Regularization and Global Optimization in Model-Based Clustering

Raphael Araujo Sampaio[a,b*]**, Joaquim Dias Garcia**[a,c]**, Marcus Poggi**[b]**, Thibaut Vidal**[b,d]

[a] PSR, Rio de Janeiro, Brazil

[b] Departamento de Informática, Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio), Brazil

[c] LAMPS, Departamento de Engenharia Elétrica, PUC-Rio, Brazil

[d] CIRRELT & SCALE-AI Chair in Data-Driven Supply Chains, Department of Mathematical and Industrial Engineering, Polytechnique Montréal, Canada

rsampaio@psr-inc.com, joaquim@psr-inc.com, poggi@inf.puc-rio.br, thibaut.vidal@polymtl.ca

Abstract. Due to their conceptual simplicity, k-means algorithm variants have been extensively used for unsupervised cluster analysis. However, one main shortcoming of these algorithms is that they essentially fit a mixture of identical spherical Gaussians to data that vastly deviates from such a distribution. In comparison, general Gaussian Mixture Models (GMMs) can fit richer structures but require estimating a quadratic number of parameters per cluster to represent the covariance matrices. This poses two main issues: (i) the underlying optimization problems are challenging due to their larger number of local minima, and (ii) their solutions can overfit the data. In this work, we design search strategies that circumvent both issues. We develop efficient global optimization algorithms for general GMMs, and we combine these algorithms with regularization strategies that avoid overfitting. Through extensive computational analyses, we observe that global optimization or regularization in isolation does not substantially improve cluster recovery. However, combining these techniques permits a completely new level of performance previously unachieved by k-means algorithm variants, unraveling vastly different cluster structures. These results shed new light on the current status quo between GMM and k-means methods and suggest the more frequent use of general GMMs for data exploration. To facilitate such applications, we provide open-source code as well as Julia packages (UnsupervisedClustering.jl and RegularizedCovarianceMatrices.jl) implementing the proposed techniques.

Keywords. Clustering; Gaussian Mixture Models; Regularization; Global Optimization; Hybrid Genetic Algorithm.

* Corresponding author

arXiv:2302.02450v1 [cs.LG] 5 Feb 2023

1. Introduction

Unsupervised clustering is the process of splitting unlabeled data into homogeneous groups with similar samples. This task is essential for data preparation and exploration. Moreover, it is linked with numerous applications in social media analysis, image processing, text analysis, and bioinformatics, among many others [28]. As a consequence of the diversity of applications, data sources, similarity measures, goals, and constraints, clustering encompasses a vast universe of methods.

Model-based clustering approaches hold a central place due to their well-grounded statistical principles and their controllable parsimony [8, 37]. These approaches assume that each cluster corresponds to a statistical process represented by a probability density function (e.g., mixtures of multivariate Gaussians). Maximum likelihood estimators of the statistical processes are used to obtain the cluster labels, which are unobserved latent parameters.

Restricting the Gaussian mixtures to be spherical (identity matrix as covariance) leads to a maximum likelihood problem almost equivalent to the minimum sum-of-squares clustering problem [6]. The classical k-means algorithm [23, 34] is a well-known local optimizer for this problem. It optimizes the $\mathcal{O}(Kd)$ model parameters representing the mean (or centroid) of K Gaussian distributions (i.e., the center of each cluster in a feature space of dimension d). The covariance matrix of these Gaussians is the identity matrix in the case of k-means. Due to their speed and simplicity, k-means and their variations have been widely used.

Assuming only spherical Gaussian mixtures is a rough approximation of real datasets. More general models with arbitrary covariance matrices have been studied to better fit the data, leading to *elliptical* variants of the k-means algorithms [9, 57, 41] and the well-known expectation-maximization (EM) clustering algorithms [12, 36] typically used to search for the maximum likelihood of the Gaussian mixture.

Although more general, the joint estimation of covariances and means within a Gaussian mixture model poses at least two fundamental issues. Firstly, the covariance matrices can become ill-conditioned during likelihood maximization without proper regularization. The ill-conditioned matrices create numerical instabilities, prevent matrix inversion, and increase the risk of overfitting the data instead of revealing useful structure data. Secondly, due to its larger number of parameters, likelihood maximization is much harder to solve in a general form, and the presence of numerous local maxima cripples solution methods. The difficulty of this optimization problem is the main reason why general EM and elliptical k-means variants have been previously qualified as a "failed opportunity" in the clustering domain [9].

In this study, we revisit some of these methods, considering different regularization techniques and

global optimization methods, to better understand to what extent and with which components general EM can be effectively used. Indeed, significant progress has been made in covariance matrix estimation, and many regularized estimators have been proposed. We attempt to use such sophisticated regularized estimators within EM-based algorithms. However, despite these regularization techniques, we observe that general EM models' performance remains underwhelming in feature spaces of dimensions 20 or more. This lack of performance apparently relates to the second issue, i.e., the presence of numerous local maxima. To circumvent this issue, we extend previous works focused on generating promising starting points for EM algorithms [4] and propose more sophisticated global optimization heuristics relying on recombination operations inspired by the study of [21]. As visible in our experiments, this second improvement counterbalances the weaknesses of the general EM, leading to high-quality solutions in a much larger regime for datasets with fewer samples relative to the feature space dimension.

To summarize, the contributions of this paper are the following.

- We test the general EM model with different regularization strategies, identifying the most successful approaches in terms of solution quality and computational effort.

- We propose advanced global optimization algorithms that permit an escape from local maxima and effectively combine them with the regularization strategies. The proposed methods can be viewed as a multi-start EM in which the initial parameters of the clusters (covariance and mean) are carefully inherited and recombined from previously-found solutions.

- We demonstrate, on controlled synthetic data, that the *joint* use of sophisticated regularization and global optimization techniques allows for the recovery of the original clusters much more systematically, outperforming classical spherical EM or k-means on a wider range of datasets.

- We compare the proposed methods in real data sets and verify that the regularized methods with global optimization achieve better cluster recovery.

This paper is structured as follows. Section 2 discusses fundamental notions and reviews the related literature on model-based clustering, focusing principally on regularization and global optimization. Section 3 presents the proposed methodology, including a global optimization technique and different regularization methods. Section 4 presents our computational experiments on synthetic and real datasets. Finally, Section 5 concludes and points to directions for future research.

2. Fundamental Notions and Related Studies

Let X be a dataset containing n data points (i.e., samples), $X = \{\boldsymbol{x}_1, ..., \boldsymbol{x}_n\}$, in which each point $\boldsymbol{x}_i = \{x_i^{(1)}, ..., x_i^{(d)}\}$ is represented as a vector of d features. We consider an unsupervised learning context in which the samples do not have known labels (i.e., groups), and the only available information is the number of these k labels.

We aim to find an underlying structure in X in the form of clusters. More precisely, we should distribute these samples into k labels so that the samples that are most similar to each other belong to the same cluster, and those with low similarity belong to different clusters.

Clustering algorithms can be broadly divided into two leading families: *soft* assignment methods output points-to-cluster assignments in the form of a probability distribution, whereas *hard* assignment methods associate each data point to a single cluster. Moreover, since clustering is a very general task, there exists, within each sub-category, various approaches, which can lead to the discovery of different solutions and underlying structures [15].

Expectation Maximization. Among soft assignment clustering methods, Expectation-Maximization (EM) algorithms based on Gaussian Mixture Models (GMM) hold a prominent place due to their generality and simplicity. Proposed by Dempster et al. [12], EM is a general solution approach for maximum likelihood estimation (MLE) problems with missing data [also see 5, 7, 36]. EM jointly estimates model parameters and missing data (the sample label in the case of clustering). In its most general form, EM searches for the model parameters θ that maximize the expected value of a log-likelihood function, $\log \mathcal{L}(\theta|X,Y)$, which depends on available information, X, and hidden information, Y, about the samples. Y has a posterior probability distribution as a function of X and θ, i.e., $Y|X,\theta$. Therefore, we take the expectation concerning Y conditioned on both X and θ:

$$\max_{\theta} \quad \mathbb{E}_{Y|X,\theta}[\log \mathcal{L}(\theta|X,Y)]. \tag{1}$$

As their name indicates, EM algorithms include two steps: *expectation* and *maximization*. In the expectation step, parameters θ are fixed to calculate a posterior of Y given X and θ. Consequently, we can obtain the expected log-likelihood function $\theta \to Q(\theta)$. In any iteration t, $Q^{(t)}(\theta)$ is obtained by fixing a previous value of the parameters $\theta^{(t-1)}$ in the expectation conditioning, and it is expressed as:

$$Q^{(t)}(\theta) = \mathbb{E}_{Y|X,\theta^{(t-1)}}\left[\log \mathcal{L}(\theta|X,Y)\right]. \tag{2}$$

4

Based on this, the maximization step optimizes the current approximation $Q^{(t)}(\theta)$ with respect to θ to obtain an improved estimate, as shown in Equation (3):

$$\theta^{(t)} = \underset{\theta}{\operatorname{argmax}} \; Q^{(t)}(\theta). \tag{3}$$

The method, therefore, iteratively attempts to estimate the model parameters θ and the labels Y. In each step, the likelihood function is improved. These steps are repeated until a local optimum is attained.

Gaussian Mixture Models. One frequently used generative model for clustering is to assume that data is distributed as a mixture of Gaussian distributions, i.e., a Gaussian Mixture Model (GMM – Banfield and Raftery 2, Moore 40, Zivkovic 60). GMM is a generative process involving two steps for each sample: i) deciding from which cluster the sample will belong; ii) generating the sample from the corresponding Gaussian distribution of the mixture. Each Gaussian is characterized by a mean vector, $\boldsymbol{\mu}$, and a covariance matrix, Σ. The vector $\boldsymbol{\pi}$ defines the weights of the Gaussians. The joint probability distribution of all samples is given by:

$$\sum_{j=1}^{k} \pi_j \; \mathcal{N}(X \mid \boldsymbol{\mu}_j, \Sigma_j) \tag{4}$$

In this model, $\boldsymbol{\mu}_j$ and Σ_j can be seen as the parameters governing the shape of the j^{th} cluster. The eigenvectors of Σ_j describe the ellipsis' main axis orientation, while the eigenvalues describe their length. Finally, π_j is the *a priori* probability that a sample belongs to a given cluster j.

Covariance estimation within EM-GMM. Estimating the covariance matrices in the EM-GMM gives greater flexibility in the families of data that can be accurately represented, but it also poses significant methodological challenges. In particular, the number of parameters of a covariance matrix grows quadratically with the number of features. Consequently, the number of samples involved in the estimate must be large enough to avoid over-fitting with ill-conditioned or singular matrices [25]. This also poses significant numerical issues in EM-type algorithms since full-rank matrices are typically needed in their intermediary steps. To circumvent this issue, several studies have proposed regularization (or shrinkage) techniques [33]. Regularization of covariance matrices can improve the estimator by reducing the condition number, which is the ratio between the largest and the smallest eigenvalue.

Most covariance regularization methods return a convex combination of the empirical matrix Σ

with a scaled identity matrix, as shown in Equation (5).

$$\Sigma_{\text{reg}} = (1 - \delta)\Sigma + \delta \frac{\text{Tr}\,\Sigma}{d}I \tag{5}$$

The constant multiplying of the identity matrix permits keeping the regularized matrix's magnitude similar to the magnitude of the original matrix. This constant is the average of the diagonal entries of Σ, equivalently written as $\frac{\text{Tr}\,\Sigma}{d}$.

The empirical covariance matrix can be seen as a random variable having its own moments (e.g., variance and expectation). If some changes in samples lead to large changes in the estimated matrix, then the estimator's variance is significant. The original empirical estimator Σ has a small bias but a large variance. On the other hand, $\frac{\text{Tr}\,\Sigma}{d}I$ usually has a much smaller variance and a large bias. The purpose of regularization is to reduce the variance of estimators (at the cost of a higher bias). Therefore, the parameter $\delta \in [0, 1]$ controls the trade-off between bias and variance.

The regularization techniques considered in this work are all variants of the above-mentioned convex combination method, but they differ in the δ constant choice. This calculation also has immediate implications on the eigenvalues: the eigenvalues of Σ_{reg} are eigenvalues of the empirical matrix Σ to which a positive constant has been added. Therefore, indirectly, the regularization imposes a lower bound on the eigenvalues of the covariance matrix.

The ill-conditioned matrices estimating problem has inspired several studies. Hastie and Tibshirani [24] developed one of the first classifications works using mixtures of Gaussians and regularization techniques. Different approaches were used to calibrate the δ parameter and achieve a good covariance estimation, such as (i) maximizing the Leave-One-Out Likelihood (LOOL) criterion for each cluster [14]; (ii) estimating the parameter from the data using the Minimum Message Length (MML) principle [30]; and (iii) using a modified Bayesian Information Criterion (BIC) as a model selection criterion to determine δ [42].

The Ledoit-Wolf (LW) approach [31] is a well-known regularization methodology that minimizes the mean squared error between the estimated and the true covariance matrix in the case of arbitrary distribution of the samples. The LW leads to well-conditioned matrices, and it is based on a simple formula that can be easily and efficiently implemented. Halbe et al. [22] used this technique on EM-GMM and noticed significant improvements in the performance of multivariate probability density estimation. Next, [10] proposed another well-known method, called Oracle Approximating Shrinkage (OAS), which reduces the LW mean squared error for the special case of Gaussian processes. This

method has a simple analytical formula and can be efficiently implemented.

[54] proposes and analyzes the selection of δ by cross-validation, leading to better results at the expense of a larger computational effort. [55] describes a convex optimization problem minimizing a Gaussian log-likelihood function with an explicit constraint on the condition number of the matrix. Although the authors present an efficient algorithm, it is still bounded by a spectral decomposition. Finally, [32] describes a non-linear regularization technique in which it is possible to select a specific regularization term for each eigenvalue leading to possibly better results than standard LW. However, like the previous method, it also requires a spectral decomposition.

Finding better local optima in EM-GMM. Finding good solutions to the likelihood maximization problem in general GMMs is challenging due to the large number of parameters that need to be estimated. Effectively, the most popular clustering algorithms are local search methods, such as k-means and EM-GMM, which can easily get trapped in local optima and depend on the initial conditions of the search. A common strategy to improve such methods consists of repeating the approach from different initial points and keeping the best solution. Although this is the first step toward methodological improvements, more efficient ways exist to improve the search using classical meta-heuristics, which are methods precisely designed to "orchestrate the local search and higher-level strategies to escape from local optimum and reach better solutions" [20].

A simple meta-heuristic approach, called Iterated Local Search (ILS) [35], consists of iteratively applying a local optimization algorithm (e.g., k-means or EM) to reach a local minimum, followed by a perturbation operator to generate a new starting point and pursue the search. This strategy effectively produces a diversity of starting points and profits from the information from previous iterations. As such, it outperforms random restarts in a variety of applications. One application of the ILS concept to clustering problems has been proposed by Fränti [18] under the name of the Random Swap (RS) approach. The design of this algorithm follows the classical ILS framework. Each time the local search converges towards a solution, the algorithm applies a SWAP which consists of moving a random cluster's position to the location of an arbitrary point in the dataset. Zhao et al. [59] combined the Random Swap and EM (RSEM) to improve EM-GMM clustering. In each RSEM iteration, a random Gaussian component is removed and relocated to a random data point location. The covariance matrix and weights are kept the same as the previous iteration in the Random Swap operation to preserve the cluster's parameters' relative magnitude.

Another way to search for the global optimum is with genetic algorithms (GA). GAs rely on the basic principles of natural evolution (selection, mutation, crossover) to improve a population of solutions.

These metaheuristics are often hybridized with local search to achieve state-of-the-art performance on challenging combinatorial optimization problems [see, e.g., 16, 52, 38, 51]. In a sense, hybrid GAs with local search extends the capabilities of ILS approaches, as they generate new starting points for the local search by *recombining* existing solutions from a *population* of high-quality local minima found in previous iterations. This way, the generation of new starting points is tightly connected to their success in the search history.

Specific to clustering problems, Pernkopf and Bouchaffra [44] proposed a genetic-based EM called GAEM, which selects the number of model components using the minimum description length (MDL) criterion and exploits a single-point crossover [39, 17]. For the minimum-sum-of-squares clustering problem (the model for which k-means is the usual choice), Gribel and Vidal [21] proposed a hybrid genetic algorithm called HG-means. This algorithm relies on a recombination operation built upon a bipartite matching problem and mutation steps to generate diverse and promising start points based on the information of the population. It achieved state-of-the-art results on a wide variety of benchmark instances compared to other popular clustering algorithms for the same model. Finally, in addition to genetic algorithms, other works have proposed split-and-merge techniques to escape from local minima. In Split-and-Merge EM (SMEM) [50], a pair of clusters is merged and another cluster is split in between two EM iterations. Stepwise SMEM (SSMEM) [56] goes further and uses the SMEM strategy to estimate the optimal number of Gaussians in the mixture. Finally, Competitive EM [58] is another SMEM variant with different criteria to select which clusters to merge and split.

The aforementioned algorithms can lead to improved solutions; however, as shown in the experiments, these algorithms do not generalize well in a high-dimensional cases with few observations, resulting in over-fitted solutions. To circumvent this issue, we introduce a global optimization algorithm using regularization techniques leading to better-conditioned matrices and better generalization as the outcome.

3. A Hybrid Genetic EM with Regularization

In this section, we introduce a global optimization algorithm for ellipsoidal clustering. This algorithm is grounded on evolutionary computation and combines the general EM local search with regularization strategies. Evolutionary algorithms apply natural selection, crossover, and mutation principles to improve a population of solutions. In such algorithms, it is also essential to balance selection pressure and diversification to search for a wide spectrum of solutions.

The hybrid genetic search (HGS) we propose follows the same basic principles as Gribel and Vidal [21] and, therefore, combines the classical variation operators (crossover and mutation) with local

improvement through EM. We use the regularization techniques presented in the previous section to avoid overfitting and bad conditioning of the covariance matrices. The general structure of the algorithm is summarized in Algorithm 1. The rest of this section details each component of the method: the solution representation and population initialization methods, the variation (crossover and mutation) operators for solution generation, the local search with a regularized EM algorithm, and the population-management methods.

Algorithm 1 General structure of HGS

1: Initialize population   ▷ Section 3.1
2: **for** T generations without improvement **do**
3:      Select parents p_1 and p_2 ▷ Section 3.2
4:      Generate an offspring θ from p_1 and p_2 (crossover)   ▷ Section 3.2
5:      Generate an individual θ' by mutating θ (mutation)   ▷ Section 3.2
6:      Apply clustering local improvement on θ' ▷ Section 3.3
7:      Add θ' to the population
8:      **if** population size exceeds maximum size $\Pi_{\max}$ **then**   ▷ Section 3.4
9:          Remove all clones in the population
10:          Select survivors until the population reach $\Pi_{\min}$
11:      **end if**
12: **end for**
13: **return** best solution

3.1. Solution Representation and Population Initialization

Each solution s is represented in the HGS as a set of three chromosomes:

- $M = \{\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_k\}$ containing the d-dimensional coordinates of the k clusters centers;

- $\Sigma = \{\Sigma_1, ..., \Sigma_k\}$ defining the $\mathbb{R}^{d \times d}$ covariance matrix of each cluster;

- $\boldsymbol{\pi} = \{\pi_1, ..., \pi_k\}$ representing the mixture weights of each cluster.

To obtain an initial population containing diverse solutions, the algorithm creates $\Pi_{\max}$ solutions by (i) randomly assigning the center of each cluster to random sample locations with uniform probability, (ii) initializing the covariance matrix of each cluster to the identity matrix and setting uniform initial mixture weights, and (iii) applying the local search algorithm (described in Section 3.3) from this starting point.

The fitness of each solution is then calculated as the log-likelihood of Equation (6):

$$\sum_{i=1}^{n} \log \left(\sum_{j=1}^{k} \pi_j \, \mathcal{N}(\boldsymbol{x}_i \mid \boldsymbol{\mu}_j, \Sigma_j) \right) \tag{6}$$

$$\mathcal{N}(\boldsymbol{x}_i \mid \boldsymbol{\mu}_j, \Sigma_j) = \frac{1}{\sqrt{(2\pi)^d \det \Sigma_j}} \exp\left(-\frac{1}{2}\left(\boldsymbol{x}_i - \boldsymbol{\mu}_j\right)^\top \Sigma_j^{-1}\left(\boldsymbol{x}_i - \boldsymbol{\mu}_j\right)\right). \tag{7}$$

3.2. Solution Generation by Crossover and Mutation

The HGS uses a binary-tournament selection to obtain two parents, P_1 and P_2. As illustrated in Figure 1, two solutions are randomly selected with uniform probability in the population during each binary tournament, and the best one is retained.
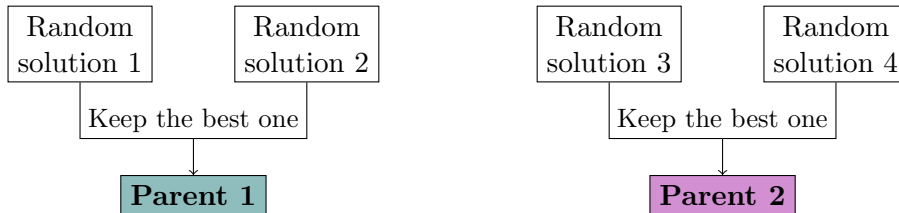

Figure 1: Diagram representing the workflow of the binary tournament

During the crossover step, the algorithm solves a matching problem in a bipartite graph $G = (V_1, V_2, E)$, in which each vertex $v \in V_1$ (respectively V_2) stands for a cluster in P_1 (respectively, P_2), therefore $|V_1| = |V_2| = k$. Each edge $(i, j) \in E$ for $i \in V_1$ and $j \in V_2$ represents the possibility of associating the i^{th} cluster of P_1 with the j^{th} cluster of P_2 for an associated cost:

$$c_{ij} = \frac{\sqrt{\left(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\right)^\top \Sigma_j^{-1}\left(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\right)} + \sqrt{\left(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i\right)^\top \Sigma_i^{-1}\left(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i\right)}}{2}. \tag{8}$$

Finding a minimum-cost bipartite matching in this graph gives us a better association between the centers of the two parents. We use the Hungarian Algorithm [29] to obtain this matching. Then, for each pair of clusters (i.e., edge) of the bipartite matching solution, we randomly retain one of the two clusters in the child along with its covariance matrix. At the same time, the mixture coefficient of this cluster is defined as the average from both matched clusters. This process effectively permits inheriting centers from both parents while keeping a good distribution of centers. Figure 2 illustrates the crossover operator.

The solution generated by the crossover operator passes through a mutation step, which is designed to diversify the search. The mutation operator randomly selects a cluster and a data sample according to a uniform probability distribution. The cluster is reallocated to the position of the chosen data sample, and the covariance matrix is re-initialized as the average of the other clusters' covariance matrices. This maintains the relative scale of the clusters' parameters and effectively prevents situations in which one specific cluster would progressively lose importance in the mixture and ultimately disappear. Figure 3

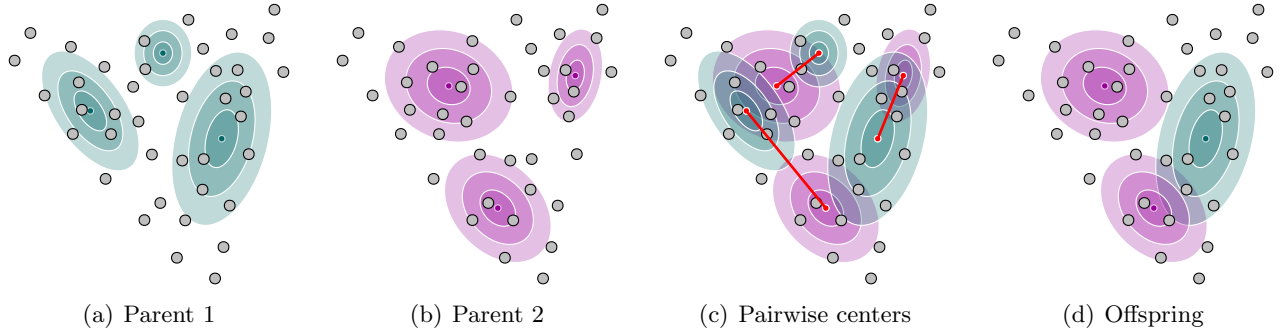(a) Parent 1   (b) Parent 2   (c) Pairwise centers   (d) Offspring

Figure 2: Crossover operator: (a) and (b) represent the parents; (c) solution of the matching step; (d) offspring obtained after retaining one cluster for each matched pair

illustrates the mutation operator as well as the final solution obtained after applying the local search (described in Section 3.3) on it.
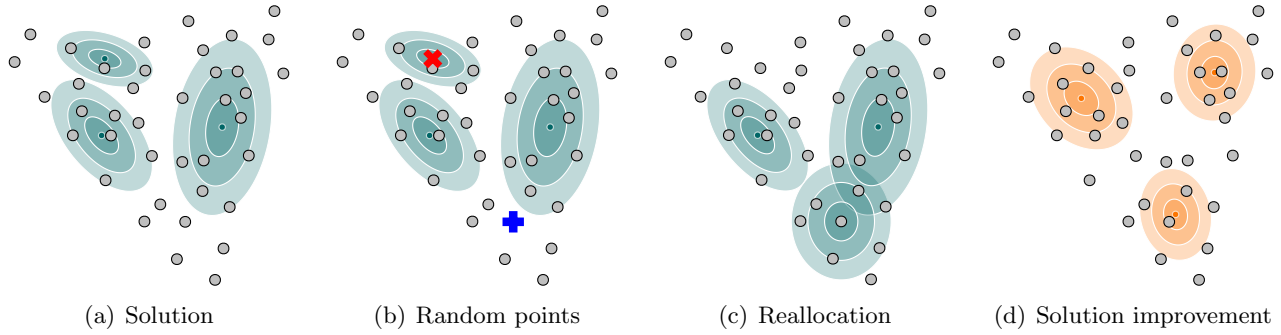

(a) Solution   (b) Random points   (c) Reallocation   (d) Solution improvement

Figure 3: Mutation operator: (a) current solution; (b) selection of a random cluster (red × mark) and random data sample (blue + mark); (c) reallocation of the cluster; (d) resulting solution after local search

3.3. Local Search with a Regularized EM

The solution produced by the crossover and mutation operators serves as the starting point for a local search using a regularized variant of the EM algorithm for GMM. Algorithm 2 describes the general structure of this algorithm.

Algorithm 2 EM-GMM Clustering

1: Initialize solution
2: **while** the stopping criterion has not been met **do**
3:     Compute γ ownership weights
4:     Recompute the mixing weights π
5:     Recompute the clusters' centers μ
6:     Recompute the covariance matrices Σ
7: **end while**
8: **return** solution

The specialization of the EM algorithm for GMM clustering is described by [6] and [19]. The EM

11

alternates between expectation and maximization steps; it optimizes the mixture's parameters along with latent variables representing the ownership weight of each sample $\in X$, i.e., the probability of attribution of this point to each cluster, also known as responsibility. The *expectation* step is described in Equation (9). Given the current solution $(\pi, \boldsymbol{\mu}, \Sigma)$, this step computes the estimated ownership weights of each sample belonging to each cluster:

$$\gamma_{i,j} = \frac{\pi_j \, \mathcal{N}(\boldsymbol{x}_i \mid \boldsymbol{\mu}_j, \Sigma_j)}{\sum_{c=1}^{k} \pi_c \, \mathcal{N}(\boldsymbol{x}_i \mid \boldsymbol{\mu}_c, \Sigma_c)}. \tag{9}$$

In contrast, the *maximization* step re-optimizes the parameters of the EM-GMM (π, $\boldsymbol{\mu}$ and Σ) based on the values of the latent variables, as depicted in Equations (10–12):

$$\pi_j = \sum_{i=1}^{n} \gamma_{i,j} \left/ \sum_{j=1}^{k}\sum_{i=1}^{n} \gamma_{i,j} \right. \tag{10}$$

$$\boldsymbol{\mu}_j = \sum_{i=1}^{n} \gamma_{i,j}\boldsymbol{x}_i \left/ \sum_{i=1}^{n} \gamma_{i,j} \right. \tag{11}$$

$$\Sigma_j = \sum_{i=1}^{n} \gamma_{i,j} \left(\boldsymbol{x}_i - \boldsymbol{\mu}_j\right) \left(\boldsymbol{x}_i - \boldsymbol{\mu}_j\right)^{\top} \left/ \sum_{i=1}^{n} \gamma_{i,j} \right. . \tag{12}$$

Each of these two steps is guaranteed to improve the likelihood, as it concentrates on re-optimizing a subset of the parameters. The EM algorithm "iterates" by alternating between these two steps until no more improvement is achievable. An example of the progression of this algorithm is depicted in Figure 4.
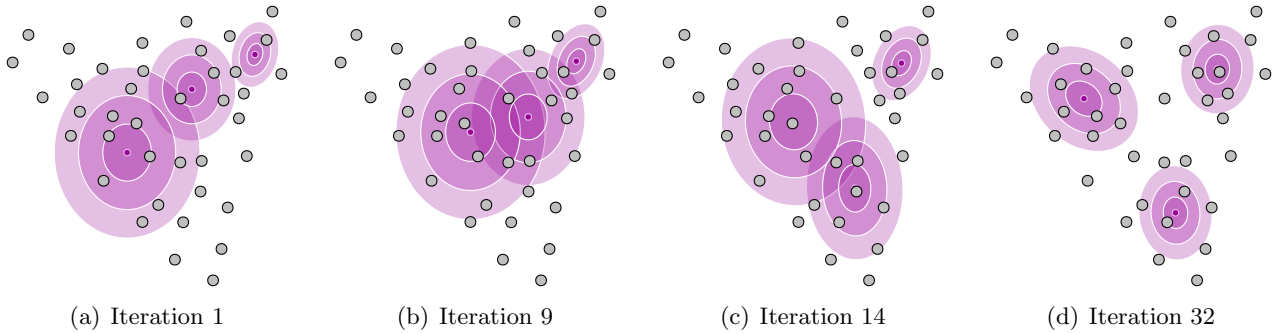

(a) Iteration 1 (b) Iteration 9 (c) Iteration 14 (d) Iteration 32

Figure 4: Successive solutions in the EM-GMM

Regularization in EM. As previously discussed, the classical EM-GMM suffers from poor conditioning of the covariance matrices. To circumvent this issue, we use regularization techniques applied at each iteration of the algorithm immediately after estimating the empirical covariance matrix in the

maximization step. We consider three alternative regularization techniques that can be efficiently implemented. These approaches impact how parameter δ of Equation (5) is selected.

1. The first regularization method, called Shrunk methodology, consists of fixing the parameter $\delta_{\text{Shrunk}} = 0.1$. This is the default approach adopted in scikit-learn [43].

2. The second regularization method, called Ledoit-Wolf Shrinkage (LW), attempts to minimize the mean squared error between the estimated and the real covariance matrix. The resulting matrix can be calculated in closed form, as seen in Ledoit and Wolf [31].

3. The last regularization method, called Oracle Approximating Shrinkage (OAS – Chen et al. 10) assumes that the data came from a Gaussian distribution. According to that study, the shrinkage coefficient δ reduces further the mean squared error compared to the LW regularization. Equation (13) is the formula that computes the δ in OAS:

$$\delta_{\text{OAS}} = \min\left(1, \frac{(1 - 2/d)\,\text{Tr}\left(\Sigma^2\right) + \text{Tr}^2\left(\Sigma\right)}{(n + 1 - 2/d)\left[\text{Tr}\left(\Sigma^2\right) + \text{Tr}^2\left(\Sigma\right)/d\right]}\right). \tag{13}$$

Figure 5 compares the results of the different covariance estimation methods. As visible in this figure, the impact of regularization is more marked in situations where few data points are available to estimate the covariance matrices. In those cases, the regularization strategies provide better-conditioned matrices.
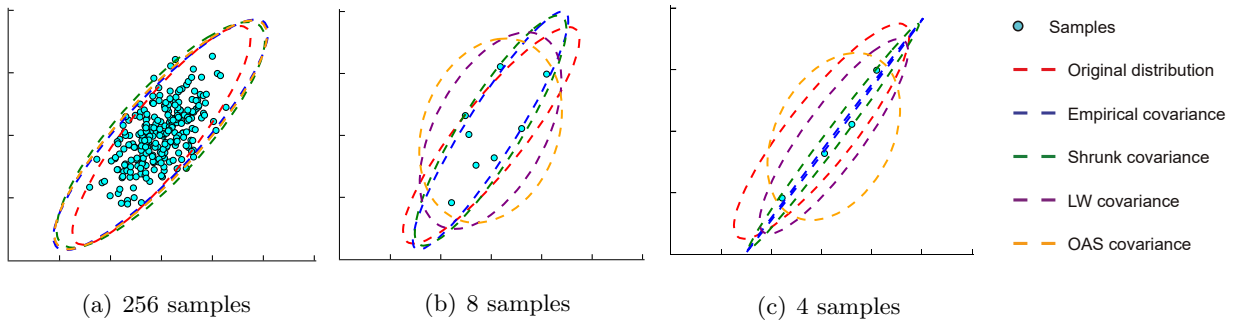

(a) 256 samples   (b) 8 samples   (c) 4 samples

Figure 5: Behavior of the empirical and the regularized covariance estimations

3.4. Population Management

Finally, whenever the population exceeds the maximum size of Π_{max}, HGS performs a survivor selection mechanism to retain only Π_{min} solutions. Survivor selection ensures a gradual selection pressure towards higher-quality solutions. To preserve diversity, we implement a clone elimination

13

mechanism, which detects individuals with the same objective value and prioritizes removing one of them. Once no such clone exists anymore and if the population still exceeds Π_{min}, the algorithm proceeds by removing the solutions with the smallest objective value (i.e., smallest likelihood) until it reaches the desired population size.

4. Computational Experiments

We conduct extensive computational experiments with two main goals: (i) to measure the impact of regularization techniques on cluster recovery and (ii) to assess the importance of more sophisticated global optimization algorithms and their interaction with the regularization strategies. To that end, we implemented the proposed methods in Julia v1.8.4 [3] and conducted our experiments on an Intel(R) Xeon(R) Platinum 8124M CPU @ 3.00 GHz. All the source code, scripts, and data needed to reproduce these experiments are openly accessible at `github.com/raphasampaio/` `RegularizationAndGlobalOptimizationInModelBasedClustering.jl`.

4.1. Datasets and Experimental Setup

We generated the benchmark instances using the *ClusterGeneration* package [47] in R programming language. The difficulty of the clustering task for these datasets is governed by the Separation Index [45, 46]. The higher this index, the more separated the clusters are (see Figure 6). We therefore generated datasets with different values of the separability index ($c \in \{-0.26, -0.10, 0.01, 0.21\}$), different number of features ($d \in \{2, 5, 10, 20, 30, 40\}$) and different number of clusters ($k \in \{3, 10, 20\}$). For each configuration, described as a tuple (c, d, k), we fixed the dataset size ($n = 100 \cdot k$) and generated 20 random datasets. In all datasets, the covariance matrix's eigenvalues describing each cluster's shape were randomly sampled from a uniform distribution over the range $[1, 200]$. Table 1 summarizes the main factors considered in our analyses and their possible levels, i.e., the parameters governing the structure of the instances.
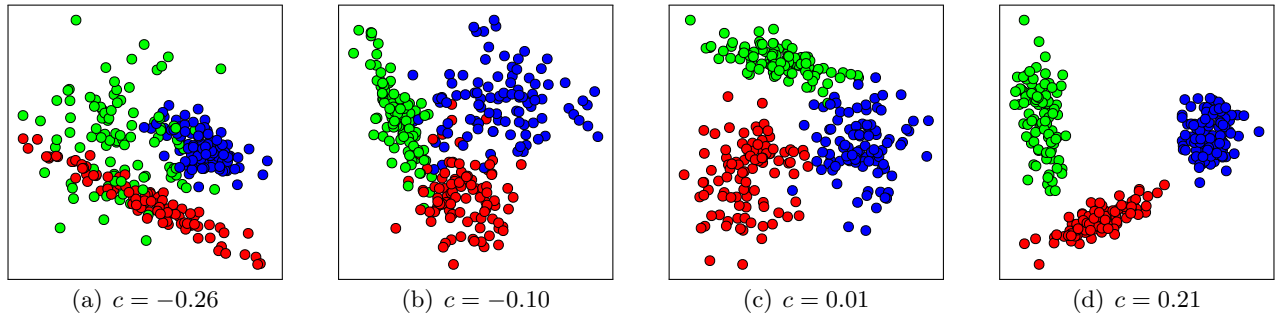

(a) $c = -0.26$ (b) $c = -0.10$ (c) $c = 0.01$ (d) $c = 0.21$

Figure 6: Examples of instances generated with different levels of the separability index c

14

Factors	Levels
Separability (c)	$\{-0.26, -0.10, 0.01, 0.21\}$
Number of Features (d)	$\{2, 5, 10, 20, 30, 40\}$
Number of Clusters (k)	$\{3, 10, 20\}$

Table 1: Main factors considered in our experiments

We adopt the adjusted Rand index (ARI) [27] as the similarity measure to compare two assignments. This score is the corrected version of the Rand index (RI) [48, 53], a widely used measure for clustering analysis [49]. The ARI establishes a baseline using the expected similarity of all pairwise comparisons between the classifications. The ARI score belongs to the range $[-1, 1]$, where 1 corresponds to perfect assignments, identical to the true values, and 0 would stand for random uniform cluster assignments.

The following sections evaluate the performance of the proposed methods in terms of ARI. We carry out our analyses in several steps, using the standard GMM algorithm (i.e., "Local Search" with "Empirical" covariance estimation) as our fundamental baseline. First, we will vary the regularization techniques within the GMM to evaluate their impact (Section 4.2). Next, we will combine the use of different regularization techniques with more sophisticated search algorithms to observe their interactions (Section 4.3). In these two analyses, we use the synthetic datasets described in Section 4.1 to measure the impact of factors related to the characteristics of the instances. Next, Section 4.4 extends our analyses to include additional methods (k-means) and practical considerations (computational time). Finally, Section 4.5 reports results on datasets issued from the UCI, whose characteristics are more diverse but less controllable.

4.2. Impact of Regularization

In this section, we study the impact of covariance estimation with regularization techniques. When the square of the number of features, d^2, grows large compared to the number of samples n, the sample covariance matrices become ill-conditioned and represent poor estimators of the sample distribution. The use of regularization techniques can help overcome these issues. To estimate the impact of regularization, we first compare the performance of EM-GMM (for short GMM), without covariance matrix regularization with its counterpart with different regularization techniques (Shrunk, OAS, and LW – defined in the previous section).

Therefore, Figure 7 compares the clustering performance of the different approaches in terms of ARI. To measure the impact of the number of features, we report the results for varying values of d while keeping the parameters governing the number of clusters and separability index fixed to $k = 10$ and $c = 0.01$. For each method and value of d, we average the results obtained from applying the method

15

in all 20 datasets and represent it as a barplot with an additional interval representing the standard deviation. In all cases, results from regularized versions of GMM are equal to or better than standard GMM. The cluster-recovery performance still seems to degrade as the number of features grows, as a consequence of possible overfitting, but the decay is slower for methods that include covariance matrix regularization.
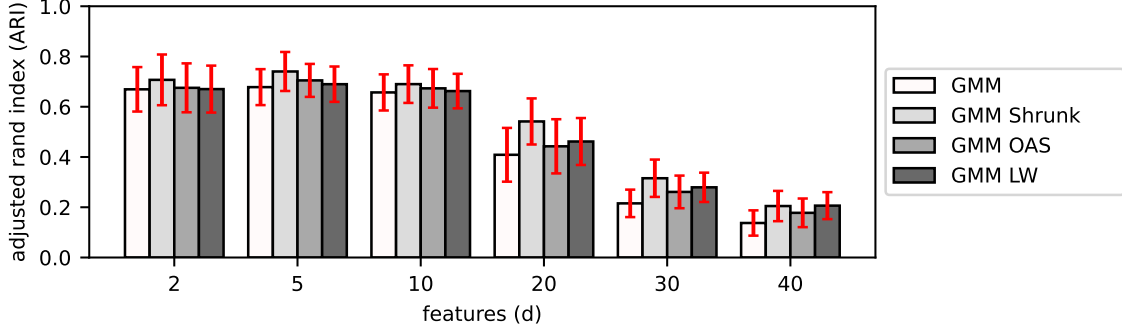

Figure 7: Performance of the GMM and its regularized variants, for datasets with $k = 10$ and $c = 0.01$

In Figure 8, we further analyze the impact of two additional factors: the dataset's separability index c and the number of clusters k. Our results are represented as four heatmaps, each corresponding to GMM algorithms with different covariance estimations: GMM (empirical estimation), GMM Shrunk, GMM OAS, and GMM LW. The rows correspond to different values of the separability index, and the columns correspond to different numbers of clusters. Each cell is the average over 120 datasets (20 different datasets for each number of features: $d = \{2, 5, 10, 20, 30, 40\}$).
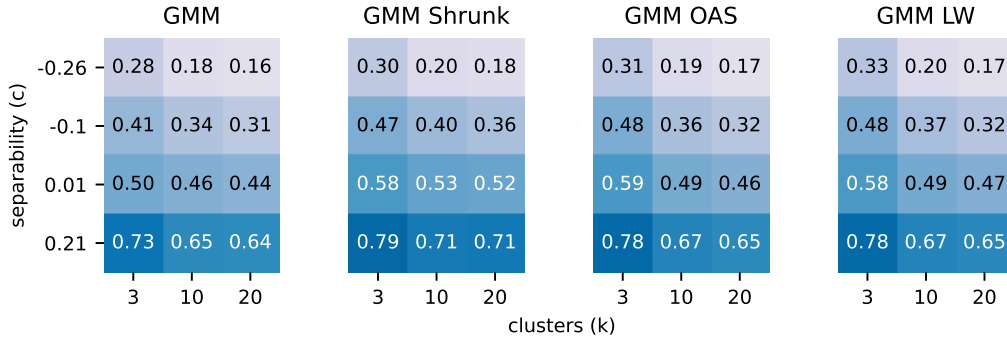

Figure 8: Average of the ARI performance of GMM and its regularized variants over the $d = \{2, 5, 10, 20, 30, 40\}$ datasets for different separability indices and number of clusters

As seen in Figure 8, the Shrunk methodology outperformed OAS, LW, and no regularization when the number of clusters increases. The significance of these performance differences is confirmed by pairwise Wilcoxon tests between the results of GMM Shrunk and those of other methods at a 0.05% significance level, presented in Table 2.

Pair of methods	p-value
GMM Shrunk – GMM	3.21×10^{-127}
GMM Shrunk – GMM OAS	7.61×10^{-57}
GMM Shrunk – GMM Ledoitwolf	1.46×10^{-42}

Table 2: Impact of regularization on GMM: Pairwise Wilcoxon tests

These differences are likely due to the fact that OAS and LW have a larger impact on the shape matrix than the Shrunk approach. Regularization methods that estimate the data shape as too spherical reduce the possible space of clustering solutions. In contrast, Shrunk ensures that the shape matrix is well conditioned but has a more limited impact on it. Due to its good results and simplicity, this regularization approach will be used in the remainder of our experiments.

4.3. Combining Regularization and Global Optimization

In this section, we evaluate whether the Shrunk regularization combined with more sophisticated search strategies can achieve even better clustering performance. We consider, to that end, three search strategies. The first is a simple Multi Start (GMM MS) approach, which repeats the GMM from n_{IT} starting points and retains the solution with maximum likelihood. The second is a Random Swap (GMM RS) method similar to Zhao et al. [59] and Fränti [18]. This method iteratively applies the GMM to find a local minimum in terms of likelihood and then randomly relocates a cluster position to generate a new starting point for the GMM. The third and final approach is the Hybrid Genetic Search (GMM HG) described in Section 3, which generates new starting points for the GMM by recombination of previous solutions. We set a maximum of $n_{\mathrm{IT}} = 100$ iterations without improvement for each meta-heuristic, where an iteration refers to one application of GMM from a starting point (obtained randomly in the case of MS, from a relocation process in RS, or the crossover and mutation operator in HG).

Figure 9 compares the ARI of the classical GMM with those achieved by the three improved search strategies, without and with Shrunk regularization. We focus this experiment on datasets for which the number of clusters and separability index are fixed to $k = 10$ and $c = 0.01$ while varying the values of d. For each method and value of d, we average the results and represent it as a barplot with an additional interval representing the standard deviation.

The results presented in Figure 9 lead to a rich set of observations. As it appears, the use of more sophisticated search approaches such as MS, RS, or HG leads to a performance deterioration when no regularization mechanism is used, especially for datasets with a large number of features. With the Shrunk regularization, however, using more sophisticated search approaches (HG, especially) leads to remarkable improvements in clustering performance. This connection between regularization and search
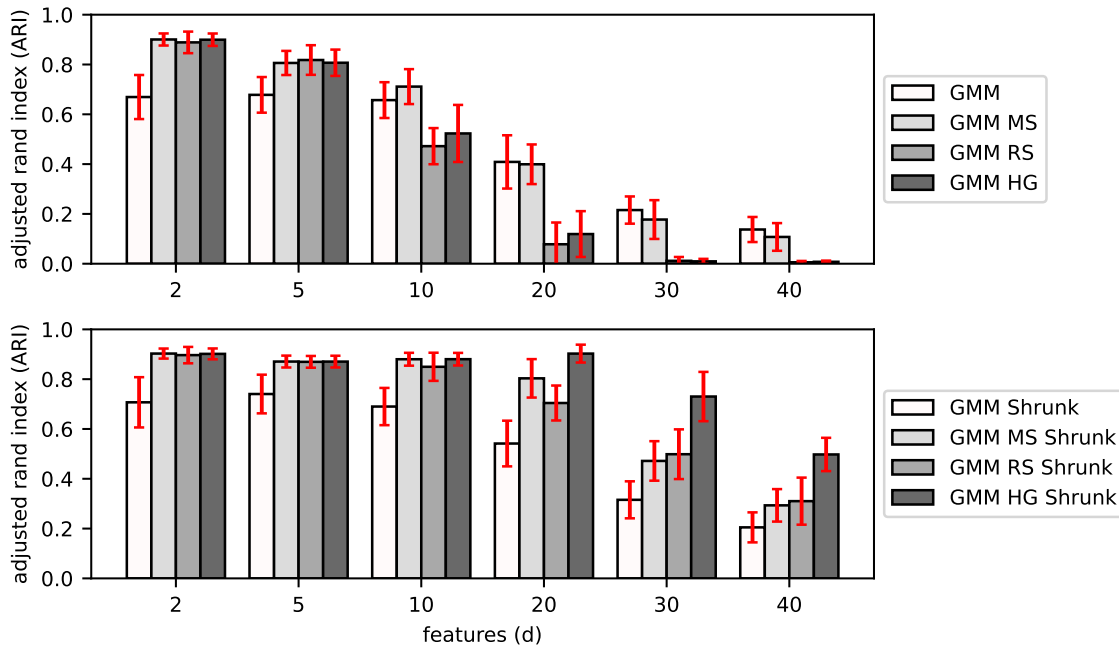
Figure 9: Average and standard deviation of the ARI obtained by GMM and its global search strategies variations over the $k = 10$ and $c = 0.01$ datasets for different numbers of features, with and without regularization

intensity is a known phenomenon: without any proper control over the number of parameters, global optima will tend to overfit the assignment of samples to clusters and covariance matrix estimation. This effect is due to the number of parameters estimated when doing ellipsoidal clustering, which is $\mathcal{O}(kd^2)$ compared to $\mathcal{O}(kd)$ for classical k-means methods.

Figure 10 and 11 additionally report the results of GMM, GMM HG, GMM Shrunk, and GMM HG Shrunk for varying separability levels c, number of clusters k, and number of features d. Each heatmap gives the results of one of the four aforementioned search strategies. In both figures, the rows correspond to different levels of separability. In Figure 10, each column corresponds to a different number of features, whereas in Figure 11 each column corresponds to a different number of clusters.

As seen in these results, the use of regularization improves the performance of the baseline GMM when the number of features increases. In contrast, the use of the HG search component for a more intensive search is, without regularization, only beneficial for problems of relatively-low dimension (e.g., it permits near-perfect cluster recovery with ARI of 0.99 compared to 0.85 when $c = 0.21$ and $d = 5$), and largely detrimental for higher-dimensional cases.

Now, as seen in the results located in the bottom-right of Figure 10 and right of Figure 11, using the Shrunk regularization and the HG search strategy jointly leads to a clear-cut improvement in all the considered situations, without any form of compromise. The regularization strategy effectively avoids overfitting in HG and permits fully harnessing its capabilities to achieve good clustering results
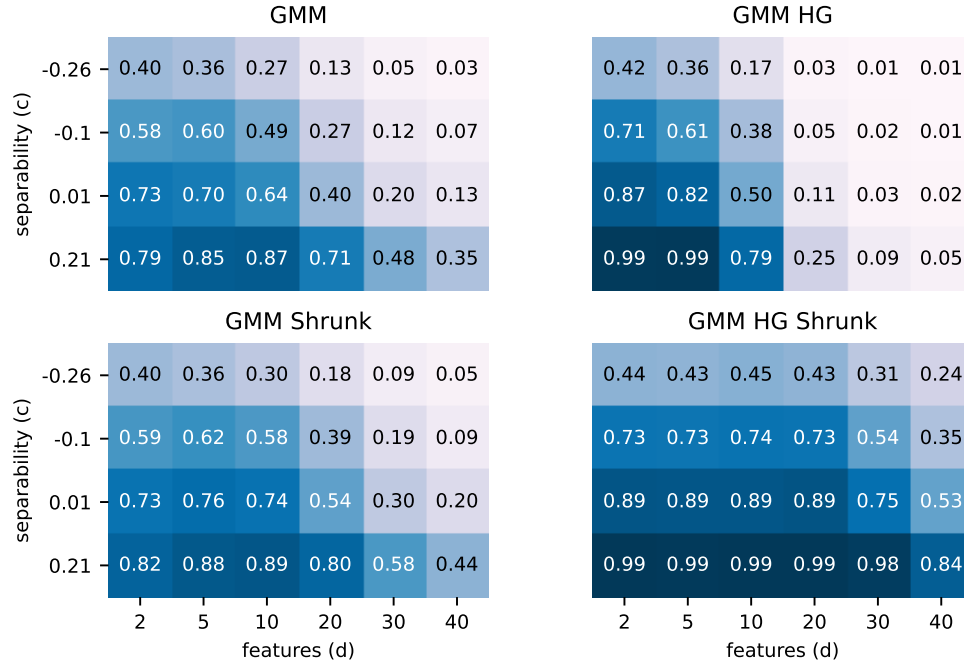
Figure 10: Average ARI performance of GMM, GMM HG, GMM Shrunk, and GMM HG Shrunk over the $k = \{3, 10, 20\}$ datasets, for different separability indices and number of features
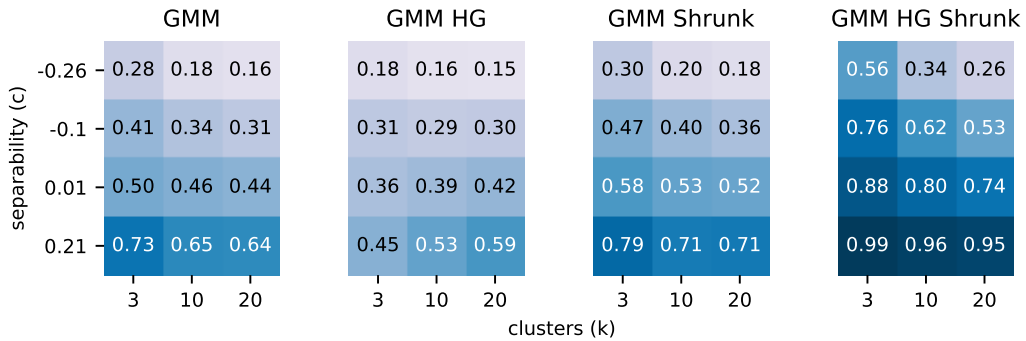

Figure 11: Average ARI performance of GMM, GMM HG, GMM Shrunk, and GMM HG Shrunk over the $d = \{2, 5, 10, 20, 30, 40\}$ datasets, for different separability indices and number of clusters

19

in all regimes. Therefore, the results of this section indicate that combining GMM, regularization, and better search procedures is a promising approach for clustering.

4.4. Comparisons with k-means and HG-means

To broaden our analyses, we now include comparisons with two variants of the k-means algorithm. Indeed, k-means naturally lead to spherical clusters (due to the use of squared Euclidean distance in the objective of the underlying optimization problem) and can be viewed as a limit case of regularization. As the combination of k-means with HG search appeared to be a promising option in previous studies [21], it is meaningful to include additional comparisons of it with the newly proposed GMM HG algorithm with Shrunk regularization, also considering clustering performance and computational time.

Figure 12 and 13 therefore compare the clustering performance of a classical k-means, the Hybrid Genetic k-means as presented in Gribel and Vidal [21] (HG-means), a simple GMM, and the Hybrid Genetic GMM with regularization (GMM HG Shrunk). Figure 12 present results for different values of d with fixed values of the number of clusters and separability index ($k = 10$ and $c = 0.01$). Next, Figure 13 provides heatmaps that describe the evolution of the clustering performance when varying the separability index (c) and the number of clusters (k). The format of these figures is the same as in Section 4.3.
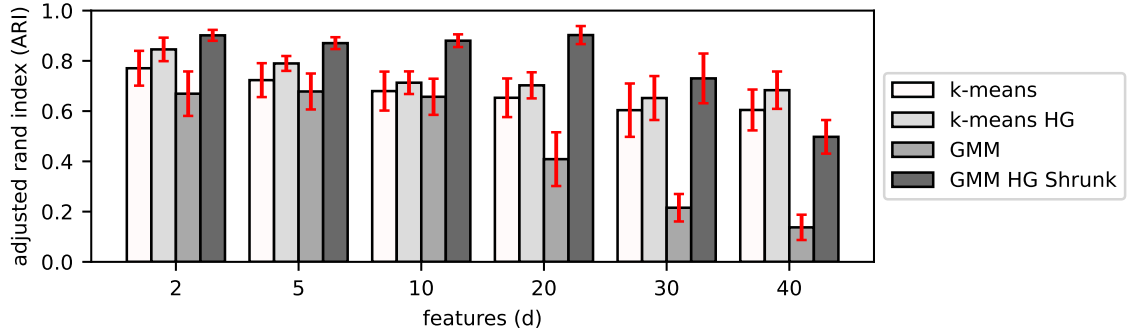
Figure 12: Average and standard deviation of the ARI obtained by k-means, GMM and its global search variations over the $k = 10$ and $c = 0.01$ datasets

Both figures lead us to similar conclusions. Standard GMM systematically appears to be the worst method in terms of clustering quality, except when $c = -0.26$ and $k = 3$. Next, k-means ranks as the second-worst option, as the HG-based algorithms systematically outperform it. Finally, GMM HG Shrunk is generally superior to all other methods. This observation is confirmed by pairwise Wilcoxon tests between the pairs of methods presented in Table 3, considering the complete sets of results on all data sets.

Interestingly, GMM HG Shrunk also achieved outstanding results compared to HG-means in

20
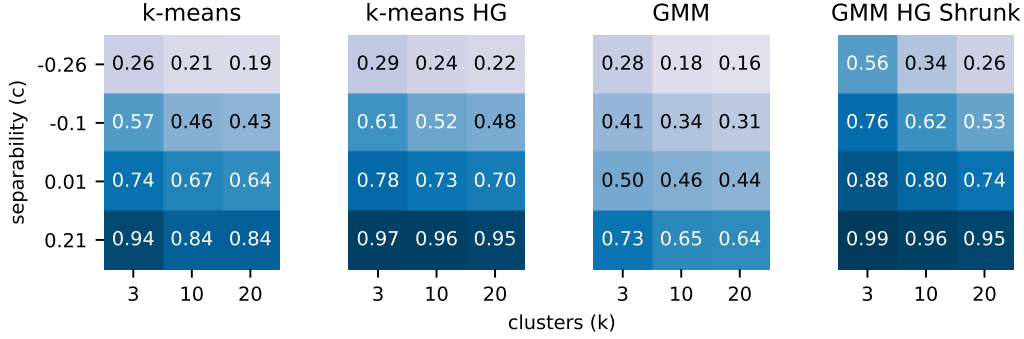
Figure 13: Average ARI performance of k-means, GMM and its global search variations over the $d = \{2, 5, 10, 20, 30, 40\}$ datasets for different separability indices and number of clusters

Algorithm	p-value
GMM HG Shrunk – k-means	4.34×10^{-156}
GMM HG Shrunk – k-means HG	2.79×10^{-81}
GMM HG Shrunk – GMM	3.41×10^{-219}

Table 3: Comparison with k-means and k-means HG: Pairwise Wilcoxon tests

instances with a low separability index.

Finally, Table 4 and 5 compare the different algorithms' computational effort for subsets of the instances with a different number of clusters or features.

k	k-means	k-means HG	GMM	GMM MS	GMM RS	GMM HG	GMM Shrunk	GMM MS Shrunk	GMM RS Shrunk	GMM HG Shrunk
3	0.01	0.80	0.03	1.76	1.36	1.50	0.01	1.30	0.80	1.16
10	0.03	4.36	0.15	24.13	7.64	19.07	0.11	20.11	6.28	14.00
20	0.10	15.28	0.54	106.19	33.68	88.99	0.54	104.17	24.01	67.23
Avg.	0.05	6.81	0.24	44.03	14.23	36.52	0.22	41.86	10.36	27.46

Table 4: Average CPU time in seconds, for a different target number of clusters

As seen in these tables, k-means remains the fastest algorithm overall. Indeed, the steps of the method are very similar to GMM, but no covariance matrix estimation takes place, significantly decreasing the needed computational effort. More sophisticated search approaches such as MS, RS, or HG require several iterations of k-means or GMM, increasing their computational effort proportionally. Therefore, these methods are recommended in situations where computational effort is affordable.

4.5. Performance on UCI Datasets

While synthetic data sets give fine-grained control over key factors (n, c, d, and k), they do not provide the full picture of practical situations. In this section, we extend our experiments to datasets from the UCI repository [13] to observe whether the previous observations still apply. We gathered

d	k-means	k-means HG	GMM	GMM MS	GMM RS	GMM HG	GMM Shrunk	GMM MS Shrunk	GMM RS Shrunk	GMM HG Shrunk
2	0.03	4.23	0.02	4.11	1.46	3.39	0.02	3.59	1.37	3.08
5	0.03	4.78	0.05	8.61	3.28	6.63	0.04	7.59	2.76	5.62
10	0.03	5.18	0.11	17.18	17.19	14.03	0.08	14.95	4.74	10.03
20	0.04	5.73	0.24	45.36	18.39	35.91	0.20	39.95	9.93	26.57
30	0.05	9.97	0.42	79.26	19.25	67.47	0.40	76.70	18.13	48.97
40	0.07	10.95	0.56	109.60	25.75	91.67	0.57	108.34	25.21	70.47
Avg.	0.04	6.80	0.24	44.02	14.22	36.52	0.22	41.85	10.36	27.46

Table 5: Average CPU time in seconds, for different number of features

classical and recent datasets widely used in clustering benchmark comparisons. We focus on datasets containing a ground truth for properly assessing the clustering results. Table 6 lists the datasets labels and their dimensions (n, d, and k). Additional references are also provided for some datasets.

#	Dataset	n	k	d
A	Facebook Live Sellers [11]	7050	4	9
B	Glass	214	6	9
C	Handwritten Digits	5620	10	62
D	HCV	589	5	12
E	Human Activity Recognition [1]	10299	6	561
F	Image Segmentation	2310	7	18
G	Ionosphere	351	2	33
H	Iris	150	3	4
I	Letter Recognition	20000	26	16
J	MAGIC Gamma Telescope	19020	2	10
K	Mice Protein Expression [26]	1047	8	74
L	Pen-Based Recognition	10992	10	16
M	Scadi	70	7	141
N	Seeds	210	3	7
O	Statlog (Shuttle)	58000	7	9
P	SPECT Heart	267	2	22
Q	Waveform	5000	3	21
R	Wholesale Customers	440	6	6
S	Wines	178	3	13
T	Yeast	1484	10	8

Table 6: UCI datasets labels and dimensions

Tables 7 and 8 report the clustering accuracy (ARI) and CPU time of each algorithm on each dataset. The following methods are compared to obtain a complete picture: standard k-means, k-means HG (HG-means), standard GMM, GMM MS, GMM RS, GMM HG, GMM Shrunk, GMM MS Shrunk, GMM RS Shrunk, and GMM HG Shrunk. To improve statistical significance, since all algorithms depend on pseudo-random choices, ten runs of each algorithm have been done with a different random seed on each instance. Therefore, each table reports an average (ARI and time) measurement over

these ten runs. We also highlight in each table the best performance in bold and use shades of grey to depict better results (darker is better).

#	k-means	k-means HG	GMM	GMM MS	GMM RS	GMM HG	GMM Shrunk	GMM MS Shrunk	GMM RS Shrunk	GMM HG Shrunk
A	0.10	0.09	0.16	0.10	0.13	0.09	**0.23**	0.21	**0.23**	**0.23**
B	0.24	0.26	0.20	0.04	0.17	0.05	**0.27**	**0.27**	**0.27**	0.26
C	0.64	0.67	0.12	0.30	0.11	0.23	0.73	**0.79**	**0.79**	**0.79**
D	0.39	0.47	0.21	0.20	0.22	0.19	**0.55**	0.26	0.47	0.32
E	0.49	0.46	0.32	0.29	0.30	0.29	0.34	**0.53**	0.38	0.51
F	0.34	0.35	0.22	0.28	0.25	0.26	0.40	**0.44**	**0.44**	0.38
G	0.18	0.18	0.32	0.68	0.60	0.60	0.63	**0.73**	**0.73**	**0.73**
H	0.60	0.73	0.66	0.75	0.76	0.73	0.68	**0.94**	**0.94**	**0.94**
I	0.13	0.13	0.16	0.09	0.09	0.07	0.22	0.24	0.23	**0.25**
J	0.06	0.06	0.06	0.06	0.06	0.06	**0.18**	0.17	0.17	0.16
K	0.51	0.57	0.32	0.45	0.33	0.62	0.44	0.57	0.78	**0.79**
L	0.54	0.53	0.30	0.25	0.23	0.23	0.60	0.68	**0.75**	0.74
M	0.30	0.31	0.23	0.27	0.23	0.22	0.24	0.31	0.27	**0.36**
N	0.71	**0.72**	0.64	0.56	0.57	0.56	0.66	**0.72**	**0.72**	**0.72**
O	0.30	**0.89**	0.19	0.15	0.15	0.10	0.51	0.41	0.52	0.42
P	0.01	-0.00	-0.05	-0.01	0.11	0.08	0.01	0.05	**0.28**	0.24
Q	0.25	0.25	0.54	0.58	0.58	0.58	0.50	**0.59**	0.58	0.58
R	0.14	0.12	0.12	0.12	0.14	0.13	**0.21**	0.15	0.20	0.14
S	0.37	0.37	**0.47**	0.44	0.22	0.23	0.30	0.33	0.34	0.33
T	0.14	0.15	0.02	0.07	0.04	0.07	0.13	**0.21**	0.20	0.20
Avg.	0.32	0.36	0.26	0.28	0.26	0.27	0.39	0.43	**0.46**	**0.46**

Table 7: Accuracy (ARI) comparison on the UCI datasets

The results of this experiment confirm the observations from the previous sections. We see that the combination of GMM, regularization and global optimization consistently achieves the best performance. Overall, the regularized GMM with a global search strategy (i.e., GMM HG Shrunk, closely followed by GMM RS Shrunk) appears to be the best method overall and achieves the best results or closely-tied results on all datasets except one (Wholesale Customers – R). As seen in Table 9, the superiority of GMM HG Shrunk in terms of accuracy is confirmed at a 5% confidence level by pairwise Wilcoxon tests (considering the individual results of all runs on all data sets) comparing the ARI of the GMM HG Shrunk method with that of the other algorithms.

In terms of computational effort, most global-search algorithms require multiple search trajectories and consequently additional computational time. Nevertheless, CPU time does not exceed a few minutes except on the Letter Recognition (I) and Pen-Based Recognition (L) datasets. In those cases, if the computational effort becomes a bottleneck, parallel implementations of the proposed methods can be developed, e.g., by performing multiple solution generation by crossover and EM local search simultaneously.

#	k-means	k-means HG	GMM	GMM MS	GMM RS	GMM HG	GMM Shrunk	GMM MS Shrunk	GMM RS Shrunk	GMM HG Shrunk
A	0.35	45.72	0.91	120.32	101.80	121.99	0.16	31.46	19.69	24.51
B	0.01	0.51	0.21	87.04	68.55	205.07	0.01	1.02	0.60	1.05
C	0.39	79.97	1.73	302.78	84.01	279.19	2.46	420.81	155.97	278.65
D	0.03	3.68	0.02	4.65	2.66	4.54	0.02	2.97	2.04	2.28
E	3.38	638.53	32.18	4561.10	1768.65	3998.46	17.11	3779.30	1631.54	3162.05
F	0.10	15.27	0.84	180.83	151.28	114.98	0.13	23.76	15.77	19.47
G	0.01	0.74	0.01	1.79	1.04	1.66	0.01	1.29	0.97	1.33
H	0.01	0.21	0.01	0.36	0.23	0.40	0.01	0.38	0.25	0.35
I	3.50	562.00	6.80	3460.18	15382.47	18090.10	5.01	945.54	318.72	693.82
J	0.54	97.61	0.13	24.58	26.89	28.38	0.14	26.45	21.88	27.46
K	0.06	7.14	0.43	434.43	213.42	277.03	0.31	72.36	22.47	45.86
L	0.67	115.91	24.86	4200.66	4754.77	5945.49	0.77	156.51	76.45	122.72
M	0.01	0.33	0.16	32.54	14.42	31.52	0.05	6.65	10.58	10.77
N	0.01	0.36	0.01	0.51	0.66	0.77	0.01	0.41	0.45	0.58
O	3.30	632.13	4.00	698.21	1663.07	1401.05	2.06	402.58	160.10	352.79
P	0.01	0.69	0.08	1.29	0.47	0.83	0.01	0.77	0.28	0.50
Q	0.08	19.53	0.18	23.22	16.57	20.28	0.12	24.07	22.76	28.16
R	0.01	1.84	0.01	2.12	4.61	2.78	0.01	1.97	1.14	1.47
S	0.01	0.32	0.05	0.74	1.56	1.00	0.03	5.56	6.36	7.13
T	0.06	9.38	0.12	16.80	9.65	23.81	0.05	9.29	3.78	7.01
Avg.	0.63	111.59	3.64	707.71	1213.34	1527.47	1.42	295.66	123.59	239.40

Table 8: CPU time (s) comparison on the UCI datasets

Algorithm	p-value
GMM HG Shrunk – k-means	4.27×10^{-23}
GMM HG Shrunk – k-means HG	8.83×10^{-14}
GMM HG Shrunk – GMM	9.13×10^{-27}
GMM HG Shrunk – GMM MS	2.18×10^{-28}
GMM HG Shrunk – GMM RS	2.26×10^{-32}
GMM HG Shrunk – GMM HG	5.13×10^{-31}
GMM HG Shrunk – GMM Shrunk	3.96×10^{-09}
GMM HG Shrunk – GMM MS Shrunk	7.39×10^{-05}
GMM HG Shrunk – GMM RS Shrunk	3.73×10^{-02}

Table 9: Pairwise Wilcoxon tests comparing the ARI of GMM HG Shrunk and that of the other algorithms

5. Conclusion

Due to their conceptual simplicity, k-means algorithm variants have been extensively used for unsupervised clustering. However, these algorithms are limited in their capacity by the fact that they essentially fit spherical Gaussian distributions to data that significantly deviates from these distributional assumptions. Against this background, GMM approaches give more modeling flexibility, but their increased number of parameters leads to more challenging underlying optimization problems and possible overfit.

In this work, we have revisited this status quo by examining the relations between regularization techniques and more sophisticated search methods within the GMM model, therefore circumventing both aforementioned weaknesses. We introduced an efficient population-based GMM (HG GMM) combined with simple regularization strategies. Through extensive numerical experiments on synthetic and real data, we demonstrate that this combination of global optimization and proper regularization permits achieving a totally new level of performance in terms of clustering accuracy (ARI). Strikingly, the use of more advanced sophisticated optimization alone was not sufficient to achieve these results; and even proved detrimental to some high-dimensional datasets. In a similar fashion, the use of regularization alone did not have a significant impact on the performance of the classic GMM. It is truly the combination of both strategies that permitted a major performance breakthrough.

This study opens many promising research avenues. First of all, we suggest pursuing the analysis of different regularization strategies in this context, e.g., to better learn or adapt the δ_{Shrunk} parameter. We also suggest further refining the solution methods to achieve faster and more accurate results. In particular, if computational time becomes a bottleneck, parallel computing may be used to generate multiple solutions simultaneously and speed up the solution process of HG or RS algorithmic variants. Finally, additional adaptations and specializations could be envisaged for different domains or data types, e.g., for computer vision or time-series analysis. To facilitate further studies, we provide all the data and analyses from our study in an open-source repository at `github.com/raphasampaio/RegularizationAndGlobalOptimizationInModelBasedClustering.jl`, as well the Julia packages UnsupervisedClustering.jl and RegularizedCovarianceMatrices.jl.

References

[1] Anguita, D., A. Ghio, L. Oneto, X. Parra, J.L. Reyes-Ortiz, et al. 2013. A public domain dataset for human activity recognition using smartphones. *Esann*, vol. 3. 3.

[2] Banfield, J.D., A.E. Raftery. 1993. Model-based Gaussian and non-Gaussian clustering. *Biometrics* 803–821.

[3] Bezanson, J., A. Edelman, S. Karpinski, V.B. Shah. 2017. Julia: a fresh approach to numerical computing. *SIAM Review* **59**(1) 65–98.

[4] Biernacki, C., G. Celeux, G. Govaert. 2003. Choosing starting values for the EM algorithm for getting the highest likehood in multivariate Gaussian mixture models. *Computational Statistics and Data Analysis* **41**(3-4) 561–575.

[5] Bilmes, J.A., et al. 1998. A gentle tutorial of the EM algorithm and its application to parameter estimation for Gaussian mixture and hidden Markov models. *International Computer Science Institute* **4**(510) 126.

[6] Bishop, C.M. 2006. *Pattern Recognition and Machine Learning*. Springer, New York.

[7] Blum, A., J. Hopcroft, R. Kannan. 2020. *Foundations of data science*. Cambridge University Press.

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

[9] Cerioli, A. 2005. K-means cluster analysis and Mahalanobis metrics: A problematic match or an overlooked opportunity. *Statistica Applicata* **17**(1) 61–73.

[10] Chen, Y., A. Wiesel, Y.C. Eldar, A.O. Hero. 2010. Shrinkage algorithms for MMSE covariance estimation. *IEEE Transactions on Signal Processing* **58**(10) 5016–5029.

[11] Dehouche, Nassim. 2020. Dataset on usage and engagement patterns for Facebook Live sellers in Thailand. *Data in brief* **30** 105661.

[12] Dempster, A.P., N.M. Laird, D.B. Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)* **39**(1) 1–22.

[13] Dua, D., C. Graff. 2017. UCI Machine Learning Repository.

[14] Dundar, M.M., D. Landgrebe. 2002. A model-based mixture-supervised classification approach in hyperspectral data analysis. *IEEE Transactions on Geoscience and Remote Sensing* **40**(12) 2692–2699.

[15] Estivill-Castro, V. 2002. Why so many clustering algorithms: a position paper. *ACM SIGKDD Explorations Newsletter* **4**(1) 65–75.

[16] Falkenauer, E. 1996. A hybrid grouping genetic algorithm for bin packing. *Journal of Heuristics* **2**(1) 5–30.

[17] Fogel, D.B. 1997. Evolutionary algorithms in theory and practice. *Complexity* **2**(4) 26–27.

[18] Fränti, P. 2018. Efficiency of random swap clustering. *Journal of Big Data* **5**(1) 13.

[19] Friedman, J., T. Hastie, R. Tibshirani, et al. 2001. *The elements of statistical learning*, vol. 1. Springer Series in Statistics New York.

[20] Glover, F.W., G.A. Kochenberger. 2006. *Handbook of metaheuristics*, vol. 57. Springer Science & Business Media.

[21] Gribel, D., T. Vidal. 2019. HG-MEANS: A scalable hybrid genetic algorithm for minimum sum-of-squares clustering. *Pattern Recognition* **88** 569–583.

[22] Halbe, Z., M. Bortman, M. Aladjem. 2013. Regularized mixture density estimation with an analytical setting of shrinkage intensities. *IEEE Transactions on Neural Networks and Learning Systems* **24**(3) 460–470.

[23] Hartigan, J.A., M.A. Wong. 1979. Algorithm AS 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **28**(1) 100–108.

[24] Hastie, T., R. Tibshirani. 1996. Discriminant adaptive nearest neighbor classification and regression. *Advances in Neural Information Processing Systems*. 409–415.

[25] Hawkins, D.M. 2004. The problem of overfitting. *Journal of Chemical Information and Computer Sciences* **44**(1) 1–12.

[26] Higuera, C., K.J. Gardiner, K.J. Cios. 2015. Self-organizing feature maps identify proteins critical to learning in a mouse model of down syndrome. *PloS One* **10**(6).

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

[28] Jain, A.K. 2010. Data clustering: 50 years beyond k-means. *Pattern Recognition Letters* **31**(8) 651–666.

[29] Kuhn, H.W. 1955. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly* **2**(1-2) 83–97.

[30] Law, M.H.C., M.A.T. Figueiredo, A.K. Jain. 2004. Simultaneous feature selection and clustering using mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **26**(9) 1154–1166.

[31] Ledoit, O., M. Wolf. 2004. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis* **88**(2) 365–411.

[32] Ledoit, O., M. Wolf. 2019. The power of (non-)linear shrinking: a review and guide to covariance matrix estimation. *Journal of Financial Econometrics* .

[33] Ledoit, Olivier, M. Wolf. 2004. Honey, I shrunk the sample covariance matrix. *The Journal of Portfolio Management* **30**(4) 110–119.

[34] Lloyd, S. 1982. Least squares quantization in PCM. *IEEE Transactions on Information Theory* **28**(2) 129–137.

[35] Lourenço, H.R., O.C. Martin, T. Stützle. 2003. Iterated local search. *Handbook of metaheuristics*. Springer, 320–353.

[36] McLachlan, G., T. Krishnan. 2007. *The EM algorithm and extensions*, vol. 382. John Wiley & Sons.

[37] McNicholas, P.D. 2016. Model-based clustering. *Journal of Classification* **373** 331–373.

[38] Mecler, J., A. Subramanian, T. Vidal. 2021. A simple and effective hybrid genetic search for the job sequencing and tool switching problem. *Computers & Operations Research* **127** 105153.

[39] Michalewicz, Z., D.B. Fogel. 2013. *How to solve it: Modern heuristics*. Springer Science & Business Media.

[40] Moore, A.W. 1999. Very fast EM-based mixture model clustering using multiresolution kd-trees. *Advances in Neural information processing systems*. 543–549.

[41] Morales-Esteban, A., F. Martínez-Álvarez, S. Scitovski, R. Scitovski. 2014. A fast partitioning algorithm using adaptive Mahalanobis clustering with application to seismic zoning. *Computers & Geosciences* **73** 132–141.

[42] Pan, W., X. Shen. 2007. Penalized model-based clustering with application to variable selection. *Journal of Machine Learning Research* **8** 1145–1164.

[43] Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **12** 2825–2830.

[44] Pernkopf, F., D. Bouchaffra. 2005. Genetic-based EM algorithm for learning Gaussian mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **27**(8) 1344–1348.

[45] Qiu, W., H. Joe. 2006. Generation of random clusters with specified degree of separation. *Journal of Classification* **23**(2) 315–334.

[46] Qiu, W., H. Joe. 2006. Separation index and partial membership for clustering. *Computational Statistics & Data Analysis* **50**(3) 585–603.

[47] Qiu, W., H. Joe. 2015. *clusterGeneration: Random Cluster Generation (with Specified Degree of Separation)*. R package version 1.3.4.

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

[49] Steinley, D. 2004. Properties of the Hubert-Arable Adjusted Rand Index. *Psychological Methods* **9**(3) 386.

[50] Ueda, N., R. Nakano, Z. Ghahramani, G.E. Hinton. 2000. SMEM algorithm for mixture models. *Neural Computation* **12**(9) 2109–2128.

[51] Vidal, T. 2022. Hybrid genetic search for the CVRP: Open-source implementation and SWAP* neighborhood. *Computers and Operations Research* **140** 105643.

[52] Vidal, T., T.G. Crainic, M. Gendreau, N. Lahrichi, W. Rei. 2012. A hybrid genetic algorithm for multidepot and periodic vehicle routing problems. *Operations Research* **60**(3) 611–624.

[53] Vinh, N.X., J. Epps, J. Bailey. 2010. Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *Journal of Machine Learning Research* **11**(Oct) 2837–2854.

[54] Warton, D.I. 2008. Penalized normal likelihood and ridge regularization of correlation and covariance matrices. *Journal of the American Statistical Association* **103**(481) 340–349.

[55] Won, J.H., J. Lim, S.J. Kim, B. Rajaratnam. 2013. Condition-number-regularized covariance estimation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **75**(3) 427–450.

[56] Xian Wang, H., B. Luo, Q. Bing Zhang, S. Wei. 2004. Estimation for the number of components in a mixture model using stepwise split-and-merge EM algorithm. *Pattern Recognition Letters* **25**(16) 1799–1809.

[57] Xiang, S., F. Nie, C. Zhang. 2008. Learning a Mahalanobis distance metric for data clustering and classification. *Pattern Recognition* **41**(12) 3600–3612.

[58] Zhang, B., C. Zhang, X. Yi. 2004. Competitive EM algorithm for finite mixture models. *Pattern Recognition* **37**(1) 131–144.

[59] Zhao, Q., V. Hautamäki, I. Kärkkäinen, P. Fränti. 2012. Random Swap EM algorithm for Gaussian Mixture Models. *Pattern Recognition Letters* **33**(16) 2120–2126.

[60] Zivkovic, Z. 2004. Improved adaptive Gaussian mixture model for background subtraction. *Pattern Recognition, 2004. ICPR 2004. Proceedings of the 17th International Conference on*, vol. 2. IEEE, 28–31.