




# Modal clustering on PPGMMGA projection subspace

Luca Scrucca<sup>1</sup> 

*Università degli Studi di Perugia*

## Summary

PPGMMGA is a projection pursuit (PP) algorithm aimed at detecting and visualising clustering structures in multivariate data. The algorithm uses the negentropy as PP index obtained by fitting Gaussian mixture models (GMMs) for density estimation and, then, exploits genetic algorithms (GAs) for its optimisation. Since the PPGMMGA algorithm is a dimension reduction technique specifically introduced for visualisation purposes, cluster memberships are not explicitly provided. In this paper a modal clustering approach is proposed for estimating clusters of projected data points. In particular, a modal EM algorithm is employed to estimate the modes corresponding to the local maxima in the projection subspace of the underlying density estimated using parsimonious GMMs. Data points are then clustered according to the domain of attraction of the identified modes. Simulated and real data are discussed to illustrate the proposed method and evaluate the clustering performance.

*Key words:* Gaussian mixture models; genetic algorithms; modal clustering; modal EM; negentropy index; projection pursuit.

## 1. Introduction

Finite mixture models are an important tool in statistical data analysis. Because of their flexibility, they have been successfully used for different purposes, from modelling of the heterogeneity to clustering, from supervised classification to semiparametric density estimation (McLachlan & Peel 2000a; McLachlan, Lee & Rathnayake 2019). Model-based clustering is often used as a synonym to indicate the use of mixture models to investigate the presence of clusters in the data.

Projection pursuit (PP) refers to a collection of algorithms aimed at finding interesting structures in the data (Friedman & Tukey 1974; Friedman 1987). Often, the main interest consists in detecting the presence of homogeneous and well separated groups of data. Recently, a projection pursuit algorithm based on Gaussian mixture models (GMMs) has been proposed by Scrucca & Serafini (2019). This PP algorithm employs the negentropy obtained from a multivariate density estimated by GMMs as the PP index to be maximised. Optimisation of the negentropy index is performed using genetic algorithms (GAs). The resulting PPGMMGA algorithm estimates the optimal orthogonal projection subspace which is able to show the directions of maximal non-normality, and for this reason it can be effectively employed to detect and visualise the presence of clustering structures in multivariate data.

---

<sup>1</sup>Dipartimento di Economia, Università degli Studi di Perugia, Via A. Pascoli 20, Perugia 06123, Italy. e-mail: luca.scrucca@unipg.it

[Correction added on 18 May 2022, after first online publication: CRUI-CARE funding statement has been added.]

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

However, being PPGMMGA a dimension reduction method, no clustering memberships are assigned to the projected data points.

The main goal of this contribution is to propose a clustering approach based on the modal EM algorithm which allows to estimate the modes corresponding to the local maxima of the underlying density of projected data. Thus, groups are obtained according to the mode to which each data point converges, thereby providing a clustering procedure for the PPGMMGA algorithm. The modal clustering procedure discussed in this paper follows the definition of clusters proposed by Hartigan (1975 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’.

Section 2 provides a brief review of parsimonious GMMs and of the model-based projection pursuit algorithm used to maximise the negentropy of a GMM. Section 3 contains the modal clustering proposal for applying the modal EM algorithm on the estimated projection subspace. Section 4 describes the empirical results obtained from the application of the proposed modal EM algorithm to examples using both synthetic and real datasets. The final section provides some concluding remarks.

## 2. Background

### 2.1. Parsimonious Gaussian mixtures

Let  $\mathbf{x}$  be a random vector in  $\mathbb{R}^p$  with density  $f(\mathbf{x})$ , which we assume it can be described by a finite mixture of Gaussian distributions:

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

where  $\boldsymbol{\theta} = \{\pi_1, \pi_2, \dots, \pi_{G-1}, \boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_G, \text{vech}(\boldsymbol{\Sigma}_1), \dots, \text{vech}(\boldsymbol{\Sigma}_G)\}$  are the parameters of the mixture model with  $G$  components,  $(\pi_1, \pi_2, \dots, \pi_G)$  the mixing weights, so that  $\pi_g > 0$  and  $\sum_{g=1}^G \pi_g = 1$ , and  $\phi(\mathbf{x}; \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g)$  the underlying multivariate Gaussian density function of the  $g$ th component having mean  $\boldsymbol{\mu}_g$  and covariance matrix  $\boldsymbol{\Sigma}_g$ . Note that  $\text{vech}(\cdot)$  is the vector operator, which forms a vector by extracting unique elements of a symmetric matrix.

Parsimonious parametrisation of covariance matrices for GMMs is available through the eigen-decomposition  $\boldsymbol{\Sigma}_g = \lambda_g \mathbf{U}_g \boldsymbol{\Delta}_g \mathbf{U}_g^\top$ , where  $\lambda_g = |\boldsymbol{\Sigma}_g|^{1/p}$  is a scalar which controls the volume,  $\boldsymbol{\Delta}_g$  is the diagonal matrix of normalised eigenvalues of  $\boldsymbol{\Sigma}_g$ , such that  $|\boldsymbol{\Delta}_g| = 1$ , which controls the shape, and  $\mathbf{U}_g$  is the orthogonal matrix of eigenvectors of  $\boldsymbol{\Sigma}_g$  controlling the orientation of the ellipsoid (Banfield & Raftery 1993; Celeux & Govaert 1995). Maximum likelihood estimates of the mixture parameters can be obtained using the expectation–maximisation (EM) algorithm (Dempster, Laird & Rubin 1977; McLachlan & Krishnan 2008), whereas information criteria, such as the Bayesian information criterion (BIC; Schwartz 1978), are typically used for model selection. The latter concerns the decision on both the number of mixture components and on which parsimonious parametrisation of covariances to adopt. Other approaches are also available in the literature to reduce the number of unknown parameters. For instance, mixture of factor analysers, proposed in McLachlan & Peel (2000b) and based on the work of Ghahramani & Hinton (1997), could be considered here by allowing each Gaussian component in the mixture to be represented in a different lower-dimensional manifold.

## 2.2 Model-based projection pursuit

Consider the orthogonal linear projection  $\mathbf{z} = \mathbf{B}^\top \mathbf{x} \in \mathbb{R}^d$ , with  $d < p$ , and the basis  $\mathbf{B} \in \mathbb{R}^{p \times d}$ , a matrix spanning the dimension reduction subspace  $\mathcal{S}(\mathbf{B})$ . A remarkable result is the *linear transformation property of GMMs* (see Appendix A in Scrucca & Serafini 2019), which shows that if  $\mathbf{x}$  follows a Gaussian mixture distribution as in (1), then any linear projection also follows a Gaussian mixture distribution, so we can write

$$f(\mathbf{z}) = \sum_{g=1}^G \pi_g \phi(\mathbf{z}; \mathbf{B}^\top \boldsymbol{\mu}_g, \mathbf{B}^\top \boldsymbol{\Sigma}_g \mathbf{B}). \quad (2)$$

Thus, the distribution of the projected data can still be described by a GMM with a suitable transformation of the parameters of the Gaussian components. The basis  $\mathbf{B}$  of the subspace can be estimated by maximising the *negentropy*, a PP index that summarises deviations from the least interesting case represented by the Gaussian distribution. The negentropy is defined as

$$J(\mathbf{z}) = h(\phi(\boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z)) - h(\mathbf{z}) = \frac{1}{2} \log((2\pi e)^d |\boldsymbol{\Sigma}_z|) + \int \log(f(\mathbf{z}))f(\mathbf{z})d\mathbf{z},$$

where  $h(\phi(\boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z))$  is the entropy of the marginal multivariate Gaussian distribution with mean  $\boldsymbol{\mu}_z$  and covariance  $\boldsymbol{\Sigma}_z$ , and  $h(\mathbf{z})$  is the entropy of the mixture density of the projected data. Efficient approximations for  $h(\mathbf{z})$  are available, such as the *unscented transformation* (UT), so the PP index can be maximised over the set of all orthonormal bases  $\mathbf{B}$  using genetic algorithms. For further details see Scrucca & Serafini (2019). This briefly described algorithm is named PPGMGA, and it is implemented in the R package `ppgmmga` (Serafini & Scrucca 2019).

## 3. Modal clustering on PPGMGA projection subspace

In the modal approach to clustering, clusters are defined as the local maxima of the underlying probability density function. For a recent review of different approaches in modal clustering see Menardi (2016), while for a more general discussion on the recent renewed interest on modal procedures see Chacón (2020). The Modal EM (MEM) algorithm is an iterative algorithm aimed at identifying the local maxima of a finite mixture density function (Li, Ray & Lindsay 2007). Once the modes have been estimated, clustering of data points is obtained according to the mode to which they converge. Recently, Scrucca (2021) generalised the MEM algorithm to efficiently deal with Gaussian mixtures having any parsimonious covariance matrices of the type described in Section 2.2.

To effectively apply the MEM approach on the PPGMGA projection subspace a density estimate is needed. In principle, the mixture density in (2) could be used as the underlying density to base the MEM procedure. However, if this is reasonable at the population level, in practice only an estimate of the density (2) from sample data is available. Even more important, this estimate is typically chosen among a set of plausible models according to a model selection criterion, such as the BIC (Fraley & Raftery 1998). Since the latter penalises model complexity, more parsimonious models tend to be selected when the sample size of available data is limited compared to the dimensionality of the problem. However, the adopted parameterisation may prove to be inadequate to represent the density of the data on the projection subspace.

A limiting case of this situation is presented in Section 4.1, where a simulated dataset of 100 observations on a 50-dimensional feature space is considered. In this case the Gaussian mixture density selected by BIC has only one component, and if this poses no problem for estimating the PPGMGA subspace, it clearly fails if used for estimating the modes of the underlying density on the two-dimensional projection subspace. For this reason, a different route is needed and, again, a parsimonious Gaussian mixture is called to the rescue.

### 3.1. MEM algorithm

Consider the PPGMGA basis  $\mathbf{B}$  estimated for a fixed dimensionality  $d$ . The projected data can be computed as  $\mathbf{z}_i = \mathbf{B}^\top \mathbf{x}_i$  for all  $i = 1, \dots, n$ , or equivalently in matrix form as  $\mathbf{Z} = \mathbf{X}\mathbf{B}$ . The Gaussian mixture density on the projection subspace can be written as

$$f(\mathbf{z}; \Psi) = \sum_{k=1}^K \tau_k \phi(\mathbf{z}; \boldsymbol{\eta}_k, \boldsymbol{\Gamma}_k), \quad (3)$$

where  $\Psi = \{\tau_1, \tau_2, \dots, \tau_{K-1}, \boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_K, \text{vech}(\boldsymbol{\Gamma}_1), \dots, \text{vech}(\boldsymbol{\Gamma}_K)\}$  are the parameters of the mixture model with  $K$  components,  $(\tau_1, \tau_2, \dots, \tau_K)$  the mixing weights, so that  $\tau_k > 0$  and  $\sum_{k=1}^K \tau_k = 1$ , and  $\phi(\mathbf{z}; \boldsymbol{\eta}_k, \boldsymbol{\Gamma}_k)$  the multivariate Gaussian density function of the  $k$ th component having mean  $\boldsymbol{\eta}_k$  and covariance matrix  $\boldsymbol{\Gamma}_k$ . Following the details provided in Section 2.1, the best fitting parsimonious mixture density (3) can be selected by BIC with respect to both the number of mixture components  $K$  and the parsimonious covariance matrices  $\{\boldsymbol{\Gamma}_k\}_{k=1}^K$ . By using the estimated density it is possible to identify the modes in the dimension reduction subspace  $\mathcal{S}(\mathbf{B})$ , thus enabling a modal clustering approach for the identification of homogeneous and separated groups. The MEM algorithm is briefly described below, while further details for interested readers, including a discussion on the number of modes in Gaussian mixtures, are provided in Scrucca (2021).

The MEM algorithm starts at initial data points  $\mathbf{z}_i^{(0)} = \mathbf{z}_i$  for  $i = 1, \dots, n$ , then at each iteration  $t$ , assuming the mixture parameters  $\{\tau_k, \boldsymbol{\eta}_k, \boldsymbol{\Gamma}_k\}_{k=1}^K$  as known and fixed, performs the following steps:

*E-step*—update the posterior conditional probability of the projected data point  $\mathbf{z}_i$  to belong to the  $k$ th mixture component using

$$\zeta_{ik}^{(t)} = \frac{\tau_k \phi(\mathbf{z}_i^{(t-1)}; \boldsymbol{\eta}_k, \boldsymbol{\Gamma}_k)}{\sum_{j=1}^K \tau_j \phi(\mathbf{z}_i^{(t-1)}; \boldsymbol{\eta}_j, \boldsymbol{\Gamma}_j)}, \quad \text{for all } k = 1, \dots, K.$$

*M-step*—update the current value of  $\mathbf{z}_i$  by solving the optimisation problem

$$\mathbf{z}_i^{(t)} = \arg \max_{\mathbf{z}_i \in \mathbb{R}^d} \sum_{k=1}^K \zeta_{ik}^{(t)} \log \phi(\mathbf{z}_i^{(t-1)}; \boldsymbol{\eta}_k, \boldsymbol{\Gamma}_k).$$

The procedure above is iterated until a convergence criterion is satisfied, or a pre-specified maximum number of iterations is reached.

Note that the objective function in the M-step can be written as

$$Q(\mathbf{z}_i) = \sum_{k=1}^K \zeta_{ik} \log \phi(\mathbf{z}_i; \boldsymbol{\eta}_k, \boldsymbol{\Gamma}_k),$$

so the gradient and Hessian of this function with respect to the projected vector  $\mathbf{z}_i$  are, respectively,

$$\begin{aligned}\nabla Q(\mathbf{z}_i) &= -\sum_{k=1}^K \zeta_{ik} \Gamma_k^{-1} (\mathbf{z}_i - \boldsymbol{\eta}_k), \quad \text{and} \\ \nabla^2 Q(\mathbf{z}_i) &= -\sum_{k=1}^K \zeta_{ik} \Gamma_k^{-1}.\end{aligned}$$

Therefore, the M-step presents a closed-form solution given by

$$\mathbf{z}_i^* = \left( \sum_{k=1}^K \zeta_{ik} \Gamma_k^{-1} \right)^{-1} \sum_{k=1}^K \zeta_{ik} \Gamma_k^{-1} \boldsymbol{\eta}_k. \quad (4)$$

A direct application of the update equation above may yield large jumps during the initial iterations of the algorithm for data points lying in low-density regions, thus escaping from the region of the attracting mode to the domain of attraction of a different mode. In other word, a data point located in a low-density region might see large jumps at the beginning of the iterative algorithm, hence converging to a mode further from its domain of attraction. For this reason, Scrucca (2021, section 3.1) introduced a step size to get more reliable update at iteration  $t$  using a convex linear combination of the solution at previous iteration,  $\mathbf{z}_i^{(t-1)}$ , and the new proposal  $\mathbf{z}_i^*$  from (4), as follows

$$\mathbf{z}_i^{(t)} = (1 - \omega) \mathbf{z}_i^{(t-1)} + \omega \mathbf{z}_i^*, \quad (5)$$

where  $\omega = 1 - \exp\{-0.1t\}$  is a parameter that controls the step size. Following this approach, at earlier iterations the step size is small, hence updates are done in small steps, but as the number of iterations increase the step size converges to one and the updates become essentially equal to the proposal  $\mathbf{z}_i^*$ . The proposed modification does not break the monotonicity property of EM (see Appendix for the proof), although it slows the convergence somewhat by increasing the number of iterations required.

Finally, we note that the execution time can be greatly reduced by batch updating all the data points at once instead of working one point at time, as detailed in Scrucca (2021).

### 3.2. Some final comments

In the model-based clustering approach, clusters are defined by the underlying distribution used for the mixture components. For instance, GMMs assume ellipsoidal group structures. Then, the so-called maximum a posteriori (MAP) principle is used for clustering. On the contrary, as discussed in Section 1, modal clustering is an attempt at addressing the definition of clusters proposed by Hartigan (1975). If data points projected onto the PPGMGA projection subspace show ellipsoidal structures and are well-separated then the model-based clustering approach using GMMs and the modal-clustering via MEM algorithm produce equivalent results. However, in case the groups show non-Gaussian structures or substantially overlap on the projection subspace, the final clustering might be different. This mostly reflects the different definition adopted by the two approaches regarding what a cluster is. An example of such a case is discussed in Section 4.4.

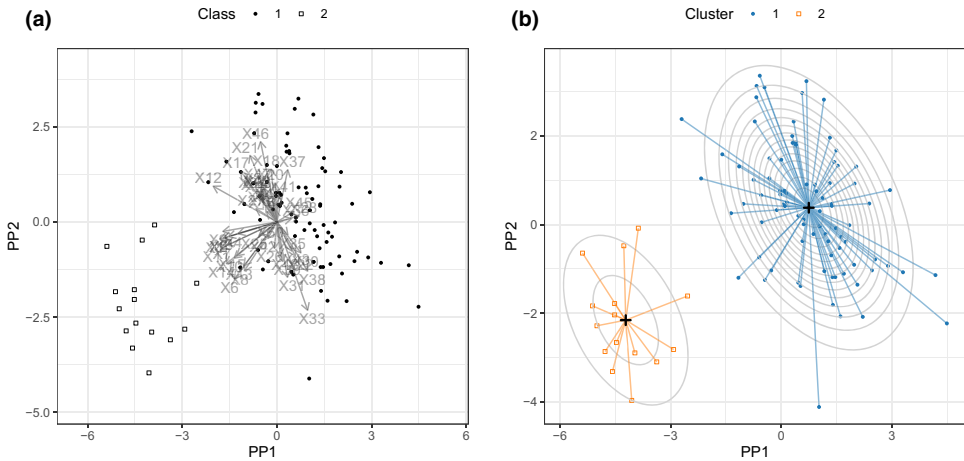


Figure 1. Plot of two-group synthetic data points projected along the PPGMMGA directions with points marked according to the known class in panel (a), and using the modal clustering solution on panel (b). The last graph also displays the modal expectation–maximisation iteration paths towards the attracting modes represented using the + symbol.

#### 4. Data analysis examples

In this section we apply the proposed methodology to simulated and real datasets.

##### 4.1. Simulated data with clustering on a subset of attributes

We consider a reduced version of the two-group synthetic dataset discussed in Friedman & Meulman (2004). For the first group 85 observations are randomly generated from a 50-dimensional standard Gaussian distribution. The second group of 15 observations are also drawn from a 50-dimensional Gaussian distribution, but its first  $p_0 = 15$  attributes have mean 1.5 and standard deviation 0.2, whereas the remaining  $50 - p_0$  attributes have zero mean and unit standard deviation. Thus, the two groups differ only on the first  $p_0$  attributes.

Figure 1a shows the data projected onto the two-dimensional projection subspace estimated by the PPGMMGA algorithm, which clearly reveals the two-group clustering structure. It is interesting to note that the GMM selected by BIC on the original 50-dimensional features space has a single Gaussian component. In this case, an attempt at clustering on the original features space would produce a solution consisting of a single cluster. Nonetheless, the estimated directions that maximise the negentropy are able to show the underlying structure as shown in Figure 1a.

Iteration paths of MEM algorithm for each data point towards the attracting mode are depicted in Figure 1b. The resulting modal clustering solution obtained on the projection subspace is perfectly able to recover the two groups of observations.

##### 4.2. Rectangular Gaussian block clusters in three dimensions with noise attributes

In this second example, we consider a simulated dataset made up of 400 observations on eight variables. The first three variables are generated from an equal weight mixture

Table 1. PPGMMGA negentropy, number of modes ( $m$ ) and adjusted Rand index (ARI) for increasing dimension ( $d$ ) of the projection subspace.

$d$	Negentropy	$m$	ARI
1	0.5609	3	0.4244
2	1.0832	7	0.8653
3	1.5075	8	0.9942
4	1.5271	8	0.9942
5	1.5351	8	0.9942

of eight Gaussian distributions with mean vectors at the corners of a rectangular block in three dimensions. The remaining five variables only contain some Gaussian noise. Details about the mean vector for each mixture component and the common covariance matrix are provided in Bolton & Krzanowski (2003 Appendix A).

We consider the PPGMMGA estimation of the projection subspace for varying dimensions. The negentropy values for each solution with  $d$  from 1 to 5 suggest a three-dimensional projection subspace as a reasonable choice (see Table 1). Figure 2 shows some views of the data points projected along the estimated PPGMMGA directions with modes and data points marked according to the cluster memberships estimated by the MEM algorithm. The two-dimensional scatterplots in panels (a)–(c) show pairwise projections for the estimated PPGMMGA directions, with arrows indicating the biplot vectors corresponding to the original features. As it can be seen, only the first three variables contribute to the definition of the PPGMMGA directions (see also the coefficients of the estimated projection basis reported in Table 2). Finally, Figure 2d shows a static view of the data points projected onto the estimated three-dimensional projection subspace.

Table 1 contains, in addition to the negentropy for increasing dimensionality of the projection subspace, the number of modes found by the MEM algorithm and the accuracy of the proposed modal clustering procedure. This is measured by the Adjusted Rand Index (ARI; Hubert & Arabie 1985), a measure of agreement between two partitions, which attains unit value when the two partitions perfectly agree. Overall, the MEM algorithm is able to recover almost perfectly the true clusters for  $d \geq 3$ .

### 4.3. Chironomus larvae data

This dataset comes from a study on morphometric attributes of Chironomus larvae (Atchley & Martin 1971). A sample of 149 larvae was collected from three different species: cloacalis, frommeri and staegeri. For each larva, 17 features of the larval head capsula were measured; for a detailed description of the dataset we refer the reader to Montanari & Viroli (2010).

Figure 3a shows the data projected from the original 17-dimensional feature space to the two-dimensional projection subspace estimated by the PPGMMGA algorithm. The estimated modes and the corresponding modal clustering solutions are shown in Figure 3b, together with the paths induced by the MEM algorithm. The 3-mode clustering structure

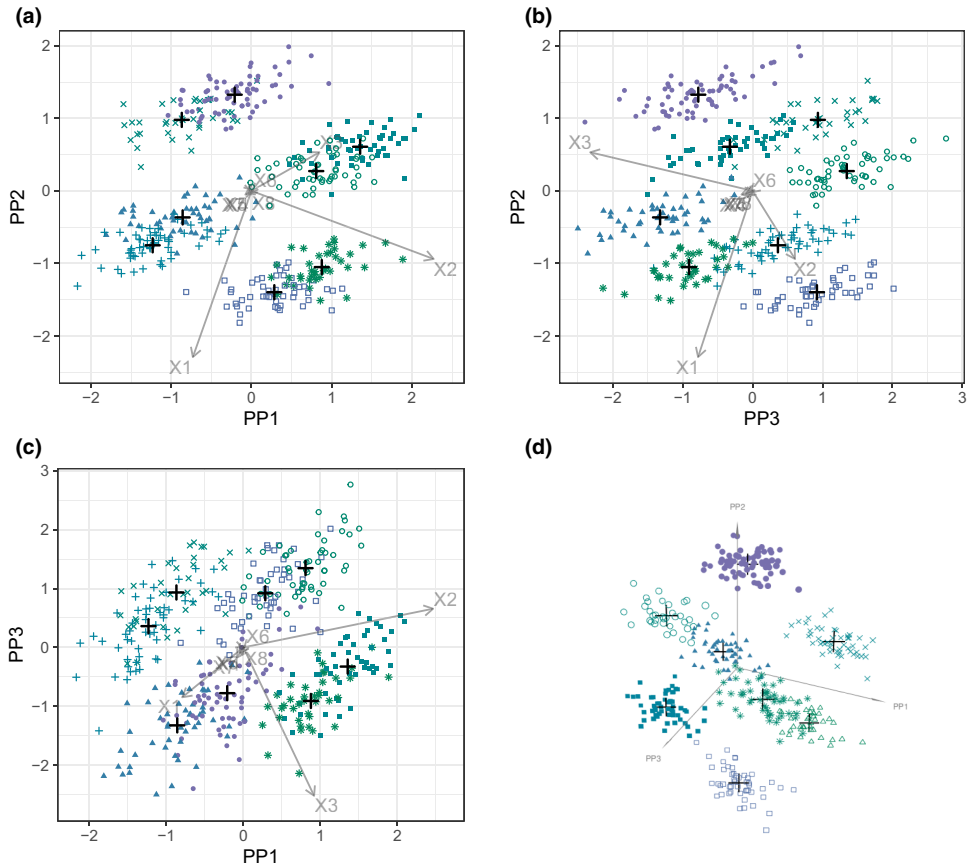


Figure 2. Pairwise scatterplots (panels (a) to (c)) and a static view of 3D scatterplot (panel (d)) of PPGMGA projections with data points marked according to the clusters assigned by the modal expectation–maximisation procedure and the corresponding modes, represented using + symbol.

identified is then compared with the true larvae species in Figure 3c, showing a quite good agreement with only one misclassified observation and an adjusted Rand index of 0.981 (slightly larger than the ARI reported in Montanari & Viroli (2010)).

Table 2. Estimated PPGMGA projection basis for the  $d = 3$  solution.

	PP1	PP2	PP3
$X_1$	-0.2885	-0.9041	-0.3080
$X_2$	0.8964	-0.3706	0.2394
$X_3$	0.3343	0.2101	-0.9168
$X_4$	-0.0097	-0.0203	-0.0425
$X_5$	-0.0184	-0.0149	-0.0361
$X_6$	0.0099	0.0026	0.0111
$X_7$	-0.0297	-0.0203	-0.0626
$X_8$	0.0020	-0.0097	-0.0109



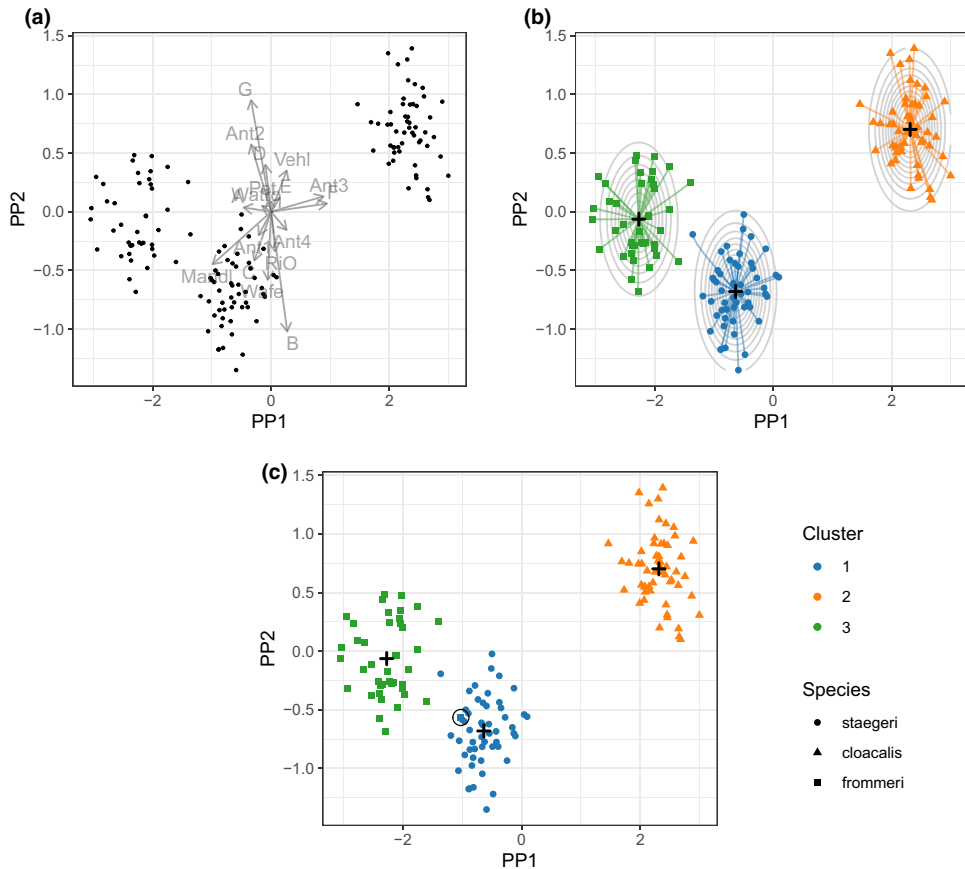


Figure 3. Chironomus larvae data graphs: (a) plot of data points projected along the PPGMGA directions; (b) plot of the modal clustering solution identified by the MEM algorithm with the estimated modes represented by a + symbol; (c) plot of data points marked according to the true species of the larvae with different symbols, and different colours for the identified clusters (circled point represents the misclassified observation).

#### 4.4. Coffee data

The coffee dataset (Streuli 1973) contains 12 chemical measurements on 43 samples of Arabica or Robusta coffee from 29 countries. Scrucca & Serafini (2019) analysed the coffee data by fitting a one-dimensional PPGMGA model. Figure 4a shows the data points projected along the estimated direction, clearly suggesting that coffee varieties can be clustered on this one-dimensional subspace. Graph on Figure 4b reports the underlying density estimated by GMM and used by the MEM algorithm, together with the ascending paths for each projected data point towards the closest mode. The clustering obtained perfectly recovers the latent coffee varieties.

In Section 3.2 the modal clustering approach via MEM algorithm and the model-based clustering approach using the MAP principle were discussed and contrasted. The considerations made therein can now be illustrated in practice using the coffee dataset. In particular, Figure 5 reports the clustering results obtained using the two approaches on the

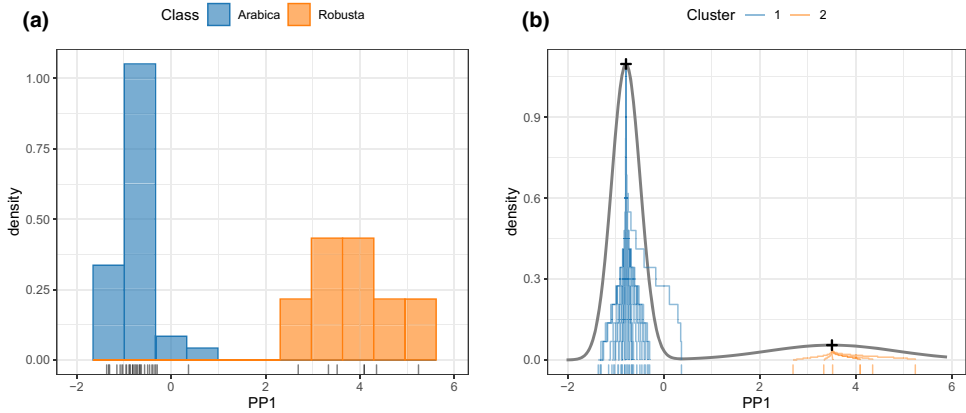


Figure 4. Histogram of coffee data points projected along the first PPGMGA direction marked by the true variety of coffee on panel (a). Underlying density estimated by Gaussian mixture models employed by the modal expectation–maximisation algorithm on panel (b), with step lines showing the iteration paths towards the attracting modes represented by a + symbol. Rug points along the x-axis on both graphs represent the projected data points.

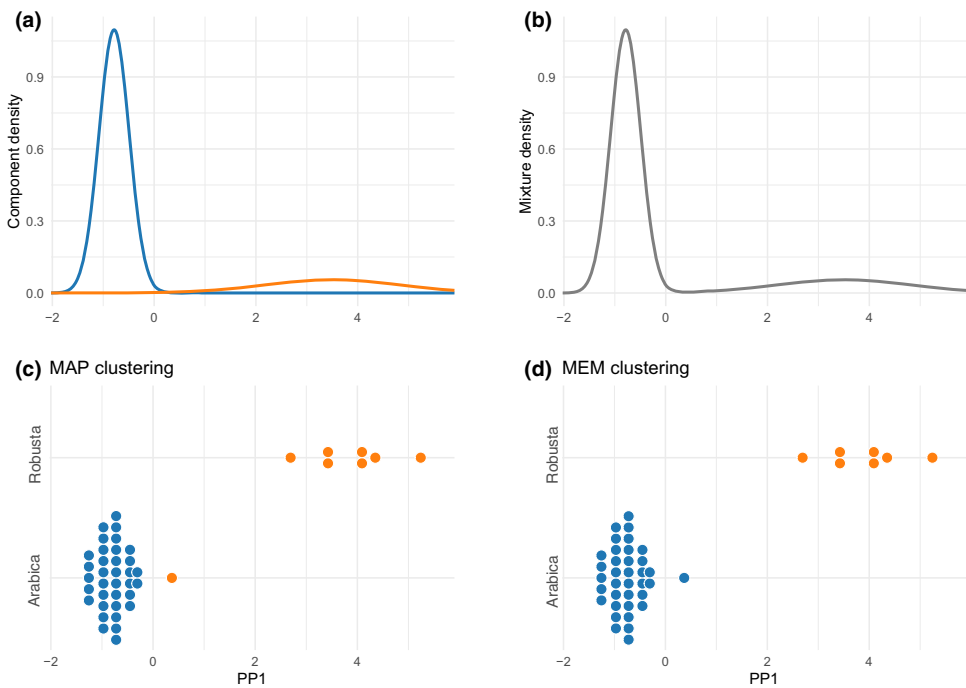


Figure 5. A graphical comparison of model-based clustering versus modal clustering on the one-dimensional PPGMGA projection for the coffee dataset. The two Gaussian components of the mixture (weighted by the mixing proportion) are shown in panel (a), whereas the corresponding mixture density is reported in panel (b). The clustering solutions obtained by applying the maximum a posteriori principle and the modal expectation–maximisation algorithm are shown, respectively, in panels (c) and (d).

one-dimensional PPGMMGA projection. Panel (a) shows the two Gaussian components of the mixture estimated on the first PPGMMGA direction, with the corresponding MAP clustering on panel (c). On the contrary, panel (b) contains the estimated mixture density, with the corresponding MEM clustering on panel (d). According to the MAP principle, data points are assigned to the cluster whose corresponding component density (weighted by the mixing proportion) is the largest. Note that one data point is assigned to the group of Robusta despite being of type Arabica and quite far from the other data points assigned to the same cluster. However, this appears reasonable under the model-based perspective because the estimated conditional probability for that point is larger for Robusta than for Arabica (and this is due to the large estimated variance for the second component). By adopting, instead, a modal clustering perspective, the mixture density must be used to identify the modes, and data points are assigned to the closest clusters represented by such modes. In this case, the classification obtained perfectly matches the true coffee varieties. Indeed, the previously misclassified data point is now assigned to the Arabica cluster because its domain of attraction is the same as for the other Arabica coffees.

## 5. Concluding comments

PPGMMGA is a linear projection method introduced by Scrucca & Serafini (2019) with the aim of visualising multivariate data when a clustering structure is sought. The PPGMMGA algorithm estimates the directions of maximal non-normality as measured by the negentropy PP index. Being a linear mapping from a potential high-dimensional features space to a subspace of reduced dimensionality, it involves linear transformations, such as rotation and reflection, of the original data. Thus, if a clustering structure is present in the original features space, PPGMMGA should preserve it and, hopefully, magnify it.

However, if visualisation is an important step in understanding the underlying structure of the data, an automatic procedure for the identification of homogenous groups would be desirable. In this paper we introduce a modal clustering approach to address this issue. Working on the projection subspace estimated by the PPGMMGA method, a modal EM algorithm is proposed for identifying the modes of the density estimated by fitting parsimonious Gaussian mixture models.

The proposed approach can, in principle, be applied to any dimension reduction method, such as local linear embedding (LLE; Roweis & Saul 2000), t-distributed stochastic neighbour embedding (t-SNE; Van Der Maaten & Hinton 2008) and uniform manifold approximation and projection (UMAP; McInnes, Healy & Melville 2018). We defer the investigation of these extensions to future work.

## Appendix

We prove here that the inclusion of the step size in (5) does not break the monotonicity property of MEM, that is,  $f(\mathbf{z}_i^{(t)}) \geq f(\mathbf{z}_i^{(t-1)})$ , for any iteration  $t$ . First, we note that by the ascending property of the MEM algorithm for any mixture we have that  $f(\mathbf{z}_i^*) \geq f(\mathbf{z}_i^{(t-1)})$  (see Li, Ray & Lindsay 2007, Appendix A), where  $\mathbf{z}_i^*$  is the proposal from (4) at current iteration  $t$ ,  $\mathbf{z}_i^{(t-1)}$  is the solution at iteration  $t - 1$ , and  $f(\cdot)$  is the mixture density (3). Then, at iteration  $t$  of the algorithm we have

$$\begin{aligned}
f(\mathbf{z}_i^{(t)}) &= f((1-\omega)\mathbf{z}_i^{(t-1)} + \omega\mathbf{z}_i^*) \\
&\geq (1-\omega)f(\mathbf{z}_i^{(t-1)}) + \omega f(\mathbf{z}_i^*) && \text{(by the monotonicity along the learning path)} \\
&= f(\mathbf{z}_i^{(t-1)}) + \omega(f(\mathbf{z}_i^*) - f(\mathbf{z}_i^{(t-1)})) \\
&\geq f(\mathbf{z}_i^{(t-1)}). && \text{(since } f(\mathbf{z}_i^*) - f(\mathbf{z}_i^{(t-1)}) \geq 0 \text{ and } \omega \geq 0)
\end{aligned}$$

## References

- ATCHLEY, W.R. & MARTIN, J. (1971). A morphometric analysis of differential sexual dimorphism in larvae of chironomus. *The Canadian Entomologist* **103**, 319–327.
- BANFIELD, J. & RAFTERY, A.E. (1993). Model-based Gaussian and non-Gaussian clustering. *Biometrics* **49**, 803–821.
- BOLTON, R.J. & KRZANOWSKI, W.J. (2003). Projection pursuit clustering for exploratory data analysis. *Journal of Computational and Graphical Statistics* **12**, 121–142.
- CELEUX, G. & GOVAERT, G. (1995). Gaussian parsimonious clustering models. *Pattern Recognition* **28**, 781–793.
- CHACÓN, J.E. (2020). The modal age of statistics. *International Statistical Review* **88**, 122–141. doi: <https://doi.org/10.1111/insr.12340>.
- DEMPSTER, A.P., LAIRD, N.M. & RUBIN, D.B. (1977). Maximum likelihood from incomplete data via the EM algorithm (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **39**, 1–38.
- FRALEY, C. & RAFTERY, A.E. (1998). How many clusters? Which clustering method? Answers via model based cluster analysis. *The Computer Journal* **41**, 578–588.
- FRIEDMAN, J.H. (1987). Exploratory projection pursuit. *Journal of the American Statistical Association* **82**, 249–266.
- FRIEDMAN, J.H. & MEULMAN, J.J. (2004). Clustering objects on subsets of attributes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **66**, 815–849.
- FRIEDMAN, J.H. & TUKEY, J.W. (1974). A projection pursuit algorithm for exploratory data analysis. *IEEE Transactions on Computers* **100**, 881–890.
- GHAHRAMANI, Z. & HINTON, G.E. (1997). The EM algorithm for factor analyzers. In *Technical Report CRG-TR-96-1*. University of Toronto.
- HARTIGAN, J.A. (1975). *Clustering Algorithms*. New York: John Wiley & Sons.
- HUBERT, L. & ARABIE, P. (1985). Comparing partitions. *Journal of Classification* **2**, 193–218.
- LI, J., RAY, S. & LINDSAY, B.G. (2007). A nonparametric statistical approach to clustering via mode identification. *Journal of Machine Learning Research* **8**, 1687–1723.
- MCINNES, L., HEALY, J. & MELVILLE, J. (2018). UMAP: Uniform manifold approximation and projection for dimension reduction. [arXiv:1802.03426](https://arxiv.org/abs/1802.03426) [stat.ML].
- McLACHLAN, G. & KRISHNAN, T. (2008). *The EM Algorithm and Extensions*, 2nd edn. Hoboken, New Jersey: Wiley-Interscience.
- McLACHLAN, G.J., LEE, S.X. & RATHNAYAKE, S.I. (2019). Finite mixture models. *Annual Review of Statistics and its Application* **6**, 355–378.
- McLACHLAN, G.J. & PEEL, D. (2000a). *Finite Mixture Models*. New York: Wiley.
- McLACHLAN, G.J. & PEEL, D. (2000b). Mixtures of factor analyzers. In *Proceedings of the Seventeenth International Conference on Machine Learning*, pp. 599–606. Burlington, MA: Morgan Kaufmann.
- MENARDI, G. (2016). A review on modal clustering. *International Statistical Review* **84**, 413–433.
- MONTANARI, A. & VIROLI, C. (2010). Heteroscedastic factor mixture analysis. *Statistical Modelling* **10**, 441–460.
- ROWEIS, S.T. & SAUL, L.K. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science* **290**, 2323–2326.
- SCHWARTZ, G. (1978). Estimating the dimension of a model. *Annals of Statistics* **6**, 31–38.
- SCRUCCA, L. (2021). A fast and efficient modal EM algorithm for Gaussian mixtures. *Statistical Analysis and Data Mining* <https://doi.org/10.1002/sam.11527>
- SCRUCCA, L. & SERAFINI, A. (2019). Projection pursuit based on Gaussian mixtures and evolutionary algorithms. *Journal of Computational and Graphical Statistics* **28**, 847–860.

- SERAFINI, A. & SCRUCCA, L. (2019). ppgmmga: Projection Pursuit Based on Gaussian Mixtures and Evolutionary Algorithms. Available from URL: <https://CRAN.R-project.org/package=ppgmmga>. R package version 1.2.
- STREULI, H. (1973). Der heutige stand der kaffeechemie. In *Association Scientifique Internationale du Cafe, 6th International Colloquium on Coffee Chemistry, Bogota, Columbia*, vol. 61. pp. 61–72.
- VAN DER MAATEN, L. & HINTON, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research* **9**, 2579–2605.