giving contradictory indications about the number of groups with respect to the models that have been mixed together. This issue shares strong contact points with the situations where the number of components exceeds the number of groups; for example a two components Gaussian mixture may result in a unimodal density leading to a counterintuitive partition with no clear separation between the two groups. In the model-based clustering framework the problem has been addressed by resorting to merging procedures (Baudry et al. 2010; Hennig 2010) where mixture components are combined together and their union seen as a single cluster.

In this work we take a different path and naturally circumvent the problem by shifting the concept of cluster to its modal formulation. Consistently, the grouping structure is searched in the modality of an estimate of the density (1). Hence, each cluster is built by gathering all those observations ascending to the same mode of the density. Operationally, a partition is obtained by means of a mode-searching algorithm. Our proposal shares some conceptual connections with the work by Hennig (2010) where the author proposes merging methods aimed at finding unimodal clusters.

The proposed solution, staying in the realm of density-based clustering, inherits and enjoys its relevant strengths as the chance to frame the problem in a standard inferential setting where proper statistical tools can be employed for evaluation, and to obtain whole sample space partitions whose features are inferentially explorable. Moreover it has been shown already (see Scrucca 2016; Chacón 2019) that blending together parametric and nonparametric approaches to clustering can lead to some relevant improvements in some, otherwise troublesome, situations.

2.2 Model estimation

The procedure outlined in Sect. 2.1 requires a practical way to estimate the density in (1). Note that, since $\hat{\Psi}_m$ has been previously estimated, the only unknown parameters involved are the α_m s. These parameters represent the weights to be assigned at every single model in the ensemble, hence their estimation is crucial in governing the resulting shape of the density, its modal structure and consequently the final partition. A reasonable estimation procedure would result in giving nearly zero weights to those models in the ensemble which do not suitably capture the features of the underlying density, while weighting more the adequate ones.

In order to obtain an estimate for the weight vector $\alpha = (\alpha_1, \dots, \alpha_M)$, based on the sample \mathcal{X} , we aim to maximize the log-likelihood of the model (1), defined as

$$\ell(\alpha; \mathcal{X}) = \sum_{i=1}^n \log \sum_{m=1}^M \alpha_m f_m(x_i | \hat{\Psi}_m). \tag{2}$$

However, if (2) is considered as the objective function to maximize, the procedure will incur in the overfitting problem since the most complex models in the ensemble, which provide a better fit by construction, will weight more. This behaviour will commonly result in wiggler estimates not appropriately seizing the relevant features of the density, hence some regularization has to be considered in the estimation.

A tentative solution has been proposed by Smyth and Wolpert (1999) where a *stacking* procedure is adapted to the density estimation framework. The authors avoid to fall into the overfitting trap by exploiting a cross-validation scheme when combining the candidate models to obtain ensemble density estimates.

We take a different path by replacing the log-likelihood in (2) with a penalized version, generally defined as

$$\ell_P(\alpha; \mathcal{X}) = \ell(\alpha; \mathcal{X}) - \lambda g(\alpha, \nu). \tag{3}$$

Here $g(\cdot)$ is a penalty function to be specified, $v = (v_1, \dots, v_M)$ is a vector measuring the complexity of the models in the ensemble, while λ is a parameter controlling for the strength of the penalization. Within this general framework, we set v_m to be the cardinality of $\hat{\Psi}_m$, as it appears a sensible proxy of the complexity of the m th model. Additionally, we consider $g(\alpha, v) = \sum_m \alpha_m v_m$ as a simple choice which guarantees a stronger penalization to the most complex models. Note that, since both α_m and v_m are positive numbers, $\sum_m \alpha_m v_m = \sum_m |\alpha_m v_m|$. As a consequence our penalty might be seen as a generalization of the LASSO one where v_m is introduced to account for the complexities of the models in the ensemble.

For a given value of λ , the parameters is then estimated to maximize the penalized log-likelihood

$$\hat{\alpha} = \hat{\alpha}(\lambda) = \arg \max_{\alpha} \ell_P(\alpha; \mathcal{X}).$$

To this end, due to the mixture structure easily recognizable in (1), we can resort to a slightly simplified version of the EM-algorithm in order to maximize the penalized log-likelihood (3); note that the number of free parameters to estimate is equal to $M - 1$ since $\hat{\Psi}_m$ with $m = 1, \dots, M$ have to be considered as fixed and $\sum_{m=1}^M \alpha_m = 1$. At iteration t of the *E-step*, conditionally to an estimate $\hat{\alpha}^{(t-1)}$ for the vector α at the previous iteration, we compute

$$\tau_{mi}^{(t)} = \frac{\hat{\alpha}_m^{(t-1)} f_m(x_i | \hat{\Psi}_m)}{\sum_{m'=1}^M \hat{\alpha}_{m'}^{(t-1)} f_{m'}(x_i | \hat{\Psi}_{m'})}. \tag{4}$$

Then the *M-step* will consist in maximizing, with respect to α , the expected value of the complete-data penalized log-likelihood, in our setting expressed as

$$Q_P(\alpha; \hat{\alpha}^{(t-1)}) = \sum_{m=1}^M \sum_{i=1}^n \tau_{mi}^{(t)} [\log \alpha_m + \log f_m(x_i | \hat{\Psi}_m)] - \lambda \sum_{m=1}^M \alpha_m v_m, \tag{5}$$

under the constraint $\sum_m \alpha_m = 1$ with $\alpha_m > 0, \forall m = 1, \dots, M$. Since closed form solutions are not available, $\hat{\alpha}^{(t)}$ is obtained by maximizing (5) numerically. As usual, the two steps will be iterated until a convergence criterion is met. Several initialization strategies might be adopted in order to obtain $\hat{\alpha}^{(0)}$: in this work we consider a uniform initialization where, in the first step of the EM-algorithm, all the models in the ensemble are equally weighted.

Regarding the choice of λ , some more caution is needed, since an accurate selection turns out to be essential in order to obtain a meaningful estimate which properly reflects the geometrical structure of the underlying density. In this work some different options have been taken into consideration. A first, possible, strategy consists in estimating λ by means of the observed data resorting to a cross-validation scheme defined as follows:

- Randomly split the set $\{1, \dots, n\}$ into V equally-sized subsets $\mathcal{F}_1, \dots, \mathcal{F}_V$;
- For $v = 1, \dots, V$:

- Consider as a training sample $\mathcal{X}_{\text{train}}(v) = \{x_i\}_{i \notin \mathcal{F}_v}$ and as a test sample $\mathcal{X}_{\text{test}}(v) = \{x_i\}_{i \in \mathcal{F}_v}$;
 - For varying λ in a reasonable grid Λ , maximize $\ell_P(\alpha; \mathcal{X}_{\text{train}}(v))$ and obtain $\hat{\alpha}_v(\lambda)$;
 - For each $x \in \mathcal{X}_{\text{test}}(v)$ predict the density $\tilde{f}(x; \hat{\alpha}_v(\lambda))$;
- Define a *test log-likelihood*

$$\ell_{\text{test}}(\lambda) = \sum_{v=1}^V \sum_{x \in \mathcal{X}_{\text{test}}(v)} \log \tilde{f}(x; \hat{\alpha}_v(\lambda))$$

and select

$$\lambda_{CV} = \arg \max_{\lambda \in \Lambda} \ell_{\text{test}}(\lambda)$$

The selected λ_{CV} is finally used to obtain an estimate of α based on the whole sample.

Another reasonable approach consists in taking inspiration from the formulations of some information criteria which may be seen, in all respects, as penalized likelihood. Therefore we introduce the *AIC-type* and the *BIC-type* penalizations, stemming directly from the definitions of AIC and BIC, that induce penalized log-likelihoods defined as

$$\ell_{P,AIC}(\alpha; \mathcal{X}) = 2\ell(\alpha; \mathcal{X}) - 2 \sum_{m=1}^M \alpha_m \nu_m \tag{6}$$

$$\ell_{P,BIC}(\alpha; \mathcal{X}) = 2\ell(\alpha; \mathcal{X}) - \log(n) \sum_{m=1}^M \alpha_m \nu_m, \tag{7}$$

hence implying $\lambda_{AIC} = 1$ and $\lambda_{BIC} = \log(n)/2$ according to the formulation in (3). Although requiring an higher computational effort, it stands to reasons that the cross-validation based approach has some relevant advantages in the regularization process. By resorting to a fully data-driven selection of λ , we end up with a more adaptive parameter than λ_{BIC} and λ_{AIC} , both in terms of sample size and features of the observed data. However, the latter penalties are computationally faster and simple rules of thumbs enable, in practice, to produce satisfactory results, as will be discussed in Sect. 4.

Once the density (1) is estimated, a partition is operationally obtained by identifying its modal regions. To this aim we consider the so called Modal EM (MEM; Li et al. 2007) in the modified version proposed by Scrucca (2020). Designed for densities which are built as mixtures, this technique alternates two iterative steps in the guise of the EM, but unlike the EM its goal is to find the local maxima of the density. Since the density (1) may still be seen as a peculiar mixture model, MEM can be fruitfully adapted to our situation where, given an estimate $\hat{\alpha}$, a starting value $x^{(0)}$ and setting initially $r = 0$, the iterative steps are defined as follows

– Let

$$p_m = \frac{\hat{\alpha}_m f_m(x^{(r)} | \hat{\Psi}_m)}{\sum_{m'=1}^M \hat{\alpha}_{m'} f_{m'}(x^{(r)} | \hat{\Psi}_{m'})}$$

– Update

$$x^{(r+1)} = \arg \max_x \sum_{m=1}^M p_m \log f_m(x | \hat{\Psi}_m) .$$

The algorithm iteratively performs these steps until a convergence criterion is met. The outlined iterative procedure draws a path leading to a local maximum of the density (see Li et al. 2007, for a proof of the ascending property of the algorithm). Lastly, a partition is operationally obtained by using each observation $\{x_i\}_{i=1, \dots, n}$ in the sample as an initial value in the MEM and by grouping together those observations converging to the same mode.

3 Some remarks

In this section we discuss further the procedure introduced so far by pointing out some practical considerations and highlighting its properties along with some links with other existing methods.

Remark 1 Estimator (1) has been introduced by considering the models to be mixed $f_m(\cdot | \hat{\Psi}_m)$, as well as their number M , as given. In fact, a virtually huge number of models could be estimated, and selecting which ones should enter in the ensemble could have some impact on the resulting partitions. In practice, when choosing the ensemble size, different paths might be considered.

A first possible approach consists in populating the ensemble with all the models estimated in the previous step of the procedure, being reasonable candidates and representing a wide batch of alternatives recording a general uncertainty. In such a way the possibly troublesome selection of M is somehow circumvented by letting the penalty term to do the job. In practical applications the penalized estimation strategy would indeed shrink towards zero the weights of the models considered as irrelevant hence somehow automatically selecting M , here defined as the number of models considerably weighted in the ensemble.

Another alternative may consist in considering an *Occam's window* to choose a set of models as proposed by Madigan and Raftery (1994). The main idea is to discard those models providing estimates being qualitatively too far from the ones provided by the best model. A rule of thumb would be to discard the m th model if $|\text{BIC}_{\text{best}} - \text{BIC}_m| > 10$, where BIC_{best} and BIC_m represent respectively the values of the BIC for the best model and for the m th one. Lastly another viable approach consists in choosing M subjectively and picking those models, among the estimated ones, resulting in a good fitting of the data. In this case M should vary also reflecting the case-specific uncertainty witnessed in the modelling process.

Finally a word of caution; finding substantial arguments that motivate some general recommendations is challenging and cannot leave aside the specificities of the data and of the problem at hand.

Remark 2 The estimation procedure outlined in Sect. 2.2 is fully frequentist in nature. Alternatively, a Bayesian approach could be an interesting development claiming some advantages. The work by Malsiner-Walli et al. (2017) faces, from a Bayesian perspective, the estimation of mixtures of mixture models. Even if the underlying motivation is different some ideas could be fruitfully borrowed and exploited in order to average different mixture models. As an example, the consideration of a shrinkage prior on the weights of the models in the ensemble could practically overcome the previously discussed issue of selecting M .

Remark 3 Model selection often precedes inference that is usually conducted considering the chosen model as fixed. However, since the selection is itself data-dependent, it possesses its own variability. Drawing inference without accounting for the selection of the model corresponds to neglect completely a source of uncertainty and usually results in anti-conservative statements (Leeb and Pötscher 2005). Even in the full awareness of the fact that, in parametric clustering, the main focus typically lies on obtaining partitions rather than on inference or uncertainty quantification, we believe that a model averaging approach can entail better estimation properties and more informative confidence intervals for the parameters when needed.

Remark 4 In the supervised framework ensemble approaches have been found tremendously effective in improving predictions of a plethora of different models. For those techniques it has been frequently noticed (see, e.g. Dietterich 2000) how the concept of *diversity* is a key factor in increasing classification performances of the *base learners* that are combined. As a consequence, often weak learners are considered in the supervised context. These classifiers are highly unstable, consequently different one from the others, as they possibly focus on distinct features of the observed data. Even in a clustering framework the impact of the diversity among the combined partitions has been empirically studied and proved to be impactful by Fern and Brodley (2003) and Kuncheva and Hadjitodorov (2004).

We are aware that, when the proposed method is used to go beyond the *single best model* paradigm, the models in the ensemble cannot be considered as weak and consequently diversity among them is not achieved. Nonetheless, even if introduced with a specific aim, the proposal can in principle be exploited in all those cases where averaging multiple density-induced clusterings could be fruitful. As a consequence, the diversity can be somehow determined for example averaging densities computed on bootstrap samples or on general subsamples of the observed data. Since initialization plays a crucial role when resorting to the EM algorithm (see, e.g. Scrucca and Raftery 2015), another appealing application consists in combining models estimated using different starting values. As a consequence of the estimation instability these models would probably be more heterogeneous hence entailing greater diversity.

Remark 5 The model introduced so far, despite being based on a different rationale, shares some connections with the general framework of *Deep Gaussian Mixture Models* investigated by Viroli and McLachlan (2019). Deep Gaussian Mixture Models are

networks of multiple layers of latent variables distributed as a mixture of Gaussian densities. Since the outlined representation encompasses the specification of a mixture of mixtures (Li 2005), model (1) can be seen as a two layers Deep Gaussian Mixture Model where the parameters involved in the inner layer are fixed.

4 Empirical analysis

4.1 Synthetic data

The idea of averaging together different densities to obtain a more informative summary for clustering purposes is explored in this section via simulations. The simulation study has multiple goals. On one side we want to evaluate the performances of our proposal in terms of the quality of the produced density estimates. These performances are studied with respect to the true and known density function considering the MISE as evaluating criterion. On the other hand the clustering performances of the proposed method are investigated. As an assessment criterion we employ the *Adjusted Rand Index* (ARI, Hubert and Arabie 1985) between the obtained partitions and the true component memberships of the observations. An additional aim consists in evaluating how the sample size impacts on these comparisons.

As a side goal of the numerical explorations we want to study which penalization strategy introduced in Sect. 2.2 produces more satisfactory results. In particular, we evaluate whether the increased computational costs implied by the cross-validation are worth the effort or if less intensive strategies such as the *BIC*- and *AIC*-type penalties produce comparable results. Lastly, we want to compare our proposals with some reasonable competitors. We consider a fully parametric approach, using the single best model chosen among a set of Gaussian mixture models corresponding to combinations of the number of components and of different covariance matrix parametrizations. Moreover, we consider a nonparametric clustering method where the density is estimated using a kernel estimator with unconstrained gradient as bandwidth matrix, a standard choice in nonparametric density literature (see Chacón and Duong 2018, for a detailed tractation). The partition is afterwards practically obtained resorting to the *mean-shift algorithm* (Fukunaga and Hostetler 1975; Cheng 1995). Furthermore, we examine also an hybrid approach consisting in finding the modes, via Modal EM algorithm, of the density estimated by the single best model. The possible improvements introduced by our proposal may be due to two different motivations: the first related to a better estimation of the underlying density while the second is concerned with the modal-inspired allocation procedure. Considering an hybrid approach as a competitor can help disentangling properly these distinct sources.

A total of $B = 200$ samples have been drawn, with sizes $n \in \{500, 5000\}$, for each of the bivariate densities depicted in Fig. 2 and whose parameters are reported in “Appendix”. These densities have been considered to encompass different situations which pose different challenges from a model-based clustering perspective. The densities on the top panels of Fig. 2 represent indeed settings where the single best model is expected to achieve satisfactory results, being the data generated from Gaussian mixtures. On the other hand the densities on the bottom panels, showing

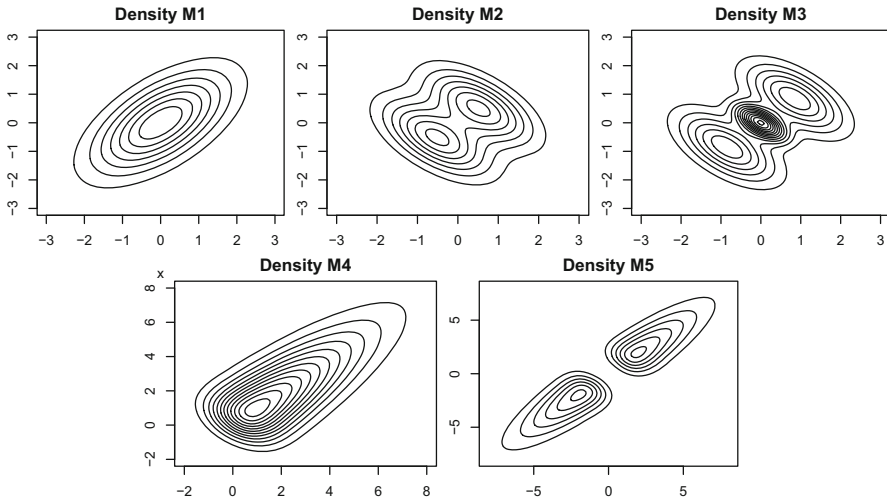


Fig. 2 Bivariate density functions used in the simulation study

strong asymmetric behaviours, constitute more challenging settings where Gaussian mixture models generally produce inadequate partitions. All the reported analyses have been conducted in the R environment (R Core Team 2019) with the aid of the `mclust` (Scrucca et al. 2016), `ks` (Duong 2019) and `EMMIXskew` (Wang et al. 2018) packages.

Throughout the simulations we estimated a total of 126 models, corresponding to the default setting in `mclust`, where the 14 different parametrizations of the component covariance matrices (see Celeux and Govaert 1995; Scrucca et al. 2016, for more details) are combined with varying number of mixture components $K = 1, \dots, 9$. Afterwards we have considered $M = 30$ best models ranked according to their BIC values, coherently with Remark 1 in Sect. 3; this choice moves towards the direction of retaining a large number of models, letting the estimation procedure to select the most relevant ones, while keeping the computations feasible. Note that some additional analyses have shown that clustering performances are not strongly influenced by a further increase in the ensemble size. Moreover the choice of the information criterion to rank the models and to select the best M ones is neither constraining nor strongly influential; here the BIC has been considered in order to be consistent with the standard practice in the model-based clustering framework. We also explored the option of selecting M by the *Occam's window* to build the ensemble as discussed in Remark 1; nonetheless results, not reported here, indicate that this strategy often leads to the selection of a too small subset of models due to the strong reliance on BIC. The three options λ_{AIC} , λ_{BIC} and λ_{CV} discussed in Sect. 2.2 are evaluated, the last one resorting to a V -fold cross-validation scheme with $V = 5$.

Results are reported in Tables 1, 2, 3, 4 and 5. A first expected behaviour indicates that the performances of the considered methods tend to improve with the increase of the sample size, both from a clustering and from a density estimation point of view.

Table 1 Top panel: the MISE ($\times 1000$) and the ARI (black lines) as functions of λ for $n = 500, 5000$. Red, purple and brown horizontal lines represent the same quantities respectively for the single best model (SB), the nonparametric approach (NP) and the hybrid approach (SB-NP). The vertical lines represent the mean values over the B samples of λ_{CV} (in green), λ_{AIC} (in light blue) and λ_{BIC} (in orange). Bottom panel: numerical values of the MISE ($\times 1000$) and average ARI (and their standard errors) over the Monte Carlo samples for the competing considered methods. Results refer to density M1 (color table online)

	$n = 500$		$n = 5000$	
	MISE	ARI	MISE	ARI
SB	0.238 (0.162)	1.000 (0.000)	0.022 (0.015)	1.000 (0.000)
NP	1.246 (0.392)	0.500 (0.501)	0.270 (0.076)	0.455 (0.499)
SB-NP	-	1.000 (0.000)	-	1.000 (0.000)
λ_{CV}	0.241 (0.164)	0.995 (0.007)	0.024 (0.016)	1.000 (0.000)
λ_{AIC}	0.465 (0.380)	0.915 (0.280)	0.029 (0.019)	1.000 (0.000)
λ_{BIC}	0.244 (0.167)	0.990 (0.099)	0.022 (0.015)	1.000 (0.000)

Generally speaking our proposal, regardless of the penalization used, produces satisfactory density estimates and partitions of the datasets. The first three scenarios have been considered to see how the ensemble approach behaves in situations where the *single best model* has a head start; in these cases the true generative model is indeed among the ones estimated in the model-based clustering routine. Even in these somewhat unfavourable settings, where in some sense an ensemble approach is not strictly needed, the proposed method behaves well producing overall comparable results with respect to the parametric ones.

In the skewed scenarios M4 and M5, where Gaussian mixture models are known to be less effective as a clustering tool, the ensemble approach induces remarkable improvements in the performances, both in terms of MISE and ARI. Note that, regarding the relation between performances and sample size, we are witnessing some results constituting an exception with respect to what we pointed out before. Indeed, especially for the setting M5, the increased availability of data points forces Gaussian mixture models to resort to an higher number of components, even in the presence of two groups, to properly model the asymmetry thus deteriorating the clustering results. In commenting these results some words of caution are needed since obtaining the allocation according to the modal concept of groups can have a strong impact in these two settings. Nonetheless comparisons with the hybrid approach help shedding light

Table 2 Cf. Table 1. Results refer to density M2 (color table online)

	$n = 500$		$n = 5000$	
	MISE	ARI	MISE	ARI
SB	0.666 (0.714)	0.677 (0.125)	0.057 (0.035)	0.680 (0.067)
NP	1.652 (0.434)	0.535 (0.186)	0.412 (0.094)	0.661 (0.083)
SB-NP	-	0.690 (0.119)	-	0.720 (0.012)
λ_{CV}	0.687 (0.402)	0.688 (0.064)	0.058 (0.036)	0.720 (0.012)
λ_{AIC}	0.809 (0.435)	0.687 (0.063)	0.072 (0.044)	0.719 (0.013)
λ_{BIC}	0.714 (0.522)	0.683 (0.129)	0.057 (0.035)	0.720 (0.012)

on this and to study further the improvements intrinsically introduced by averaging together distinct densities. The method proposed, despite showing comparable results when $n = 500$, attains notable enhancements when $n = 5000$ along with decreased standard errors. This could constitute, from a clustering standpoint, an indication of the improved quality of the density estimates produced considering model (1) with respect to the ones produced by a single mixture model; better ARI values could indeed indicate smoother estimates, being easier to be explored when searching for the modes.

The aforementioned decrease in the variability of the results of the proposal with respect to the competitors is witnessed across all the scenarios. This represents a substantial and somewhat expected advantage of the ensemble approach, since a gain in robustness and stability moves towards the desired direction when mixing models together.

With regard to the choice of the penalization scheme some different considerations arise. As expected, building on a data-based rationale, λ_{CV} seems to be more reliable when the aim is to obtain an accurate estimate of the density. Choosing the amount of the penalization via cross-validation appears to be particularly suitable especially when $n = 500$ while, with increasing sample size, the performances of the three considered schemes tend to be more similar. However, when clustering is the final aim of the analysis λ_{BIC} turns out to be a serious candidate as it often produces better results with respect to λ_{CV} and λ_{AIC} ; this constitutes a notable result since the *BIC-type* penalization, unlike the cross-validation based one, requires a null computational cost when dealing with the selection of λ . On the other hand, not even depending on

Table 3 Cf. Table 1. Results refer to density M3 (color table online)

	$n = 500$		$n = 5000$	
	MISE	ARI	MISE	ARI
SB	1.751 (1.387)	0.809 (0.099)	0.128 (0.063)	0.865 (0.008)
NP	3.753 (0.701)	0.580 (0.151)	1.071 (0.191)	0.715 (0.097)
SB-NP	-	0.804 (0.078)	-	0.831 (0.012)
λ_{CV}	1.586 (0.855)	0.790 (0.075)	0.129 (0.065)	0.830 (0.016)
λ_{AIC}	1.718 (0.837)	0.790 (0.071)	0.140 (0.064)	0.829 (0.017)
λ_{BIC}	1.759 (1.083)	0.809 (0.063)	0.128 (0.065)	0.830 (0.012)

the sample size, λ_{AIC} tends to produce the most unsatisfactory results among the three as expected.

Lastly note that the performances of the fully nonparametric approach appear not to be competitive with the other approaches considered. Nonetheless we believe that some tuning in choosing the smoothing parameters used could lead to an improvement in the results. This is not explored in our numerical experiments since appropriate bandwidth selection is not the aim of the present study, hence it appears reasonable to resort to a standard selector as we did.

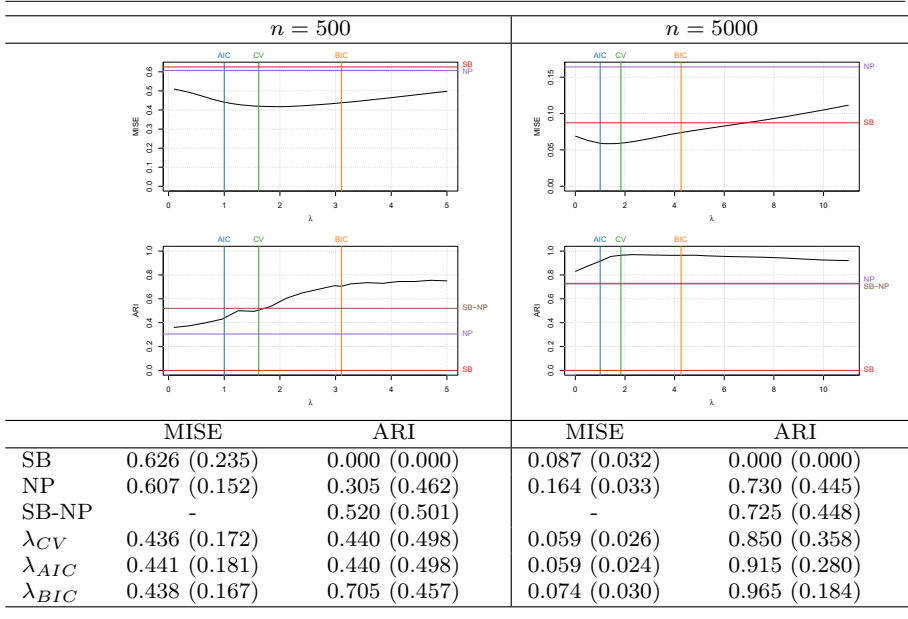
4.2 Real data

In this section we consider three illustrative examples on real datasets. As in the previous section, we fit our proposed model considering the three different penalization schemes introduced in Sect. 2.2 and we use as competitors the parametric, the nonparametric and the hybrid approaches. The number of models in the ensemble is set to $M = 30$ following the same rationale as the one discussed in the simulated examples. Not having a real density to refer to, the analyses focus on the quality of the partitions obtained, evaluated via the ARI.

4.2.1 Iris data

The *Iris* dataset (available at <https://archive.ics.uci.edu/ml/datasets/Iris>), already mentioned in Sect. 2.1 to motivate our proposal, have been thoroughly studied since the

Table 4 Cf. Table 1. Results refer to density M4 (color table online)



seminal paper by Fisher (1936) and it consists in $d = 4$ variables (sepal length and width, petal length and width) measured on $n = 150$ iris plants with $K_{true} = 3$ classes equally sized. A widely known characteristic of these data consists in having a class being linearly separable from the other two, in turn hardly to detect as separate groups. Results are shown in Table 6. The method proposed here clearly outperforms all the considered competitors. As seen in Sect. 2.1 the BIC select a two-component model hence giving wrong indications about the number of groups. As a consequence, both the parametric and the hybrid approaches, relying on the single best model, tend to produce unsatisfactory results. On the other hand the detection of 7 groups, via modal clustering based on kernel density estimation, is a symptom of an undersmoothed density estimate with the selected bandwidth matrix. Note that the high degree of rounding in the dataset could affect nonparametric performances since the estimator is built to work with continuous data, hence without duplicated values. Our method, regardless of the penalization scheme, produces strong improvements in the clustering results. The *AIC-type* and the CV-based penalties wrongly find 4 clusters with one spurious, yet small, group detected. On the contrary, a closer examination of the results reveals that λ_{BIC} assumes roughly twice the value of λ_{AIC} and λ_{CV} and leads to the correct identification of 3 groups.

4.2.2 DLBCL data

The *Diffuse Large B-cell Lymphoma* (DLBCL) dataset is provided by the British Columbia Cancer Agency (Spidlen et al. 2012; Aghaeepour et al. 2013). The sample

Table 5 Cf. Table 1. Results refer to density M5 (color table online)

$n = 500$			$n = 5000$		
	MISE	ARI	MISE	ARI	
SB	0.727 (0.236)	0.784 (0.193)	0.081 (0.030)	0.367 (0.032)	
NP	0.697 (0.128)	0.862 (0.143)	0.187 (0.025)	0.862 (0.164)	
SB-NP	-	0.910 (0.158)	-	0.829 (0.177)	
λ_{CV}	0.385 (0.124)	0.789 (0.178)	0.058 (0.014)	0.934 (0.114)	
λ_{AIC}	0.382 (0.122)	0.798 (0.175)	0.057 (0.013)	0.961 (0.091)	
λ_{BIC}	0.459 (0.139)	0.889 (0.159)	0.074 (0.019)	0.986 (0.052)	

Table 6 Results obtained on the Iris dataset. The true number of cluster is $K_{true} = 3$

	SB	NP	SB-NP	λ_{CV}	λ_{AIC}	λ_{BIC}
λ	-	-	-	1.449	1.000	2.505
ARI	0.568	0.556	0.568	0.869	0.845	0.941
\hat{K}	2	7	2	4	4	3

consists in fluorescent intensities of $d = 3$ markers, namely CD3, CD5 and CD19, measured on $n = 8183$ lymph nodes cells from subjects with a DLBCL diagnosis. A scatter plot of the data is shown in Fig. 3. In flow cytometry analysis these measurements are used to study normal and abnormal cell structures and to monitor human diseases and response to therapies. An essential step in this framework consists in obtaining a grouping of the cells according to their fluorescences. This task is usually accomplished via the so called *gating* process: the experts obtain a partition manually by visually inspecting the data. This approach is usually time-consuming and infeasible in high-dimensional situations, therefore clustering tools could come in aid to automate the gating process. The 3-dimensional structure of the data, illustrated in Fig. 3, allows us to visually inspect the true cluster configuration, displaying elongated and skewed group shapes. As noted in the simulated scenarios, results in Table 7 show how the model-based approach by using symmetric components tends to perform badly when dealing with such situations, since it detects an higher number of groups with respect to the true one. In this setting, building mixtures on more flexible, possibly skew component densities could help in improving the fit by means of a




Fig. 3 3D scatter plot of the DLBCL data with colors representing the true clustering labels

Table 7 Results obtained on the DLBCL dataset. The true number of cluster is $K_{true} = 4$

	SB	NP	SB-NP	λ_{CV}	λ_{AIC}	λ_{BIC}
λ	–	–	–	0.001	1.000	4.505
ARI	0.401	0.857	0.867	0.912	0.909	0.910
\hat{K}	7	5	4	4	4	4

single model. Conversely, the nonparametric and the hybrid approaches, which search for the modes of the density, do not suffer of the same drawbacks and outperform the parametric strategy. Nonetheless, while the former appears to undersmooth again the density, the latter detects the true number of clusters, yet with improved performance in the allocation of units.

Our proposal, regardless of the penalization scheme adopted, enjoys the very same advantage of modal clustering methods when dealing with asymmetric shapes. In fact the results obtained improve with respect to the hybrid approach thus indicating that our model produces a density estimate better tailored for the clustering scope. In this case different penalization schemes lead to irrelevant changes in the ARI values, and indicate a weaker dependency on the selected penalty value.

4.2.3 Olive oil data

As a last example we consider the *Olive oil* dataset, originally introduced in Forina et al. (1986). The data consist of $d = 8$ chemical measurements on $n = 572$ olive oils produced in 9 regions of Italy (North and South Apulia, Calabria, Sicily, Sardinia coast and inland, Umbria, East and West Liguria) that can be further aggregated in three macro-areas (Centre-North, South and Sardinia island). Clustering tools may come in aid in reconstructing the geographical origin of the oils on the basis of their chemical compositions.

Table 8 Results obtained on the Olive oil dataset. The unaggregated regions have been considered as true labels hence $K_{true} = 9$

	SB	NP	SB-NP	λ_{CV}	λ_{AIC}	λ_{BIC}
λ	–	–	–	1.155	1.000	3.175
ARI	0.782	0.604	0.792	0.902	0.902	0.892
\hat{K}	6	20	6	8	8	8

Table 9 Olive oil results, partition obtained with penalization parameter λ_{AIC}

	1	2	3	4	5	6	7	8
South								
Apulia north	24	1	0	0	0	0	0	0
Apulia south	0	6	200	0	0	0	0	0
Calabria	0	56	0	0	0	0	0	0
Sicily	6	30	0	0	0	0	0	0
Sardinia								
Sardinia inland	0	0	0	65	0	0	0	0
Sardinia coast	0	0	0	0	33	0	0	0
Centre-North								
Liguria east	0	0	0	0	0	1	42	7
Liguria west	0	0	0	0	0	0	0	50
Umbria	0	0	0	0	0	48	3	0

Compared to the cases considered previously, this example allows exploring the performances of the proposal in a moderately higher dimensional setting. Results in Table 8 show how our proposal outperforms the competitors, regardless of the penalization adopted, and how it yields a more faithful partition of the data into the 9 considered regions. The parametric and the hybrid approaches detect 6 groups, aggregating Sardinia coast and inland oils and highlighting some issues concerning the correct classification of oils produced in South macro-area. On the other hand, probably suffering of the higher dimensionality of the data, the fully nonparametric approach clearly produces a partition based on a severely undersmoothed density with 20 modes.

As it happened in Sect. 4.2.2 the clustering performances of our proposal appear to be quite insensitive to the specific penalization adopted. In Table 9 we report the partition induced by considering λ_{AIC} as penalizing parameter. Again it appears harder to discriminate the oils produced in the southern macro-area, with calabrian and sicilian ones assigned mainly to the same cluster, while oils in the other two macro-areas are substantially correctly identified.

5 Conclusions

In this work we have addressed the issue of overcoming the strong reliance of model-based clustering on a single best model, selected according to some information

criterion. Making reference to a single model may be suboptimal both for clustering and for density estimation, since alternative well-fitted models may provide useful information by uncovering different and complementary features which are otherwise discarded. It has been pointed out that possible solutions may be found in the ensemble learning literature. In this setting, we have proposed a clustering method building on a density function which averages different estimated models, and whose modal regions are then associated to the groups. The introduced density estimator is defined as a convex linear combination of the estimates of the models in the ensemble, with weights estimated via penalized maximum likelihood. This choice allows assigning relevance to the only models which better fit the data while avoiding the risk of overfitting.

The proposed methodology finds a relevant strength in the coherency not to resort to distance-based approaches to practically identify a grouping of the data. While it exhibits an hybrid taxonomy which borrows both concepts and operational tools from the parametric and nonparametric settings, it enjoys some of the relevant properties of any formulation of density-based clustering, as for example its mathematically sound formalization. In fact, it additionally inherits the strengths of both approaches. On one hand, by resorting to parametric tools and to model average, density estimation is strengthened and allows obtaining more accurate results of both nonparametric tools and single parametric models. Additionally, the introduced penalization scheme allows us to end up with a single parsimonious mixture model when the true generative mechanism underlying the data is among the candidates to be averaged. On the other hand, relying on the modal concept of clusters, groups are not constrained by any shape induced by distributional assumptions, and they can arise with arbitrary shapes which naturally comply with the geometric intuition.

The introduced method finds its original and main motivation in the need of proposing a viable alternative to the single best model paradigm in the model-based clustering framework. Nonetheless a different reading key may be given if the proposal is considered from a modal perspective. Nonparametric clustering is known to have its Achilles' heel in the density estimation step, since nonparametric estimators tend not to be reliable in moderately high-dimensional situations (Scott 2015). It has been shown (Fraley and Raftery 2002) that Gaussian mixture models, used with density estimation purposes, usually outperform kernel density estimator even in low-dimensional spaces. Our proposed method may therefore be thought as a way to strengthen the estimation process, by using an ensemble of Gaussian mixtures, with a modal clustering aims in mind.

It is worth noting that the proposed density estimator lies itself half-away between parametric and nonparametric methods. In fact, when considering the number of components as an unknown parameter, mixture models can be seen as a semi-parametric compromise between classical parametric models and nonparametric methods as kernel density estimators. The model we introduced has an increased number of components with respect to a single mixture, inherited by the averaging procedure, hence it takes another step forwards the nonparametric approach to density estimation.

The performances of the proposal have been investigated both on simulated and on real data, selected to encompass different situations and confirm the considerations above. The method produces satisfactory results both from a density estimation and from a clustering perspective, and it compares favorably with the considered com-

petitors. A deeper examination of the results leads to disentangle the reasons of the improvements into two different sources: on one side partitioning the data according to the modal formulation produces promising results in some specific scenarios, on the other hand several clues have been obtained which highlight enhancements in the density estimation process. Moreover, since nonparametric density estimation performances are known to deteriorate in high-dimensional settings, our proposal is expected to produce more pronounced improvements in these scenarios. Even if in the simulation study we explored only two-dimensional situations due to the computational burden, some analysis we have conducted, not reported here, appears to confirm this conjecture.

Concerning the introduced penalization schemes, the results seem to suggest the use of the *BIC-type* penalization, being more suitable for clustering, or of the cross-validation-based one, being able to adapt more to the features of the considered dataset. Lastly note that the use of a penalized log-likelihood, as well as the ways we considered to select the penalization parameter, targets an improvement in density estimation. While the aforementioned considerations motivate the soundness of this choice, a possible interesting research direction may consist in proposing some clustering-oriented penalization schemes and comparing them with the one introduced in this work.

Funding Open access funding provided by Università degli Studi di Padova within the CRUI-CARE Agreement.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix: Parameter settings

In the following the parameter settings of the densities selected for the simulations in Sect. 4.1 are presented. For Density M1, M2 and M3, being Gaussian mixture models, we adopt the usual notation where, for a given k component, π_k represents the k th mixture weight, μ_k and Σ_k the corresponding mean vector and covariance matrix. On the other hand, for Density M4 and M5 we consider multivariate skew normal distributions (or mixture of) hence the additional parameter δ_k regulates the skewness of the k th component (for details on the parametrization readers can refer to Azzalini and Dalla Valle 1996).

Density M1

Component	π_k	μ_k	Σ_k
1	1	$\begin{pmatrix} 0 \\ 0 \end{pmatrix}$	$\begin{pmatrix} 1.25 & 0.75 \\ 0.75 & 1.25 \end{pmatrix}$

Density M2

Component	π_k	μ_k	Σ_k
1	0.5	$\begin{pmatrix} -0.53 \\ -0.53 \end{pmatrix}$	$\begin{pmatrix} 0.68 & -0.41 \\ -0.41 & 0.68 \end{pmatrix}$
2	0.5	$\begin{pmatrix} 0.53 \\ 0.53 \end{pmatrix}$	$\begin{pmatrix} 0.68 & -0.41 \\ -0.41 & 0.68 \end{pmatrix}$

Density M3

Component	π_k	μ_k	Σ_k
1	0.4	$\begin{pmatrix} -0.85 \\ -0.85 \end{pmatrix}$	$\begin{pmatrix} 0.58 & -0.35 \\ -0.35 & 0.58 \end{pmatrix}$
2	0.4	$\begin{pmatrix} 0.85 \\ 0.85 \end{pmatrix}$	$\begin{pmatrix} 0.58 & -0.35 \\ -0.35 & 0.58 \end{pmatrix}$
3	0.2	$\begin{pmatrix} 0 \\ 0 \end{pmatrix}$	$\begin{pmatrix} 0.16 & -0.09 \\ -0.09 & 0.16 \end{pmatrix}$

Density M4

Component	π_k	μ_k	Σ_k	δ_k
1	1	$\begin{pmatrix} 0 \\ 0 \end{pmatrix}$	$\begin{pmatrix} 0.8 & -0.4 \\ -0.4 & 0.8 \end{pmatrix}$	$\begin{pmatrix} 3 \\ 3 \end{pmatrix}$

Density M5

Component	π_k	μ_k	Σ_k	δ_k
1	0.5	$\begin{pmatrix} 1 \\ 1 \end{pmatrix}$	$\begin{pmatrix} 0.8 & -0.4 \\ -0.4 & 0.8 \end{pmatrix}$	$\begin{pmatrix} 3 \\ 3 \end{pmatrix}$
2	0.5	$\begin{pmatrix} -1 \\ -1 \end{pmatrix}$	$\begin{pmatrix} 0.8 & -0.4 \\ -0.4 & 0.8 \end{pmatrix}$	$\begin{pmatrix} -3 \\ -3 \end{pmatrix}$

References

- Aghaeepour N, Finak G, Hoos H, Mosmann T, Brinkman R, Gottardo R, Scheuermann R, FlowCAP Consortium, DREAM Consortium (2013) Critical assessment of automated flow cytometry data analysis techniques. *Nat Methods* 10(3):228
- Azzalini A, Dalla Valle A (1996) The multivariate skew-normal distribution. *Biometrika* 83(4):715–726
- Banfield J, Raftery AE (1993) Model-based Gaussian and non-Gaussian clustering. *Biometrics* 49(3):803–821
- Baudry JP, Raftery AE, Celeux G, Lo K, Gottardo R (2010) Combining mixture components for clustering. *J Comput Graph Stat* 19(2):332–353
- Biernacki C, Celeux G, Govaert G (2000) Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE T Pattern Anal* 22(7):719–725
- Celeux G, Govaert G (1995) Gaussian parsimonious clustering models. *Pattern Recognit* 28(5):781–793
- Chacón JE (2019) Mixture model modal clustering. *Adv Data Anal Classif* 13(2):379–404
- Chacón JE, Duong T (2018) Multivariate kernel smoothing and its applications. Chapman and Hall/CRC, London
- Cheng Y (1995) Mean shift, mode seeking, and clustering. *IEEE Trans Pattern Anal* 17(8):790–799
- Claeskens G, Hjort N (2008) Model selection and model averaging. Cambridge University Press, Cambridge
- Dempster A, Laird N, Rubin D (1977) Maximum likelihood from incomplete data via the EM algorithm. *J R Stat Soc Ser B Stat Methodol* 39(1):1–22
- Dietterich T (2000) An experimental comparison of three methods for constructing ensembles of decision trees: bagging, boosting, and randomization. *Mach Learn* 40(2):139–157
- Duong T (2019) ks: Kernel Smoothing. R package version 1.11.4. <https://CRAN.R-project.org/package=ks>. Accessed Aug 2019
- Fern XZ, Brodley CE (2003) Random projection for high dimensional data clustering: a cluster ensemble approach. In: Proceedings of the 20th international conference on machine learning, pp 186–193
- Fisher R (1936) The use of multiple measurements in taxonomic problems. *Ann Eugen* 7(2):179–188
- Forina M, Armanino C, Castino M, Ubigli M (1986) Multivariate data analysis as a discriminating method of the origin of wines. *Vitis* 25(3):189–201
- Fraley C, Raftery AE (2002) Model-based clustering, discriminant analysis, and density estimation. *J Am Stat Assoc* 97(458):611–631
- Friedman J, Hastie T, Tibshirani R (2001) The elements of statistical learning. Springer, New York
- Fukunaga K, Hostetler L (1975) The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Trans Inform Theory* 21(1):32–40
- Glodek M, Schels M, Schwenker F (2013) Ensemble Gaussian mixture models for probability density estimation. *Comput Stat* 28(1):127–138
- Hennig C (2010) Methods for merging Gaussian mixture components. *Adv Data Anal Classif* 4(1):3–34
- Hubert L, Arabie P (1985) Comparing partitions. *J Classif* 2(1):193–218
- Kuncheva L, Hadjitodorov S (2004) Using diversity in cluster ensembles. In: 2004 IEEE international conference on systems, man and cybernetics, vol 2. IEEE, pp 1214–1219
- Leeb H, Pötscher B (2005) Model selection and inference: facts and fiction. *Econom Theory* 21(1):21–59
- Li J (2005) Clustering based on a multilayer mixture model. *J Comput Graph Stat* 14(3):547–568
- Li J, Ray S, Lindsay B (2007) A nonparametric statistical approach to clustering via mode identification. *J Mach Learn Res* 8:1687–1723
- Madigan D, Raftery AE (1994) Model selection and accounting for model uncertainty in graphical models using Occam's window. *J Am Stat Assoc* 89(428):1535–1546
- Malsiner-Walli G, Frühwirth-Schnatter S, Grün B (2017) Identifying mixtures of mixtures using Bayesian estimation. *J Comput Graph Stat* 26(2):285–295
- Menardi G (2016) A review on modal clustering. *Int Stat Rev* 84(3):413–433
- Monti S, Tamayo P, Mesirov J, Golub T (2003) Consensus clustering: a resampling-based method for class discovery and visualization of gene expression microarray data. *Mach Learn* 52(1–2):91–118
- R Core Team (2019) R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>. Accessed Aug 2019
- Rigollet P, Tsybakov A (2007) Linear and convex aggregation of density estimators. *Math Methods Stat* 16(3):260–280
- Russell N, Murphy TB, Raftery AE (2015) Bayesian model averaging in model-based clustering and density estimation. arXiv preprint [arXiv:1506.09035](https://arxiv.org/abs/1506.09035)

- Schwarz G (1978) Estimating the dimension of a model. *Ann Stat* 6(2):461–464
- Scott D (2015) *Multivariate density estimation: theory, practice, and visualization*. Wiley, New York
- Scrucca L (2016) Identifying connected components in Gaussian finite mixture models for clustering. *Comput Stat Data Anal* 93:5–17
- Scrucca L (2020) A fast and efficient modal EM algorithm for Gaussian mixtures. arXiv preprint [arXiv:2002.03600](https://arxiv.org/abs/2002.03600)
- Scrucca L, Raftery AE (2015) Improved initialisation of model-based clustering using Gaussian hierarchical partitions. *Adv Data Anal Classif* 9(4):447–460
- Scrucca L, Fop M, Murphy TB, Raftery AE (2016) mclust 5: clustering, classification and density estimation using Gaussian finite mixture models. *R J* 8(1):289
- Smyth P, Wolpert D (1999) Linearly combining density estimators via stacking. *Mach Learn* 36(1–2):59–83
- Spidlen J, Breuer K, Rosenberg C, Kotecha N, Brinkman R (2012) Flowrepository: a resource of annotated flow cytometry datasets associated with peer-reviewed publications. *Cytom Part A* 81(9):727–731
- Strehl A, Ghosh J (2002) Cluster ensembles—a knowledge reuse framework for combining multiple partitions. *J Mach Learn Res* 3:583–617
- Stuetzle W (2003) Estimating the cluster tree of a density by analyzing the minimal spanning tree of a sample. *J Classif* 20(1):025–047
- Tibshirani R, Wainwright M, Hastie T (2015) *Statistical learning with sparsity: the lasso and generalizations*. Chapman and Hall, London
- Viroli C, McLachlan G (2019) Deep Gaussian mixture models. *Stat Comput* 29(1):43–51
- Wang K, Ng A, McLachlan G (2018) EMMIXskew: the EM algorithm and skew mixture distribution. <https://CRAN.R-project.org/package=EMMIXskew>. R package version 1.0.3
- Wei Y, McNicholas PD (2015) Mixture model averaging for clustering. *Adv Data Anal Classif* 9(2):197–217

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.