An Evolutionary Algorithm with Crossover and Mutation for Model-Based Clustering

Sharon M. McNicholas*, Paul D. McNicholas* and Daniel A. Ashlock**

*Department of Mathematics and Statistics, McMaster University, Ontario, Canada.
**Department of Mathematics and Statistics, University of Guelph, Ontario, Canada.

Abstract

An evolutionary algorithm (EA) is developed as an alternative to the EM algorithm for parameter estimation in model-based clustering. This EA facilitates a different search of the fitness landscape, i.e., the likelihood surface, utilizing both crossover and mutation. Furthermore, this EA represents an efficient approach to "hard" model-based clustering and so it can be viewed as a sort of generalization of the k-means algorithm, which is itself equivalent to a restricted Gaussian mixture model. The EA is illustrated on several datasets, and its performance is compared to other hard clustering approaches and model-based clustering via the EM algorithm.

Keywords: clustering; crossover; evolutionary algorithm; mixture models; mutation; model-based clustering.

1 Introduction

Model-based clustering is the use of mixture models for clustering. Recent reviews of model-based clustering are given by Bouveyron and Brunet-Saumard (2014) and McNicholas (2016b), while extensive details on mixture models are provided by Titterington et al. (1985) and McLachlan and Peel (2000a). The expectation-maximization (EM) algorithm (Dempster et al., 1977) is commonly used to estimate parameters in model-based clustering problems; however, it is prone to becoming stuck at local maxima (see, e.g., Titterington et al., 1985). As an alternative to using the EM algorithm to estimate parameters, an evolutionary algorithm (EA) is developed herein. This facilitates a different search of the fitness landscape (i.e., the likelihood surface) and can be viewed as an algorithm for "hard" model-based clustering. Let

$$z_{ig} = \begin{cases} 1 & \text{if observation } \mathbf{x}_i \text{ belongs to component } g, \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

so that the goal of model-based clustering is to predict z_{ig} for each (p-dimensional) observation $\mathbf{x}_i$ and each component g. In this context, "hard" means that the prediction for z_{ig} is restricted to taking a value 0 or 1 and, for convenience, we use the notation $\tilde{z}_{ig} \in \{0, 1\}$ to denote the prediction for z_{ig} arising from our EA. This differs from the typical EM approach, where the prediction is a probability $\hat{z}_{ig} \in [0, 1]$ which may be reported as is (soft) or subjected to an *a posteriori* hardening. Our EA can also be viewed as a sort of generalization of the k-means algorithm. Celeux and Govaert (1992) show that the k-means algorithm is equivalent to a classification EM (CEM) algorithm for a Gaussian mixture model with equal mixing proportions and common spherical component covariances (i.e., component covariance matrices $\boldsymbol{\Sigma}_g = \lambda \mathbf{I}_p$, where $\lambda \in \mathbb{R}^+$ and $\mathbf{I}_p$ is the p-dimensional identity matrix) — Vermunt (2011) also presents an argument in this direction. Some historical context and further details on the CEM algorithm are provided by McLachlan (1982).

2 Background

2.1 Evolutionary Algorithms

Evolutionary computation is a paradigm in which a computer algorithm incorporates some of the elements of the biological theory of evolution. Evolutionary operations are performed on members of a population, which reproduce and create new population members. The new members replace less "fit" members from the previous generation, and the process is continued until some stopping criterion is met. Evolutionary operations include actions such as crossover and mutation. The measure of "fitness" used is determined by the goal of the evolution, i.e., the function that is being optimized. The field of evolutionary computation is interdisciplinary, with practitioners approaching it from a variety of different perspectives such as computer science, biology, and statistics. This leads to a situation where terminology is often ambiguous; the following explanation of the basic terminology used herein is based on the concepts as outlined in Ashlock (2010).

Evolution occurs when a population is subject to change over time. Population members from each generation undergo a selection process. Crossover is the combining of two data structures to produce at least one new structure. Two-point crossover occurs when two data structures are selected to be parents. Then, two positions (the same for each structure) are selected at random and the genetic material between these two positions is exchanged between the parents (see Table 1). The resulting structures are called offspring, or children. Mutation is the process by which random changes are made to a population member's structure and can be used to produce a constant supply of minor variation in the population over time. In fitness-based reproduction, solutions that are deemed to be fitter, based on a predetermined fitness function, are preferentially selected to reproduce.

Table 1: Illustration of 2-point crossover, similar to that used by Ashlock (2010).

Parent 1	xxxxxxxxxxxxxxxxxxxxxx
Parent 2	yyyyyyyyyyyyyyyyyyyyyy
Child 1	xxxxxxxyyyyyyyxxxxxxx
Child 2	yyyyyyyyxxxxxxxyyyyyyy

2.2 Classification

Classification is a mechanism by which group membership labels are assigned to unlabelled observations (the group may be a class or cluster). Unsupervised classification, or clustering, is the special case where all observations are *a priori* unlabelled or treated as such. A common definition of a cluster suggests that it occurs when observations are grouped together in such a way that members of one cluster are more similar to each other than to observations in other clusters. As McNicholas (2016a) mentions, such a definition might be considered flawed because it is satisfied by a solution that places each observation into its own cluster and a definition that casts a cluster as a component in a suitable finite mixture model may be preferable.

Finite mixture models lend themselves well to classification problems. Consider a finite mixture distribution so that $\mathbf{X}$ is a p-dimensional random vector that, for all $\mathbf{x} \subset \mathbf{X}$, has density of the form

$$f(\mathbf{x} \mid \boldsymbol{\vartheta}) = \sum_{g=1}^{G} \pi_g f_g(\mathbf{x} \mid \boldsymbol{\theta}_g), \tag{2}$$

where $\pi_g > 0$, such that $\sum_{g=1}^{G} \pi_g = 1$, are the mixing proportions, $f_1(\mathbf{x} \mid \boldsymbol{\theta}_g), \ldots, f_G(\mathbf{x} \mid \boldsymbol{\theta}_g)$ are the component densities, and $\boldsymbol{\vartheta} = (\boldsymbol{\pi}, \boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_G)$ is the vector of parameters where $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_G)$. Historically, in clustering applications, it has been customary to assume that the component densities are multivariate Gaussian (Wolfe, 1965) — as an aside, it is notable that much recent work focuses on mixtures of non-Gaussian distributions (recent examples include work by Bagnato et al., 2017; Wallace et al., 2018; Pesevski et al., 2018; Lin et al., 2018; Morris et al., 2019; Tortora et al., 2019; Wei et al., 2020). The Gaussian mixture density can be written

$$f(\mathbf{x} \mid \boldsymbol{\vartheta}) = \sum_{g=1}^{G} \pi_g \phi(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g), \tag{3}$$

where $\phi(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g)$ is the multivariate Gaussian density with mean $\boldsymbol{\mu}_g$ and covariance matrix $\boldsymbol{\Sigma}_g$.

Consider a clustering scenario so that we have unlabelled data $\mathbf{x}_1, \ldots, \mathbf{x}_n$, i.e., there are no known group membership labels. In this case, the Gaussian mixture model likelihood can be expressed as

$$\mathcal{L}(\boldsymbol{\vartheta} \mid \mathbf{x}) = \prod_{i=1}^{n} \sum_{g=1}^{G} \pi_g \phi(\mathbf{x}_i \mid \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g), \tag{4}$$

where π_g can be thought of as the *a priori* probability that observation $\mathbf{x}_i$ belongs in component g (McLachlan and Peel, 2000a; McNicholas, 2016a). To enable clustering, the notation z_{ig} is used to represent component (group) membership and has the same definition as before; see (1). Within the EM algorithm framework, at each iteration, the z_{ig} are replaced by their conditional expected values given the observed data, using the current estimates for the unknown parameters, and we denote this quantity

$$\hat{z}_{ig} := \mathbb{E}[Z_{ig} \mid \mathbf{x}_i, \hat{\boldsymbol{\vartheta}}] = \Pr[Z_{ig} = 1 \mid \mathbf{x}_i, \hat{\boldsymbol{\vartheta}}] = \frac{\hat{\pi}_g \phi(\mathbf{x}_i \mid \hat{\boldsymbol{\mu}}_g, \hat{\boldsymbol{\Sigma}}_g)}{\sum_{h=1}^{G} \hat{\pi}_h \phi(\mathbf{x}_i \mid \hat{\boldsymbol{\mu}}_h, \hat{\boldsymbol{\Sigma}}_h)}. \tag{5}$$

Once parameter estimation has been carried out, the predicted group memberships are obtained from the *a posteriori* probability that observation $\mathbf{x}_i$ belongs to component g — this is just given by $\hat{z}_{ig}$ evaluated at the (final) parameter estimates, see (5). Depending on the problem at hand, this probability may be reported as it is, in which case it is a soft classification, or the probability may be rounded to 0 or 1, i.e., hardened. Hardening is usually carried out by computing maximum *a posteriori* (MAP) classifications, i.e., $\text{MAP}\{\hat{z}_{ig}\} = 1$ if $g = \arg\max_k\{\hat{z}_{ik}\}$ and $\text{MAP}\{\hat{z}_{ig}\} = 0$ otherwise. Note that MAP classifications are reported for the EM algorithms in Section 4. Crucially, regardless of whether $\hat{z}_{ig}$ or $\text{MAP}\{\hat{z}_{ig}\}$ is ultimately returned, the parameter estimation algorithms most often used permit values $\hat{z}_{ig} \in [0, 1]$ as the algorithm iterates. This is true for the EM algorithm and some popular alternatives, including variational Bayes approximations (see, e.g., McGrory and Titterington, 2007; Subedi and McNicholas, 2014, 2019). The EA approach that is taken herein uses values $\tilde{z}_{ig} \in \{0, 1\}$ as the algorithm iterates, which is a fundamental difference between it and the more common approaches.

For model-based clustering, the Bayesian information criterion (BIC; Schwarz, 1978) is a commonly used technique for determining the number of components G (if unknown). We can write

$$\text{BIC} = 2l(\hat{\boldsymbol{\vartheta}}) - \rho \log n, \tag{6}$$

where $\hat{\boldsymbol{\vartheta}}$ is the maximum likelihood estimate of $\boldsymbol{\vartheta}$, $l(\hat{\boldsymbol{\vartheta}})$ is the maximized log-likelihood, ρ is the number of free parameters in the model, and n is the number of observations. Leroux (1992), Roeder and Wasserman (1997), Kass and Wasserman (1995) and Dasgupta and Raftery (1998), amongst others, provide arguments that support using the BIC for estimating the number of components in a mixture model. Alternatives have been suggested for model selection in model-based clustering, (see, e.g., Biernacki et al., 2000), but the BIC remains the most common approach.

2.3 Model-Based Clustering and k-Means Clustering

As mentioned in Section 1, Celeux and Govaert (1992) show that the k-means clustering algorithm is equivalent to a CEM algorithm for model-based clustering using a Gaussian mixture model with equal mixing proportions and component covariance structure

$$\boldsymbol{\Sigma}_g = \lambda \mathbf{I}_p, \tag{7}$$

where $\lambda \in \mathbb{R}^+$ and $\mathbf{I}_p$ is the p-dimensional identity matrix. The CEM algorithm uses estimates $\hat{z}_{ig} \in \{0, 1\}$ throughout the algorithm and maximizes the classification likelihood (see McLachlan and Peel, 2000a, Section 2.21 for further details). Although introduced well over a decade ago, the CEM algorithm never caught on because of the seemingly exaggerated difficulties associated with local maxima. Of particular concern is its propensity for getting "stuck" at local maxima — the EM algorithm suffers from the same problem but to a lesser extent. For a summary on the drawbacks of the CEM algorithm, see McLachlan and Peel (2000a, Section 2.21).

Note that the component covariance structure in (7) limits the associated Gaussian mixture model — and, by extension, k-means clustering — to spherical components of equal radius. Accordingly, such an approach will only be effective if the clusters are either roughly spheres of equal radius or well separated. Further to the latter situation, clustering where the clusters are well separated can be considered a trivial case and warrants no further consideration. The `x2` dataset from the `mixture` package (Browne and McNicholas, 2014b) for R (R Core Team, 2018) is an example of a seemingly easy clustering situation where k-means clustering does not work well. The results for k-means are shown in Figure 1.
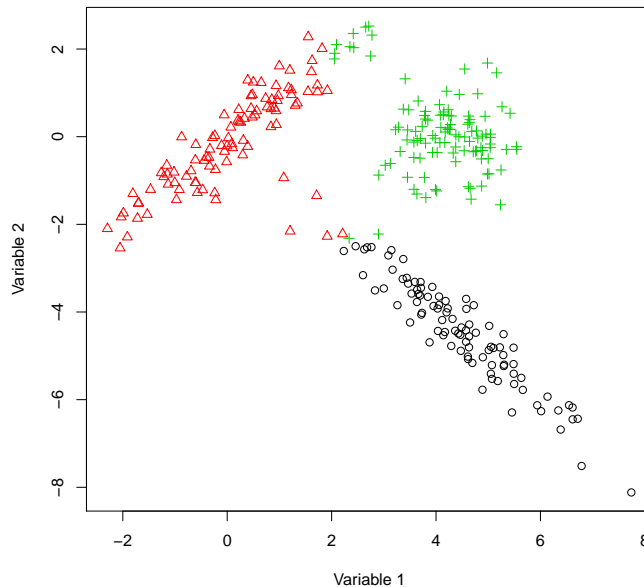

Figure 1: Scatterplot of the `x2` data, where plotting symbols and colours reflect the predicted classifications from k-means clustering.

The failure of k-means depicted in Figure 1 is directly attributable to the fact that, unless the clusters are well separated, it accommodates spherical clusters (of equal radius) and is not related to the fact that it is a hard clustering technique. Herein, an evolutionary algorithm is used to develop a hard clustering approach, based on a Gaussian mixture model, where

clusters have more flexibility in shape, volume, and orientation. Accordingly, in addition to being an alternative approach for parameter estimation in model-based clustering, this work can be viewed as an extension of k-means clustering to cases where the components need not be spherical.

3 Methodology

3.1 Model and Fitness Function

A mixture of multivariate Gaussian distributions is selected as the base model for the EA developed herein. As usual, z_{ig} is used to denote component membership labels; see (1). Because we are working within a clustering paradigm, all component memberships are *a priori* unknown or treated as such. While the EM algorithm maximizes the conditional expected value of the complete-data log-likelihood at each iteration, the EA developed herein has a fitness function based on the (observed) log-likelihood

$$l(\boldsymbol{\vartheta}) = \sum_{i=1}^{n} \log \left\{ \sum_{g=1}^{G} \pi_g \phi(\mathbf{x}_i \mid \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g) \right\}, \tag{8}$$

where $\boldsymbol{\vartheta} = (\pi_1, \ldots, \pi_G, \boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_G, \boldsymbol{\Sigma}_1, \ldots, \boldsymbol{\Sigma}_G)$ denotes the model parameters.

As our EA progresses, the estimated value of z_{ig} evolves. To avoid confusion with the expected values $\hat{z}_{ig}$ used in the EM algorithm (see Section 2.2), we continue to use $\tilde{z}_{ig}$ to denote the estimate of z_{ig} used in our EA. Accordingly, the estimated component membership of $\mathbf{x}_i$ in our EA is given by $\tilde{\mathbf{z}}_i = (\tilde{z}_{i1}, \ldots, \tilde{z}_{iG})$ for $\tilde{z}_{ig} \in \{0, 1\}$. The fitness function is just the log-likelihood (8) evaluated at the estimates

$$\tilde{\pi}_g = \frac{n_g}{n}, \qquad \tilde{\boldsymbol{\mu}}_g = \frac{1}{n_g} \sum_{i=1}^{n} \tilde{z}_{ig} \mathbf{x}_i, \qquad \tilde{\boldsymbol{\Sigma}}_g = \frac{1}{n_g} \sum_{i=1}^{n} \tilde{z}_{ig} (\mathbf{x}_i - \tilde{\boldsymbol{\mu}}_g)(\mathbf{x}_i - \tilde{\boldsymbol{\mu}}_g)', \tag{9}$$

where $n_g = \sum_{i=1}^{n} \tilde{z}_{ig}$.

3.2 Evolutionary Algorithm

In our EA, a number of single parents are used and each is cloned many times, with the cloned children reproducing as discussed here. For each child, i.e., each clone of a single parent, two observations are chosen at random and their $\tilde{\mathbf{z}}_i$ values are swapped. There is a check in the code, so that the swap only occurs if the estimated group memberships are different, i.e., if the $\tilde{\mathbf{z}}_i$ are different; if not, other observations are selected until two different $\tilde{\mathbf{z}}_i$ are found. The children of the different single parents never interbreed with each other. After one instance of crossover has been carried out on each cloned child, all of the children (plus the original few single parents) are put into one list in descending order of fitness.

The top few are selected to become the new generation of single parents. This crossover procedure will help avoid stopping at local maxima of the fitness surface, i.e., the likelihood surface. However, crossover alone will not suffice in clustering applications — to see why this is so, consider that it is not always possible to improve a clustering result by just swapping the membership label from one observation with that from another — so a mutation step is also carried out at each iteration. These iterations, of crossover followed by mutation, are repeated until our EA stagnates.

To crystallize the exact procedure followed in our EA, pseudocode is provided (Algorithm 1). Note that the code used herein was written in R and comments within the pseudocode in Algorithm 1 use the R comment style, i.e., `#`. Note also that, after the very first crossover step, the `parents` are simply the best `pars` elements in terms of fitness. Furthermore, note that `compute fitness` entails computing the estimates in (9) and then computing the log-likelihood (8). The greedy nature of the mutation step is clear; for a given parent, once a mutation increases the fitness, our EA moves on to the next parent. There is a nice general interpretation to this EA: the crossover step provides diversity while the mutation step allows fitness (clustering) improvements that cannot be facilitated by crossover alone. The effectiveness of our EA for traversing the fitness (likelihood) surface is illustrated in Section 4.

4 Illustrations

4.1 Overview and Performance Assessment

The purpose of these illustrations is to compare our EA to two well-established hard clustering approaches, i.e., k-means and k-medoids, and the EM algorithm for a Gaussian mixture model. For brevity, the latter shall be referred to as "the EM algorithm" hereafter. To facilitate comparison with our EA, the EM algorithm is run using the `gpcm()` function from the `mixture` package (Browne and McNicholas, 2014b) with `mnames="VVV"`. First, a simulated dataset is used (Section 4.2). The remaining analyses (Sections 4.3–4.5) focus on real data. For all analyses, our EAs are run with two parents: one parent is initialized using k-means and the other is initialized using k-medoids.

Although all of the illustrations in this section are carried out as real cluster analyses (i.e., the data are treated as unlabelled), the true labels are in fact known. Therefore, it is possible to compare the predicted classifications — for our EA, the final $\tilde{z}_{ig}$ values and, for the EM algorithm, the final $\hat{z}_{ig}$ values — with the true labels. We carry out this comparison using the adjusted Rand index (ARI; Hubert and Arabie, 1985), which is the Rand index (Rand, 1971) corrected for chance agreement. The Rand index is the ratio of pairwise agreements to total pairs. Arguments as to why the ARI should be used in this circumstance, as opposed to alternatives like the misclassification rate, are given by Steinley (2004).

Algorithm 1 Pseudocode for an EA for model-based clustering.

```
input: x, G, z, pars=2, stagnation, clones
# x is data matrix; G is number of components (groups);
# z is a list where each element is a matrix of tilde_z_ig values for one parent;
# pars is the number of parents (defaults to 2); stagnation is the stagnation value;
# clones is the number of clones
N = number of rows in x
stag=0
while stag < stagnation
    # First, crossover
    for a in 1 to pars
        for b in 1 to clones
            randomly select two unequal labels from parent a
            swap them to get clone (child) b from parent a
            compute fitness for clone (child) b from parent a
        end for
    end for
    sort parents plus children by descending fitness
    # The top four are now the parents
    if top four are unchanged from previous iteration
        stag ++
    else
        stag = 0
    end if
    # Now, mutation
    for a in 1 to pars
        rand = random permutation of 1,2,...,N
        for i in rand
            swap two distinct elements in label i
            if fitness increases
                break for (i in rand)
            end if
        end for
    end for
    if no mutation has increased log-likelihood
        stag = stag +1
    else
        sort parents by descending fitness
    end if
end while
return fitness values (log-likelihoods) and labels for the parents
```

4.2   x2 Data

As a first step, consider the `x2` dataset from the `mixture` package in R (Figure 2). The `x2` data are generated from a Gaussian mixture model with three components and were used by Browne and McNicholas (2014a) to illustrate their MM (majorization-minimization, in this case) algorithms — see Hunter and Lange (2004) for an overview of MM algorithms. As McNicholas (2016a, Chapter 2) points out, the `x2` data are a good illustration of data generated from a mixture model where the "best" clustering result clearly does not correspond perfectly to the labels from the generating model.
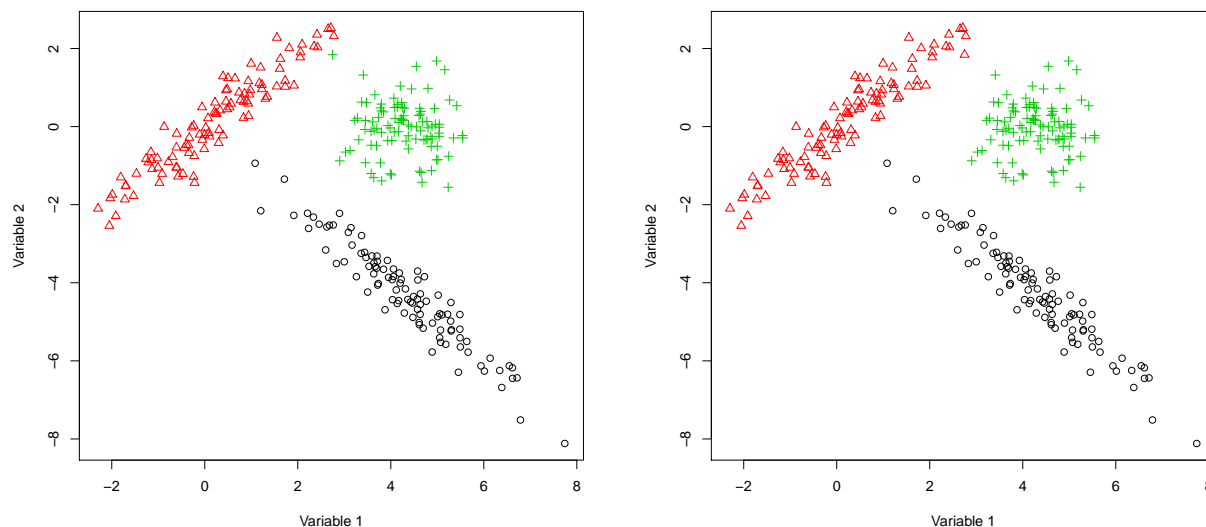

Figure 2: Scatterplots of the `x2` data, where plotting symbols and colours reflect the components from the model used to generate the data (left) or the predicted classifications from our EA (right).

Our EA approach is applied to the `x2` data with the following settings (see Algorithm 1): `stagnation=3` and `clones=10`. The result (Figure 2) indicates that "perfect" classification performance is attained — note, again, that this does not quite correspond to the generating model. It is interesting to consider the performance of other hard clustering approaches and so k-means and k-medoids are also applied to these data. The results (Figure 3) indicate that neither approach performs as well as our EA. Note that the EM algorithm gives the same classification results as our EA for the `x2` data.

4.3   Female Voles Data

The female voles (`f.voles`) data are available in the `Flury` package (Flury, 2012) for R. They contain six morphometric measurements, as well as age, for 86 female voles from two species: *Microtus californicus* and *Microtus ochrogaster* (Figure 4).
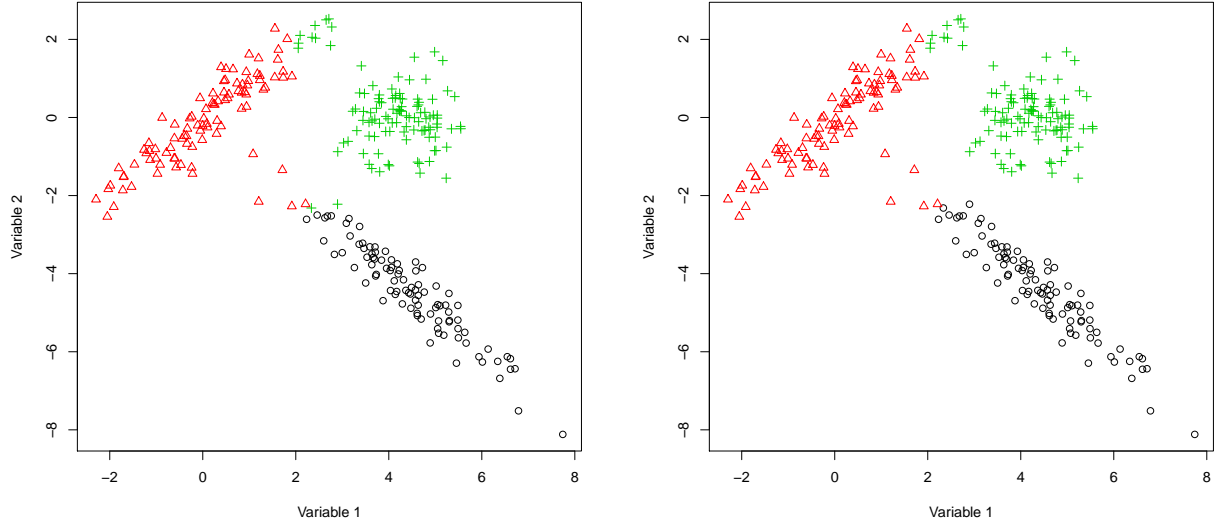
Figure 3: Scatterplots of the x2 data, where plotting symbols and colours reflect the predicted classifications from k-means clustering (left) or k-medoids clustering (right).

The EA approach introduced herein is applied to these data with stagnation $\in \{3, 4, 5\}$ and clones $\in \{10, 20, 30, 40\}$. Over all runs, identical and excellent classification performance was obtained, with just one misclassification (Table 2; ARI = 0.953). The respective classification performance of k-means (Tables 3; ARI = 0.737) and k-medoids (Tables 4; ARI = 0.658) on these data is notably inferior. The EM algorithm gives slightly inferior classification performance, misclassifying one additional vole (Table 5; ARI = 0.908) when compared to our EA (Table 2; ARI = 0.953).

Table 2: Cross-tabulation of the predicted classifications (A, B) from our EA versus species for the female voles data.

	A	B
Microtus californicus	41	0
Microtus ochrogaster	1	44

Table 3: Cross-tabulation of the predicted classifications (A, B) from k-means versus species for the female voles data.

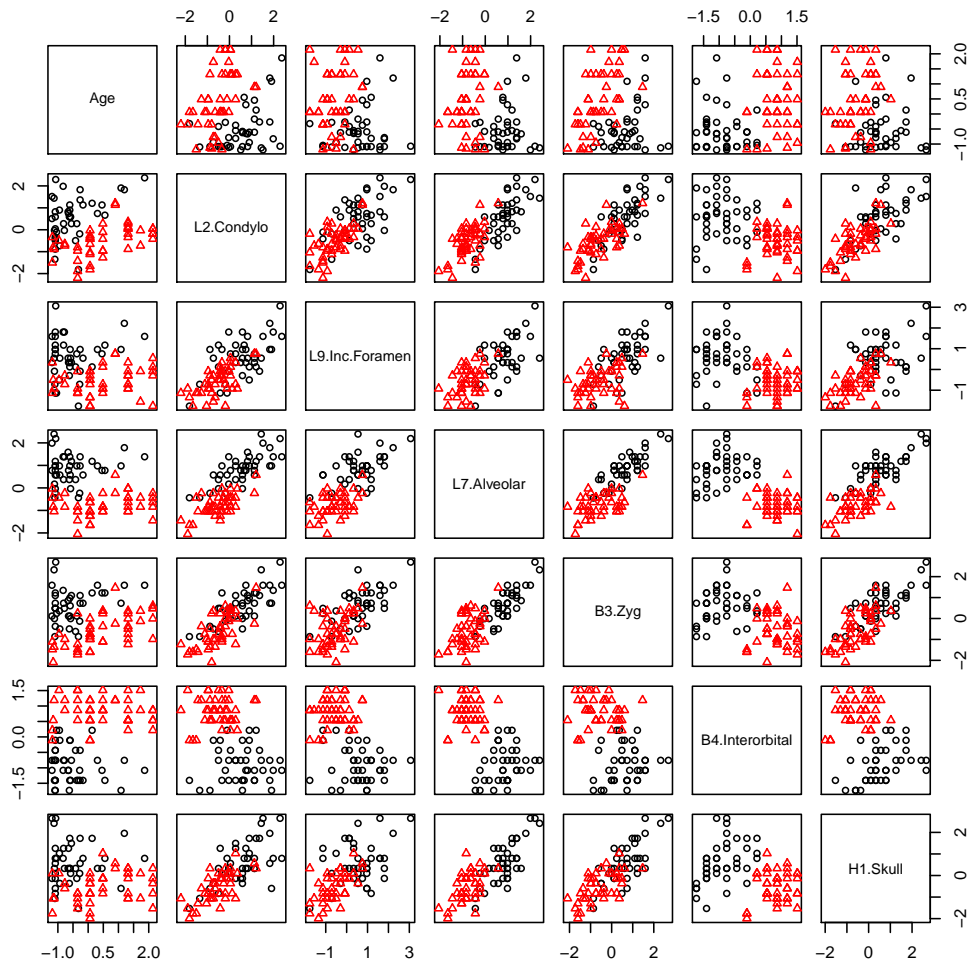|	A	B
Microtus californicus	36	5
Microtus ochrogaster	1	44

Figure 4: Scatterplot of the female voles data, where plotting symbols and colours reflect species.

Table 4: Cross-tabulation of the predicted classifications (A, B) from k-medoids versus species for the female voles data.

	A	B
Microtus californicus	34	7
Microtus ochrogaster	1	44

Table 5: Cross-tabulation of the predicted classifications (A, B) from the EM algorithm versus species for the female voles data.

	A	B
Microtus californicus	41	0
Microtus ochrogaster	2	43

4.4 Banknote Data

The `banknote` data are available from the `mclust` package (Fraley et al., 2012) in R. They contain six measurements, all in mm, on 100 genuine and 100 counterfeit Swiss 1000-franc banknotes (Figure 5).
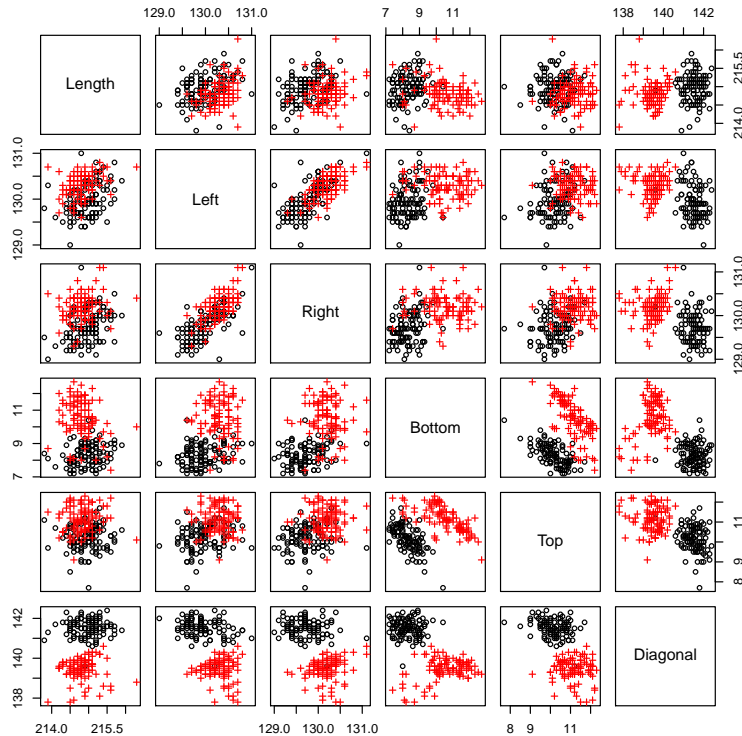

Figure 5: Scatterplot of the Swiss banknote data, where plotting symbols and colours reflect true class, i.e., genuine or counterfeit.

Our EA is applied to these data with `stagnation` $\in \{3, 4, 5\}$ and `clones` $\in \{10, 20, 30, 40\}$. Over all runs, identical and excellent classification performance was obtained, with just one misclassification (Table 6; ARI = 0.980). The respective classification performance of k-means (Table 7; ARI = 0.846) and k-medoids (Table 8; ARI = 0.941) on these data is somewhat inferior. The EM algorithm gives the same classification results as our EA on these data (Table 6; ARI = 0.980).

Table 6: Cross-tabulation of the predicted classifications (A, B) from our EA versus true class for the Swiss banknote data.

	A	B
Counterfeit	100	0
Genuine	1	99

Table 7: Cross-tabulation of the predicted classifications (A, B) from k-means versus true class for the Swiss banknote data.

	A	B
Counterfeit	100	0
Genuine	8	92

Table 8: Cross-tabulation of the predicted classifications (A, B) from k-medoids versus true class for the Swiss banknote data.

	A	B
Counterfeit	100	0
Genuine	3	97

4.5 Italian Wine Data

Thus far, the real datasets considered have contained just two classes. To move beyond this, consider the Italian wine data that were collected by Forina et al. (1986). A subset of these data, containing 13 chemical and physical properties of three cultivars (Barolo, Grignolino, Barbera) from the Piedmont region of Italy, is available in the `gclus` package (Hurley, 2004) for R.

Table 9: The thirteen chemical and physical properties of Italian wines available in `gclus`.

Magnesium	Malic acid	Total phenols
Alcohol	Ash	Alcalinity of ash
Flavonoids	Nonflavonoid phenols	Proanthocyanins
Color Intensity	Hue	Proline
OD_{280}/OD_{315} of diluted wines		

The EA approach introduced herein is applied to these data with `stagnation` $\in \{3, 4, 5\}$ and `clones` $\in \{10, 20, 30, 40\}$. Over all runs, identical and excellent classification performance was obtained, with just one misclassification (Table 10; ARI = 0.982). The respective classification performance of k-means (Tables 11; ARI = 0.897) and k-medoids (Tables 12; ARI = 0.741) on these data is notably inferior. For the Italian wine data, the EM algorithm gives slightly inferior classification performance (Table 13; ARI = 0.945) compared to our EA (Table 10; ARI = 0.982).

Table 10: Cross-tabulation of the predicted classifications (A, B, C) from our EA versus true class for the Italian wine data.

	A	B	C
Barolo	59	0	0
Grignolino	1	70	0
Barbera	0	0	48

Table 11: Cross-tabulation of the predicted classifications (A, B, C) from k-means versus true class for the Italian wine data.

	A	B	C
Barolo	59	0	0
Grignolino	3	65	3
Barbera	0	0	48

Table 12: Cross-tabulation of the predicted classifications (A, B, C) from k-medoids versus true class for the Italian wine data.

	A	B	C
Barolo	59	0	0
Grignolino	15	55	1
Barbera	0	0	48

Table 13: Cross-tabulation of the predicted classifications (A, B, C) from the EM algorithm versus true class for the Italian wine data.

	A	B	C
Barolo	59	0	0
Grignolino	3	68	0
Barbera	0	0	48

5 Discussion

An EA has been introduced for model-based clustering. Each iteration of our EA uses a crossover step followed by a (greedy) mutation step; no comparable approach has been taken for Gaussian mixture models. In fact, the closest approach uses mutations only, ignoring crossover (see Andrews and McNicholas, 2013). The clustering philosophy associated with our approach is that of hard clustering, i.e., $\tilde{z}_{ig} \in \{0,1\}$. At no point in our parameter estimation scheme do we entertain soft values. This is in contrast to the commonly used EM algorithm, which uses $\hat{z}_{ig} \in [0,1]$.

In terms of future work, there is much that could be done. For one, our EA uses an unconstrained (VVV) covariance structure and the other 11 covariance structures used by Celeux and Govaert (1995) could easily be implemented — for practical reasons, comparison with algorithms other than the EM might be preferable for two of these 11 (see Browne and McNicholas, 2014a,c, for alternative algorithms). For each of the toy examples (Sections 4.3–4.5), our EA returned the same results for the different values of `stagnation` and `clones` used; however, one would not expect that to be the case in general and so further consideration of model selection is needed. Comparison of our EA to the EM algorithm was limited in scope, i.e., we considered only classification performance, and extent; accordingly, a more in-depth study is required. Extension to higher dimensional problems should be considered and could be achieved in a number of ways, including by extending our EA to mixtures of factor analyzers (Ghahramani and Hinton, 1997; McLachlan and Peel, 2000b;

McNicholas and Murphy, 2008, 2010). It would also be of interest to consider how an EA approach similar to the one introduced herein would work within the mixture discriminant analysis (see Scott, 1992; McLachlan, 1992; Fraley and Raftery, 2002), semi-supervised (see, e.g., Dean et al., 2006; McNicholas, 2010) or fractionally supervised classification (see Vrbik and McNicholas, 2015; Gallaugher and McNicholas, 2019b) frameworks. Although our EA is based on a Gaussian mixture model, an analogous approach could be taken for mixtures of non-Gaussian distributions and this will be a topic for future work. Another avenue for future work is using an analogous EA approach within the matrix variate mixture setting, where there has been significant work of late (Gallaugher and McNicholas, 2018, 2019a, 2020; Melnykov and Zhu, 2018, 2019; Sarkar et al., 2020; Murray et al., 2020). Finally, it will be of interest to compare the EA approach developed herein to the variational approximations-deviance information criterion rubric used by Subedi and McNicholas (2019).

Acknlwledgements

References

Andrews, J. L. and P. D. McNicholas (2013). Using evolutionary algorithms for model-based clustering. *Pattern Recognition Letters 34*, 987–992.

Ashlock, D. (2010). *Evolutionary Computation for Modeling and Optimization.* New York: Springer-Verlag.

Bagnato, L., A. Punzo, and M. G. Zoia (2017). The multivariate leptokurtic-normal distribution and its application in model-based clustering. *Canadian Journal of Statistics 45*(1), 95–119.

Biernacki, C., G. Celeux, and G. Govaert (2000). Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Transactions on Pattern Analysis and Machine Intelligence 22*(7), 719–725.

Bouveyron, C. and C. Brunet-Saumard (2014). Model-based clustering of high-dimensional data: A review. *Computational Statistics and Data Analysis 71*, 52–78.

Browne, R. P. and P. D. McNicholas (2014a). Estimating common principal components in high dimensions. *Advances in Data Analysis and Classification 8*(2), 217–226.

Browne, R. P. and P. D. McNicholas (2014b). *mixture: Mixture Models for Clustering and Classification*. R package version 1.1.

Browne, R. P. and P. D. McNicholas (2014c). Orthogonal Stiefel manifold optimization for eigen-decomposed covariance parameter estimation in mixture models. *Statistics and Computing 24*(2), 203–210.

Celeux, G. and G. Govaert (1992). A classification EM algorithm for clustering and two stochastic versions. *Computational Statistics and Data Analysis 14*(3), 315–332.

Celeux, G. and G. Govaert (1995). Gaussian parsimonious clustering models. *Pattern Recognition 28*(5), 781–793.

Dasgupta, A. and A. E. Raftery (1998). Detecting features in spatial point processes with clutter via model-based clustering. *Journal of the American Statistical Association 93*, 294–302.

Dean, N., T. B. Murphy, and G. Downey (2006). Using unlabelled data to update classification rules with applications in food authenticity studies. *Journal of the Royal Statistical Society: Series C 55*(1), 1–14.

Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B 39*(1), 1–38.

Flury, B. (2012). *Flury: Data Sets from Flury, 1997*. R package version 0.1-3.

Forina, M., C. Armanino, M. Castino, and M. Ubigli (1986). Multivariate data analysis as a discriminating method of the origin of wines. *Vitis 25*, 189–201.

Fraley, C. and A. E. Raftery (2002). Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association 97*(458), 611–631.

Fraley, C., A. E. Raftery, T. B. Murphy, and L. Scrucca (2012). mclust version 4 for R: Normal mixture modeling for model-based clustering, classification, and density estimation. Technical Report 597, Department of Statistics, University of Washington, Seattle, WA.

Gallaugher, M. P. B. and P. D. McNicholas (2018). Finite mixtures of skewed matrix variate distributions. *Pattern Recognition 80*, 83–93.

Gallaugher, M. P. B. and P. D. McNicholas (2019a). Mixtures of skewed matrix variate bilinear factor analyzers. *Advances in Data Analysis and Classification*. To appear, https://doi.org/10.1007/s11634-019-00377-4.

Gallaugher, M. P. B. and P. D. McNicholas (2019b). On fractionally-supervised classification: Weight selection and extension to the multivariate t-distribution. *Journal of Classification 36*(2), 232–265.

Gallaugher, M. P. B. and P. D. McNicholas (2020). Parsimonious mixtures of matrix variate bilinear factor analyzers. In T. Imaizumi, A. Nakayama, and S. Yokoyama (Eds.), *Advanced Studies in Behaviormetrics and Data Science: Essays in Honor of Akinori Okada*, pp. 177–196. Singapore: Springer.

Ghahramani, Z. and G. E. Hinton (1997). The EM algorithm for factor analyzers. Technical Report CRG-TR-96-1, University Of Toronto, Toronto.

Hubert, L. and P. Arabie (1985). Comparing partitions. *Journal of Classification 2*(1), 193–218.

Hunter, D. L. and K. Lange (2004). A tutorial on MM algorithms. *The American Statistician 58*(1), 30–37.

Hurley, C. (2004). Clustering visualizations of multivariate data. *Journal of Computational and Graphical Statistics 13*(4), 788–806.

Kass, R. E. and L. Wasserman (1995). A reference Bayesian test for nested hypotheses and its relationship to the Schwarz criterion. *Journal of the American Statistical Association 90*(431), 928–934.

Leroux, B. G. (1992). Consistent estimation of a mixing distribution. *The Annals of Statistics 20*(3), 1350–1360.

Lin, T.-I., W.-L. Wang, G. J. McLachlan, and S. X. Lee (2018). Robust mixtures of factor analysis models using the restricted multivariate skew-t distribution. *Statistical Modelling 18*, 50–72.

McGrory, C. and D. Titterington (2007). Variational approximations in Bayesian model selection for finite mixture distributions. *Computational Statistics and Data Analysis 51*(11), 5352–5367.

McLachlan, G. J. (1982). The classification and mixture maximum likelihood approaches to cluster analysis. In P. R. Krishnaiah and L. Kanal (Eds.), *Handbook of Statistics*, Volume 2, pp. 199–208. Amsterdam: North-Holland.

McLachlan, G. J. (1992). *Discriminant Analysis and Statistical Pattern Recognition*. New Jersey: John Wiley & Sons.

McLachlan, G. J. and D. Peel (2000a). *Finite Mixture Models*. New York: John Wiley & Sons.

McLachlan, G. J. and D. Peel (2000b). Mixtures of factor analyzers. In *Proceedings of the Seventh International Conference on Machine Learning*, San Francisco, pp. 599–606. Morgan Kaufmann.

McNicholas, P. D. (2010). Model-based classification using latent Gaussian mixture models. *Journal of Statistical Planning and Inference 140*(5), 1175–1181.

McNicholas, P. D. (2016a). *Mixture Model-Based Classification.* Boca Raton: Chapman & Hall/CRC Press.

McNicholas, P. D. (2016b). Model-based clustering. *Journal of Classification 33*(3), 331–373.

McNicholas, P. D. and T. B. Murphy (2008). Parsimonious Gaussian mixture models. *Statistics and Computing 18*(3), 285–296.

McNicholas, P. D. and T. B. Murphy (2010). Model-based clustering of microarray expression data via latent Gaussian mixture models. *Bioinformatics 26*(21), 2705–2712.

Melnykov, V. and X. Zhu (2018). On model-based clustering of skewed matrix data. *Journal of Multivariate Analysis 167*, 181–194.

Melnykov, V. and X. Zhu (2019). Studying crime trends in the USA over the years 2000–2012. *Advances in Data Analysis and Classification 13*(1), 325–341.

Morris, K., A. Punzo, P. D. McNicholas, and R. P. Browne (2019). Asymmetric clusters and outliers: Mixtures of multivariate contaminated shifted asymmetric Laplace distributions. *Computational Statistics and Data Analysis 132*, 145–166.

Murray, P. M., R. P. Browne, and P. D. McNicholas (2020). Mixtures of hidden truncation hyperbolic factor analyzers. *Journal of Classification 37*(2). To appear.

Pesevski, A., B. C. Franczak, and P. D. McNicholas (2018). Subspace clustering with the multivariate-t distribution. *Pattern Recognition Letters 112*(1), 297–302.

R Core Team (2018). *R: A Language and Environment for Statistical Computing.* Vienna, Austria: R Foundation for Statistical Computing.

Rand, W. M. (1971). Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association 66*(336), 846–850.

Roeder, K. and L. Wasserman (1997). Practical Bayesian density estimation using mixtures of normals. *Journal of the American Statistical Association 92*, 894–902.

Sarkar, S., X. Zhu, V. Melnykov, and S. Ingrassia (2020). On parsimonious models for modeling matrix data. *Computational Statistics and Data Analysis 142*.

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics 6*, 461–464.

Scott, D. W. (1992). *Multivariate Density Estimation.* New York: Wiley.

Steinley, D. (2004). Properties of the Hubert-Arabie adjusted Rand index. *Psychological Methods 9*, 386–396.

Subedi, S. and P. D. McNicholas (2014). Variational Bayes approximations for clustering via mixtures of normal inverse Gaussian distributions. *Advances in Data Analysis and Classification 8*(2), 167–193.

Subedi, S. and P. D. McNicholas (2019). A variational approximations-DIC rubric for parameter estimation and mixture model selection within a family setting. *Journal of Classification*. To appear, https://doi.org/10.1007/s00357-019-09351-3.

Titterington, D. M., A. F. M. Smith, and U. E. Makov (1985). *Statistical Analysis of Finite Mixture Distributions*. Chichester: John Wiley & Sons.

Tortora, C., B. C. Franczak, R. P. Browne, and P. D. McNicholas (2019). A mixture of coalesced generalized hyperbolic distributions. *Journal of Classification 36*(1), 26–57.

Vermunt, J. K. (2011). *K*-means may perform as well as mixture model clustering but may also be much worse: Comment on Steinley and Brusco (2011). *Psychological Methods 16*(1), 82–88.

Vrbik, I. and P. D. McNicholas (2015). Fractionally-supervised classification. *Journal of Classification 32*(3), 359–381.

Wallace, M. L., D. J. Buysse, A. Germain, M. H. Hall, and S. Iyengar (2018). Variable selection for skewed model-based clustering: Application to the identification of novel sleep phenotypes. *Journal of the American Statistical Association 113*(521), 95–110.

Wei, Y., Y. Tang, and P. D. McNicholas (2020). Flexible high-dimensional unsupervised learning with missing data. *IEEE Transactions on Pattern Analysis and Machine Intelligence 42*(3), 610–621.

Wolfe, J. H. (1965). A computer program for the maximum-likelihood analysis of types. USNPRA Technical Bulletin 65-15, U.S. Naval Personal Research Activity, San Diego.