



An adequacy approach for deciding the number of clusters for OTRIMLE robust Gaussian mixture-based clustering

Christian Hennig¹ and Pietro Coretto^{2,*}

University of Bologna and University of Salerno

Summary

We introduce a new approach to deciding the number of clusters. The approach is applied to Optimally Tuned Robust Improper Maximum Likelihood Estimation (OTRIMLE; Coretto & Hennig, *Journal of the American Statistical Association* **111**, 1648–1659) of a Gaussian mixture model allowing for observations to be classified as ‘noise’, but it can be applied to other clustering methods as well. The quality of a clustering is assessed by a statistic Q that measures how close the within-cluster distributions are to elliptical unimodal distributions that have the only mode in the mean. This non-parametric measure allows for non-Gaussian clusters as long as they have a good quality according to Q . The simplicity of a model is assessed by a measure S that prefers a smaller number of clusters unless additional clusters can reduce the estimated noise proportion substantially. The simplest model is then chosen that is adequate for the data in the sense that its observed value of Q is not significantly larger than what is expected for data truly generated from the fitted model, as can be assessed by parametric bootstrap. The approach is compared with model-based clustering using the Bayesian information criterion (BIC) and the integrated complete likelihood (ICL) in a simulation study and on two real data sets.

Key words: model-based clustering; noise component; parametric bootstrap; unimodality

1. Introduction

We introduce an approach for finding a suitable number of clusters for use with Optimally Tuned Robust Improper Maximum Likelihood (OTRIMLE) clustering (Coretto & Hennig 2016, 2017), which attempts to find approximately Gaussian distributed clusters allowing for some observations to be classified as noise or outliers. The approach in its general form is very flexible and can be adapted to other clustering methods and other types of clusters, but we focus on its use with OTRIMLE here. The approach is based on adequacy testing

*Author to whom correspondence should be addressed.

¹Dipartimento di Scienze Statistiche, Università di Bologna, 40126, Bologna, Italy. e-mail: christian.hennig@unibo.it

²Department of Economics and Statistics, University of Salerno, 84084 Fisciano, Italy. e-mail: pcoretto@unisa.it

[Corrections added on 30 November 2021 after the first publication: In the sixth paragraph of Introduction, the word “mclust” is changed to “tclust” to read the sentence as “These issues are acknowledged for example by the authors of the R package tclust for robust trimmed clustering (Fritz, García-Escudero & Mayo-Iscar 2012)”. Also, the last sentence of the Abstract is changed from “on real two data sets” to “on two real data sets”]

[Correction added on 20 May 2022, after first online publication: CRUI-CARE funding statement has been added.]

Acknowledgements. Open Access Funding provided by Università degli Studi di Bologna within the CRUI-CARE Agreement.

of a fitted model by using parametric bootstrap. An early forerunner of this approach is McLachlan (1987).

A key issue with choosing the number of clusters is that model assumptions never hold precisely in reality. It is therefore important that statistical methods produce reasonable results even if the model assumptions are violated. The problem with this is that what the method tries to estimate is usually defined in terms of the nominal (assumed) model, and if the model does not hold, it is not always clear what a ‘reasonable’ result would be. If clusters are supposed to be (approximately) Gaussian, using a Gaussian mixture model for clustering (Banfield & Raftery 1993) looks attractive. Estimation of the number of clusters for this is often done using the Bayesian information criterion (BIC), for example in the R package `mclust` (Scrucca *et al.* 2016). The BIC has been proven to be consistent for estimating the number of mixture components (Keribin 2000) under some rather restrictive assumptions, and is believed to be more generally consistent. This may be seen as an advantage, but is in fact a problem if the aim is to interpret the mixture components as clusters rather than just finding a good approximating Gaussian mixture for the data. A Gaussian mixture model with a sufficiently large number of mixture components can approximate almost any distribution arbitrarily well (for a recent precise version of this statement and a discussion of some older versions, see Nguyen *et al.* 2020), and in reality, clusters are not precisely Gaussian. This means that if the number of observations n becomes larger, a consistent method for estimating the number of mixture components can be expected to add mixture components in order to fit the real distribution better, and ultimately several components will fit an approximately, but not precisely, Gaussian subset of the data that intuitively would qualify as a single cluster, in turn overestimating the number of clusters. For the same reason, a likelihood ratio test will reject a single Gaussian distribution for such clusters for large n with high probability. This has also been observed in practice for the BIC (Hennig 2010). The estimation of the number of clusters is therefore affected by violations of the model assumptions in a more critical way than most standard statistical estimation problems.

The problem of choosing the number of clusters is notoriously difficult and has been treated by many authors. In the mixture context, several alternatives to the BIC have been considered. One popular alternative is the integrated completed likelihood (ICL; Biernacki, Celeux & Govaert 2000), which as the method introduced here is meant to address the issue that the BIC can choose too many mixture components for non-Gaussian clusters. Other alternatives to the BIC, include the AICmix (Hui, Warton & Foster 2015) and MSCAD (Chen & Khalili 2008), testing different model orders against each other by means of bootstrapping the likelihood ratio (McLachlan 1987; Feng & McCulloch 1996) or theoretically (Chen, Li & Fu 2012), and Bayesian approaches (Xie & Xu 2020), with more references in Frühwirth-Schnatter, Celeux & Robert (2020). An alternative approach to fit non-Gaussian clusters with Gaussian mixtures is merging of mixture components (Baudry *et al.* 2010; Hennig 2010; Malsiner-Walli, Frühwirth-Schnatter & Grün 2017). Hennig & Lin (2015) use parametric bootstraps from a null model for homogeneous data for choosing the number of clusters. Section 4.3 of Ritter (2014) reviews versions of the BIC for robust clustering with trimming.

The precise definition of outliers/noise in cluster analysis is another issue. There is ambiguity between noise and clusters in two respects. First, it is not clear how large a group of outliers has to be in order to be interpreted as a cluster on its own, and second, there may be very widely spread observations that can be well approximated by a Gaussian distribution with a very low density everywhere, but may more appropriately be interpreted as noise

than as a cluster, depending on the subject matter and the meaning of the data. Not allowing for a noise classification does not solve these issues as long as there are observations that should appropriately be interpreted as outliers; integrating them into regular clusters affects the estimation of these clusters and can also lead to misclassification of other observations.

A further issue is that to some extent more mixture components can be traded off against more flexible covariance matrices. Covariance matrices that are too flexible are already an issue for a fixed number of mixture components, because of potential degeneracy or near-degeneracy of the likelihood (García-Escudero *et al.* 2018).

The consequence is that finding an appropriate number of clusters should not be seen as a well-defined estimation problem of a statistical model. Rather it essentially requires decisions by the user: how much better approximation of the data, how much simpler covariance matrix structure that is less prone to degeneracy and what decrease of the noise proportion, would justify adding another mixture component? A method that does not require any user input such as the BIC should not be trusted, naively. These issues are acknowledged for example by the authors of the R package *tclust* for robust trimmed clustering (Fritz, García-Escudero & Mayo-Isacar 2012), who do not offer an automatic method for choosing the number of clusters, but rather some graphical displays that allow the user to track the different aspects to be traded off against each other.

On the other hand, in many situations, users do not have sufficient background knowledge to make all the required decisions in a well-founded manner, and also, for systematically evaluating the quality of an approach, automation that does not require manual adaptation to every data set is required. For this reason, we not only offer an approach that allows the user to make the required tuning decisions, but also suggest some default choices to give the user a starting point and to enable evaluation by simulation. Ultimately, an optimal tuning should depend on knowledge about the subject matter background and the aim of clustering.

The approach is based on the concept of ‘adequacy’ introduced by Davies (1995). According to this concept, a model (Davies’ use of the term ‘model’ includes specific distributions with given parameter values) is adequate for a data set, with respect to a statistic Q , if the value of Q on the data set is ‘typical’ for data sets generated by the model. This basically means that a significance test based on Q does not reject the model. The statistic Q is chosen to reflect the sense in which the model needs to ‘fit’ the data in a given application rather than following optimality considerations such as those by Neyman–Pearson; more than one test statistic can be chosen and can be combined using Bonferroni’s correction. Unless the distribution of Q on the model can be handled analytically, the parametric bootstrap can be applied to approximate this distribution. The selection of the number of clusters is a model selection problem, and Davies recommends to select the simplest model that is adequate for the data (Davies & Kovac 2001), which could be the model with the lowest number of mixture components, but see Section 4.4. Note that whenever a mixture with a low number of mixture components fits the data adequately, the data could also be fit by a model with more mixture components (one could just add small components around single observations), which means that the data actually cannot distinguish between a model with a small number of well-fitting mixture components and a model with a larger number of components, despite the fact that automatic rules such as the BIC may be interpreted by users as if this were possible. The simplest model that fits can be preferred for reasons of parsimony, avoidance of overfitting, potentially better suitability for generalisation and ease of interpretation.

Some other work comparing a model fit on data with data generated from parametric bootstrap comprises Wagenmakers *et al.* (2004), and posterior predictive checking in a Bayesian framework (Meng 1994; Gelman, Meng & Stern 1996).

The OTRIMLE method is introduced in Section 2. Section 3 gives an outline of the approach for deciding the number of clusters. This approach requires a number of decisions by the user. Section 4 contains proposals for these decisions. Particularly, a statistic Q is proposed that measures to what extent the found clusters in a data set for a given number of clusters qualify as ‘adequate’. In Section 5, we compare the method with the BIC and ICL for Gaussian mixtures, Gaussian mixtures with noise, mixtures of t and skew t distributions. Section 6 concludes the paper.

We are very happy to be invited to contribute this paper to a Special Issue in honour of Geoff McLachlan, who is a pioneer of the use of the parametric bootstrap for estimating the number of mixture components (McLachlan 1987), mixtures of t and skew t distributions and their use for accommodating outliers (Peel & McLachlan 2000; Lee & McLachlan 2013), and who contributed to making statements about the approximation of arbitrary distributions by mixtures precise (Nguyen *et al.* 2020). He has also contributed to inspiring and improving work of ours by many valuable remarks, for which we are very grateful.

2. The OTRIMLE approach to robust clustering

When using mixture models for cluster analysis, usually mixtures of families of distributions are considered that formalise the idea of a homogeneous cluster. Every mixture component is then interpreted as modelling a cluster, and the number of mixture components corresponds to the number of clusters (there are exceptions to this, see Hennig 2010).

The most popular choice for continuous data is the family of Gaussian distributions. A standard Gaussian mixture model assumes data $\mathbf{x}_1, \dots, \mathbf{x}_n$ to be generated independently identically distributed from a distribution with density

$$f(\mathbf{x}; \boldsymbol{\theta}) = \sum_{g=1}^G \pi_g \phi_p(\mathbf{x}; \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g), \quad (1)$$

where $\phi_p(\cdot; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ is the p -variate Gaussian density with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, $\pi_g \in [0, 1]$ for $j = 1, 2, \dots, G$, $\sum_{i=1}^G \pi_g = 1$, and $\boldsymbol{\theta}$ is the parameter vector collecting all $\pi_g, \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g, j = 1, 2, \dots, G$. For given G , the parameters $\boldsymbol{\theta}$ can be estimated by maximum likelihood (ML). More precisely, a global optimum is often not available, and algorithms such as the EM-algorithm are used that find a local optimum of the likelihood. Given estimators (here denoted by $\hat{\boldsymbol{\theta}}, \hat{\pi}_g, \hat{\boldsymbol{\mu}}_g, \hat{\boldsymbol{\Sigma}}_g, j = 1, 2, \dots, G$), probabilities that observations $\mathbf{x}_i, i = 1, 2, \dots, n$, were generated by mixture component g can be estimated as

$$\hat{p}_{ig} = \frac{\hat{\pi}_g \phi_p(\mathbf{x}_i; \hat{\boldsymbol{\mu}}_g, \hat{\boldsymbol{\Sigma}}_g)}{f(\mathbf{x}_i; \hat{\boldsymbol{\theta}})}, \quad (2)$$

and observation i can be assigned to the mixture component g that maximises \hat{p}_{ig} . This is implemented in the R package `mclust` (Banfield & Raftery 1993; Scrucca *et al.* 2016), along with a number of models defined by various constraints on the within-component covariance matrices. The `mclust`-approach for deciding the number of mixture components G and the covariance matrix model is to minimise the Bayesian information criterion (BIC),

$$\text{BIC} = k \ln(n) - 2 \ln(\hat{L}_n), \quad (3)$$

where k is the number of free parameters ($k = (G - 1) + pG + p(p + 1)G/2$ for a model with fully free covariance matrices) and \hat{L}_n is the maximised likelihood for the model under investigation.

It is well known that statistical methods based on a Gaussian distributional assumption can be strongly affected by outliers, and this not different in cluster analysis. For fixed G , outliers have to be included in a cluster, in turn affecting their mean and covariance matrix estimators and often the classification of many further observations. In order to deal with this, Banfield & Raftery (1993) proposed to add a so-called ‘noise component’ to the mixture in order to collect outliers and to prevent them from affecting the Gaussian clusters. The density then becomes

$$f(\mathbf{x}; \boldsymbol{\theta}) = \pi_0 \delta + \sum_{g=1}^G \pi_g \phi_p(\mathbf{x}; \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g), \quad (4)$$

$\delta \geq 0$, $\pi_0 \in [0, 1]$, and now $\sum_{i=0}^G \pi_g = 1$. They proposed to estimate the δ as $1/M$, where M is the hypervolume of the smallest hyperrectangle to cover all data, assuming that $\delta = 0$ outside that hyperrectangle. The number of clusters is still estimated by the BIC, adding the π_0 -parameter to the parameter count. Although this method often works reasonably well, it is actually not the ML estimator for δ (Coretto & Hennig 2011), and neither is it breakdown robust, because a single extreme outlier can make M arbitrarily large, preventing any other outlier from being classified as noise (Hennig 2004). The same holds for another mixture approach that is meant to be more robust than plain Gaussian mixtures, namely mixtures of t distributions (Peel & McLachlan 2000).

Hennig (2004) noted that a method with a better breakdown point can be defined by fixing δ in (1). Allowing δ to be positive on the whole Euclidean space makes f an improper density, although a proper density can be defined that constrains the noise component to occur in an unspecified set of Lebesgue measure $1/\delta$ that is assumed to cover all actually observed data. In this way, all other parameters can still be estimated using the EM algorithm, enjoy improved robustness properties and observations can still be clustered using (2). For multivariate Gaussian mixtures, this has in detail been explored by Coretto & Hennig (2016, 2017) under the name ‘robust improper maximum likelihood estimator’ (RIMLE). Coretto & Hennig (2016) propose to choose δ as

$$\arg \min_{\delta} (D(\delta) + \beta \hat{\pi}_0(\delta)), \quad (5)$$

where $D(\delta)$ is a measure of the Kolmogorov-type difference between the distribution function of within-cluster Mahalanobis distances weighted by (2) between the observations and the cluster centre, and the χ^2 -distribution function, which should be observed for perfectly Gaussian distributed observations. The weighting assigns all observations to the clusters according to the estimated probability of being generated by that cluster, which particularly means that observations that have a high estimated probability of being ‘noise’ will be downweighted. Minimising $D(\delta)$ means that δ is chosen so that the estimated clusters will look optimally Gaussian. This happens if $\beta = 0$ is chosen. The parameter β is a tuning constant that allows for tolerating more non-normality within clusters if in turn the estimated noise probability $\hat{\pi}_0(\delta)$ is decreased. Coretto & Hennig (2016) suggest $\beta = 1/3$ as alternative to

$\beta = 0$. This is particularly useful for estimating the number of clusters with clusters that are not necessarily required to be normal – see Section 5.1.

The function $D(\delta)$ can degenerate and becomes meaningless if δ is so large that all or most observations are classified as noise. Therefore, using (5) requires that the average posterior pseudo probability of observations to have been generated by the noise component is limited, and Coretto & Hennig (2017) propose an upper bound of 0.5.

Like other methods based on Gaussian mixtures, RIMLE needs to address the issue of a potentially degenerating likelihood due to covariance matrices with very small or zero eigenvalues. This is done imposing the constraint

$$\lambda_{\max}(\boldsymbol{\theta})/\lambda_{\min}(\boldsymbol{\theta}) \leq \gamma < +\infty, \quad (6)$$

where $\lambda_{\max}(\boldsymbol{\theta})$ and $\lambda_{\min}(\boldsymbol{\theta})$ are the maximum and minimum of the eigenvalues of the covariance matrices of the different Gaussian mixture components parameterised in $\boldsymbol{\theta}$, and $\gamma \geq 1$ is a constant to be chosen by the user. Based on experiments in Coretto & Hennig (2017), $\gamma = 20$ seems to be a sensible choice for standardised data (if the measurements of different variables in the data set have different orders of magnitude, there is hardly any reasonable way to specify γ), although occasionally a user may look for either more spherical clusters (which requires smaller γ) or for even more flexibility of the covariance matrices (which requires larger γ). See García-Escudero *et al.* (2018) for a comprehensive discussion of covariance matrix constraints in Gaussian mixture modelling, and particularly Section 4.1 of Ritter (2014) for robust clustering. Cerioli *et al.* (2018) argue that the choice of γ has impact on the number of clusters, and explore this for the case of a plain Gaussian mixture model.

The resulting method is called ‘Optimally Tuned RIMLE’ (OTRIMLE), and implemented in the R package `otrimle` (Coretto & Hennig 2019). Theory including consistency for the canonical functional, a breakdown point and detailed information about computation is given in Coretto & Hennig (2017). A simulation study comparing OTRIMLE with plain Gaussian mixtures and alternative robust methods is in Coretto & Hennig (2016). Likelihood-based methods, such as BIC and ICL, should not be used for estimating the number of mixture components with OTRIMLE, at least not in their original form, because the parameter δ is not chosen by ML and affects the comparison of models fitted for different numbers of components in a way not covered by likelihood-based theory.

3. An adequacy approach to decide the number of clusters

We have argued in the Introduction that the problem of finding a suitable number of clusters is essentially different from the problem of estimating the number of mixture components. Even if a Gaussian mixture model is precisely fulfilled, a ‘submixture’ of several poorly separated Gaussian components taken together can still be unimodal and even look fairly close to a single Gaussian distribution. In most applications this would qualify as a single cluster, and the number of meaningful real clusters in such a case would be smaller than the number of Gaussian mixture components.

The problem of estimating the number of Gaussian mixture components is ill-posed because any data set generated from a Gaussian mixture with a certain number of components can be arbitrarily well approximated by a mixture with more components. This particularly means that if the Gaussian mixture model assumption is not precisely fulfilled (as is always




Figure 1. Data generated from a mixture of three multivariate t_3 -distributions with clustering by Gaussian mixture fitting (left side) and Gaussian mixture fitting with noise component (right side); the number of mixture components was estimated by the BIC.

the case in reality), with enough observations a mixture with arbitrarily many components will fit the data better than a mixture with few components, even if the latter may look like an excellent representation of the intuitive clusters in the data. This is illustrated in Figure 1, which shows data generated by a mixture of three multivariate t_3 distributions (generated by the setup ‘TGauss.3I’ in Coretto & Hennig 2016). The left side shows a clustering from a plain Gaussian mixture produced by `mclust` with default settings. Although there are three elliptical clusters clearly visible, the BIC estimates the number of Gaussian mixture components as 6, because the intuitive clusters have not been generated exactly by a Gaussian distribution. Adding a uniform noise component (right side of Figure 1) classifies some outliers appropriately as ‘noise’, but does not help with the estimation of the number of clusters, as the BIC still estimates six Gaussian components. A mixture of t distributions will fit these data well with three mixture components; however, if the underlying distributions are not exactly t distributions, it runs into similar problems – see Section 5.1. A consistent method such as the BIC has more use for picking a mixture that fits the empirical density well than for interpreting the resulting components as clusters.

This implies that the problem of deciding the number of clusters is not a well-defined statistical estimation problem. It does not only rely on parameters of an assumed underlying distribution, but also on user decisions. Even assuming that the Gaussian distribution is used as a ‘cluster prototype’, that is, a cluster should look Gaussian or similar, the user has to decide:

1. what is required of a data subset to be interpreted as cluster,
2. how far from a Gaussian distribution a within-cluster distribution is tolerated to be, and
3. in case that some observations can be classified as outliers/noise, how small and homogeneous an outlying data subset is required to be in order to be interpreted as cluster rather than a group of outliers?

These decisions cannot be made from the data alone, and therefore user tuning is essential for estimating the number of clusters. We believe that this is quite generally the case in

cluster analysis, and that the vast majority of the literature ignores this, probably because most users expect a solution without having to make decisions, and a solution that depends crucially on user tuning may not be accepted as ‘objective’; see Gelman & Hennig (2017) for a discussion of this issue.

We now introduce a general scheme for deciding the number of clusters that can be applied to general model-based clustering methods, and that can be tuned by the user addressing the issues above.

The scheme is based on a general approach to model selection proposed first in Davies (1995) and more explicitly (in the context of non-parametric regression) in Davies & Kovac (2001). The idea is that one can choose the simplest model that is adequate for the data in the sense that it produces data that cannot be distinguished from typical data generated by the model. Obviously, more complex models can be adequate as well, as is the case in mixture modelling, but a more complex model will not be chosen if a simpler one exists that is already adequate. Entry points for user tuning are:

1. the target model, that is, the model for which adequacy of the data is evaluated (in cluster analysis this will often be a mixture model; here a Gaussian mixture model, as we assume that the Gaussian distribution serves as ‘cluster prototype’),
2. the statistic, or potentially more than one statistics, that are used to distinguish the data from what is expected under the model (in cluster analysis a statistic Q is required that measures whether what is interpreted as clusters behave as clusters should behave in the application at hand),
3. how atypical data has to look like in order to decide against the model (standard significance levels such as 0.01 or 0.05 may be used), and
4. the formal definition of simplicity S (in cluster analysis the standard choice would be the number of clusters, but we will penalise this with the estimated noise proportion in order to stop the method from declaring too many observations ‘noise’).

We will work with a statistic Q that does not allow for simple analytic derivation of its distribution for data generated by a mixture, and therefore its distribution will be approximated by parametric bootstrap.

Let $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})^\top \in \mathbb{R}^p$, $i = 1, \dots, n$ be the data set and $C_G(\mathbf{X})$ be the output of the clustering method C with G clusters on \mathbf{X} . Here is the general scheme:

1. Choose a target model, a clustering method that fits the target model, a statistic Q that measures clustering quality, and a statistic S measuring the simplicity of a fit. In practice also, a maximum number G_{\max} of clusters and a number of bootstrap resamples B are required.
2. For $G = 1, \dots, G_{\max}$, compute a fit (clustering) of \mathbf{X} with G clusters.
3. For $G = 1, \dots, G_{\max}$, generate B data sets $D_{G,b}$, $b = 1, \dots, B$ from the fitted model.
4. For given G , the clustering is adequate for the data if $Q(C_G(\mathbf{X}))$ is consistent with the empirical distribution of $Q(C_G(D_{G,b}))$; see Section 4.3.
5. The final number of clusters is chosen as $\arg \min_{G \text{ adequate}} S(G)$. In the simplest case $S(G) = G$, and the scheme can be stopped once an adequate G is found.

A possible outcome of the scheme is that no clustering is adequate. This is informative for the user in its own right, and means that the data are not compatible with the target model, at least not for $G \leq G_{\max}$. There are various options to enforce a clustering if it is required anyway. One could try a larger G_{\max} , choose the best found clustering according

to $C(G)$, or $((C(\mathbf{X}) - m_{QG}))/s_{QG}$ (see Section 4.3 for the definition), or try a non-model based clustering method.

4. Key decisions and tuning

We illustrate the general approach by using it for deciding the number of clusters with OTRIMLE. The number of bootstrap replications B and the maximum number of clusters G_{\max} should optimally be as large as possible, but the method is computationally intensive, so they need to be limited for pragmatic reasons. The choice of G_{\max} should also depend on potential background information about a realistic or required number of clusters. Variable B should be at least around 20 to give the method some stability, but $B = 100$ and higher would be better. The further choices are less straightforward.

4.1. Data generation from the target model

The target model in case of OTRIMLE should be a Gaussian mixture with noise, similar to (4), but (4) in the given form is not a proper probability model without constraining the set where noise (i.e. observations from mixture component zero) can occur.

With all parameters estimated by OTRIMLE and assuming the noise to be constrained to an unspecified set of Lebesgue measure $1/\delta$, the estimated posterior probability of observation \mathbf{x}_i , $i = 1, \dots, n$, to be noise is

$$\hat{p}_{i0} = \frac{\hat{\pi}_0 \delta}{f(\mathbf{x}_i; \hat{\boldsymbol{\theta}})}.$$

For data generation from the target model for the parametric bootstrap, an observation is assigned to the noise with probability $\hat{\pi}_0$, and given that it is assigned to the noise, we propose to resample it from the existing data set with the noise distribution defined by

$$\hat{\text{Pr}}_0\{\mathbf{x}_i\} = \frac{\hat{p}_{i0}}{\sum_{h=1}^n \hat{p}_{h0}},$$

so that the probability of every observation to be drawn as noise is proportional to its estimated probability to be noise in the data set. Non-noise data are generated in a standard way from the estimated Gaussian mixture.

4.2. The clustering quality statistic

The clustering quality statistic Q is meant to formalise what a ‘good’ clustering is. We do not insist on a precisely Gaussian shape, but we assume that the clusters of interest here should be elliptical and unimodal with density decreasing from the mean symmetrically in all directions. In such a case, the use of the Gaussian distribution as a cluster prototype and the Gaussian mixture approach seem justified.

The Q proposed here measures in a non-parametric way to what extent the clusters have such a shape. We start from a one-dimensional measure for a single cluster. For $p > 1$, within-cluster principal components (PCs) are computed, and the values of the one-dimensional measure are then aggregated over all PCs and over all clusters to compute the overall Q . The definition is not motivated by any model-based optimality theory, but rather custom-made

in order to express exactly what is required. It is based on a test for unimodality by Pons (2013, p. 79).

Assuming one-dimensional data standardised to have mean zero and variance one in cluster $g = 1, \dots, G$, we use the following definition:

1. Choose a kernel density estimator and q points $z_1 < z_2 < \dots < z_q$ symmetrically around the mean. Our software uses the default of the R function `density`, $q = 100$, and the 100 points are chosen as p -quantiles of the standard Gaussian distribution with p ranging from 0.005 to 0.995 in equidistant manner.
2. Compute kernel density estimators at the quantiles $\hat{f}(z_1), \dots, \hat{f}(z_q)$ based on a weighted sample in which x_{ij} has a weight according to (2).
3. Let $\hat{f}^{(1)} \geq \hat{f}^{(2)} \geq \dots \geq \hat{f}^{(q)}$ be the sorted version of $\hat{f}(z_1), \dots, \hat{f}(z_q)$.
4. For $h = 1, \dots, q/2$, let $\hat{f}^{*h} = (\hat{f}^{2h-1} + \hat{f}^{2h})/2$. This implies that $f^{*1}, f^{*2}, \dots, f^{*(q/2)}, f^{*(q/2)}, \dots, f^{*1}$ are a symmetric versions of the original $\hat{f}(z_1), \dots, \hat{f}(z_q)$.
5. Compare the symmetrised kernel density with the mean (q_l and q_r refer to the left and right side of the mean respectively):

$$q_l = \sum_{i=1}^{q/2} (\hat{f}(z_{q/2+1-i}) - \hat{f}^{*i})^2, \quad q_r = \sum_{i=1}^{q/2} (\hat{f}(z_{q/2+i}) - \hat{f}^{*i})^2.$$

$$\text{Aggregating: } \tilde{Q}_g = \sqrt{\frac{1}{q}(q_l + q_r)}.$$

The process is illustrated in Figure 2. If the estimated density already decreases monotonically and symmetrically from the mean, we have $\tilde{Q}_g = 0$, which is the best possible value.

For aggregating \tilde{Q}_g -values over different clusters, it is important to take the size of the estimated clusters, that is, \hat{n}_g , $g = 1, \dots, G$, into account in order to avoid the overall measure being dominated by a highly unreliable value from a small clusters. The rationale is not to give bigger clusters more weight, because this is about estimating the number of clusters, so small clusters that are bad should not be tolerated. However, \tilde{Q}_g can also be expected to be more variable for even valid small clusters, and this needs to be accounted for. Therefore we use

$$Q_g^* = \frac{\tilde{Q}_g - E_{\hat{n}_g} \tilde{Q}_g}{\sqrt{\text{var}_{\hat{n}_g}(\tilde{Q}_g)}},$$

where the expectation E_m and variance var_m are computed assuming m i.i.d. observations from an $N(0, 1)$ -distribution in the corresponding cluster. These values can be simulated to very high precision and interpolated to allow for non-integer m . This idea is similar to ‘pivoting’ in bootstrap inference (Hall 1992).

For p -dimensional clusters with $p > 1$, within-cluster PCs are computed first, based on the weighted within cluster data with weights according to (2) again. For $j = 1, \dots, p$, let Q_{jg} be Q_g^* computed on the j th standardised within-cluster PC of cluster g . Aggregating information from the PCs,

$$Q_g = \frac{1}{p} \sum_{j=1}^p (Q_{jg}^2) \mathbf{1}(Q_{jg} > 0),$$

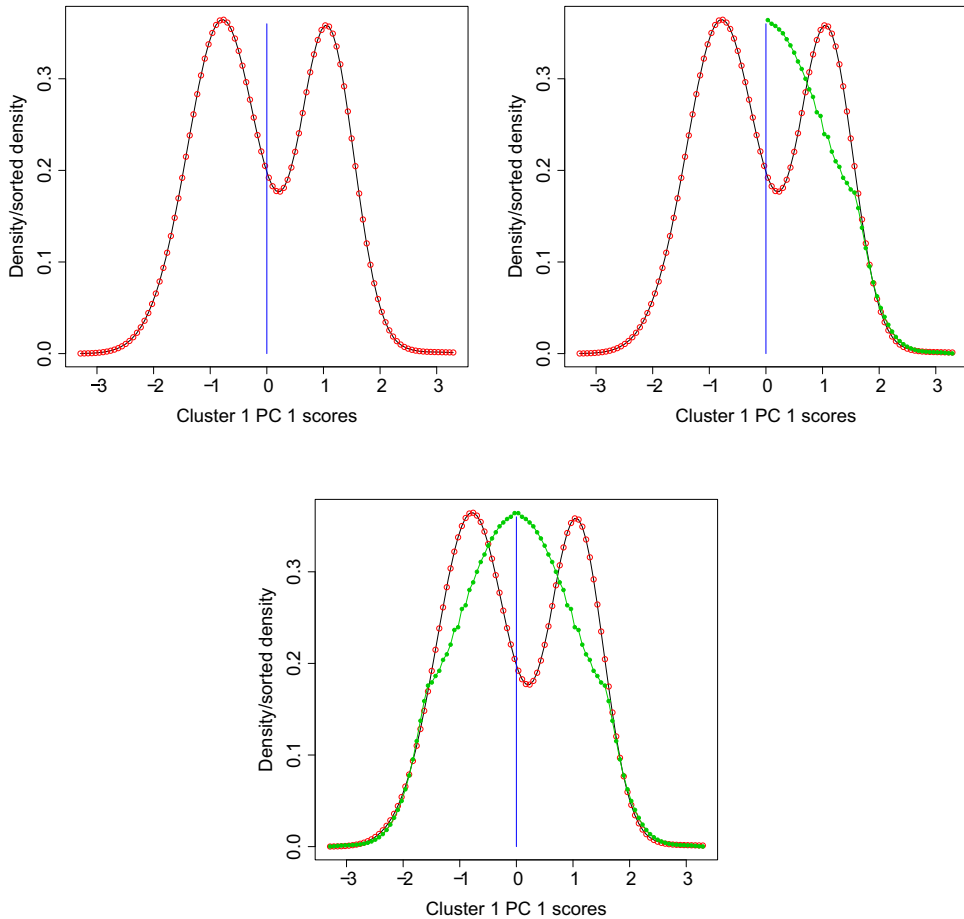


Figure 2. Illustration of the one-dimensional measurement of cluster quality. Left side: Suppose this is the kernel-estimated density for the weighted data set within the first estimated cluster at z_1, \dots, z_q , obviously not looking unimodal. Middle: Density values at z_1, \dots, z_q are ordered from the largest to the smallest. Pairs of density values (the two largest ones, then the third and fourth largest and so on) are averaged, and the resulting density values are shown on the right side of the mean at $z_{q/2+1}, \dots, z_q$ from largest to smallest. Right side: The same values are also put on the left side of the mean in descending order from the mean to the outskirts, producing a density symmetric about the mean. \tilde{Q}_g is the root of the averaged squared difference between these.

where $\mathbf{1}(\cdot)$ denotes the indicator function. The rationale here is that (a) if $Q_{jg} \leq 0$, it means that on the j th PC, the symmetric unimodality statistic behaves as expected under a Gaussian distribution or even better, so there is no indication whatsoever against this being a cluster, and (b) squaring positive Q_{jg} will emphasise problematic issues in certain PCs. The contribution of the first PCs is not upweighted in the definition of Q_g , because potential issues with unimodality are of interest along all PCs in the same way, although one could intuitively expect that issues occur more often along the first PC.

Finally, for the same reason, squares are applied when aggregating over the clusters in order to make q sensitive against substantial issues in any cluster:

$$Q(G) = \sqrt{\sum_{g=1}^G (Q_g^2)}.$$

A number of alternative choices of Q could be considered. Other tests for unimodality have been proposed (Silverman 1981; Hartigan & Hartigan 1985; Siffer *et al.* 2018). The statistic Q as defined above was chosen because it does not only measure unimodality but also symmetry (elliptic shape is measured to some extent by assessing the symmetric shape in all PC directions), and it allows for relatively straightforward aggregation over different mixture components and dimensions, because it measures deviations from the symmetric unimodal shape directly. Measuring unimodality for multivariate data is hard, and it cannot be ruled out that violations are only apparent in directions other than the PCs. Multimodality can often be expected to lead to increased variance (Siffer *et al.* 2018), and therefore the first PCs are good candidates for detecting it, but exceptions exist.

It may be possible to adapt other unimodality statistics to our approach as well. Also, Q could be defined in different ways to measure different cluster characteristics of interest. As a simple variant the symmetry requirement can be dropped by centring step 5 at the maximum estimated density mode rather than the mean, and compare the estimated data density with the un-symmetrised sorted density values on the left and right side of the mode, respectively, which is the original proposal by Pons (2013, p. 79). The statistic Q could also take into account classification entropy as the ICL does. More than one statistic can be employed at the same time to measure multiple features of the clusters (Akhanli & Hennig 2020). Elaboration of these ideas is left to future work. Section 6 has some more discussion.

4.3. Bootstrap adequacy

Because the method is computer intensive and precise quantiles may require a too large B , G will be defined to be adequate if

$$\frac{Q(G) - m_{QG}}{s_{QG}} \leq c, \quad (7)$$

where m_{QG} and s_{QG} are location and scatter statistics of the empirical distribution of $Q(G)$ for data generated from the fitted model. We have observed that with OTRIMLE (as potentially with other clustering methods) $Q(G)$ may produce outlying values. Certain fitted distributions may generate data sets that are quite ambiguous regarding the optimal clustering and the number of clusters. Such outlying values normally indicate a very bad clustering, and $Q(G)$ on the original data set should not be assessed as adequate just because certain $Q(G)$ on bootstrapped data are even worse. For this reason, m_{QG} and s_{QG} should be chosen robustly. We suggest the robust τ -estimator for location and scale (Maronna & Zamar 2002). With appropriate consistency factor, this is consistent if the parametric bootstrap distribution of $Q(G)$ is Gaussian (we currently do not have a proof for this, so this is just heuristic; Chebyshev's inequality can

be used for guidance with general distributions), allowing for a standard interpretation of the constant c .

The statistic Q is assumed to be defined so that lower values imply a better clustering quality, and adequacy will only be rejected if $Q(G)$ is too large. Choosing, for example, $c = 2$ then means that if $Q(G)$ on bootstrapped data follows a Gaussian distribution, the probability that adequacy is rejected is about 0.023.

4.4. The simplicity measure

The simplest choice for the simplicity measure S is $S(G) = G$; a model is seen as simpler if it has fewer clusters. This is appropriate for standard non-robust clustering, but it is problematic if it is allowed to classify a number of observations as ‘noise’. With OTRIMLE, as well as with trimmed clustering and the noise component in `mclust`, it would be possible to declare all observations ‘noise’ that make clustering ambiguous or belong to small clusters, in which case a high-quality clustering with small G for the remaining observations could be found easily. For this reason, and because it is generally ambiguous whether observations that belong to small groups in some distance from the bigger clusters should be declared noise or clusters on their own, too much noise should be penalised. We propose

$$S(G) = G + \frac{\hat{\pi}_0}{p_0}, \quad (8)$$

where p_0 is a constant chosen by the user. It specifies the smallest percentage of additional noise that the user is willing to trade in for adding another cluster, that is, if $p_0 = 0.05$ (which we use as a default), it means, say, that a clustering with $G = 6$ and $\hat{\pi}_0 = 0.04$ is assessed as ‘simpler’ as a clustering with $G = 5$ and $\hat{\pi}_0 = 0.1$. The former clustering will then be preferred by our method if both clusterings are adequate. Particularly, this will normally imply that clusters with $\hat{\pi}_0 < p_0$ are not found, because they could simply be declared noise and the resulting clustering would be ‘simpler’ and as adequate, although there may be exceptions in case that the smallest cluster has a high-quality Q_g compared to the other clusters.

5. Experiments

The adequacy approach to choose the number of clusters with OTRIMLE (called AOTRI in the following) is compared to different mixture model-based methods in a simulation study and on two data sets of scientific interest, one with and the other one without given true G . There is always a tension between stating that a method requires user tuning dependent on the specific situation, and running it in a default fashion on artificial data sets, but we think that both of these have their justification. Where user decisions can be used with convincing justification to adapt the method to what is required in a given application, this is certainly recommended. However, in many situations, the user does not have a clear idea how to make some or all of these choices, and therefore defaults are often useful. They are also required in order to compare the method in a ‘neutral’ fashion with others. In the following, we choose $p_0 = 0.05$ in (8), that is, we prefer a solution with one cluster more if that reduces the estimated noise by 0.05 or more. We did some experiments with $p_0 = 0.02$ (not shown), but results were rarely different. We choose $c = 2$ in (7) as maximum value of the standardised clustering quality for the model to still count as ‘adequate’. The maximum

eigenvalue ratio for covariance matrices was chosen as $\gamma = 20$. Variables were standardised before clustering in order to allow for a scale-independent interpretation of γ except where explicitly mentioned.

We looked at both $\beta = 0$ and $\beta = 1/3$ in (5), the latter called AOTRIB, and meaning that for fixed G more non-Gaussianity within clusters is tolerated if that reduced noise. Results were occasionally different. Note that β trades non-Gaussianity against noise for fixed G , whereas c tunes trading non-Gaussianity against non-adequacy of the non-noise, usually leading to a larger G (if anything changes at all, which it often does not).

We chose $G_{\max} = 10$ in the simulations. This choice does not matter, however, as long as the finally chosen G has a value of $S(G) < G_{\max} + 1$ in (8), because then it will be chosen regardless of results for higher G . As far we have seen, for all data sets, larger G_{\max} could not have changed results for this reason; for the BIC and ICL this can never be known, which is an advantage of our approach.

The number of bootstrap replicates is chosen as $B = 30$ in the computer intensive simulations, but $B = 100$ in Sections 5.2 and 5.3.

Declaration of selection bias. As this paper introduces a new method, as a proof of concept, we need to show some situations in which it works well. We looked at some other data sets and data generating mechanisms (although usually with a very small number of test runs). In many cases, there was no big difference between the different methods, and sometimes `mclust` with or without noise, or a mixture of t distributions or skew t distributions worked better, though never all of them. Sometimes nothing worked well. So we do not claim that AOTRI/AOTRIB is universally the best, just where we show it is.

5.1. Simulation study

In this study, we compare AOTRI and AOTRIB with some mixture model-based clustering methods that estimate the number of clusters using the BIC or the ICL. More precisely, we use the `R` package `mclust` with default settings for fitting a Gaussian mixture with BIC and ICL (GBIC, GICL), and with noise component (GNBIC); the noise component is initialised by the `R` function `NNClean` in package `prabclus` with parameter `nnk=5` (Byers & Raftery 1998). We use the `R` package `teigen` (Andrews *et al.* 2018) for fitting mixtures of multivariate t distributions using the BIC and ICL (TBIC, TICL). We use the `R` package `EMMIXskew` for fitting mixtures of skew t distributions (Wang, Ng & McLachlan 2009; Lee & McLachlan 2013; SKTBIC). We use fully flexible covariance matrices and degrees of freedom if possible, but sometimes `EMMIXskew` does not deliver a solution with the default settings, in which case we try out more constrained covariance matrix models as offered by `EMMIXskew` until a valid solution is found, which in the simulations ultimately always was the case. 100 data sets have been generated from each DGP.

We consider the chosen number of clusters and the adjusted Rand index comparing the resulting clustering with the true clustering (ARI; Hubert & Arabie 1985). This becomes 1 for perfect correspondence, and 0 is its expected value for comparing two random clusterings. For the AOTRI variants and GNBIC, the noise component is included as a cluster in the computation of the ARI with one exception. In a real situation, classifying observations as ‘noise’ indicates that cluster membership is unclear, and on this basis these observations could be excluded from the computation of the ARI, but this could be seen as an unfair




Figure 3. First two dimensions of data from simulated DGP 1, generated from a mixture of three multivariate Gaussian distributions.

advantage for these methods, because the other methods are assessed based on all observations including those that are hardest to classify. Therefore, we decided to include the noise in the ARI computation, although we show both results for DGP 4. A number of alternatives to the ARI have been proposed in the literature. For some experiments we also evaluated VI (Meila 2007) and BCubed (Amigo *et al.* 2009), but both Pearson and Spearman correlations between any two of the three indexes were larger than 0.96, so we present results with the ARI only here.

We simulated data from four DGPs. DGP 1 is a plain Gaussian mixture with $G = 3$ components, all with probability $\frac{1}{3}$. There are $n = 1000$ observations in $p = 10$ dimensions. The means and covariance matrices of the first two variables (see Figure 3) are: $\mu_1 = (-3, 0)^\top$, $\Sigma_1[1, 1] = \Sigma_1[2, 2] = 1$, $\Sigma_1[1, 2] = 0.5$; $\mu_2 = (8, 0)^\top$, $\Sigma_2[1, 1] = \Sigma_2[2, 2] = 2$, $\Sigma_2[1, 2] = -1.5$; $\mu_3 = (5, 9)^\top$, $\Sigma_3[1, 1] = \Sigma_3[2, 2] = 2$, $\Sigma_3[1, 2] = 1.3$. The third to tenth variable are generated from $N(0, 1)$ independently of the others. Standardising the variables implies downweighting of the cluster structure compared to the non-informative third to tenth variables, causing problems for all methods including GBIC and GICL of which the model assumptions are fulfilled here, and therefore for DGP 1 and DGP 2 variables were not standardised before clustering.

The results for DGP 1 are shown in Figure 4 (see Table 1 for the mean ARI values). GICL produces perfect results here, as does, somewhat surprisingly, SKTBIC. TBIC and TICL perform substantially worse, TICL being the better of the two. What happens here is analogous to the problem shown in Figure 1; as the clusters are not t distributed, the methods often add mixture components to approximate the Gaussian distributions better by t distributions. The remaining methods including the two AOTRI methods almost always find the correct clustering, with a few exceptions.

DGP 1 serves as a baseline for DGP 2, which is identical to DGP 1, except that one observation from cluster 1 has its value in the third variable replaced by 1000, and therefore is now a gross outlier. The results for DGP 2 are shown in Figure 5 and Table 1. The AOTRI variants produce even slightly better results than in DGP 1; the added extreme outlier may occasionally stop truly Gaussian observations from being classified as noise. The outlier also seems to regularise TBIC and TICL to some extent, although their results are still not

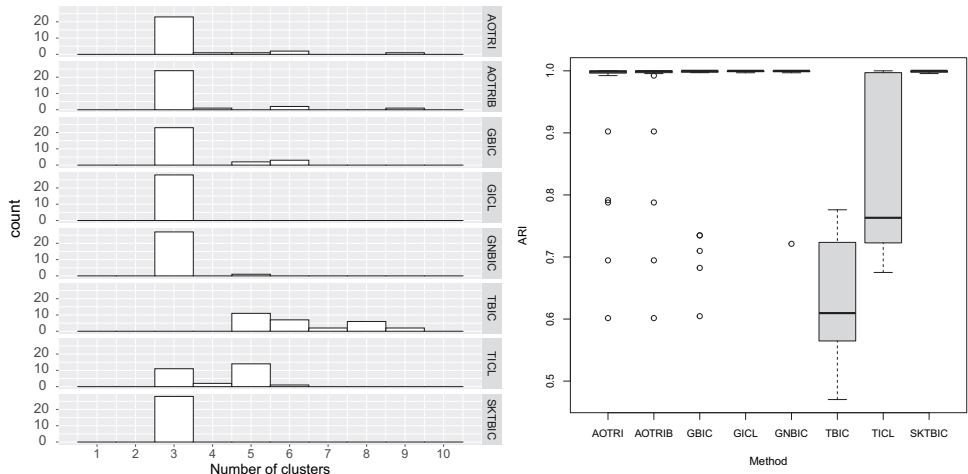


Figure 4. Left side: Distribution of numbers of clusters by method for DGP 1 (true $G=3$) over 100 simulation runs. Right side: Corresponding distributions of adjusted Rand index values.

excellent. The result of SKTBIC is still very good, typically adding a cluster for the outlier alone, and the robustness of GNBIC looks satisfactory. On the other hand, GBIC deteriorates strongly, and GICL worsens significantly as well. The effect of the outlier on these methods is in many cases that the choice of the covariance matrix model is affected and a model is used that requires more mixture components to fit the clusters.

DGP 3 with $n=2000$, $p=20$ was designed to deviate from the model assumptions in a way that does not make the clusters look strikingly different from Gaussian ones, but with some heavier tails. Again the clustering structure is present only in the first two variables, but these are now t_3 distributed; the third to twentieth variable are again standard Gaussian; outliers as occasionally generated by t_3 distributions are now in the same variables that also have the clustering structure, as opposed to DGP 1 and 2. See the supplement of Coretto & Hennig (2016) for full details. Figure 1 shows the first two variables generated by this DGP. Results are shown in Figure 6 and Table 1.

AOTRI estimates $G=3$ correctly for 88 data sets (AOTRIB for 87), and gets the clustering almost completely right in these cases, which does not hold for any of the other methods. In the other cases, they estimate either $G=2$ or $G=4$. At least in the latter case, the ARI-values are still very high. As this DGP has t as well as Gaussian distributions, the model assumptions of none of these methods is perfectly fulfilled. They estimate $G=4$ in almost all cases, adding a fourth mixture component that collects observations so that the

Table 1. Average adjusted Rand index values over 100 simulation runs. The last line gives the values for DGP 4 excluding the observations that were classified as noise.

Method	ARI	AOTRI	AOTRIB	GBIC	GICL	GNBIC	TBIC	TICL	SKTBIC
DGP 1		0.955	0.963	0.945	0.999	0.989	0.631	0.842	0.999
DGP 2		0.967	0.976	0.632	0.844	0.918	0.901	0.906	0.968
DGP 3		0.947	0.957	0.861	0.864	0.845	0.853	0.862	0.636
DGP 4		0.718	0.801	0.662	0.675	0.699	0.654	0.662	0.838
DGP 4b		0.919	0.822	0.662	0.675	0.716	0.654	0.662	0.838

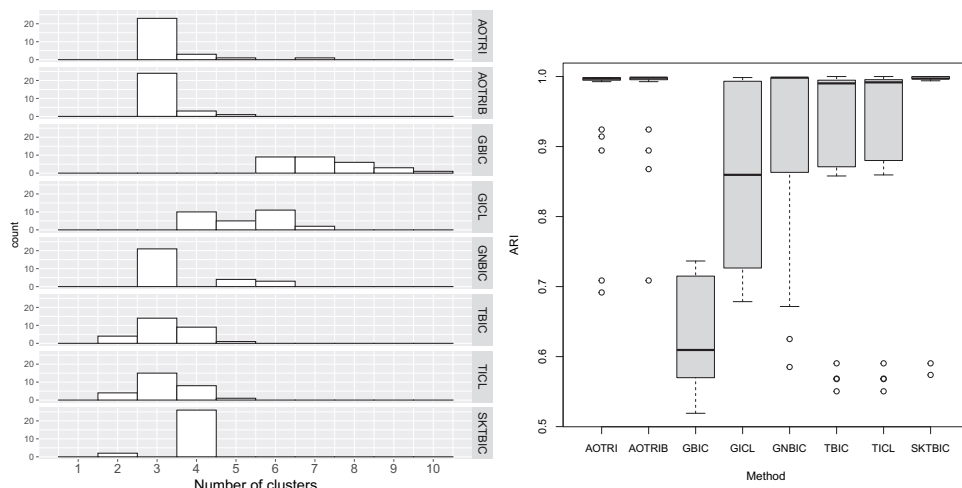


Figure 5. Left side: Distribution of numbers of clusters by method for DGP 2 (true $G=3$) over 100 simulation runs. Right side: Corresponding distributions of adjusted Rand index values.

three main components look closer to the assumptions. SKTBIC estimates $G=2$ with a worse ARI in most cases.

DGP 4 with $n=660$, $p=6$ brings together different shapes of distributions in the same data set, as is the case in some real applications. Cluster structure occurs on the first four variables, the fifth variable is standard Gaussian, the sixth is t_2 , generating some outliers. There are two Gaussian clusters with sizes 250 and 150, an independent product of exponential variables with 70 observations, a shifted multivariate t_2 distribution with 70 observations and a tight uniform with 100 observations, therefore $G=5$. There are 20 ‘true’ noise points, 10 of which are generated by a wide uniform distribution and 10 by a wider




Figure 6. Left side: Distribution of numbers of clusters by method for DGP 3 (true $G=3$) over 100 simulation runs. Right side: Corresponding distribution of adjusted Rand index values.

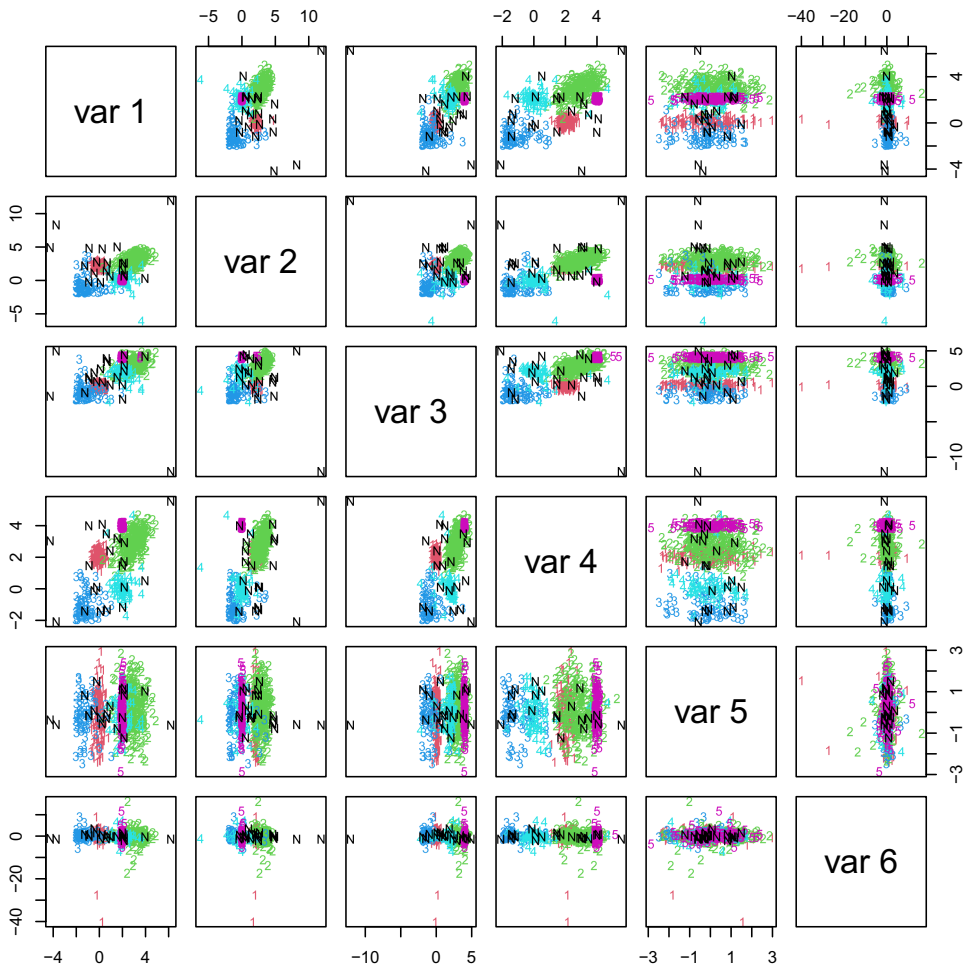


Figure 7. Data simulated from DGP 4 with true clustering. ‘N’ denotes noise, half of which was generated by a uniform and half by a t_3 , see Hennig (2007).

spread t_2 (see Figure 7). This was taken from Hennig (2007), where details are given. Only the uniform cluster was added, centred at (2, 0, 4, 4) with range 0.4 on the first four variables.

For the results, see Figure 8 and Table 1. The best performance is shown by SKTBIC regarding choosing G and also regarding the plain ARI. AOTRI and AOTRIB have a tendency to underestimate the number of clusters. This can mainly be explained by the fact that the strongly asymmetric exponential cluster is not well represented by a mode at the mean, and therefore the Q -criterion will prefer solutions that classify this as noise. This is not a proper cluster in the sense defined by Q (as long as its asymmetric version is used) and should arguably not be counted when operating with a symmetric prototype idea of a cluster. We also give ARI-results not involving the observations classified as noise in Table 1 (DGP 4b). Regarding these, AOTRI and AOTRIB perform better than the SKTBIC; if in a real application it is acceptable to not classify and interpret the observations classified as ‘noise’, AOTRI classifies the remaining observations very reliably. Regarding the number of clusters

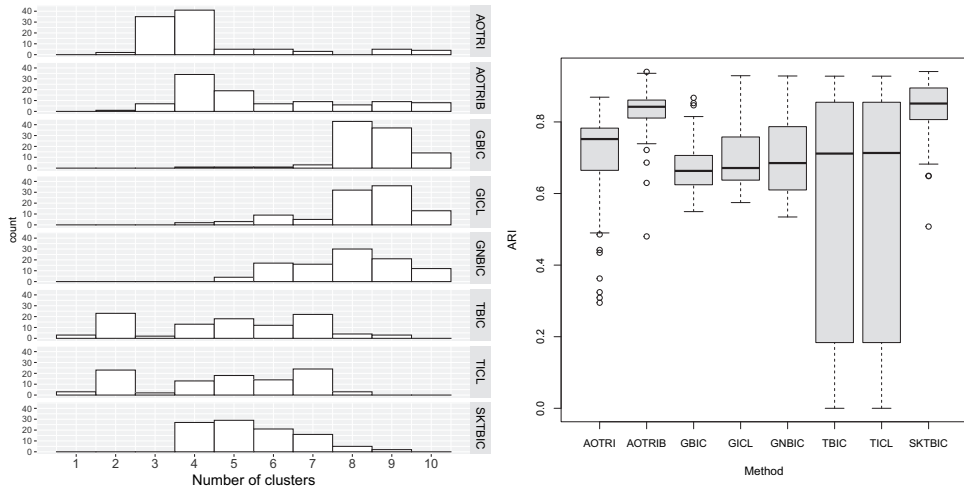


Figure 8. Left side: Distribution of numbers of clusters by method for DGP 4 (true $G=5$) over 100 simulation runs. Right side: Corresponding distribution of adjusted Rand index values.

and raw ARI, AOTRIB with $\beta = 1/3$ is almost as good as SKTBIC and clearly better than AOTRI. The latter is better when estimated noise is discounted. This is largely due to the larger estimated noise proportion: Observations that are not identified as noise by AOTRI are those that are easier to classify. GBIC, GICL and GNBIC tend to fit some non-Gaussian clusters with more than one Gaussian component, and overestimate G in this way. The ICL does not help much here. TBIC and TICL produce a large variance of the ARI and the estimated G , sometimes over- and sometimes underestimating it.

Overall AOTRI and AOTRIB show the best performance, although they are not best for every single DGP and run into occasional problems. The remaining methods all have some strength and some weaknesses; SKTBIC does very well except in DGP 3. GBIC and GNBIC overestimate G in case of non-Gaussian clusters, with deteriorating effect on the ARI. GICL does somewhat better, but does not solve the issue completely. TBIC and TICL do fairly well for DGPs 2 and 3 but are much weaker in DGPs 1 and 4.

5.2. Olive oil data

The first real data set is from Forina *et al.* (1983). The data set contains $p=8$ chemical measurements on $n=572$ different specimen of olive oil produced in $G=9$ regions in Italy (northern Apulia, southern Apulia, Calabria, Sicily, inland Sardinia and coast Sardinia, eastern and western Liguria, Umbria). It has been used several times for benchmarking of supervised and unsupervised classification methods. Interpreting the regions as the true clusters, some of them have a clearly non-Gaussian shape, and there are some outliers.

Assuming that a researcher analysing these data knows that clusters can be rather small (the smallest true cluster contains 4.4% of the observations) we analyse this using $p_0=0.02$. The AOTRI results are visualised in Figure 9. This shows that $G=3$ and $G \geq 6$ are adequate. $G=9$ has a substantially lower noise proportion than the lower numbers of clusters. It is therefore the G with the smallest $S(G)$ out of the adequate ones, and is chosen as optimal, so that G is estimated correctly. It still classifies more than 22% of the observations as noise.

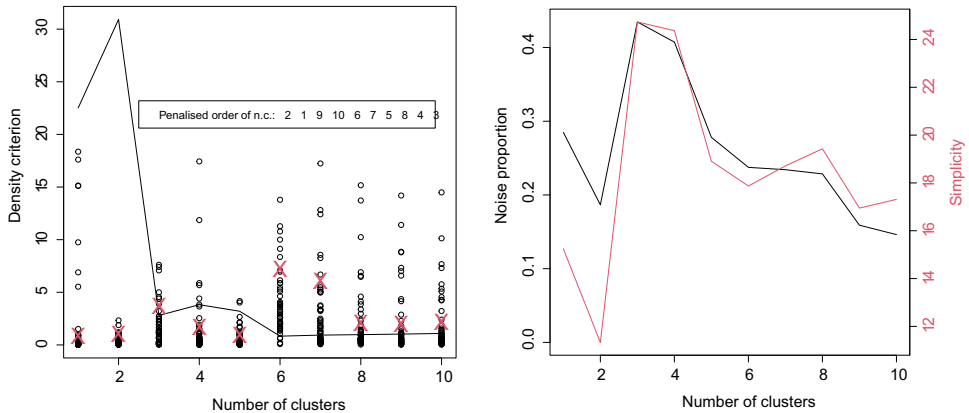


Figure 9. AOTRI results for olive oil data. Left side: Density-based clustering quality criterion $Q(G)$ for the different numbers of clusters. The connected lines refer to the clustering of the original data set, the circles to the clustering of bootstrapped data sets. The red 'X' denotes the cutoff point for a clustering to be adequate. Right side: Noise proportions, and ordering of numbers of clusters according to $S(G)$.

This seems high, but other methods have issues with correctly classifying these observations as well. The ARI between this solution and the true regions is 0.762, including the observations classified as noise; without them it is 0.930. AOTRIB estimates $G=8$ with only 5.6% of the observations classified as noise, and it achieves a better ARI of 0.808 including the noise.

GBIC and GICL both estimate the same model with $G=10$ and an ARI of 0.552 at the upper end of the range of tried out G -values. Introducing a noise component as GNBIC does improves this to 0.599 with $G=8$ including the noise, or 0.607 without it. TBIC and TICL agree on $G=5$. The ARI is rather good at 0.773. SKTBIC estimates $G=7$ and an ARI of 0.548.

5.3. Districts of the city of Dortmund

A data set characterising the 170 districts of the German city of Dortmund is presented in Sommerer & Weihs (2005). This data set does not come with true cluster labels. We used a version consisting of five sociological key variables and transformed them in such a way that fitting Gaussian distributions within clusters makes sense. The resulting variables are the logarithm of the unemployment rate ('unemployment'), the birth/death balance divided by number of inhabitants ('birth.death'), the migration balance divided by number of inhabitants ('moves.in.out'), the logarithm of the rate of employees paying social insurance ('soc.ins.emp'), and the percentage of foreigners among all unemployed and dependently employed persons ('foreigners').

Figure 10 shows that there is an extreme outlier in the scatterplot of 'birth.death' and 'moves.in.out' (called 'Romberg Park', a park district with a clinic), with some more outlier candidates. The scatterplots of 'soc.ins.emp' with 'unemployment' and 'foreigners', respectively, show some potential non-homogeneous structure.

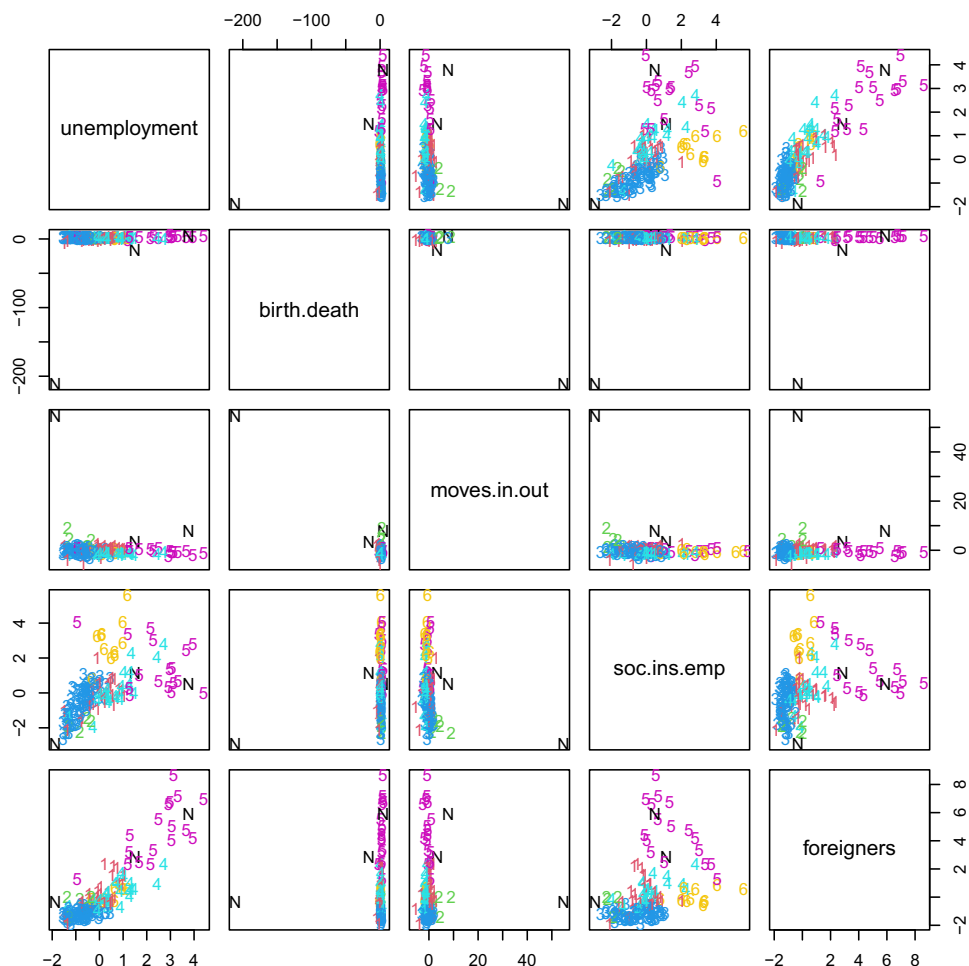


Figure 10. District of Dortmund data with clustering by AOTRI.

Figure 11 shows that the clusterings for G between 2 and 7, 9, and 10 are assessed as adequate, but that there are very high noise proportions for $G < 6$. Therefore AOTRI assesses $G = 6$ as optimal. As far as the clusters can be assessed from Figure 10, they seem sensible, with clusters 5 and 6 showing a particularly clear profile, although one may wonder whether the data could be represented by a lower number of clusters. The outliers seem well detected. Figure 10 and cluster-wise boxplots (not shown) show that the different clusters have distinct profiles in terms of the five variables. Cluster 5 has the highest values of 'unemployment', 'foreigners' and 'birth.death'. Cluster 6 has high 'social.ins.emp' and is quite homogeneous on a number of other variables. Cluster 2 is highest on 'moves.in.out' and lowest on 'social.ins.emp'. Cluster 3 is lowest on 'foreigners' and joint lowest on 'unemployment'. Cluster 4 is lowest on 'moves.in.out' and has otherwise fairly homogeneous values in the middle of the range. Cluster 1 is more difficult to interpret,

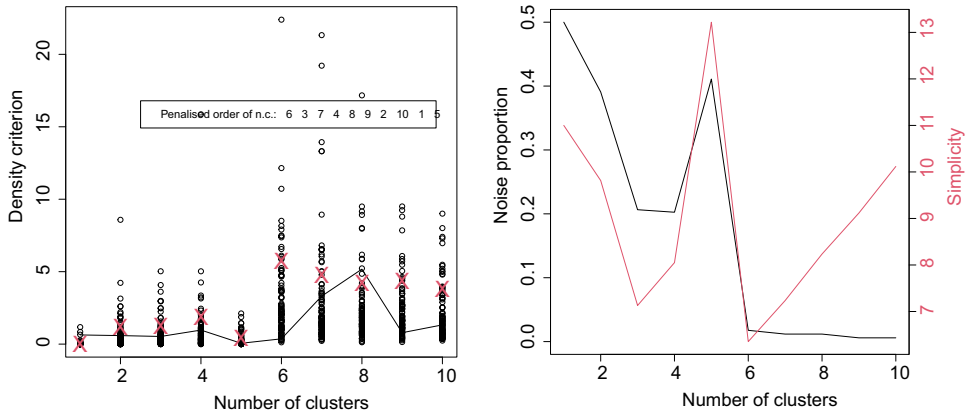


Figure 11. AOTRI results for Dortmund data. Left side: Density-based clustering quality criterion $Q(G)$ for the different numbers of clusters. The connected lines refer to the clustering of the original data set, the circles to the clustering of bootstrapped data sets. The red 'X' denotes the cutoff point for a clustering to be adequate. Right side: Noise proportions, and ordering of numbers of clusters according to $S(G)$.

with homogeneous mid-range values for 'social.ins.emp' and 'foreigners' and a large variance including the lowest values of 'birth.death'.

AOTRIB finds the same outliers as AOTRI, but prefers $G=3$ (with noise proportion 0.018, whereas AOTRI estimates this as 0.206 for $G=3$), putting the observations in the lower density region in the upper right of 'unemployment' vs. 'soc.ins.emp', together, and splits the rest into two bigger clusters, one of which is almost identical to AOTRI's cluster 3. This seems data analytically reasonable with the clusters more visibly distinct, although it encodes rougher information on the structure of the districts. GBIC and GICL choose the same model with $G=4$ as optimal. This includes a large variance cluster joining the outlier Romberg Park with some districts that have high 'foreigners' values and does not make much sense. TBIC and TICL choose $G=2$, just distinguishing the main bulk of the data from a group collecting atypical observations in various directions. SKTBIC fits the whole data set by $G=1$. All of these seem to be of little use for the understanding of the city districts. GNBIC with noise components selects a reasonable solution with $G=4$.

Overall the simulations and data examples illustrate that AOTRI/AOTRIB can give reasonable and useful results in a variety of situations in which several competitors have difficulties.

6. Conclusion

The problem of choosing the number of clusters is very difficult, particularly in applications in which observations occur that do not belong to any cluster. It is often treated as an estimation problem regarding the true number of mixture components in a parametric mixture distribution, for example a Gaussian mixture, but then clusters that make interpretative sense and are even slightly non-Gaussian are often fitted by more than one mixture component.

An appropriate decision rule for the number of clusters in a Gaussian mixture context involves a decision about what kind of non-Gaussian data subset still qualifies as a cluster.

This is formalised by our clustering quality statistic Q . The observed value of Q is compared to what is expected if data are indeed generated by a Gaussian mixture with the estimated parameter values. If an underlying distribution of a cluster has a tendency to produce better clusters than a Gaussian according to Q (which is the case for distributions such as the t distribution, for which the density goes down faster from the mean than for the Gaussian), the procedure will accept such clusters. Some users may be willing to accept certain potentially unimodal clusters even though they look somewhat worse than what is expected from the Gaussian. This could be achieved by changing the cutoff value c for adequacy to something larger, say from 2 to 3 or 4. However, this would allow for clusters that look less unimodal. Another possible modification is to re-define Q in order to allow for asymmetric clusters, although it may then be better to start with a mixture of skew distributions.

Readers may wonder whether the Gaussian mixture model is a good starting point if it is of interest to fit non-Gaussian clusters by a single mixture component. The answer is that this is appropriate if the interest is in finding clusters that are roughly Gaussian-shaped, which we define here as unimodal, and approximately elliptical. We want to avoid modelling clusters that share enough key characteristics with the Gaussian, at least approximately, by more than one mixture component, which is the reason why we do not choose Q as a likelihood ratio or a goodness-of-fit statistic for a Gaussian distribution. Furthermore in many applications, it is desirable to have a distribution with light tails as a cluster prototype distribution, because distributions with heavier tails generate observations with larger probability that are far from the main bulk of the data, and are therefore often more appropriately interpreted as outliers rather than cluster members – see McLachlan & Peel (2000, p. 231 ff.) for mixtures of t distributions.

It is ultimately up to the user to decide what kind of clusters is required in a given application. Without such decisions, the data on their own do not provide sufficient information about the clustering structure required to fit them; there are severe identifiability problems when choosing a mixture model. Most (if admittedly not all) of the required tuning of the method proposed here can be directly related to such decisions and can therefore be seen as a feature rather than a drawback.

The general adequacy approach presented here can be used for choosing the number of clusters for other clustering methods, as long as a model is given that formalises a prototype clustering structure of interest to which parametric bootstrap can be applied. The clustering quality statistic Q can be chosen in different ways, formalising other concepts of admissible clusters, or even the same concept in alternative ways. One could also attempt to select parameters such as δ and γ in this way, although this is probably more difficult due to the continuous nature of these parameters. This is left to future work.

The approach as presented here along with the accompanying plots shown in Sections 5.2 and 5.3 is implemented in the R package `otrimle`.

References

- AKHANLI, S.E. & HENNIG, C. (2020). Comparing clusterings and numbers of clusters by aggregation of calibrated clustering validity indexes. *Statistics and Computing* **30**, 1523–1544.
- AMIGO, E., GONZALO, J., ARTILES, J. & VERDEJO, F. (2009). A comparison of extrinsic clustering evaluation metrics based on formal constraints. *Information Retrieval* **12**, 461–486.
- ANDREWS, J.L., WICKINS, J.R., BOERS, N.M. & McNICHOLAS, P.D. (2018). `teigen`: an R package for model-based clustering and classification via the multivariate t distribution. *Journal of Statistical Software* **83**, 1–32.

- BANFIELD, J.D. & RAFTERY, A.E. (1993). Model-based Gaussian and non-Gaussian clustering. *Biometrics* **49**, 803–821.
- BAUDRY, J.P., RAFTERY, A., CELEUX, G., LO, K. & GOTTARDO, R. (2010). Combining mixture components for clustering. *Journal of Computational and Graphical Statistics* **19**, 332–353.
- BIERNACKI, C., CELEUX, G. & 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.
- BYERS, S. & RAFTERY, A.E. (1998). Nearest-neighbor clutter removal for estimating features in spatial point processes. *Journal of the American Statistical Association* **93**, 577–584.
- CERIOLO, A., GARCÍA-ESCUADERO, L.A., MAYO-ISCAR, A. & RIANI, M. (2018). Finding the number of normal groups in model-based clustering via constrained likelihoods. *Journal of Computational and Graphical Statistics* **27**, 404–416.
- CHEN, J. & KHALILI, A. (2008). Order selection in finite mixture models with a nonsmooth penalty. *Journal of the American Statistical Association* **103**, 1674–1683.
- CHEN, J., LI, P. & FU, Y. (2012). Inference on the order of a normal mixture. *Journal of the American Statistical Association* **107**, 1096–1105.
- CORETTO, P. & HENNIG, C. (2011). Maximum likelihood estimation of heterogeneous mixtures of Gaussian and uniform distributions. *Journal of Statistical Planning and Inference* **141**, 462–473.
- CORETTO, P. & HENNIG, C. (2016). Robust improper maximum likelihood: tuning, computation, and a comparison with other methods for robust Gaussian clustering. *Journal of the American Statistical Association* **111**, 1648–1659.
- CORETTO, P. & HENNIG, C. (2017). Consistency, breakdown robustness, and algorithms for robust improper maximum likelihood clustering. *Journal of Machine Learning Research* **18**, 1–39.
- CORETTO, P. & HENNIG, C. (2019). otrimle: Robust model-based clustering. R package version 1.3. Available from URL: <https://CRAN.R-project.org/package=otrimle>.
- DAVIES, P.L. (1995). Data features. *Statistica Neerlandica* **49**, 185–245.
- DAVIES, P.L. & KOVAC, A. (2001). Local extremes, runs, strings and multiresolution. *The Annals of Statistics* **29**, 1–65.
- FENG, Z.D. & MCCULLOCH, C.E. (1996). Using bootstrap likelihood ratios in finite mixture models. *Journal of the Royal Statistical Society: Series B (Methodological)* **58**, 609–617.
- FORINA, M., ARMANINO, C., LANTERI, S. & TISCORNIA, E. (1983). Classification of olive oils from their fatty acid composition. In *Food Research and Data Analysis*, eds. H. Martens and H. Russwurm, pp. 189–214. Barking: Applied Science Publishers.
- FRITZ, H., GARCÍA-ESCUADERO, L.A. & MAYO-ISCAR, A. (2012). tclust: An R package for a trimming approach to cluster analysis. *Journal of Statistical Software*, **47**, 1–26.
- FRÜHWIRTH-SCHNATTER, S., CELEUX, G. & ROBERT, C.P. (Eds.) (2020). *Handbook of Mixture Analysis*. Boca Raton, FL: CRC/Chapman & Hall.
- GARCÍA-ESCUADERO, L.A., GORDALIZA, A., GRESELIN, F., INGRASSIA, S. & MAYO-ISCAR, A. (2018). Eigenvalues and constraints in mixture modeling: geometric and computational issues. *Advances in Data Analysis and Classification* **12**, 203–233.
- GELMAN, A. & HENNIG, C. (2017). Beyond objective and subjective in statistics (with discussion). *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **180**, 967–1033.
- GELMAN, A., MENG, X.L. & STERN, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica* **6**, 733–807.
- HALL, P. (1992). On bootstrap confidence intervals in nonparametric regression. *The Annals of Statistics* **20**, 695–711.
- HARTIGAN, J.A. & HARTIGAN, P.M. (1985). The dip test of unimodality. *The Annals of Statistics* **13**, 70–84.
- HENNIG, C. (2004). Breakdown points for maximum likelihood estimators of location-scale mixtures. *The Annals of Statistics* **32**, 1313–1340.
- HENNIG, C. (2007). Cluster-wise assessment of cluster stability. *Computational Statistics & Data Analysis* **52**, 258–271.
- HENNIG, C. (2010). Methods for merging Gaussian mixture components. *Advances in Data Analysis and Classification* **4**, 3–34.
- HENNIG, C. & LIN, C.J. (2015). Flexible parametric bootstrap for testing homogeneity against clustering and assessing the number of clusters. *Statistics and Computing* **25**, 821–833.
- HUBERT, L. & ARABIE, P. (1985). Comparing partitions. *Journal of Classification* **2**, 193–218.

- HUI, F.K.C., WARTON, D.I. & FOSTER, S.D. (2015). Order selection in finite mixture models: complete or observed likelihood information criteria? *Biometrika* **102**, 724–730.
- KERIBIN, C. (2000). Consistent estimation of the order of mixture models. *Sankhyā: The Indian Journal of Statistics, Series A* **62**, 49–66.
- LEE, S.X. & MCLACHLAN, G.J. (2013). On mixtures of skew normal and skew t -distributions. *Advances in Data Analysis and Classification* **7**, 241–266.
- MALSINER-WALLI, G., FRÜHWIRTH-SCHNATTER, S. & GRÜN, B. (2017). Identifying mixtures of mixtures using Bayesian estimation. *Journal of Computational and Graphical Statistics* **26**, 285–295.
- MARONNA, R.A. & ZAMAR, R.H. (2002). Robust estimates of location and dispersion of high-dimensional datasets. *Technometrics* **44**, 307–317.
- MCLACHLAN, G.J. (1987). On bootstrapping the likelihood ratio test statistic for the number of components in a normal mixture. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **36**, 318–324.
- MCLACHLAN, G.J. & PEEL, D. (2000). *Finite Mixture Models*. New York: Wiley.
- MEILA, M. (2007). Comparing clusterings – an information based distance. *Journal of Multivariate Analysis* **98**, 873–895.
- MENG, X.L. (1994). Posterior predictive p -values. *The Annals of Statistics* **22**, 1142–1160.
- NGUYEN, T.T., NGUYEN, H.D., CHAMROUKHI, F. & MCLACHLAN, G.J. (2020). Approximation by finite mixtures of continuous density functions that vanish at infinity. *Cogent Mathematics & Statistics* **7**, 1750861.
- PEEL, D. & MCLACHLAN, G.J. (2000). Robust mixture modelling using the t distribution. *Statistics and Computing* **10**, 339–348.
- PONS, O. (2013). *Statistical Tests of Nonparametric Hypotheses: Asymptotic Theory*. Singapore: World Scientific.
- RITTER, G. (2014). *Robust Cluster Analysis and Variable Selection*. Monographs on Statistics and Applied Probability. Boca Raton, FL: Chapman and Hall/CRC.
- SCRUCCA, L., FOP, M., MURPHY, T.B. & RAFTERY, A.E. (2016). mclust 5: clustering, classification and density estimation using Gaussian finite mixture models. *The R Journal* **8**, 289–317.
- SIFFER, A., FOUQUE, P.A., TERMIER, A. & LARGOUËT, C. (2018). Are your data gathered? The folding test of unimodality. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. London: ACM. pp. 2210–2218.
- SILVERMAN, B.W. (1981). Using kernel density estimates to investigate multi-modality. *Journal of the Royal Statistical Society: Series B (Methodological)* **43**, 97–99.
- SOMMERER, E.O. & WEIHS, C. (2005). Introduction to the contest ‘social milieus in Dortmund’. In *Classification—The Ubiquitous Challenge*. eds. C. Weihs and W. Gaul, pp. 667–673. Berlin: Springer.
- WAGENMAKERS, E.J., RATCLIFF, R., GOMEZ, P. & IVERSON, G.J. (2004). Assessing model mimicry using the parametric bootstrap. *Journal of Mathematical Psychology* **48**, 28–50.
- WANG, K., NG, S.K. & MCLACHLAN, G.J. (2009). Multivariate skew t mixture models: applications to fluorescence-activated cell sorting data. In *2009 Digital Image Computing: Techniques and Applications*, eds. H. Shi, Y.-C. Zhang, M. Bottema, B. Lovell and A. Maeder, pp. 526–531. New York: IEEE.
- XIE, F. & XU, Y. (2020). Bayesian repulsive Gaussian mixture model. *Journal of the American Statistical Association* **115**, 187–203.