

RESEARCH ARTICLE

A fast and efficient Modal EM algorithm for Gaussian mixtures

Luca Scrucca 

Department of Economics, Università degli Studi di Perugia, Perugia

Correspondence

Luca Scrucca, Department of Economics, Università degli Studi di Perugia, Via A. Pascoli 20, 06123 Perugia, Italy.
Email: luca.scrucca@unipg.it

Funding information

Università degli Studi di Perugia, Fondo Ricerca di Base, Grant/Award Number: 2019

Abstract

In the modal approach to clustering, clusters are defined as the local maxima of the underlying probability density function, where the latter can be estimated either nonparametrically or using finite mixture models. Thus, clusters are closely related to certain regions around the density modes, and every cluster corresponds to a bump of the density. The Modal Expectation-Maximization (MEM) algorithm is an iterative procedure that can identify the local maxima of any density function. In this contribution, we propose a fast and efficient MEM algorithm to be used when the density function is estimated through a finite mixture of Gaussian distributions with parsimonious component-covariance structures. After describing the procedure, we apply the proposed MEM algorithm on both simulated and real data examples, showing its high flexibility in several contexts.

KEYWORDS

cluster analysis, density modes, finite mixture of Gaussians, Modal EM algorithm, model-based density estimation

1 | INTRODUCTION

The term cluster analysis encompasses a large set of methods and algorithms that aim at partitioning a set of data into some meaningful groups of homogeneous data points called *clusters*. The presence of such clusters is neither known a priori, sometimes even their number is unknown, nor is case labeling available. For this reason, cluster analysis is considered an instance of so-called *unsupervised learning*.

Several approaches and methods are available in the literature to explore the clustering structure of a data set [1]. Among these, density-based approaches have been proposed to exploit the relationship between the underlying density of a data set and the presence of clusters. In the parametric or *model-based clustering* approach, each

component of a mixture distribution is associated to a cluster [2, 3]. Thus, observations are allocated to the cluster with maximal weighted component density.

However, there may be situations where more than a single component is required to represent the shape of a cluster. Merging of mixture components is a possible answer to this problem. Baudry et al. [4] proposed a merging method based on an entropy criterion, while Hennig [5] discussed several methods based on unimodality and misclassification probabilities. All these methods are hierarchical in nature, so clusters can only be obtained by merging two or more mixture components. This indeed may constitute a limitation because data points assigned to a single Gaussian component cannot be subsequently allocated to different clusters. A different approach to tackle this problem was proposed by Scrucca [6] based on the

identification of connected components from high-density regions of the underlying density function.

Modal clustering is another density-based approach to clustering where clusters are taken as the “domains of attraction” of the density modes [7]. This follows the definition proposed by Hartigan [8, p. 205], according to which “clusters may be thought of as regions of high density separated from other such regions by regions of low density.” This definition of cluster is the one adopted in the paper.

Modal EM (MEM) is an iterative algorithm aimed at identifying the local maxima of a density function [9]. Let $f(\mathbf{x}) = \sum_{k=1}^G \pi_k f_k(\mathbf{x})$ be a finite mixture density for $\mathbf{x} \in \mathbb{R}^d$, where π_k is the mixing probability of component k with density function $f_k(\mathbf{x})$, under the constraints $\pi_k > 0$ for all $k = 1, \dots, G$, and $\sum_{k=1}^G \pi_k = 1$. Given an initial starting point $\mathbf{x}^{(0)}$, the following steps are iteratively executed until a stopping criterion is met:

$$\text{E-step: } p_k^{(t)} = \frac{\pi_k f_k(\mathbf{x}^{(t-1)})}{f(\mathbf{x}^{(t-1)})}, \quad \text{for } k = 1, \dots, G,$$

$$\text{M-step: } \mathbf{x}^{(t)} = \arg \max_{\mathbf{x}} \sum_{k=1}^G p_k^{(t)} \log f_k(\mathbf{x}^{(t-1)}).$$

Li et al. [9] showed that the objective function in the M-step has a unique maximum if the $f_k(\mathbf{x})$ are Gaussian densities. They also reported a closed-form solution in case of Gaussian mixtures with common covariance matrix. This is a fairly strong assumption that rarely occurs in practice, so it would be interesting to address the general case, which is not only more complex to deal with, but also much more interesting from a practical point of view.

Regarding the question of how many modes a Gaussian mixture can have, we note that Carreira-Perpiñán and Williams [10, 11] conjectured that the number of modes cannot exceed the number of components when the components of the mixture have the same covariance matrix (homoscedastic mixture), whereas if the components are allowed to have arbitrary and different covariance matrices (isotropic and full heteroscedastic mixtures) then the number of modes can be larger than the number of components. However, the first conjecture turns out to be wrong, so in general it is not possible to know a priori the number of modes of a multivariate Gaussian mixture. For a recent contribution on this issue see Améndola et al. [12], where lower and upper bounds on the maximum number of modes of a Gaussian mixture are derived under the assumption they are finite.

Modal clustering plays a central role in the nonparametric approach to cluster analysis. Several mode-seeking algorithms have been proposed in the literature, such

as the mean-shift algorithm of Fukunaga and Hostetler [13] and its many extensions [14]. However, regardless of the algorithm adopted, detection of high-density regions requires the choice of a density estimator, typically a kernel density estimator. The latter requires the selection of an appropriate kernel bandwidth, and extension to high dimensions is known to be somewhat problematic [15, ch. 9]. Interestingly, connections exist between the MEM algorithm and the mean shift algorithm. In fact, Carreira-Perpiñán [16] showed that the mean-shift algorithm is a generalized EM algorithm when the kernel of a nonparametric kernel density estimate is Gaussian. More recently, Chacón [17] extended the use of the mean shift algorithm to nonisotropic Gaussian components. For a review on nonparametric modal clustering, see Menardi [18].

1.1 | A motivating example

Consider the data shown in Figure 1(A). They represent a sample of $n = 500$ observations drawn from the following bivariate two-component mixture:

$$f(\mathbf{x}) = \pi N(\boldsymbol{\mu}_1, \Sigma_1) + (1 - \pi) \text{Skew } N(\boldsymbol{\mu}_2, \Sigma_2, \lambda_2),$$

where $\pi = 1/3$ is the mixing weight of the first Gaussian component with mean $\boldsymbol{\mu}_1 = [5 \ -2]^\top$ and covariance matrix $\Sigma_1 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$, whereas the second component is a skew-normal distribution [19] with location $\boldsymbol{\mu}_2 = [0 \ 0]^\top$, scale matrix $\Sigma_2 = \begin{bmatrix} 1 & 0.5 \\ 0.5 & 1 \end{bmatrix}$, and skew parameter $\lambda_2 = [5 \ 1]^\top$. Figure 1(B) shows the density estimate corresponding to the “best” Gaussian finite mixture model according to the Bayesian Information Criterion (BIC). The selected model is a mixture of three components with ellipsoidal covariance matrices having common orientation (VVE in `mclust` nomenclature [20]). Figure 1(C) shows the corresponding clustering partition. Clearly, observations coming from the skewed component are not correctly identified by the estimated clustering partition. Indeed, two Gaussian components are needed to adequately represent this group of observations. However, note that the corresponding density estimate seems to correctly suggest a bimodal distribution. By exploiting this fact a better partition could be obtained, and the method discussed in this paper aims to deal with similar situations.

In this contribution, we propose a fast and efficient MEM algorithm for identifying the modes of a density estimated by finite mixture of multivariate Gaussians having any of the parsimonious covariance structures available in the `mclust` R package [20]. The outline of this article is

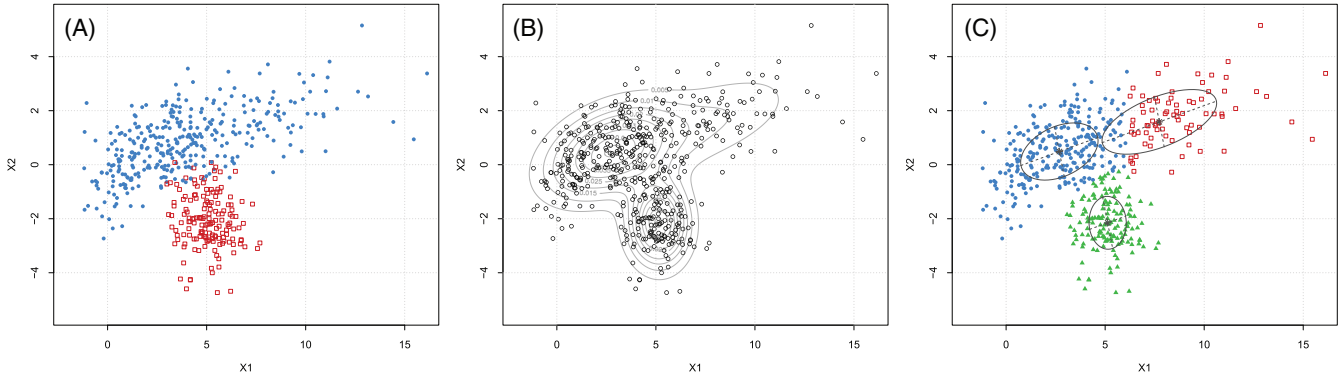


FIGURE 1 Plots of a two-component simulated data: (A) data points marked according to the true component memberships; (B) density contours obtained from the estimated Gaussian mixture model; (C) clustering and ellipses corresponding to the estimated Gaussian components of the mixture

as follows. Section 2 provides a brief review of the MEM approach for Gaussian mixtures available in the literature. Section 3 contains the proposal for extending the Modal clustering approach to any density estimated by fitting a finite mixture of Gaussian distributions with parsimonious component-covariance structures, and details on how to improve the computational efficiency of this approach. Section 4 describes the empirical results deriving from the application of the proposed MEM algorithm to examples using both synthetic and real data sets. The final section provides some concluding remarks.

2 | MEM ALGORITHM FOR GAUSSIAN MIXTURES

Gaussian mixture models (GMMs) assume that the mixture components are all multivariate Gaussians with mean μ_k and covariance Σ_k , that is, $f_k(\mathbf{x}) \equiv \phi(\mathbf{x}; \mu_k, \Sigma_k)$. Therefore, the mixture density for any data point \mathbf{x}_i can be written as

$$f(\mathbf{x}_i) = \sum_{k=1}^G \pi_k \phi(\mathbf{x}_i; \mu_k, \Sigma_k).$$

Clusters described by a GMM are centered at the means μ_k , and with other geometric characteristics (such as volume, shape, and orientation) determined by the covariance matrices Σ_k . These can be controlled by introducing some constraints on the covariance matrices through the following eigen-decomposition [21, 22]:

$$\Sigma_k = \lambda_k \mathbf{U}_k \Delta_k \mathbf{U}_k^\top, \quad (1)$$

where $\lambda_k = |\Sigma_k|^{1/d}$ is a scalar which controls the *volume*, Δ_k is a diagonal matrix, such that $|\Delta_k| = 1$ and with the normalized eigenvalues of Σ_k in decreasing order, which controls the *shape*, \mathbf{U}_k is an orthogonal matrix of

eigenvectors of Σ_k which controls the *orientation*. In this way, a total of 14 GMMs are obtained [20].

It is important to note that in this paper we shall consider the mixing proportions π_k , the mean vectors μ_k , and the covariance matrices Σ_k as fixed (either estimated or known a priori) for all $k = 1, \dots, G$.

The MEM algorithm starts with $t = 0$ and initial data point $\mathbf{x}_i^{(0)} = \mathbf{x}_i$. At iteration t , MEM performs the following steps:

- Set $t = t + 1$.
- E-step—Update the posterior conditional probability of the current data point \mathbf{x}_i to belong to the k th mixture component:

$$z_{ik}^{(t)} = \frac{\pi_k \phi(\mathbf{x}_i^{(t-1)}; \mu_k, \Sigma_k)}{\sum_{g=1}^G \pi_g \phi(\mathbf{x}_i^{(t-1)}; \mu_g, \Sigma_g)},$$

for all $k = 1, \dots, G$.

- M-step—Update the current value of \mathbf{x}_i by solving the optimization problem:

$$\mathbf{x}_i^{(t)} = \arg \max_{\mathbf{x}_i} \sum_{k=1}^G z_{ik}^{(t)} \log \phi(\mathbf{x}_i^{(t-1)}; \mu_k, \Sigma_k).$$

- Iterate the above steps until a stopping criterion is satisfied, for instance, $\max\{|\mathbf{x}_i^{(t)} - \mathbf{x}_i^{(t-1)}| / (1 + |\mathbf{x}_i^{(t-1)}|)\} < \epsilon$, where ϵ is a tolerance value, say $\epsilon = 1e-5$, or a pre-specified maximum number of iterations is reached.

By the ascending property of the MEM algorithm [9, appendix A], at convergence the value $\mathbf{x}_i^{(t)}$ is the mode associated with data point \mathbf{x}_i . Li et al. [9] presented a closed-form solution only in the specific case of Gaussian mixtures with common covariance matrix, and reported that numerical procedures are required for the M-step if

the covariance matrices are different across components. By replicating the above algorithm for all data points, it is possible to identify the modes associated with any \mathbf{x}_i ($i = 1, \dots, n$), but this process is time-consuming even for moderately large data sets. In the next section, we present an approach aimed at accelerating the MEM algorithm by iterating simultaneously for all data points and for any parsimonious covariance matrix decomposition.

3 | PROPOSAL

In this section, we detail our proposal to speed up the MEM algorithm for Gaussian mixtures having any of the parsimonious component-covariance matrix eigen-decomposition proposed by Banfield and Raftery [21], Celeux and Govaert [22], and implemented in the `mclust` package [20] for R [23].

To this end, we start by noting that, the objective function in the M-step presented in Section 2 can be written as

$$Q(\mathbf{x}_i) = \sum_{k=1}^G z_{ik} \log \phi(\mathbf{x}_i; \boldsymbol{\mu}_k, \Sigma_k).$$

The gradient and Hessian of this function with respect to the observed vector \mathbf{x}_i (again, assuming the mixture parameters $\{\pi_k, \boldsymbol{\mu}_k, \Sigma_k\}_{k=1}^G$ as known and fixed) are, respectively,

$$\nabla Q(\mathbf{x}_i) = - \sum_{k=1}^G z_{ik} \Sigma_k^{-1} (\mathbf{x}_i - \boldsymbol{\mu}_k),$$

and

$$\nabla^2 Q(\mathbf{x}_i) = - \sum_{k=1}^G z_{ik} \Sigma_k^{-1}.$$

Because all covariance matrices Σ_k are positive definite by definition, and $z_{ik} > 0$ for all k and i , the Hessian is negative definite. Thus, maximization of the Q -function can be pursued by equating the gradient to zero, and then solving for \mathbf{x}_i we obtain

$$\mathbf{x}_i^* = \left(\sum_{k=1}^G z_{ik} \Sigma_k^{-1} \right)^{-1} \sum_{k=1}^G z_{ik} \Sigma_k^{-1} \boldsymbol{\mu}_k. \quad (2)$$

Note that the last equation also arises in other mode-seeking procedures, such as in the gradient-quadratic algorithm and fixed-point iterative algorithm proposed by Carreira-Perpiñán [24], and the mean shift algorithm proposed by Chacón [17].

A straightforward application of Equation (2) requires to replicate the procedure for all data points. This can be time-consuming because it repeatedly involves calculating

matrix products and inversion of matrices. However, these objects can be efficiently computed in a single pass for all data points through the use of the Kronecker product.

Let \mathbf{z}_k be the vector of length n containing the posterior probabilities of all data points $\{\mathbf{x}_i\}_{i=1}^n$ to belong to the k th mixture component, and $\boldsymbol{\mu}_k$ be the vector of length d of component means ($k = 1, \dots, G$). Define the $(nd \times d)$ matrix

$$\mathbf{A} = \sum_{k=1}^G \mathbf{z}_k \otimes \Sigma_k^{-1},$$

and the $(nd \times 1)$ vector

$$\mathbf{b} = \sum_{k=1}^G \mathbf{z}_k \otimes \Sigma_k^{-1} \boldsymbol{\mu}_k.$$

Solutions for each $\{\mathbf{x}_i\}_{i=1}^n$ can be obtained by solving the linear systems

$$\mathbf{A}_I \mathbf{x}_i^* = \mathbf{b}_I,$$

where $I \equiv \{(i-1)d+1, \dots, id\}$ is the set containing the indices used to select the rows of matrix \mathbf{A} and the elements of vector \mathbf{b} . Equivalently, solutions of the linear systems can be written as $\mathbf{x}_i^* = \mathbf{A}_I^{-1} \mathbf{b}_I$.

Compared to the approach based on the calculation of the solution for each data point as in Equation (2), the main advantage of our proposal is that computing the matrix \mathbf{A} and vector \mathbf{b} is performed in a single step for all the observations. Then, the use of the indices I allows us to select the relevant parts of \mathbf{A} and \mathbf{b} for computing the solutions. Although algebraically equivalent, this approach turns out to be three times faster in our experiments under different settings.

The above algorithm is fast and efficient, but in practice MEM can suffer from some drawbacks which can be easily addressed as discussed below.

3.1 | Setting the step size

Large jumps can occur during the initial iterations of the algorithm for those data points \mathbf{x}_i 's located in low-density regions. In these cases, since most z_{ik} 's are very small, the inverse of $\sum_{i=1}^G z_{ik} \Sigma_k^{-1}$ in Equation (2) will contain large values (in magnitude). As a consequence, an initial data point would shift from the region of the attracting mode to the domain of attraction of a different mode, and therefore to converge to a mode different from the attractor of the data point. As an example, consider the data point at the bottom-right of Figure 2(A), and the corresponding path of MEM iterations (see the red arrows). In this case, after an initial large jump, the algorithm converges to a mode further from the domain of attraction of the data point.

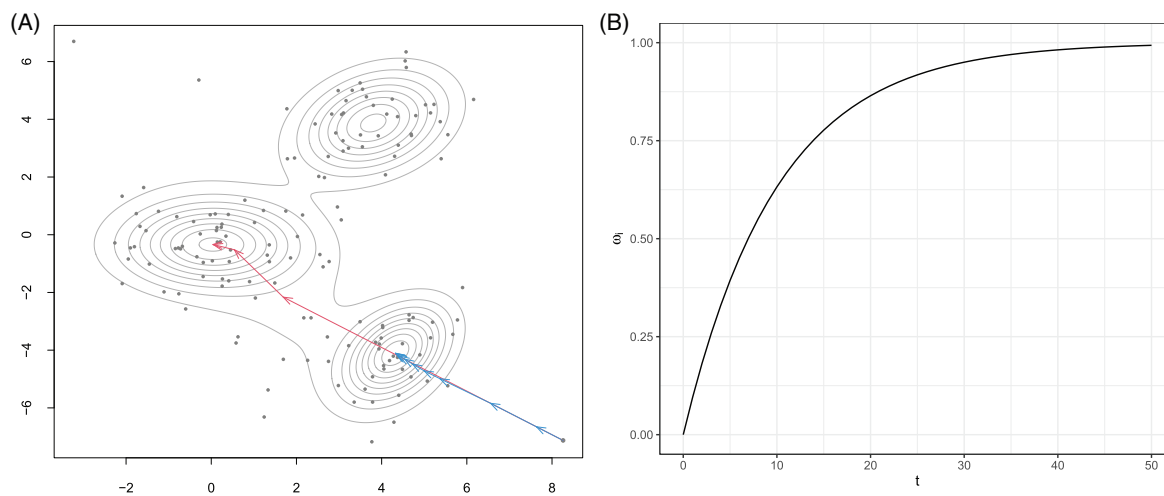


FIGURE 2 Synthetic example with paths of Modal EM (MEM) iterations for a low-density data point with (blue arrows) and without (red arrows) the proposed step size applied (A). Plot of step size as a function of the number of MEM iterations (B)

To avoid these situations, we may compute the update at iteration t as the convex linear combination of the solution at previous step and the proposed value as follows:

$$\mathbf{x}_i^{(t)} = (1 - \omega_i)\mathbf{x}_i^{(t-1)} + \tilde{\omega}_i\mathbf{x}_i^*,$$

where ω_i is a tuning parameter that controls the step size. The definition of ω_i should consider whether or not a data point lies in a low-density region, and update this value at each iteration. For instance, we could define a function that depends on $\delta_i = \left| \sum_{k=1}^G z_{ik} \Sigma_k^{-1} \right|$ for $i = 1, \dots, n$, so when the determinant δ_i is small, that is, the corresponding data point lies in a low-density region, the value of ω_i should be close to zero and the value $\mathbf{x}_i^{(t)}$ should be updated by small steps, whereas for relatively large values of δ_i the associated weights ω_i converge to one, so essentially setting $\mathbf{x}_i^{(t)} = \mathbf{x}_i^*$. However, implementing such strategy would require the computation of δ_i at each iteration for all the data points, and this may result in a significant increase of the execution time.

For this reason, in practice, we suggest to compute the step size as $\omega_i = 1 - \exp\{-0.1t\}$ (see the function drawn in Figure 2(B)). The idea is that at earlier iterations, the step size is small and $\mathbf{x}_i^{(t)}$ must be updated by small steps, but as the number of iterations increase the step size converges to one, so the updated value becomes almost equivalent to \mathbf{x}_i^* . Returning to the example given in Figure 2(A), smaller initial steps (shown as blue arrows) allow to converge to the correct density mode. Finally, we note that such a strategy inevitably increases the number of steps performed by the MEM algorithm. However, since each step is quite fast to perform, overall the algorithm is not significantly affected. For instance, in the previous example, the MEM iterations for all data points increase from 7 to 18, while the execution time from 0.04 to 0.08 s.

3.2 | Connected-components algorithm for “tight clusters”

After the final iteration of the MEM algorithm a set of points $\{\mathbf{x}_i^*\}_{i=1}^n$ are obtained. These represent the modes to which each of the data points converge. However, in the limit, points that would converge to the same mode may be numerically different from each other by a small amount, whose magnitude depends on the tolerance value used for checking the convergence of the algorithm. Thus, solutions $\{\mathbf{x}_i^*\}_{i=1}^n$ form tight clusters around the corresponding modes, widely separated from other tight clusters corresponding to different modes. The connected-components algorithm described in Carreira-Perpiñá [14] allows for the merging of those points that ideally would be identical. This can be applied as a postprocessing step to obtain the final estimated modes $\{\hat{\mathbf{x}}_m\}_{m=1}^M$.

3.3 | Denoising low-density modes

Certain regions of the features space may lack of sufficient data points to obtain reliable density estimates, particularly in high-dimensional features space. As a consequence, modes located in such regions might be spurious and, in these cases, it may be convenient to filter out these modes [24]. We consider a simple rule to drop modes associated with regions of relatively low probability. Following the approach of Banfield and Raftery [21], we postulate the presence of a noise component uniformly distributed over the data region. Let V be the hypervolume of the data region, so each log-density value of a mode not exceeding $-\log(V)$ can be considered as a noisy artifact of the density estimation process. Here, the logarithmic scale is used to improve stability and numerical accuracy.

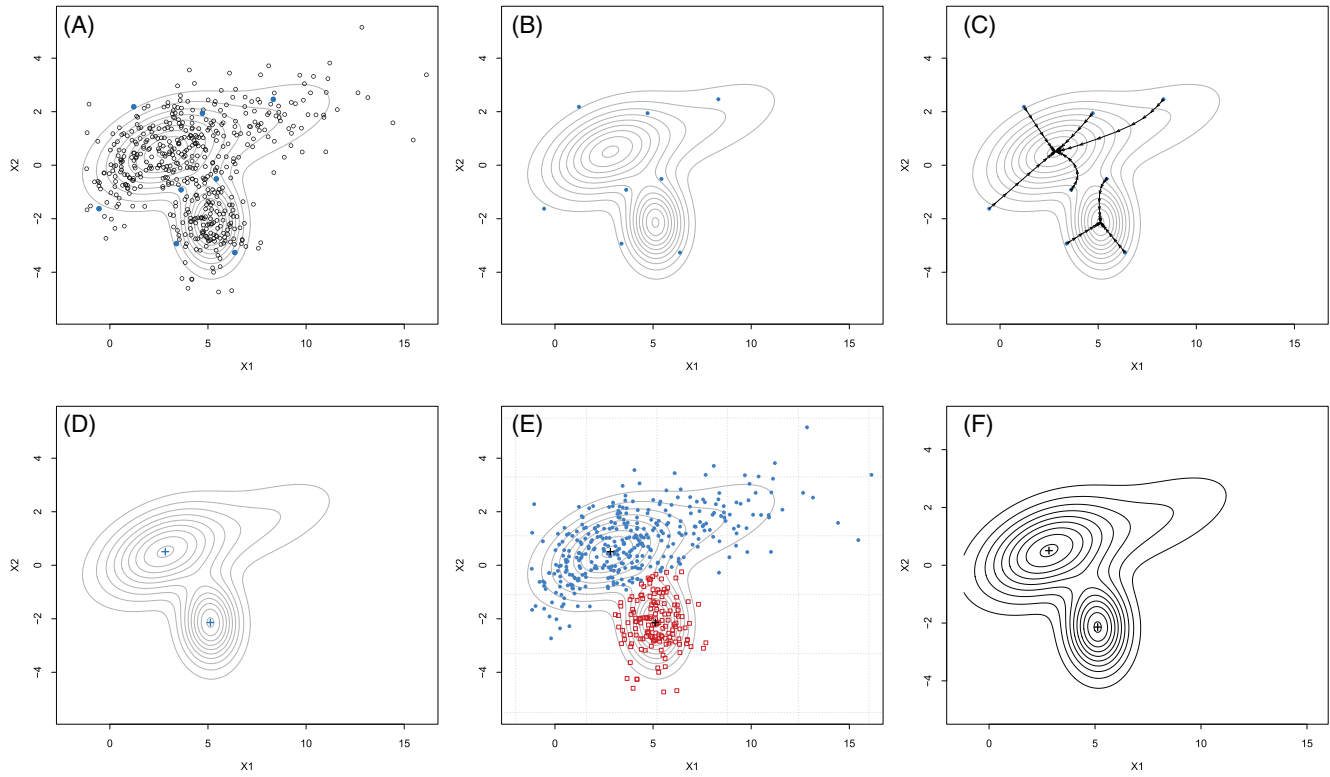


FIGURE 3 Plots for the simulated data: (A,B) density contours obtained from the select Gaussian mixture model (GMM), with some data points highlighted; (C) paths of Modal EM (MEM) algorithm for the selected data points; (D) modes found by the MEM algorithm; (E) modal clustering solution; (F) modal clustering partition of the feature space

In practice, we need to compute $\log(V)$, and a simple approximation could be obtained by taking the minimum among: (i) the volume of hyperbox containing the observed data; (ii) the volume of the hyperbox obtained from principal component scores; (iii) the volume of ellipsoid hull, that is, the ellipsoid of minimal volume such that all data points lie inside or on the boundary of the ellipsoid. Alternatively, the central $(1 - \alpha)100\%$ region of a multivariate Gaussian distribution, that is, the smallest region such that an observation falls in this region with probability $(1 - \alpha)$, can be computed. This region is an ellipsoid in d dimensions, with log-hypervolume equal to

$$\log(V) = \log(2) + \frac{d}{2} \log(\pi) - \log(d) - \log \Gamma\left(\frac{d}{2}\right) + \frac{d}{2} \log(\chi_{1-\alpha}^2(d)) + \frac{1}{2} \log |\Sigma|,$$

where $\chi_{1-\alpha}^2(d)$ is the $(1 - \alpha)100\%$ quantile of a chi-squared distribution with d degrees of freedom, and $\Gamma(\cdot)$ the gamma function. The covariance matrix Σ can be estimated as the marginal covariance matrix using the well-known relationship between the parameters of a multivariate mixture distribution and the marginal parameters [25, section

6.1.1], that is,

$$\Sigma = \sum_{k=1}^G \pi_k \Sigma_k + \sum_{k=1}^G \pi_k (\boldsymbol{\mu}_k - \boldsymbol{\mu})(\boldsymbol{\mu}_k - \boldsymbol{\mu})^\top$$

where $\boldsymbol{\mu} = \sum_{k=1}^G \pi_k \boldsymbol{\mu}_k$ is the vector of marginal means.

Thus, modes whose \log -density is smaller than $-\log(V)$ or, equivalently, with density smaller than $\exp(-\log(V)) = 1/V$, can be dropped, and points associated with them are reassigned to the remaining modes with additional few steps of the MEM algorithm. This is coherent with Hartigan's definition of cluster adopted in the paper, because the groups associated with low-density modes lack the main requirement of being a high-density cluster. The data example in Section 4.3 illustrates this approach.

4 | DATA ANALYSIS EXAMPLES

4.1 | Simulated data example

Recalling the bivariate Gaussian-skew-normal mixture distribution described in Section 1, Figure 3(A) shows the

estimated density obtained by the selected GMM, namely, model VVE with three mixture components selected by BIC. To illustrate the procedure, some points are marked as blue filled points, and they are also reported in isolation in Figure 3(B). For the selected points, the paths produced by the MEM algorithm described in Section 2 are shown as arrows in Figure 3(C). At each step of the algorithm, points move up-hill toward the density modes. The estimated modes are shown in Figure 3(D). Figure 3(E) shows the modal clustering for all the data points obtained according to the mode to which they converge. The MEM algorithm required 21 iterations and 0.23 s to run on an iMac with four-core i5 Intel CPU running at 2.8 GHz and with 16 GB of RAM. Finally, Figure 3(F) illustrates the partition of the feature space that defines the “domains of attraction” of the estimated density modes.

4.2 | Mass cytometry data

Mass cytometry is a recent technology that couples flow cytometry with mass spectrometry. It allows to simultaneously measure several features of a cell. The biological question of interest is the identification of subpopulations of cells. We consider two protein markers, CD4 and CD3all, from a mass cytometry experiment [26] to find latent classes in single-cell measurements. Data are preprocessed using the hyperbolic arcsin transformation [27], that is, $\text{asinh}(x) = \log(x + \sqrt{x^2 + 1})$. A random sample of 10,000 cells (out of 91,392) is shown in Figure 4(A). The clustering obtained with the “best” GMM selected by BIC—namely, VVV in the `mclust` nomenclature [20, table 3] with nine components—is shown in Figure 4(B). The corresponding density estimate and the modes estimated via the MEM algorithm are reported in Figure 4(C), while Figure 4(D) shows the corresponding modal clustering partition. There appears to be five clusters, one for each combination of high/low values of the CD4 and CD3all markers, and an additional cluster formed by the highest values of both the CD4 and CD3all markers. This result can be contrasted with those obtained using approaches based on merging mixture components. Figure 4(E) shows the partition derived from the entropy-based approach of Baudry et al. [4], while Figure 4(F) shows the clusters obtained using the unimodal ridgeline approach of Hennig [5]. Both partitions are clearly different from the one obtained using the MEM algorithm, with the latter producing more compact and easily interpretable clusters.

Finally, note that with 10,000 observations the MEM algorithm required 24 iterations and 5.4 seconds to run on an iMac with four-core i5 Intel CPU running at 2.8 GHz and with 16 GB of RAM.

4.3 | Bankruptcy data set

Altman [28] presented a study on financial ratios to predict corporate bankruptcy. The data set provides the ratio of retained earnings (RE) to total assets, and the ratio of earnings before interests and taxes (EBIT) to total assets, for a sample of 66 manufacturing US firms, of which 33 had filed for bankruptcy in the following 2 years. Data are shown in Figure 5(A).

The best GMM selected by BIC is the model VEI (diagonal, varying volume and equal shape) with three components. The contour plot of the corresponding density estimate is shown in Figure 5(B), with points marked according to the implied maximum a posteriori classification. There appears to be two prominent clusters, consisting mainly of solvent and bankrupt companies, but also a spread out group of firms with very low values for either financial ratios. Figure 5(C) shows the density estimate via a 3D perspective plot. Following the approach described in Section 3.3, a plane is included at the uniform density level corresponding to $\exp(-11.17492) = 1/71$, $319.39 = 1.402 \times 10^{-5}$, where $V = 71$, 319.39 is the hypervolume of the 99% central region. As it can be seen, only two bumps of densities emerge, namely, those corresponding to the main groups in the data. A very low-density mode is also present in the region corresponding to small financial ratios; specifically, the density is equal to 4.661×10^{-6} , a value approximately equal to one third the density threshold computed above. For this reason, the low-density mode is filtered out by the denoising procedure.

The modes estimated by the MEM algorithm are shown in Figure 5(D), with data points marked according to the clusters assigned by the modal clustering procedure. Only four companies are misclassified, as indicated by the circled data points. This result can be compared with those reported by Lo and Gottardo [29, table 1], where the best model (a mixture of t distributions on Box–Cox transformed data) misclassified 10 observations.

5 | CONCLUSIONS

This paper addresses the problem of computing the modes of a density estimated by fitting a GMM. The proposed approach is based on the MEM algorithm, an iterative procedure aimed at identifying the local maxima of a density function. By exploiting specific characteristics of the underlying GMM, we extend the MEM algorithm to deal with any parsimonious component-covariance matrix decomposition. Furthermore, we discuss a fast implementation of the algorithm that allows to perform the M-step simultaneously for all data points. Once the modes of the

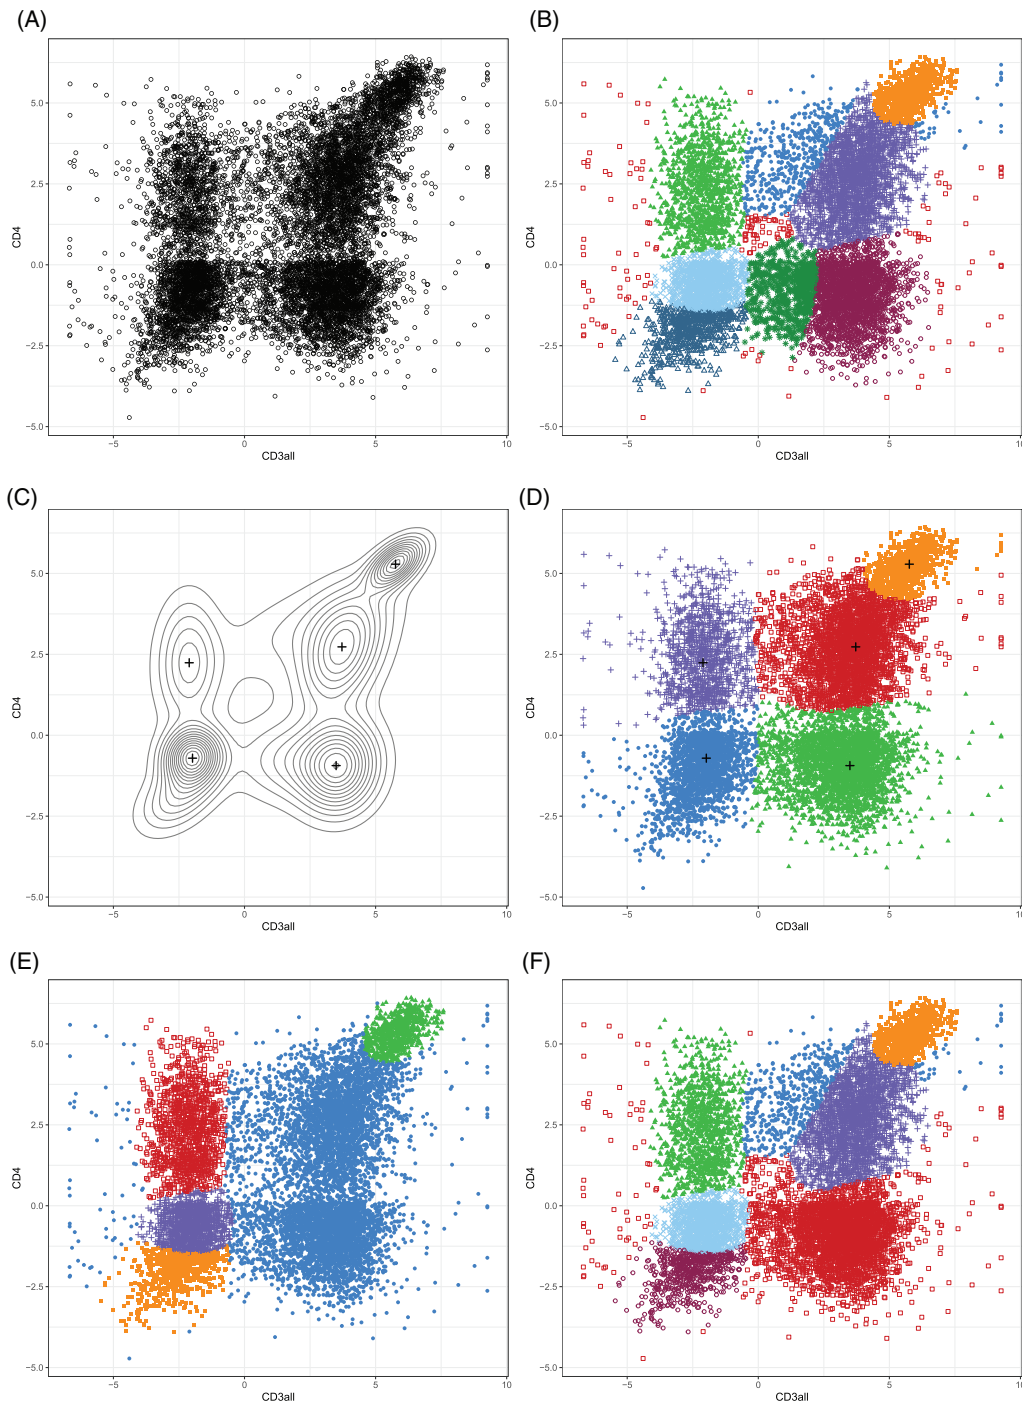


FIGURE 4 Mass cytometry experiment data: (A) plot of CD4 and CD3all markers; (B) clustering from the unrestricted Gaussian mixture model (GMM) with nine components; (C) contour plot of density estimate and modes (+) estimated by the Modal EM (MEM) algorithm; (D) modal clustering classification of cells; (E) clustering obtained by merging mixture components using the Baudry et al. [4] approach; (F) clustering obtained by merging mixture components using the unimodal ridgeline approach of Hennig [5]

underlying density are estimated, a modal clustering partition can be obtained by associating each observation to the pertaining mode.

The MEM algorithm discussed, as any other mode-seeking procedure, relies on the quality of the underlying density estimate. Clearly, if the parameters of

the mixture model are not well estimated, some issues could arise. However, provided that the general form of the density estimate is not overly biased, the proposed method should not be significantly affected.

The proposed approach seems to be very promising and, in principle, it could be extended to mixtures of

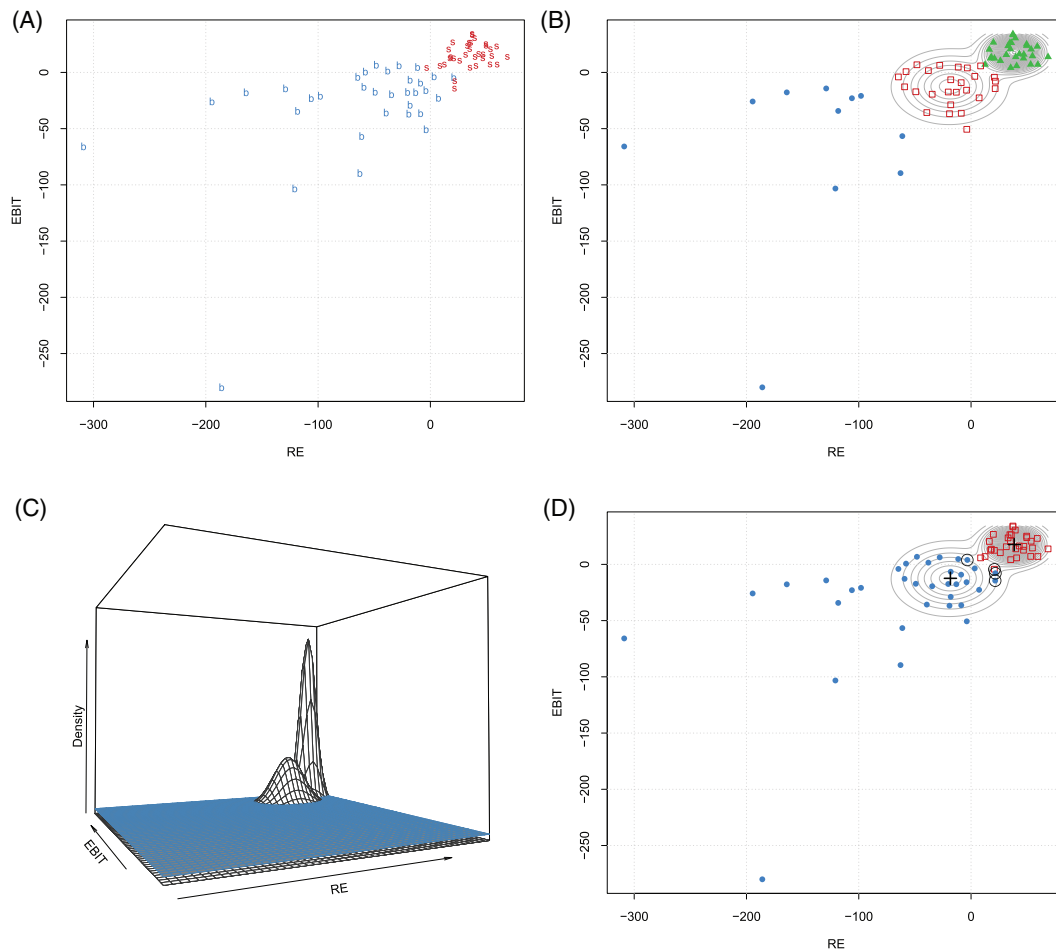


FIGURE 5 Altman's bankruptcy data set: (A) financial ratios of companies with points marked as solvent or bankruptcy; (B) contour plot of density estimate and classification obtained by the best fitting Gaussian mixture model (GMM) as selected by BIC; (C) perspective plot of density estimate and a plane drawn at the uniform density level for denoising; (D) modes and classification obtained from applying the Modal EM (MEM) algorithm, with circled points corresponding to misclassified observations

non-Gaussian distributions (e.g. t , skew-normal, skew- t , shifted asymmetric Laplace, ...). However, it is necessary to investigate the potential benefits obtained by adopting more complex probability models. Recently, an adaptation of the proposed algorithm has been used for clustering from an ensemble of Gaussian mixtures [30].

Another area of future research involves the use of the MEM algorithm in high-dimensional data settings. In this regard, we plan to study the effectiveness of the MEM algorithm applied to the subspace estimated by the GMM-based projection pursuit method proposed by Scrucca and Serafini [31].

ACKNOWLEDGMENTS

We thank the editor and the anonymous referees for their valuable comments, which substantially improved the quality of the work. Luca Scrucca received the support of "Fondo Ricerca di Base, 2019" from Università degli Studi di Perugia for the project "A fast and efficient Modal

EM algorithm for Gaussian mixtures". Open Access Funding provided by Università degli Studi di Perugia within the CRUI-CARE Agreement.

[Correction added on May 17, 2022, after first online publication: CRUI-CARE funding statement has been added.]

CONFLICT OF INTEREST

The author declares no potential conflict of interests.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.

ORCID

Luca Scrucca  <https://orcid.org/0000-0003-3826-0484>

REFERENCES

1. B. Everitt, S. Landau, M. Leese, and D. Stahl, *Cluster analysis*, 5th ed., John Wiley & Sons Ltd, Chichester, UK, 2011.

2. G. J. McLachlan and D. Peel, *Finite Mixture Models*, Wiley, New York, 2000.
3. C. Fraley and A. E. Raftery, *Model-based clustering, discriminant analysis, and density estimation*, *J. Am. Stat. Assoc.* 97(458) (2002), 611–631.
4. J. P. Baudry, A. E. Raftery, G. Celeux, K. Lo, and R. Gottardo, *Combining mixture components for clustering*, *J. Comput. Graph. Stat.* 19(2) (2010), 332–353.
5. C. Hennig, *Methods for merging Gaussian mixture components*, *ADAC* 4(1) (2010), 3–34.
6. L. Scrucca, *Identifying connected components in Gaussian finite mixture models for clustering*, *Computational Statistics & Data Analysis* 93 (2016), 5–17. <https://doi.org/10.1016/j.csda.2015.01.006>.
7. W. Stuetzle, *Estimating the cluster tree of a density by analyzing the minimal spanning tree of a sample*, *J. Classif.* 20(1) (2003), 25–47.
8. J. A. Hartigan, *Clustering algorithms*, John Wiley & Sons, New York, 1975.
9. Li, J., S. Ray, and B. G. Lindsay, 2007: *A nonparametric statistical approach to clustering via mode identification*. *Journal of Machine Learning Research*, 8, no. Aug, 1687–1723.
10. Carreira-Perpiñán, M. Á. and C. K. I. Williams On the number of modes of a Gaussian mixture. *Scale Space Methods in Computer Vision, 4th International Conference, Scale-Space 2003*, Isle of Skye, UK, June 10–12, 2003, pp. 625–640.
11. M. Á. Carreira-Perpiñán and C. K. I. Williams, *An Isotropic Gaussian Mixture Can Have More Modes Than Components*. Institute for Adaptive and Neural Computation Informatics Research Report EDI-INF-RR-0185, School of Informatics, University of Edinburgh, Scotland, 2003.
12. C. Améndola, A. Engström, and C. Haase, *Maximum number of modes of Gaussian mixtures*, *Inf. Infer. J. IMA* 9(3) (2019), 587–600. <https://doi.org/10.1093/imaia/iaz013>.
13. K. Fukunaga and L. Hostetler, *The estimation of the gradient of a density function, with applications in pattern recognition*, *IEEE Trans. Inf. Theory* 21(1) (1975), 32–40.
14. M. Á. Carreira-Perpiñán, *Clustering methods based on kernel density estimators: mean-shift algorithms*, in *Handbook of cluster analysis*, Chapman and Hall/CRC, New York, NY, 2016, 383–417.
15. D. W. Scott, *Multivariate density estimation: Theory, practice, and visualization*, 2nd ed., John Wiley & Sons, Hoboken, NJ, 2009.
16. M. Á. Carreira-Perpiñán, *Gaussian mean-shift is an EM algorithm*, *IEEE Trans. Pattern Anal. Mach. Intell.* 29(5) (2007), 767–776.
17. J. E. Chacón, *Mixture model modal clustering*, *ADAC* 13(2) (2019), 379–404. <https://doi.org/10.1007/s11634-018-0308-3>.
18. G. Menardi, *A review on modal clustering*, *Int. Stat. Rev.* 84(3) (2016), 413–433. <https://doi.org/10.1111/insr.12109>.
19. A. Azzalini, *The skew-normal and related families*, Cambridge University Press, Cambridge, 2013.
20. L. Scrucca, M. Fop, T. B. Murphy, and A. E. Raftery, *mclust 5: Clustering, classification and density estimation using Gaussian finite mixture models*, *R J.* 8(1) (2016), 205–233.
21. J. Banfield and A. E. Raftery, *Model-based Gaussian and non-Gaussian clustering*, *Biometrics* 49 (1993), 803–821.
22. G. Celeux and G. Govaert, *Gaussian parsimonious clustering models*, *Pattern Recogn.* 28 (1995), 781–793.
23. R Core Team, *R: A language and environment for statistical computing*, R Foundation for Statistical Computing, Vienna, Austria. Available at, 2019. <https://www.R-project.org/>.
24. M. Á. Carreira-Perpiñán, *Mode-finding for mixtures of Gaussian distributions*, *IEEE Trans. Pattern Anal. Mach. Intell.* 22(11) (2000), 1318–1323.
25. S. Frühwirth-Schnatter, *Finite mixture and Markov switching models*, Springer, New York, NY, 2006.
26. S. C. Bendall, E. F. Simonds, P. Qiu, D. A. El-ad, P. O. Krutzik, R. Finck, R. V. Bruggner, R. Melamed, A. Trejo, O. I. Ornatsky, R. S. Balderas, S. K. Plevritis, K. Sachs, D. Pe'er, S. D. Tanner, and G. P. Nolan, *Single-cell mass cytometry of differential immune and drug responses across a human hematopoietic continuum*, *Science* 332(6030) (2011), 687–696. <https://doi.org/10.1126/science.1198704>.
27. S. Holmes and W. Huber, *Modern statistics for modern biology*, Cambridge University Press, Cambridge, England, 2018.
28. E. I. Altman, *Financial ratios, discriminant analysis and the prediction of corporate bankruptcy*, *J. Financ.* 23(4) (1968), 589–609.
29. K. Lo and R. Gottardo, *Flexible mixture modeling via the multivariate t distribution with the box-cox transformation: An alternative to the skew-t distribution*, *Stat. Comput.* 22(1) (2012), 33–52.
30. A. Casa, L. Scrucca, and G. Menardi, *Better than the best? Answers via model ensemble in density-based clustering*, *Advances in Data Analysis and Classification* (2020). <https://doi.org/10.1007/s11634-020-00423-6>.
31. L. Scrucca and A. Serafini, *Projection pursuit based on Gaussian mixtures and evolutionary algorithms*, *J. Comput. Graph. Stat.* 28 (2019), 847–860. <https://doi.org/10.1080/10618600.2019.1598871>.

How to cite this article: L. Scrucca, *A fast and efficient Modal EM algorithm for Gaussian mixtures*, *Stat Anal Data Min: The ASA Data Sci Journal*. **14** (2021), 305–314. <https://doi.org/10.1002/sam.11527>