



Group-Wise Shrinkage Estimation in Penalized Model-Based Clustering

Alessandro Casa¹ · Andrea Cappelletto² · Michael Fop³

Accepted: 12 September 2022 / Published online: 11 October 2022
© The Author(s) 2022

Abstract

Finite Gaussian mixture models provide a powerful and widely employed probabilistic approach for clustering multivariate continuous data. However, the practical usefulness of these models is jeopardized in high-dimensional spaces, where they tend to be over-parameterized. As a consequence, different solutions have been proposed, often relying on matrix decompositions or variable selection strategies. Recently, a methodological link between Gaussian graphical models and finite mixtures has been established, paving the way for penalized model-based clustering in the presence of large precision matrices. Notwithstanding, current methodologies implicitly assume similar levels of sparsity across the classes, not accounting for different degrees of association between the variables across groups. We overcome this limitation by deriving group-wise penalty factors, which automatically enforce under or over-connectivity in the estimated graphs. The approach is entirely data-driven and does not require additional hyper-parameter specification. Analyses on synthetic and real data showcase the validity of our proposal.

Keywords Model-based clustering · Penalized likelihood · Sparse precision matrices · Gaussian graphical models · Graphical lasso · EM algorithm

1 Introduction

In their recent work, Gelman and Vehtari (2021) include regularized estimation procedures among the most important contributions to the statistical literature of the last fifty years. Technological advancements and the booming of data complexity, both from a dimensional and structural perspective, have fostered the development of complex models, often involving an increasingly large number of parameters. Different regularization strategies have

Alessandro Casa and Andrea Cappelletto contributed equally to this work.

✉ Alessandro Casa
alessandro.casa@unibz.it

¹ Faculty of Economics and Management, Free University of Bozen-Bolzano, Piazza Università 1, 39100 Bolzano, Italy

² MOX - Laboratory for Modeling and Scientific Computing, Politecnico di Milano, Milan, Italy

³ School of Mathematics & Statistics, University College Dublin, Dublin, Ireland

been proposed to obtain good estimates and predictions in these otherwise troublesome settings. In this framework, a considerable amount of effort has been put into the estimation of sparse structures (see Hastie et al., 2015, for a review). The rationale underlying the sparsity concept assumes that only a small subset of parameters of a given statistical model is truly relevant. As a consequence, sparse procedures usually include penalization terms in the objective function to be optimized, forcing the estimates of some parameters to be equal to zero. These machineries generally lead to an improvement in terms of interpretability and stability of the results, as well as to advantages from a computational perspective, while reducing the risk of overfitting.

Sparse modeling has been successfully applied in regression and in supervised classification contexts. Furthermore, these strategies have been recently employed also in the model-based clustering framework, where Gaussian mixture models are usually considered to group multivariate continuous data. As a matter of fact, these models tend to be over-parameterized in high-dimensional scenarios (see Bouveyron & Brunet-Saumard, 2014, for a discussion), where the detection of meaningful partitions becomes more troublesome. For this reason, penalized likelihood methods have been considered, inducing sparsity in the resulting parameter estimates, and possibly performing variable selection (see, e.g., Pan & Shen, 2007; Xie et al., 2008; Zhou et al., 2009). In particular Zhou et al. (2009) propose a penalized approach which drastically reduces the number of parameters to be estimated in the component inverse covariance, or precision, matrices. This method exploits the connection between Gaussian mixture models and Gaussian graphical models (GGM; Whittaker, 1990), which provides a convenient way to graphically represent the conditional dependencies encoded in the precision matrices. The estimation of such matrices is difficult when the number of variables is comparable to or greater than the sample size. For this reason, a fruitful line of research has focused on sparsity-inducing procedures, which allows to obtain estimates in high-dimensional scenarios: readers may refer to Pourahmadi (2013) for a detailed treatment of the topic.

The approach by Zhou et al. (2009) induces the intensity of the penalization imposed to be common for all the component precision matrices, thus implicitly assuming that the conditional dependence structure among the variables is similar across classes. This assumption can be harmful and too restrictive in those settings where the association patterns are cluster-dependent. For instance, the method can be inappropriate to classify subjects affected by autism spectrum disorder, which might present under or over-connected fMRI networks with respect to control individuals (see Hull et al., 2017, for a review on the topic). Another relevant example can be found in the field of digits recognition, for which dependence structures between pixels may vary greatly across digits: a comprehensive analysis is reported in Section 5.2.

In order to circumvent this drawback, a possible solution consists in considering class-specific penalization intensities. While reasonable, this approach implies a rapidly increasing computational burden and it substantially becomes impractical even with a moderate number of classes. Other viable strategies would resort to procedures that deal with the estimation of GGMs in the multi-class framework (see, e.g., Danaher et al., 2014, and references therein). Nonetheless, most of these proposals adopt a borrowing-strength strategy, encouraging the estimated precision matrices to be similar across classes. This behavior may be inappropriate in a clustering context, since it might hinder groups discrimination and jeopardize the output of the analysis.

In this work, taking our step from Friedman et al. (2008) where single class inverse covariance estimation is considered, we introduce a generalization of the method by Zhou et al. (2009), which may be consequently seen as a particular case of our proposal. Even

if considering a single penalization parameter, thus avoiding the troublesome selection of more shrinkage terms, our approach turns out to be more flexible and adaptive since it penalizes a class-specific transformation of the precision matrices rather than the matrices themselves. In such a way, we are able to encompass under or over-connectivity situations as well as scenarios where the GGMs share similar structures among the groups.

The rest of the paper is structured as follows. Section 2 briefly recalls the model-based clustering framework, with a specific focus on the strategies proposed to deal with over-parameterized mixture models. In Section 3 we motivate and present our proposal, both in terms of model specification and estimation. In Sections 4 and 5 the performances and the applicability of the proposed approach are tested on synthetic and real data, respectively. Lastly, the paper ends with a brief discussion in Section 6.

2 Preliminaries and Related Work

Model-based clustering (Fraley and Raftery, 2002; McNicholas, 2016; Bouveyron et al., 2019) represents a well established and probabilistic-based approach to account for possible heterogeneity in a population. In this framework, the data generating mechanism is assumed to be adequately described by means of a finite mixture of probability distributions, with a one-to-one correspondence between the mixture components and the unknown groups. More specifically, let $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ be the set of observed data with $\mathbf{x}_i \in \mathbb{R}^p$, for $i = 1, \dots, n$, and n denoting the sample size. The density of a generic data point \mathbf{x}_i is given by

$$f(\mathbf{x}_i; \Psi) = \sum_{k=1}^K \pi_k f_k(\mathbf{x}_i; \Theta_k) \quad (1)$$

where K is the number of mixture components, π_k are the mixing proportions with $\pi_k > 0$ and $\sum_k \pi_k = 1$, and $\Psi = \{\pi_1, \dots, \pi_{K-1}, \Theta_1, \dots, \Theta_K\}$ is the vector of model parameters.

In (1), $f_k(\cdot; \Theta_k)$ represents the generic k -th component density; even if other flexible choices have been proposed (see, e.g., McLachlan & Peel, 1998; Lin, 2009, 2010; Vrbik & McNicholas, 2014), when dealing with continuous data, Gaussian densities are commonly employed. Therefore, we assume that $f_k(\cdot; \Theta_k) = \phi(\cdot; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$, where $\phi(\cdot; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ denotes the density of a multivariate Gaussian distribution with mean vector $\boldsymbol{\mu}_k = (\mu_{1k}, \dots, \mu_{pk})$, covariance matrix $\boldsymbol{\Sigma}_k$, and with $\Theta_k = \{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\}$, for $k = 1, \dots, K$.

Operationally, maximum likelihood estimation of Ψ is carried out by means of the EM algorithm (Dempster et al., 1977). This is achieved by resorting to the missing data representation of model (1), with $\mathbf{y}_i = (\mathbf{x}_i, \mathbf{z}_i)$ denoting the complete data, with $\mathbf{z}_i = (z_{i1}, \dots, z_{iK})$ latent group indicators where $z_{ik} = 1$ if the i -th observation belongs to the k -th cluster and $z_{ik} = 0$ otherwise. Considering a one-to-one correspondence between clusters and mixture components, as it is common in the general model-based clustering framework, the partition is obtained assigning the i -th observation to cluster k^* if

$$k^* = \arg \max_{k=1, \dots, K} \frac{\pi_k \phi(\mathbf{x}_i; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{v=1}^K \pi_v \phi(\mathbf{x}_i; \boldsymbol{\mu}_v, \boldsymbol{\Sigma}_v)},$$

according to the so-called maximum a posteriori (MAP) classification rule (see Ch. 2.3 in Bouveyron et al., 2019, for details).

One of the major limitations of Gaussian mixture models is given by their tendency to be over-parameterized in high-dimensional scenarios. In fact, the cardinality of Ψ is of order $\mathcal{O}(Kp^2)$, thus scaling quadratically with the number of the observed variables and often being larger than the sample size. In order to mitigate this issue, several different approaches have been studied, and readers may refer to Bouveyron and Brunet-Saumard (2014) and Fop and Murphy (2018) for exhaustive surveys on the topic. Hereafter, we outline some of the proposals introduced to deal with over-parameterized mixture models. Roughly speaking, we might identify three different types of approaches, namely constrained modeling, sparse estimation, and variable selection.

The first strategy relies on constrained parameterizations of the component covariance matrices. The proposals by Banfield and Raftery (1993) and Celeux and Govaert (1995) aim to reduce the number of free parameters by considering an eigen decomposition of Σ_k , which allows to control the shape, the orientation, and the volume of the clusters. Other works falling within this framework are, to mention a few, the ones by McLachlan et al. (2003), McNicholas and Murphy (2008), and Bouveyron et al. (2007) and Biernacki and Lourme (2014). Most of these methodologies do not directly account for the associations between the observed variables, resorting to matrix decompositions and focusing on the geometric characteristics of the component densities. As a consequence, parsimony is induced in a rigid way and the interpretation in some cases is not straightforward.

The second class of approaches employs flexible sparsity-inducing procedures, to overcome the limitations of constrained modeling. As an example, we mention the methodology recently proposed by Fop et al. (2019), where a mixture of Gaussian covariance graph models is devised, coupled with a penalized likelihood estimation strategy. This approach eases the interpretation of the results in terms of marginal independence among the variables and allows for cluster-wise different association structures, by obtaining sparse estimates of the covariance matrices.

Lastly, variable selection has been explored in this context, following two distinct paths. On one hand, the problem has been recast in terms of model selection, with models defined considering different sets of variables being compared by means of information criteria (Raftery and Dean, 2006; Maugis et al., 2009a; b). On the other hand, the second class of approaches lies in between the variable selection and the sparse estimation methodologies. In fact, in Pan and Shen (2007), Xie et al. (2008), and Zhou et al. (2009) a penalty term is considered in the Gaussian mixture model log-likelihood to induce sparsity in the resulting estimates and thus possibly identifying a subset of irrelevant variables.

In the following, we focus specifically on the work by Zhou et al. (2009), where the penalty is placed on the inverse covariance parameters of the Gaussian mixture components. Such a penalty is considered for regularization and for obtaining sparse estimates of the association matrices in high-dimensional settings. Parameters estimation, and the subsequent clustering step, are carried out by maximizing the following penalized log-likelihood

$$\tilde{\ell}_p(\Psi) = \sum_{i=1}^n \log \sum_{k=1}^K \pi_k \phi(\mathbf{x}_i; \boldsymbol{\mu}_k, \boldsymbol{\Omega}_k) - \lambda \sum_{k=1}^K \|\boldsymbol{\Omega}_k\|_1. \quad (2)$$

The first term is the log-likelihood of a Gaussian mixture model, parametrized in terms of the component precision matrices $\boldsymbol{\Omega}_k = \boldsymbol{\Sigma}_k^{-1}$, for $k = 1, \dots, K$. The second term corresponds to the graphical lasso penalty (see, e.g., Banerjee et al., 2008; Friedman et al., 2008; Scheinberg et al., 2010; Witten et al., 2011) applied to the component-specific precision matrices, with the L_1 norm taken element-wise, i.e., $\|A\|_1 = \sum_{ij} |A_{ij}|$; in the following,

we do not apply the penalty on the diagonal elements of the precision matrices, even if it is possible in principle.

Note that Zhou et al. (2009) consider an additional penalty term $\lambda_2 \sum_{k=1}^K \sum_{j=1}^p |\mu_{jk}|$ in (2): this corresponds to the lasso penalty function (Tibshirani, 1996) applied element-wise to the mean component vectors employed to perform variable selection. Since our primary focus is in uncovering the conditional dependence structure enclosed in Ω_k , we are not concerned in providing penalized estimators for the component mean vectors, we therefore do not include this term in (2).

The graphical lasso penalty allows to induce sparsity in the precision matrices, which eases the interpretation of the model. In fact, Ω_k embeds the conditional dependencies among the variables for the k -th component, whereby in the Gaussian case zero entries between pair of variables imply that they are conditionally independent given all the others. Moreover, a convenient way to visually represent the dependence structure among the features is given by the graph of the associated Gaussian graphical model. Here, as already mentioned in the introduction, a correspondence between a sparse precision matrix and a graph is defined, with nodes representing the variables while the edges connect only those features being conditionally dependent. A recent and interesting extension is represented by the *colored graphical models*, where symmetry restrictions are added to the precision matrix thus offering a more parsimonious representation and possibly highlighting commonalities among the variables; readers may refer to Højsgaard and Lauritzen (2008), Gao and Massam (2015), and Li et al. (2021) and references therein for a more detailed discussion.

Sparse precision matrix estimation via the graphical lasso algorithm is routinely employed assuming that observations arise from the same population, adequately described by a single multivariate Gaussian distribution, which is indexed by a single GGM. However, this assumption does not hold in the cluster analysis framework where, as in (1), the observed data are assumed to arise from K different sub-populations, which might be characterized by different association patterns. Some modifications of the standard graphical lasso have been proposed, in order to make it applicable also in a multi-class setting (see, e.g., Guo et al., 2011; Mohan et al., 2014; Danaher et al., 2014; Lyu et al., 2018). Nonetheless, these approaches usually consider the matrices Ω_k to have possible commonalities and shared sparsity patterns; as a consequence, they modify the graphical lasso penalty term in order to induce the estimated GGMs to be similar to each other, while allowing for structural differences. These strategies have been usually considered with an exploratory aim in mind, in order to obtain parsimonious and interpretable characterization of the relationships among the variables within and between the classes. Undoubtedly, they might be fruitfully embedded also in a probabilistic unsupervised classification context. However, to some extent, by borrowing strength and encouraging similarity among groups, these methods may be inappropriate, if not harmful, as they might hinder the classification task itself. This particularly holds in the case of clustering, where the classes are not readily available and need to be inferred from the data. For this reason, in the next section we focus on how to obtain cluster-specific sparse precision matrices to account for cluster-wise distinct degrees of sparsity.

3 Proposal

All the multi-class GGM estimation strategies reviewed in Section 2 assume that different classes are characterized by a similar structure in the precision matrices, either explicitly or implicitly. This assumption is explicit for those approaches where similarities among the precision matrices are encouraged by the considered penalty term. Similarly, in the

approach proposed by Zhou et al. (2009) the assumption is implicitly entailed by the use of a single penalization parameter λ . The adopted penalization scheme, even if somehow weighted by the clusters sample sizes, as per (12), can be profitably considered only in those situations where the number of non-zero entries in the precision matrices is similar across classes. Therefore, these approaches do not contemplate under or over-connectivity scenarios, where the groups are characterized by significantly different amounts of sparsity. This constitutes a serious limitation in those applications where different degrees of connectivity could ultimately characterize the resulting data partition. That is, while approaches such as the one by Zhou et al. (2009) well encompass scenarios in which connected nodes are group-wise different, we aim at defining a data-driven strategy specifically designed for addressing also those situations in which groups differ in the amount of sparsity (i.e., in the number of non-zero entries in the precision matrices).

3.1 Motivating Example

To better explain this issue and to justify our solution, we provide a motivating example where we simulate $n = 200$ p -dimensional observations with $p = 20$ from a mixture with $K = 2$ Gaussian components and mixing proportions $\pi = (0.5, 0.5)$. The considered component precision matrices are associated with the two sparse at random structures depicted in the top panel of Fig. 1. The first component is characterized by an almost diagonal precision matrix while the second presents a dense structure, thus mimicking a scenario where the degree of sparsity is drastically different among the two classes. We estimate these matrices employing the penalization scheme in (2), with λ ranging over a suitable interval. The ability in recovering the association patterns inherent to the two clusters is evaluated by means of the F_1 score:

$$F_1 = \frac{tp}{tp + 0.5(fp + fn)}, \quad (3)$$

where tp denotes the correctly identified edges (i.e., the number of non-zero entries in the precision matrix correctly estimated as such), while fp and fn represent respectively the number of incorrectly identified edges and the number of missed edges (i.e., the number of non-zero entries wrongly shrunk to 0). Line plots displaying the F_1 patterns for the two components are reported in the bottom panel of Fig. 1. A trade-off is clearly visible, indicating how a common penalty term prevents the possibility to obtain a proper estimation of both the precision matrices. In fact, if for the second component a mild penalization might be adequate to estimate the dense dependence structure, for the first component the high degree of sparsity is recovered only when a stronger penalty is considered.

3.2 Model Specification

The illustrative example clearly shows how the penalization strategy proposed by Zhou et al. (2009) does not represent a suitable solution when dealing with unbalanced class-specific degrees of sparsity among the variables. As briefly mentioned in the introduction, a possible alternative would consist in uncoupling the precision matrices estimation by considering component-specific penalization coefficients. That is, λ in (2) would be substituted by K different shrinkage terms. While in principle reasonable, the increased flexibility induced by introducing K different penalties may be problematic, since in the graphical lasso framework tuning

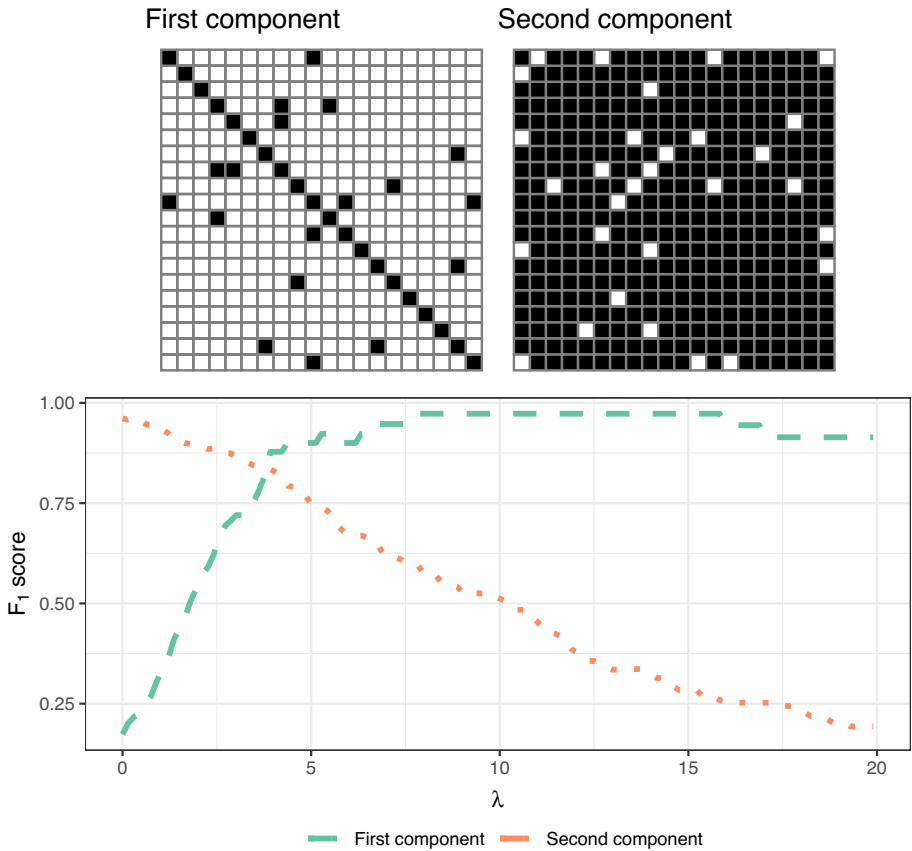


Fig. 1 On the top: true data generating component-specific precision matrices. Black squares denote the presence of an edge between the two variables in the corresponding GGM. On the bottom: F_1 score as a function of the parameter λ

these hyper-parameters is a difficult task. Time consuming grid searches are generally considered, with the optimal penalty factor selected either according to some information criteria or resorting to computationally intensive cross-validation strategies. The simultaneous presence of K penalty terms would make this approach much more complex, also from a computational perspective.

In this work we propose instead to carry out parameter estimation by maximizing a penalized log-likelihood function defined as follows:

$$\ell_P(\Psi) = \sum_{i=1}^n \log \sum_{k=1}^K \pi_k \phi(\mathbf{x}_i; \boldsymbol{\mu}_k, \boldsymbol{\Omega}_k) - \lambda \sum_{k=1}^K \|\mathbf{P}_k \circ \boldsymbol{\Omega}_k\|_1, \tag{4}$$

where \circ denotes the Hadamard product between two matrices and \mathbf{P}_k are weighting matrices employed to scale the effect of the penalty. The focus is then shifted towards the specification of $\mathbf{P}_1, \dots, \mathbf{P}_K$, which, when properly encoding information about class-specific sparsity patterns, introduces a degree of flexibility that accounts for under or over-connectivity

scenarios. In (4) a single penalization parameter for the precision matrices is considered. As a consequence other than the model selection problem concerning the selection of the number of components K , we only need to carefully tune a single penalization hyperparameter λ . In Section 3.8 we outline a standard technique to choose the number of components and the penalty parameter λ according to a model selection criterion.

Hereafter, we describe a data-driven procedure for inferring \mathbf{P}_k by means of carefully initialized sample precision matrices $\hat{\mathbf{\Omega}}_1^{(0)}, \dots, \hat{\mathbf{\Omega}}_K^{(0)}$. Our proposals rely on the definition of a function $f: \mathbb{S}_+^p \rightarrow \mathbb{S}^p$, where \mathbb{S}_+^p and \mathbb{S}^p respectively denote the space of positive semi-definite and symmetric matrices of dimension p , such that $\mathbf{P}_k = f(\hat{\mathbf{\Omega}}_k^{(0)})$.

Recommendations on how to compute $\hat{\mathbf{\Omega}}_k^{(0)}$, and how to subsequently define \mathbf{P}_k , $k = 1, \dots, K$ will be the object of the next subsections.

3.3 Initializing the Sample Precision Matrices $\hat{\mathbf{\Omega}}_k^{(0)}$

The definition of a proper strategy to initialize the matrices $\hat{\mathbf{\Omega}}_k^{(0)}$, for $k = 1, \dots, K$, strongly depends on the framework and on the specific task of the analysis. In fact, in supervised and semi-supervised scenarios, where the class labels are known for at least a subset of observations, the initialization step turns out to be straightforward. More formally, let us denote with n the number of observations in the training set, with m the number of observations with known labels and $m = \sum_{k=1}^K m_k$, with m_1, \dots, m_K denoting the class-specific sample sizes with $m_k > 0$, for $k = 1, \dots, K$. In the supervised setting, where $m = n$, and in the semi-supervised one, where $0 < m < n$, we simply define $\hat{\mathbf{\Omega}}_k^{(0)}$ to be the k -th class sample precision matrix estimated on the pertaining m_k observations. Note that, if $p > m_k$, the initialized sample precision matrix might be obtained by means of K distinct graphical lasso algorithms.

On the other hand, in a clustering framework where $m = 0$, the specification of $\hat{\mathbf{\Omega}}_k^{(0)}$ is more tricky. The absence of clear indications about the group memberships makes the approach introduced for the supervised and semi-supervised context impractical. Nonetheless, it is possible to find suitable workarounds in order to employ our proposal even in an unsupervised scenario. From a practical point of view, we consider a general multi-step procedure as follows:

1. Run any clustering algorithm on the observed data $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, to obtain an initial partition of the observations into K groups;
2. Given the obtained initial partition, estimate the cluster-specific precision matrices $\hat{\mathbf{\Omega}}_1^{(0)}, \dots, \hat{\mathbf{\Omega}}_K^{(0)}$ using only those observations assigned to that specific group; again $\hat{\mathbf{\Omega}}_k^{(0)}$ might be obtained as the sample estimate when $p < n_k$ or as the graphical lasso solution if $p \geq n_k$.

When resorting to this procedure, the choice of which clustering strategy to adopt for obtaining the initial partition is subjective and needs to be carefully taken. In fact, inadequate choices might provide incoherent indications of the true clustering structure and hinder the possibility to obtain an accurate reconstruction of the dependencies among the variables when maximizing (4). In principle, different clustering strategies may be adopted, and providing specific suggestions about the more adequate ones is beyond the scope of this work. Nonetheless, model-based techniques (see, e.g., Fraley and Raftery, 2002) might constitute a clever choice, being them coherent with the considered framework. Also ensemble strategies can be adequate, as they

aim to combine the strengths of different algorithms and lessen the impact of some otherwise cumbersome choices (see Russell et al., 2015; Wei & McNicholas, 2015; Casa et al., 2021), for some proposals from a model-based standpoint). In addition, powerful initialization strategies for partitioning the data into K groups can as well appropriately serve the purpose (Scrucca & Raftery, 2015). Lastly, we remark that the use of standard methods to obtain an appropriate initial clustering partition can cause difficulties in high-dimensional settings where $p > n$. In these scenarios, subspace clustering methods specifically designed for high-dimensional data can be employed. Examples are mixtures of factor analyzers (McNicholas & Murphy, 2008) and model-based discriminant subspace clustering (Bouveyron & Brunet, 2012); we point the reader to Section 5 of Bouveyron and Brunet-Saumard (2014) for a comprehensive overview.

3.4 Obtaining \mathbf{P}_k via Inversely Weighted Sample Precision Matrices

Once the sample precision matrices $\hat{\mathbf{\Omega}}_k^{(0)}$ have been initialized, the first viable proposal for defining \mathbf{P}_k reads as follows:

$$P_{k,ij} = 1 / \left(|\hat{\Omega}_{k,ij}^{(0)}| \right), \quad \forall i, j = 1, \dots, p, \quad (5)$$

where $P_{k,ij}$, $\hat{\Omega}_{k,ij}^{(0)}$ are respectively the (i,j) -th elements of the matrices \mathbf{P}_k and $\hat{\mathbf{\Omega}}_k^{(0)}$. Notice that it is sufficient to set $P_{k,ii} = 0$, $\forall i = 1, \dots, p$, whenever the diagonal entries of $\mathbf{\Omega}_k$ shall not be penalized. Intuitively, with (5) we are inflating/deflating the penalty enforced on the (i,j) -th element of the matrix $\mathbf{\Omega}_k$ according to the magnitude of $\hat{\Omega}_{k,ij}^{(0)}$. Clearly, values of $|\hat{\Omega}_{k,ij}^{(0)}|$ close to 0 induce a higher penalty on $\mathbf{\Omega}_{k,ij}$. Should $\hat{\mathbf{\Omega}}_k^{(0)}$ be estimated via the graphical lasso, e.g., in those situations where $p \geq n_k$, a small positive constant is added in the denominator of (5) to avoid having an undefined \mathbf{P}_k . This strategy shares connections with the proposal by Bien and Tibshirani (2011) developed in a covariance estimation context, and it might be seen as a multi-class extension of the approach proposed in Fan et al. (2009), where the *adaptive lasso* (Zou, 2006) is generalized to the estimation of sparse precision matrices.

A hard-thresholding version of (5) may also be considered, in which entries of $\mathbf{\Omega}_k$ are not shrunk if their magnitude exceeds a given constant. A sensible way to do so would be to examine the initialized partial correlation matrix for the k -th class, and to fix a value $\gamma \in (0,1)$ that acts as a user-defined threshold. This approach is related to the fixed-zero problem of Chaudhuri et al. (2007): when λ is sufficiently large, it leads to an estimate equivalent to the one obtained from a given association graph where the zero entries correspond to partial correlations of magnitude lower than the specified threshold γ . The idea is connected to the thresholding operator for sparse covariance matrix estimation (Bickel and Levina, 2008; Pourahmadi, 2013) and it has been explored in the `covglasso` R package (Fop, 2020). The hard-thresholded approach is not further considered in the following as it requires a sensitive choice of γ , whereas the other suggested proposals do not rely on any hyper-parameter specification.

3.5 Obtaining \mathbf{P}_k via Distance Measures in the \mathbb{S}_+^p Space

An alternative approach consists in setting the elements of \mathbf{P}_k proportional to the distance between $\hat{\mathbf{\Omega}}_k^{(0)}$ and $\text{diag}(\hat{\mathbf{\Omega}}_k^{(0)})$, where $\text{diag}(\hat{\mathbf{\Omega}}_k^{(0)})$ indicates a diagonal matrix whose diagonal elements are equal to the ones in $\hat{\mathbf{\Omega}}_k^{(0)}$. We propose to compute $P_{k,ij}$ as follows:

$$P_{k,ij} = \frac{1}{\mathcal{D}(\hat{\Omega}_k^{(0)}, \text{diag}(\hat{\Omega}_k^{(0)}))}, \quad \forall i, j = 1, \dots, p \quad \text{and} \quad i \neq j, \tag{6}$$

with $\mathcal{D}(\cdot, \cdot)$ being a suitably chosen measure of distance between positive semi-definite matrices. Since \mathbb{S}_+^p is a non-Euclidean space, the type of distance needs to be carefully selected: the reader is referred to Dryden et al. (2009) for a thorough dissertation on the topic. Alternatively, to account for the diagonal elements in the definition of \mathbf{P}_k , one may employ the following quantity:

$$P_{k,ij} = \frac{1}{\mathcal{D}(\hat{\Omega}_k^{(0)}, \mathbf{I}_p)}, \quad \forall i, j = 1, \dots, p \quad \text{and} \quad i \neq j, \tag{7}$$

where \mathbf{I}_p denotes the identity matrix of dimension p . The definitions of (6) and (7) simply stem from the conjecture that the “closer” $\hat{\Omega}_k^{(0)}$ is to a diagonal matrix, the higher the group-wise penalty shall be, thus forcing some of the entries in Ω_k to be shrunk to 0.

3.6 Obtaining \mathbf{P}_k : Comparison of Methods

The strategies mentioned above share the same underlying rationale as they aim to impose stronger penalization on those entries corresponding to weaker sample conditional dependencies among variables. Moreover, being the specification class specific, they fruitfully encompass situations where one or more groups present different sparsity patterns and magnitudes. While the solution in Section 3.4 induces an entry-wise different penalty, it heavily depends on the estimation of $\hat{\Omega}_k^{(0)}$. Therefore, in those situations where the reliability of the sample estimates is difficult to assess, it might be convenient to let \mathbf{P}_k depend on a group-specific constant, as for the strategy in Section 3.5.

The proposed approaches generalize the one by Zhou et al. (2009), as the strategies coincide when \mathbf{P}_k is set to be equal to a matrix of ones for all $k = 1, \dots, K$. Once the definition of \mathbf{P}_k has been established, coherently to Zhou et al. (2009), the model is estimated employing an EM algorithm: details are given in the next section.

3.7 Model Estimation

For a fixed number of components K and penalty terms λ and \mathbf{P}_k , model estimation deals with the maximization of (4) with respect to Ψ . Within the EM framework, a *penalized complete-data log-likelihood* is naturally defined as follows:

$$\ell_C(\Psi) = \sum_{i=1}^n \sum_{k=1}^K z_{ik} \log \pi_k \phi(\mathbf{x}_i; \boldsymbol{\mu}_k, \Omega_k) - \lambda \sum_{k=1}^K \|\mathbf{P}_k \circ \Omega_k\|_1, \tag{8}$$

where as usual z_{ik} is the categorical latent variable indicating the component which observation \mathbf{x}_i belongs to. The algorithm alternates between two steps. At the t -th iteration, the E-step provides the expected value $\hat{z}_{ik}^{(t)}$ for the unknown labels z_{ik} given $\hat{\Psi}^{(t-1)}$, while in the M-step (8) is maximized to determine $\hat{\Psi}^{(t)}$, conditioning on $\hat{z}_{ik}^{(t)}$.

In detail, in the E-step the posterior probability of \mathbf{x}_i belonging to component k is updated as follows:

$$\hat{z}_{ik}^{(t)} = \frac{\hat{\pi}_k^{(t-1)} \phi(\mathbf{x}_i; \hat{\boldsymbol{\mu}}_k^{(t-1)}, \hat{\boldsymbol{\Omega}}_k^{(t-1)})}{\sum_{v=1}^K \hat{\pi}_v^{(t-1)} \phi(\mathbf{x}_i; \hat{\boldsymbol{\mu}}_v^{(t-1)}, \hat{\boldsymbol{\Omega}}_v^{(t-1)})}, \quad i = 1, \dots, n. \tag{9}$$

In the M-step, the updating formulas for mixing proportions and cluster means are readily given by:

$$\hat{\pi}_k^{(t)} = \frac{n_k^{(t)}}{n}, \quad \hat{\boldsymbol{\mu}}_k^{(t)} = \frac{1}{n_k^{(t)}} \sum_{i=1}^n \hat{z}_{ik}^{(t)} \mathbf{x}_i \tag{10}$$

where $n_k^{(t)} = \sum_{i=1}^n \hat{z}_{ik}^{(t)}$. Notice that, as mentioned in Section 2, we are not concerned in providing penalized estimators for $\boldsymbol{\mu}_k$. At any rate, should sparse mean vectors be of interest, an extra penalty can be promptly accommodated by adding the term $\lambda_2 \sum_{k=1}^K \sum_{j=1}^p |\mu_{jk}|$ in (4). In this case, estimation of sparse $\boldsymbol{\mu}_k$ follows directly the steps outlined in Section 2.3.1 of Zhou et al. (2009).

When (8) is maximized with respect to $\boldsymbol{\Omega}_k$, the penalized complete log-likelihood simplifies as follows:

$$Q_{\Omega}(\boldsymbol{\Omega}_k) = \sum_{i=1}^n \hat{z}_{ik}^{(t)} \left\{ \frac{1}{2} \log \det(\boldsymbol{\Omega}_k) - \frac{1}{2} (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_k^{(t)})' \boldsymbol{\Omega}_k (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_k^{(t)}) \right\} - \lambda \|\mathbf{P}_k \circ \boldsymbol{\Omega}_k\|_1. \tag{11}$$

By rearranging terms in (11), the following optimization problem needs to be solved to obtain the estimate $\hat{\boldsymbol{\Omega}}_k^{(t)}$:

$$\arg \max_{\boldsymbol{\Omega}_k} \log \det(\boldsymbol{\Omega}_k) - \text{tr}(\mathbf{S}_k \boldsymbol{\Omega}_k) - \frac{2\lambda}{n_k^{(t)}} \|\mathbf{P}_k \circ \boldsymbol{\Omega}_k\|_1, \tag{12}$$

with the constraint that $\boldsymbol{\Omega}_k$ must be positive definite, $\boldsymbol{\Omega}_k > 0$, and \mathbf{S}_k denoting the weighted sample covariance matrix for cluster k :

$$\mathbf{S}_k = \sum_{i=1}^n \hat{z}_{ik}^{(t)} \frac{(\mathbf{x}_i - \hat{\boldsymbol{\mu}}_k^{(t)})' (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_k^{(t)})}{n_k^{(t)}}.$$

Following Zhou et al. (2009), a coordinate descent graphical lasso algorithm by (Friedman et al., 2008) is employed for solving the maximization problem in (12), where in our context the penalty is equal to $2\lambda \mathbf{P}_k / n_k^{(t)}$.

3.8 Further Aspects

Hereafter, we discuss some practical considerations related to the algorithm devised for maximizing (4) and described in the previous section.

Convergence The EM algorithm is considered to have converged once the relative difference in the objective function for two subsequent iterations is smaller than ϵ , for a given $\epsilon > 0$:

$$\frac{|\ell_P(\hat{\Psi}^{(t+1)}) - \ell_P(\hat{\Psi}^{(t)})|}{|\ell_P(\hat{\Psi}^{(t)})|} < \varepsilon.$$

In our analyses, ε is set equal to 10^{-5} .

Model Selection While the determination of \mathbf{P}_k is entirely data-driven and does not require any external tuning, model selection still needs to be performed when it comes to identify the best number of components K and the common penalty term λ . We rely on previous results (Pan and Shen, 2007; Zou et al., 2007; Lian, 2011) which propose to select λ and K by maximizing a modified version of the Bayesian Information Criterion (BIC, Schwarz, 1978):

$$BIC = 2 \log L(\hat{\Psi}) - d_0 \log(n), \quad (13)$$

where $\log L(\hat{\Psi})$ is the log-likelihood evaluated at $\hat{\Psi}$, obtained maximizing (4), and d_0 is the number of parameters that are not shrunk to 0 by the penalized estimation.

Implementation Routines for fitting the group-wise shrinkage method for model-based clustering have been implemented in R (R Core Team, 2022), and the source code is freely available at <https://github.com/AndreaCappozzo/sparsemix> in the form of an R package. Despite not being explicitly used in the present manuscript, the `sparsemix` software allows for penalizing the mean vectors μ_k along the lines of Zhou et al. (2009), thus providing a complete generalization of the methodology described therein.

Promising results are obtained when performing penalized model-based clustering with \mathbf{P}_k defined as in (6) and (7), as reported in the next section.

4 Simulated Data Experiment

4.1 Experimental Setup

We illustrate, via numerical experiments, the effectiveness of the proposed procedures in recovering the true group-wise conditional structure in a multi-class population. We generate $n = 1500$ observations from a Gaussian mixture distribution with $K = 3$ components, with the precision matrices Ω_k having various sparsity patterns, embedding different association structures. Three different scenarios are considered:

- *Equal proportion of edges in Ω_k* : for each replication of the simulated experiment, the precision matrices Ω_k are generated according to a sparse at random Erdős-Rényi graph structure (Erdős and Rényi, 1960) characterized by the same probabilities of connection, equal to 0.5. The number of variables is $p = 20$.
- *Different proportion of edges in Ω_k* : for each replication of the simulated experiment, the precision matrices Ω_k are generated according to a sparse at random Erdős-Rényi graph structure, with probabilities of connection equal to 0.1, 0.8 and 0.4 for Ω_1 , Ω_2 and Ω_3 , respectively. The number of variables is $p = 20$.
- *High-dimensional and different proportion of edges in Ω_k* : for each replication of the simulated experiment, the precision matrices Ω_k are generated according to a sparse at




Fig. 2 Example of simulation setting *Equal proportion of edges in Ω_k* . Black squares denote the presence of an edge between the two variables

random Erdős-Rényi graph structure with different probabilities of connection as per the previous scenario. The number of variables is $p = 100$.

In all scenarios, we take equal mixing proportions $\pi_k = 1/3, k = 1,2,3$ and mean vectors equal to:

$$\mu_1 = -1.5\mathbf{e}_{20}, \quad \mu_2 = 0\mathbf{e}_{20}, \quad \mu_3 = 1.5\mathbf{e}_{20},$$

for the first two scenarios, while

$$\mu_1 = 5\mathbf{e}_{100}, \quad \mu_2 = 0\mathbf{e}_{100}, \quad \mu_3 = 5\mathbf{e}_{100},$$

for the high-dimensional case, with \mathbf{e}_{20} and \mathbf{e}_{100} identifying the all-ones vector in \mathbb{R}^{20} and \mathbb{R}^{100} , respectively. Such parameters induce a moderate degree of overlapping between components in the lower-dimensional case, while producing well-separated clusters in the high-dimensional scenario. We point out again that the primary objective of the study is assessing the recovering of the true underlying sparse precision matrices, and so we do not impose any regularization on the mean parameters. An example of the graph structures resulting from the first two scenarios are reported in Figs. 2 and 3 respectively.

We repeat the experiment $B = 100$ times, and for each replication we fit the model in (4), computing \mathbf{P}_k as follows:

- Zhou et al. (2009): \mathbf{P}_k set equal to the all-one matrix for $k = 1,2,3$
- \mathbf{P}_k via Frobenius distance in \mathbb{S}_+^p : \mathbf{P}_k computed as in (6), with $\mathcal{D}(\cdot, \cdot)$ the Frobenius distance in the \mathbb{S}_+^p space,
- \mathbf{P}_k via Riemannian distance in \mathbb{S}_+^p : \mathbf{P}_k computed as in (6), with $\mathcal{D}(\cdot, \cdot)$ the Riemannian distance in the \mathbb{S}_+^p space,
- \mathbf{P}_k via inversely weighted $|\hat{\Omega}_k^{(0)}|$: \mathbf{P}_k computed as in (5).

A grid of equispaced 100 elements for the penalty term λ is considered, with lower and upper extremes set to:

$$\left[0, \max_k \left\{ \max \left\{ |\mathbf{S}_k^{(0)} - \mathbf{I}_p| \right\} \frac{n_k^{(0)}}{2} \right\} \right], \tag{14}$$

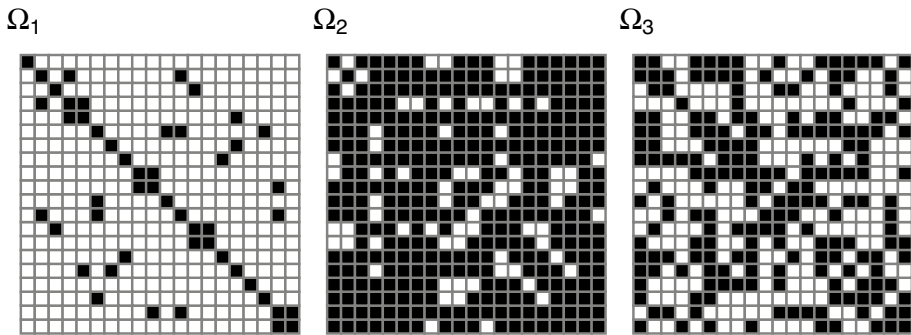


Fig. 3 Example of simulation setting *Different proportion of edges in Ω_k* . Black squares denote the presence of an edge between the two variables

where the inner max operation is taken element-wise, and with $S_k^{(0)}$ and $n_k^{(0)}$ the starting estimates of the sample covariance matrices and their associated sample sizes, initialized via Gaussian mixture models provided by the `mclust` software (Scrucca et al., 2016). Other initialization strategies are clearly possible, as described in Section 3.3. We note that the upper term in (14) forces the final estimates $\hat{\Omega}_k$ to be approximately diagonal when P_k are equal to all-ones matrices (Zhao et al., 2012).

The performance of the methods, in relation to the different specification of the P_k matrices, is evaluated via component-wise F_1 scores defined as in (3), where the problem of matching the estimated clustering to the actual classification is dealt with by means of the `matchClasses` routine of the `e1071` R package (Meyer et al., 2020). Simulation results are reported in the next subsection.

4.2 Simulation Results: Recovering the Association Structure

Results for the simulated experiments are summarized in Fig. 4, depicting smoothed lines plots of the F_1 scores for the estimated Ω_k , $k = 1, 2, 3$, under the three considered scenarios, for different specifications of the matrices P_k and shrinkage factor λ . By visually exploring Fig. 4, several interesting patterns emerge.

First off, it is immediately noticed that the methods performance in the first scenario does not vary across components, with strong similarities between F_1 score trajectories for Ω_1 , Ω_2 and Ω_3 . This is expected, as each precision matrix is generated with a probability of connection equal to 0.5. This is the “gold-standard” scenario for the method described in Zhou et al. (2009), since in principle a common λ should be sufficient for achieving the same group-wise degree of sparsity. Indeed, the highest F_1 values (around 0.7) are achieved by all proposals when small penalty values are considered. Notwithstanding, multiplying the common shrinkage term by a group-specific factor P_k moderates the rapid decline in performance when λ increases. Particularly, computing the P_k as a function of the Frobenius distance between $\hat{\Omega}_k^{(0)}$ and $\text{diag}(\hat{\Omega}_k^{(0)})$ greatly downweights the impact of the common penalty term, making the procedure less sensitive to the selection of λ .

The beneficial effect of group-wise P_k becomes apparent in the second scenario, where the true precision matrices have dissimilar a priori probability of connection. F_1 trajectories are component-wise different: the almost diagonal Ω_1 would require a higher shrinkage

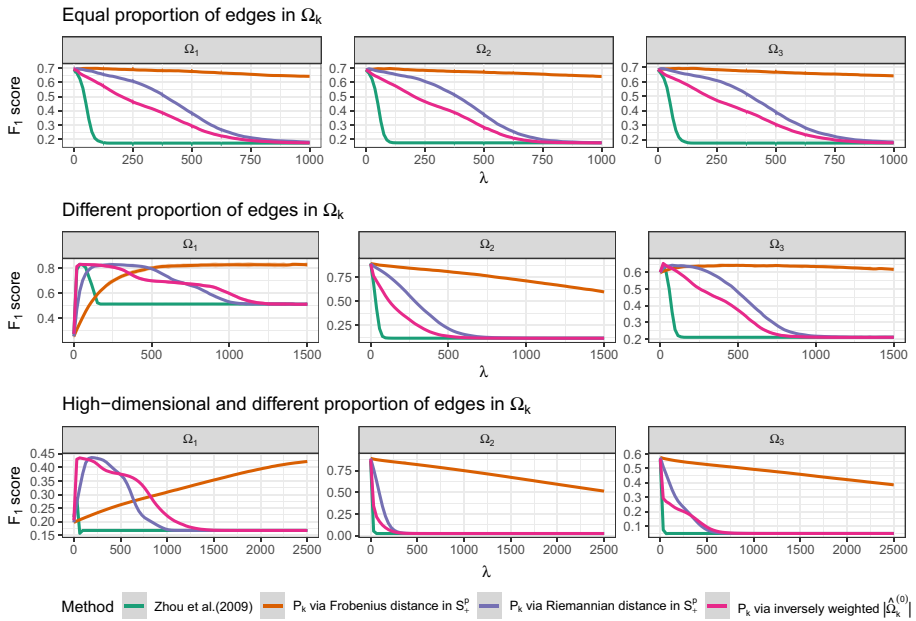


Fig. 4 Smoothed lines plots of the F_1 score for $B = 100$ repetitions of the simulated scenarios, varying method and shrinkage factor λ

for recovering the very sparse underlying graph structure, whereas the highly connected Ω_2 is well-estimated when low values of λ are considered. This trade-off is mitigated by the \mathbf{P}_k factor, which adjusts for under or over-connectivity within the estimation process. In particular, every suggested approach succeeds in improving the strategy of Zhou et al. (2009), with F_1 patterns for our proposals almost always dominating the common shrinkage method. This behavior is intensified even further in the scenario with a larger number of variables, where as soon as λ increases, the proportion of incorrectly missed edges produces a huge drop in the F_1 score for the second component. As previously highlighted, the \mathbf{P}_k via Frobenius distance in \mathbb{S}_+^p is the solution enforcing the greatest discount on the common λ , greatly improving the results for Ω_2 at the expense of overestimating the true number of edges for Ω_1 .

Figure 4 shows the overall superiority of our proposals with respect to a common penalty framework in group-wise recovering of sparse precision matrices. Nevertheless, when it comes to performing the analysis, a single value of λ must be chosen. We make use of the BIC criterion defined in Section 3.8 to determine the best λ for each method and instance of the simulated experiments. For the $B = 100$ simulations in the three scenarios, the resulting empirical distribution of the mean F_1 score averaged over $\hat{\Omega}_1$, $\hat{\Omega}_2$, and $\hat{\Omega}_3$ is reported in Fig. 5. As expected, the overall results do not dramatically change in the equal proportion of edges in Ω_k case. On the other hand, for the other two scenarios, our proposals perform substantially better compared to Zhou et al. (2009), especially in the case with larger number of variables. None of the introduced methods for computing \mathbf{P}_k seems to outperform the others; nonetheless, it is clear that, whenever the degree of conditional dependence varies greatly across components, weighting the common penalty λ by a group-specific factor improves the recovering of the group-wise different sparse structures.




Fig. 5 Boxplots of the mean F_1 score, averaged over $\hat{\Omega}_1, \hat{\Omega}_2$ and $\hat{\Omega}_3$, for the $B = 100$ simulations in the three considered scenarios. For each simulation and method, the shrinkage parameter λ is selected according to the modified BIC defined in (13)

4.3 Simulation Results: Clustering Performance

The previous section showcases the ability of our proposals in learning group-wise different sparse structures in the components precision matrices. In doing so, we did not directly assess the obtained clustering, as the mean vectors induced adequately well-separated components. We hereafter evaluate the recovering of the underlying data partition by generating further $B = 100$ samples from the *Different proportion of edges in Ω_k* scenario. Contrarily to the previous study, we fix $\mu_1 = \mu_2 = \mu_3 = \mathbf{0}$. That is, components differ only on the basis of their conditional dependence structures, and thus the final allocation is entirely driven by the estimated precision matrices $\hat{\Omega}_k$. In this case, where both the clustering and the association structures are of interest, a comprehensive model selection strategy is needed to choose K and to properly tune λ . More specifically, coherently with both the model-based clustering literature and the penalized estimation schemes setting, we evaluate each model on a grid of λ values, whose range is computed as in (14), and for different mixture components $K \in \{2, 3, 4, 5\}$. For each model, we select λ and K according to the BIC criterion defined in (13). The adjusted Rand index (ARI, Hubert and Arabie, 1985) is employed for comparing the estimated classification with the true data partition. The results are reported in Fig. 6. While the resulting clustering is satisfactorily close to the true one for all models, including in the penalty specification the group-wise shrinkage matrices \mathbf{P}_k seems to improve the overall performance. Careful analysis of the results demonstrates that the model with common penalty struggles in separating the components with high and medium degree of connectivity. Such a behavior is exacerbated even further when the data dimensionality is equal or even larger than the sample size, as demonstrated in the next section.

4.4 Simulation Results: $p \geq n$ Scenarios

In this section we further explore, via simulation, the applicability of our proposals when the data dimension is equal or even larger than the sample size. In detail, the following data generating process is considered: we sample $n = 100$ observations from a $K = 2$ Gaussian mixture, with precision matrices generated according to a sparse at random Erdős-Rényi graph structure with probabilities of connection equal to 0.1 and 0.8

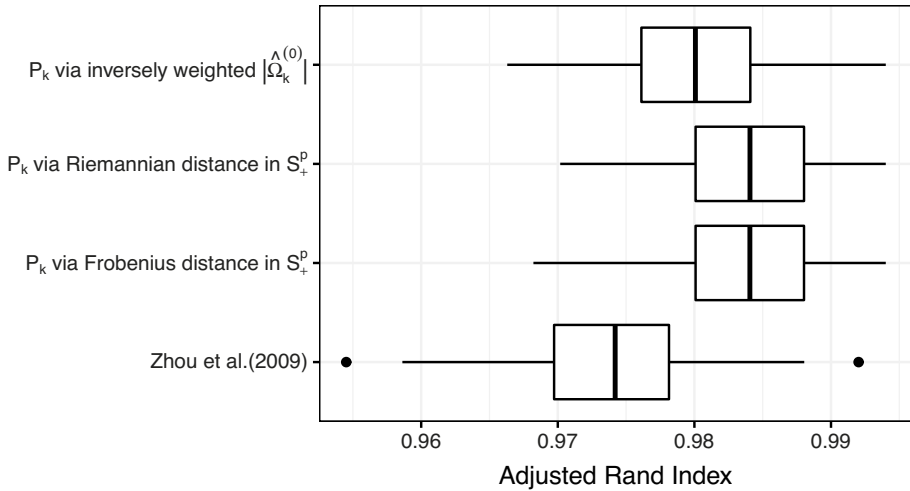


Fig. 6 Boxplots of the ARI for the $B = 100$ simulations of the scenario described in Section 4.3. For each method, the shrinkage parameter λ and the number of components K are selected according to the modified BIC defined in (13)

(the same structure as for the first two components in the *Different proportion of edges in Ω_k* scenario, see Section 4.1). We contemplate two different cases, setting the number of variables equal to $p = 100$ and $p = 200$. For each scenario, we replicate the experiments $B = 100$ times, monitoring the resulting mean F_1 scores and ARI: Figures 7 and 8 report the resulting boxplots for the former and latter metric, respectively.

Similarly to what observed in the right-most panel of Fig. 5, the inclusion of data-driven \mathbf{P}_k assures a better recovery of the true conditional association structure: the mean F_1 scores displayed by our proposals are significantly higher than the one displayed by Zhou et al. (2009). Particularly, employing the strategy described in Section 3.5, coupled with a Frobenius distance, seems to outperform all the other procedures. Interestingly, our proposals showcase fairly good results even in the challenging $n = 100$, $K = 2$, $p = 200$ scenario, demonstrating the effective applicability of such criteria even when the feature space is bigger than the sample size. The same holds only partially true when we monitor the ARI (Fig. 8). As it may be expected, the high dimensionality deteriorates the recovery of the true data partition for all penalized models. Nonetheless, the ARI for methods with group-wise different \mathbf{P}_k is never lower than the one of Zhou et al. (2009). Moreover, the clustering retrieved by the \mathbf{P}_k via *inversely weighted* $|\hat{\Omega}_k^{(0)}|$ procedure is significantly better than all the other alternatives, particularly for the $n = 100$, $K = 2$, $p = 100$ scenario.

4.5 A Note on Computing Times

All the penalized methods considered in the aforementioned simulation studies rely on an iterative algorithm when performing parameters estimation. To this extent, it is of interest to investigate the required elapsed time to fit the models. Table 1 reports the average computing times and associated standard deviations for different methods and simulated scenarios. All the simulated experiments were run on a computer cluster with 32 processors Intel Xeon E5-4610 @2.3GHz.




Fig. 7 Boxplots of the mean F_1 score, averaged over $\hat{\Omega}_1, \hat{\Omega}_2$, for the $B = 100$ simulations in the two $p \geq n$ scenarios. For each simulation and method, the shrinkage parameter λ is selected according to the modified BIC defined in (13)

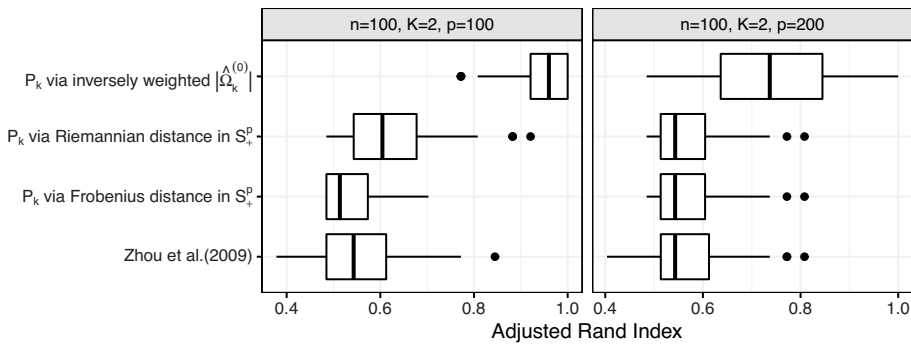


Fig. 8 Boxplots of the ARI for the $B = 100$ simulations in the two $p \geq n$ scenarios. For each simulation and method, the shrinkage parameter λ is selected according to the modified BIC defined in (13)

First off, recall that the calculation of the group-wise different \mathbf{P}_k in our proposals is performed only once prior to start the EM algorithm. Therefore, at least in principle, the extra computational effort required by our methods with respect to the one by Zhou et al. (2009) amounts only to compute the \mathbf{P}_k at the beginning of the iterative procedure. By looking at Table 1 we notice that, irrespective of the methods and quite naturally, a higher dimensionality is associated with a longer computational time. At any rate, in low dimensional settings (*Equal proportion of edges in Ω_k* and *Different proportion of edges in Ω_k*) our methods seems to be even (slightly) faster than Zhou et al. (2009). Conversely, when the dimensionality increases, Zhou et al. (2009) showcase, as expected, the shortest computing times. This is particularly true if compared with the *P_k via Frobenius distance in S_+^p* strategy, for which it appears a larger number of EM iterations are necessary to reach convergence in high-dimensional settings. Notwithstanding we argue that, as extensively demonstrated in the previous sections, the small extra price to pay in terms of computing time tends to be well worth when it comes to detect clusters with diverse degrees of sparsity between the variables.

All in all, considering a group-wise penalty not only improves the estimation of the component precision matrices, but it also enhances the quality of the resulting data partition in all the simulated scenarios, even in the most challenging ones in which the data dimension is equal or

Table 1 Average computing time (in seconds) over $B = 100$ runs for the simulation studies reported in Sections 4.2 (first three columns) and 4.4 (last two columns). The metric refers to the average time required in fitting a model with a single shrinkage term λ . Standard errors are reported in parentheses

	Equal prop of edges in Ω_k	Diff prop of edges in Ω_k	High dim and diff prop of edges in Ω_k	Section 4.4 $p = 100$	Section 4.4 $p = 200$
Zhou et al. (2009)	0.0918 (0.073)	0.0892 (0.106)	0.1461 (0.055)	1.5895 (0.092)	10.5871 (0.495)
P_k via Frobenius distance in S_+^p	0.0352 (0.016)	0.0458 (0.098)	2.7869 (0.723)	3.024 (2.514)	22.2367 (18.184)
P_k via Riemannian distance in S_+^p	0.0325 (0.014)	0.0378 (0.05)	0.4767 (0.532)	1.7297 (0.559)	11.7203 (4.547)
P_k via inversely weighted $ \hat{\Omega}_k^{(0)} $	0.0294 (0.015)	0.0347 (0.045)	0.2065 (0.156)	1.5943 (0.118)	10.7254 (0.915)

even greater than the sample size. The same happens in the illustrative data examples, as it is reported in the next section.

5 Illustrative Datasets

This section presents two illustrative data examples. In this case, contrarily to the synthetic experiments reported in Section 4, the true underlying graph structure is unknown and thus we employ several metrics to assess the model performance under the different definitions of the \mathbf{P}_k matrices. Classification accuracy is evaluated as usual via the ARI, while the number of non-zero estimated parameters in the precision matrices (indicated by $d_{\hat{\Omega}}$) is taken as a proxy of model complexity. The identification of the underlying conditional association structure is evaluated by implementing the following approach. Using the true class labels we compute the data class-specific precision matrices; then, for each method, we measure the median distance (in terms of the Frobenius norm) between the empirical and the estimated sparse precision matrices. In doing so, we make again use of the `matchClasses` routine to match the empirical class-specific precision matrix with its corresponding sparse estimate. In detail, the *Median Frobenius Distance* metric (MFD) is computed as follows:

$$\text{median}_{k \in 1, \dots, K} \left(\|\hat{\Omega}_k - \bar{\Omega}_k\|_F \right) \tag{15}$$

with $\|\cdot\|_F$ denoting the Frobenius norm, while $\hat{\Omega}_k$ is the estimated precision matrix of the k -th component for a given method and $\bar{\Omega}_k$ is the empirical k -th class precision matrix, computed using the true labels.

Similarly to the first simulation study, we fix the number of clusters to the number of classes available in the data; we do so in order to focus the attention on the model selection aspect concerning the recovery of the conditional association structure, rather than on the selection of the number of components in the mixture. Lastly note that data have been standardized before applying any modeling procedure, as it is customarily done with penalized estimation. Nonetheless, we acknowledge that standardization can have an impact on the results and we refer to the recent work by Carter et al. (2021) for a thorough discussion.

5.1 Olive Oil

The first dataset reports the percentage composition of $p = 8$ fatty acids in $n = 572$ units of olive oil. The oil samples come from $K = 9$ Italian regions: the aim is to recover the geographical partition of the oils by means of their lipidic features. This dataset was firstly described in Forina et al. (1983) and it is available in the R package `pgmm` (McNicholas et al., 2019). Results for the considered methods are reported in Table 2.

Together with the different specification of \mathbf{P}_k for sparse estimation, we include in the comparison the standard model-based clustering approach with eigen-decomposed covariance matrices selected using BIC, fitted via the `mclust` software (Scrucca et al., 2016).

For all penalized methods, the selection criterion defined in (13) is used to identify the best λ in a data-driven fashion. In general, penalized models outperform `mclust` VVE (different volume and shape but same clusters orientation) in recovering the true data partition. This might be due to the rigid dependence structure imposed by such model, where the association among variables is forced to be equal across all components. Notice that including a data-dependent specification for \mathbf{P}_k slightly improves the clustering accuracy with respect to the all-one matrix (Zhou et al., 2009). Moreover, the overall model complexity is reduced: the method with common penalty selects a λ that induces a mild sparsity, as a total of $Kp(p + 1)/2 = 324$ parameters would be considered in a fully unconstrained estimation. On the other hand, for our proposals the number of non-zero inverse covariance parameters d_Ω is lower than for the full model, and particularly the \mathbf{P}_k via *inversely weighted* $|\hat{\Omega}_k^{(0)}|$ approach substantially reduces the number of estimated parameters, while showcasing the highest ARI and the lowest Median Frobenius distance. The corresponding graphs for the 9 different clusters are reported in Fig. 9, in which we see that the conditional dependence structure appreciably varies across regions, with our proposal taking advantage of it in the estimation phase.

5.2 Handwritten Digits Recognition

The second dataset, publicly available in the University of California Irvine Machine Learning data repository (<http://archive.ics.uci.edu/ml/datasets/optical+recognition+of+handwritten+digits>), contains $n = 5620$ samples of handwritten digits represented by 64 features. Each variable counts the pixels of an 8×8 grid in which the original images were divided. The aim is to recognize the $K = 10$ digits by means of the penalized procedures introduced in the paper. This clustering problem is more challenging than the one presented in Section 5.1, due to both the higher dimensionality and the narrower separation

Table 2 BIC, ARI, number of estimated parameters and Median Frobenius Distance, as defined in (15), for different model-based clustering methods. Olive oil dataset

	BIC	ARI	d_Ω	MFD
<code>mclust</code> VVE	-4790	0.6586	100	758
Zhou et al.(2009)	-5302	0.6724	320	830
\mathbf{P}_k via inversely weighted $ \hat{\Omega}_k^{(0)} $	-5058	0.7199	242	421
\mathbf{P}_k via Frobenius distance in \mathbb{S}_+^p	-5286	0.6875	312	701
\mathbf{P}_k via Riemannian distance in \mathbb{S}_+^p	-5282	0.6812	314	798

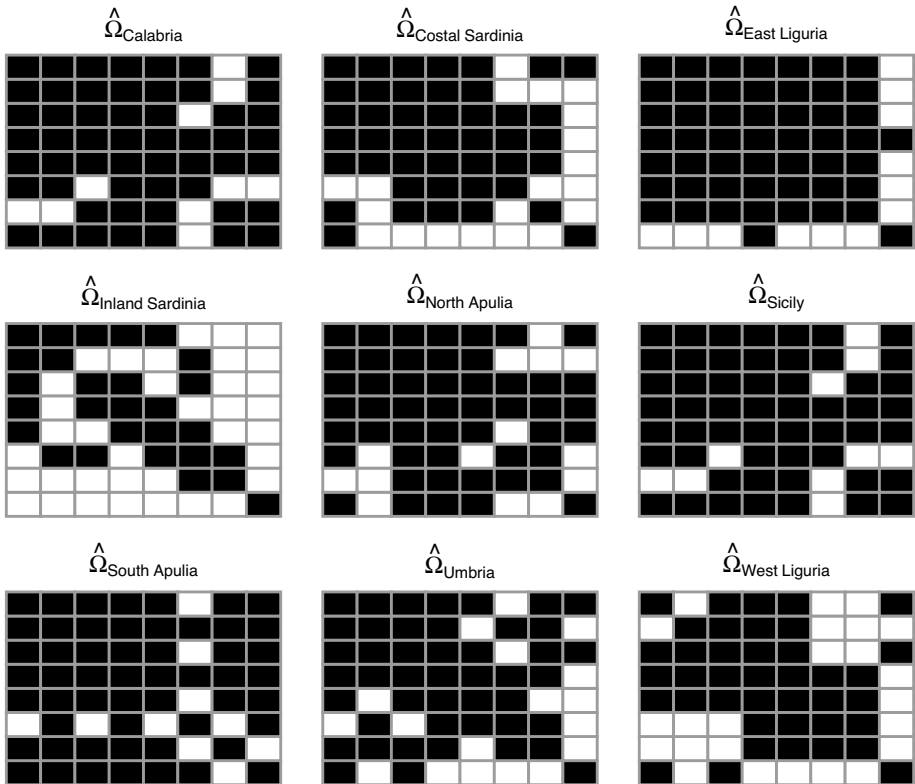


Fig. 9 Estimated graphs in the precision matrices for the \mathbf{P}_k via *inversely weighted* $|\hat{\Omega}_k^{(0)}|$ approach. Black squares denote the presence of an edge between the two variables. Olive oil dataset

between classes (see Fig. 10). Before applying the different clustering methods, we employ a preprocessing step, excluding from the subsequent analysis all predictors with near zero variance. This boils down to essentially remove the left-most and right-most pixels in each image, as being mostly white they contain no separating information. To this task, we use the default routines available in the R package *caret* (Kuhn, 2021). After having eliminated these variables, we are left with $p = 47$ features, which are then considered to perform model-based clustering. Results are reported in Table 3. The parsimonious structure selected by *mclust* forces the precision matrices to be all equal across groups. This rigid constraint undermines the classification accuracy and the uncovering of the conditional dependence structure, resulting in the worst ARI and Median Frobenius distance metrics. Conversely, the penalized methods are able to shrink the estimates in a group-wise manner. This is especially true in our proposals for which, even though the resulting classification accuracy is not dramatically affected, the Median Frobenius distance is always smaller than Zhou et al. (2009). In Fig. 11 we report the estimated graphs in the precision matrices for the \mathbf{P}_k via *Riemannian distance in S_+^p* approach which results in the highest ARI. Lastly note how the numbers of estimated edges appreciably differ between digits, an aspect that is implicitly taken into account in our data-driven specification of the \mathbf{P}_k matrices.

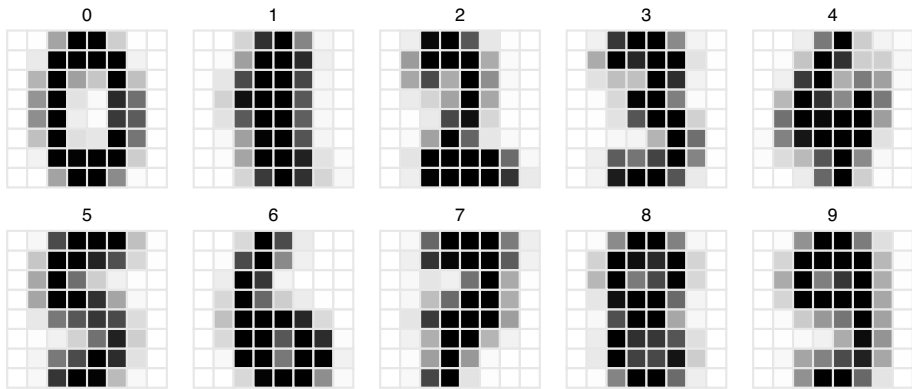


Fig. 10 Image representation of the means of digits 0 to 9. Handwritten digits dataset

Table 3 BIC, ARI, number of estimated parameters and Median Frobenius Distance, as defined in (15), for different model-based clustering methods. Handwritten digits dataset

	BIC	ARI	d_{Ω}	MFD
mclust EEE	-521220	0.6489	1128	172
Zhou et al.(2009)	-388862	0.6837	4914	148
\mathbf{P}_k via inversely weighted $ \hat{\Omega}_k^{(0)} $	-368604	0.6820	3436	104
\mathbf{P}_k via Frobenius distance in \mathbb{S}_+^p	-391359	0.6827	6066	146
\mathbf{P}_k via Riemannian distance in \mathbb{S}_+^p	-388902	0.6841	5206	147

6 Discussion

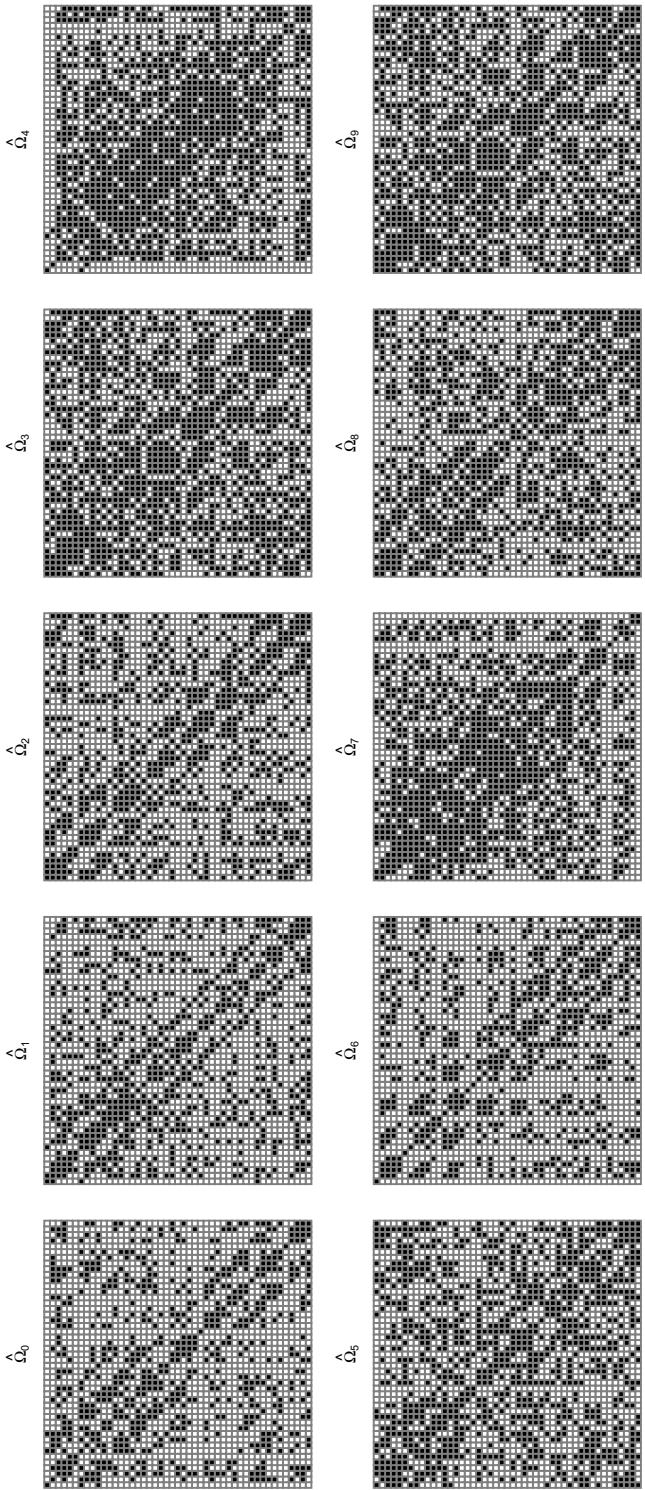
The present paper has highlighted the limitations of imposing a single penalty when performing sparse estimation of component precision matrices in a multi-class setting. We have argued that methods enforcing similarities in the graphical models across groups may not be adequate for classification, since they have detrimental effects when it comes to groups discrimination, in particular in the case of clustering. Thus, we have focused our attention on the penalized model-based method with sparse precision matrices framework of Zhou et al. (2009), where class-specific differences are preserved. Nonetheless, this methodology does not account for situations in which a component displays under or over-connectivity with respect to the remaining ones. To this extent, we have proposed some procedures to incorporate group-specific differences in the estimation, enforcing a carefully initialized solution to drive the algorithm in under or over penalizing specific components. Numerical illustrations and analyses on real data have confirmed the validity of our proposals. By means of our solutions we have achieved both group-wise flexibility in the precision matrices reconstruction and we have mitigated the impact the common shrinkage factor has in the overall sparse estimation.

The present paper opens up a quite natural direction for future research: the penalized approach could be adapted to estimate sparse covariance matrices, rather than precision matrices. In the Gaussian case, a missing edge between two nodes in the *Gaussian covariance graph model* corresponds to two variables being marginally independent, and

the so-called *covariance graph* (Chaudhuri et al., 2007) allows to represent the pattern of zeroes in the covariance matrices. A related methodology based on cluster-specific penalties has been recently introduced by Fop et al. (2019), unfortunately, such an approach relies on a time-consuming graph structure search, making it less attractive in high-dimensional problems. On the other hand, the definition of a penalized likelihood that incorporates a covariance graphical lasso term (Bien and Tibshirani, 2011; Wang, 2014) can be effectively employed in these scenarios: model definitions are being explored and they will be the object of future work.

The framework proposed here has also interesting connections with the notion of global-local shrinkage developed in the Bayesian literature for sparsity-inducing priors (Bhattacharya et al., 2015; Polson & Scott, 2010). The general formulation of these priors is based on a normal scale mixture representation, where the mean is zero and the variance is expressed as the product of two nonnegative parameters: one scaling parameter pulls the global shrinkage towards zero, while the other allows for modifications in the amount of shrinkage (Bhattacharya et al., 2015). Global-local shrinkage priors for Gaussian graphical models have been employed in Leday et al. (2017) for gene network inference in the case of a homogeneous population. The authors develop a simultaneous equations modeling approach for graph inference, where the regression parameters are given Gaussian scale mixture priors for local and global shrinkage, which allows borrowing of information among the regressions and encourages the posterior expectation of the corresponding entries of the precision matrix to be shrunk towards zero. As pointed out by Leday et al. (2017), compared to Meinshausen et al. (2006), a disadvantage of these priors is that they do not automatically perform variable selection, hence the graph structure needs to be recovered by thresholding of the posterior means of the regression coefficients. An alternative use of global-local shrinkage priors is in the Bayesian graphical lasso of Wang (2012). Here, it is shown that the graphical lasso estimator is the maximum a posteriori of a Bayesian hierarchical model where the entries of the precision matrix have exponential and double exponential prior distributions, which can be represented as a scale mixture of normals. We note that the graphical modeling Bayesian frameworks of Leday et al. (2017) and Wang (2012) are developed for the case of a homogeneous sample, and that global-local shrinkage is intended only in terms of joint shrinkage of all the entries of Ω , allowing only for variable and scale-specific adjustments. In contrast, in our proposed approach, global-local shrinkage would be intended in terms of joint shrinkage of the component precision matrices Ω_k towards a common level of sparsity, with cluster related adaptations. The penalty term λ could be considered the global shrinkage factor, which equally shrinks the entries of the precision matrices across the mixture components, while the weighting matrices \mathbf{P}_k allow for local cluster-specific adjustments. Following the literature on penalized model-based clustering, we devised our proposal under a penalized likelihood framework, which has computational advantages especially in high-dimensional scenarios. However, with the purpose of sparse Bayesian model-based clustering, carefully defined prior distributions could be defined for global-local shrinkage across mixture components and within clusters; these considerations open a path for future developments of our proposal in a Bayesian context and are currently under exploration.

As a last worthy note, even if the proposed procedure is applicable in a general setting, we believe that the definition of group-specific penalties should never leave aside prior information and subject-matter knowledge whenever available, as their incorporation in the methodology can be strongly beneficial for the analysis.



◀ **Fig. 11** Estimated graphs in the precision matrices for the \mathbf{P}_k via Riemannian distance in S_+^p approach. Black squares denote the presence of an edge between the two variables. Handwritten digits dataset

Funding Open access funding provided by Libera Università di Bolzano within the CRUI-CARE Agreement.

Code Availability The R code implementing the procedure is available at <https://github.com/AndreaCappozzo/parsemix>.

Declarations

Conflict of Interest The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Banerjee, O., Ghaoui, L.E., & d'Aspremont, A. (2008). Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *Journal of Machine Learning Research*, 9, 485–516.
- Banfield, J.D., & Raftery, A.E. (1993). Model-based gaussian and non-gaussian clustering. *Biometrics*, 49(3), 803–821.
- Bhattacharya, A., Pati, D., Pillai, N.S., & Dunson, D.B. (2015). Dirichlet–Laplace priors for optimal shrinkage. *Journal of the American Statistical Association*, 110(512), 1479–1490.
- Bickel, P.J., & Levina, E. (2008). Covariance regularization by thresholding. *The Annals of Statistics*, 36(6), 2577–2604.
- Bien, J., & Tibshirani, R.J. (2011). Sparse estimation of a covariance matrix. *Biometrika*, 98(4), 807–820.
- Biernacki, C., & Lourme, A. (2014). Stable and visualizable gaussian parsimonious clustering models. *Statistics and Computing*, 24(6), 953–969.
- Bouveyron, C., & Brunet, C. (2012). Simultaneous model-based clustering and visualization in the fisher discriminative subspace. *Statistics and Computing*, 22(1), 301–324.
- Bouveyron, C., & Brunet-Saumard, C. (2014). Model-based clustering of high-dimensional data: A review. *Computational Statistics & Data Analysis*, 71, 52–78.
- Bouveyron, C., Celeux, G., Murphy, T.B., & Raftery, A.E. (2019). *Model-based clustering and classification for data science: with applications in R*. Cambridge: Cambridge University Press.
- Bouveyron, C., Girard, S., & Schmid, C. (2007). High-dimensional data clustering. *Computational Statistics & Data Analysis*, 52(1), 502–519.
- Carter, J.S., Rossell, D., & Smith, J.Q. (2021). Partial correlation graphical lasso. arXiv:2104.10099.
- Casa, A., Scrucca, L., & Menardi, G. (2021). Better than the best? Answers via model ensemble in density-based clustering. *Advances in Data Analysis and Classification*, 15(3), 599–623.
- Celeux, G., & Govaert, G. (1995). Gaussian parsimonious clustering models. *Pattern Recognition*, 28(5), 781–793.
- Chaudhuri, S., Drton, M., & Richardson, T.S. (2007). Estimation of a covariance matrix with zeros. *Biometrika*, 94(1), 199–216.
- Danaher, P., Wang, P., & Witten, D.M. (2014). The joint graphical lasso for inverse covariance estimation across multiple classes. *Journal of the Royal Statistical Society: Series B (Methodological)*, 76(2), 373.
- Dempster, A.P., Laird, N.M., & Rubin, D.B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1), 1–22.
- Dryden, I.L., Koloydenko, A., & Zhou, D. (2009). Non-Euclidean statistics for covariance matrices, with applications to diffusion tensor imaging. *The Annals of Applied Statistics*, 3(3), 1102–1123.

- Erdős, P., & Rényi, A. (1960). On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5(1), 17–60.
- Fan, J., Feng, Y., & Wu, Y. (2009). Network exploration via the adaptive lasso and scad penalties. *The Annals of Applied Statistics*, 3(2), 521.
- Fop, M. (2020). covglasso: Sparse covariance matrix estimation. R package version 2.0. <https://CRAN.R-project.org/package=covglasso>
- Fop, M., & Murphy, T.B. (2018). Variable selection methods for model-based clustering. *Statistics Surveys*, 12, 18–65.
- Fop, M., Murphy, T.B., & Scrucca, L. (2019). Model-based clustering with sparse covariance matrices. *Statistics and Computing*, 29(4), 791–819.
- Forina, M., Armanino, C., Lanteri, S., & Tiscornia, E. (1983). Classification of olive oils from their fatty acid composition. In *Food research and data analysis: proceedings from the IUFOST Symposium September 20-23 1982, Oslo, Norway* edited by H. Martens and H. Russwurm, Jr. London: Applied Science Publishers.
- Fraley, C., & Raftery, A.E. (2002). Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association*, 97(458), 611–631.
- Friedman, J., Hastie, T., & Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3), 432–441.
- Gao, X., & Massam, H. (2015). Estimation of symmetry-constrained gaussian graphical models: application to clustered dense networks. *Journal of Computational and Graphical Statistics*, 24(4), 909–929.
- Gelman, A., & Vehtari, A. (2021). What are the most important statistical ideas of the past 50 years? *Journal of the American Statistical Association*, 116 (536), 2087–2097.
- Guo, J., Levina, E., Michailidis, G., & Zhu, J. (2011). Joint estimation of multiple graphical models. *Biometrika*, 98(1), 1–15.
- Hastie, T., Tibshirani, R., & Wainwright, M. (2015). *Statistical learning with sparsity: the lasso and generalizations*. Boca Raton: CRC Press.
- Højsgaard, S., & Lauritzen, S.L. (2008). Graphical gaussian models with edge and vertex symmetries. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(5), 1005–1027.
- Hubert, L., & Arabie, P. (1985). Comparing partitions. *Journal of Classification*, 2, 193–218.
- Hull, J.V., Dokovna, L.B., Jacokes, Z.J., Torgerson, C.M., Irimia, A., & Van Horn, J.D. (2017). Resting-state functional connectivity in autism spectrum disorders: A review. *Frontiers in Psychiatry*, 7, 205.
- Kuhn, M. (2021). caret: Classification and Regression Training. R package version 6.0-86. <https://CRAN.R-project.org/package=caret>
- Leday, G.G., de Gunst, M.C., Kpogbezan, G.B., van der Vaart, A.W., van Wieringen, W.N., & van de Wiel, M.A. (2017). Gene network reconstruction using global-local shrinkage priors. *The Annals of Applied Statistics*, 11 (1), 41–68.
- Li, Q., Sun, X., Wang, N., & Gao, X. (2021). Penalized composite likelihood for colored graphical gaussian models. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 14(4), 366–378.
- Lian, H. (2011). Shrinkage tuning parameter selection in precision matrices estimation. *Journal of Statistical Planning and Inference*, 141(8), 2839–2848.
- Lin, T.I. (2009). Maximum likelihood estimation for multivariate skew normal mixture models. *Journal of Multivariate Analysis*, 100(2), 257–265.
- Lin, T.I. (2010). Robust mixture modeling using multivariate skew t distributions. *Statistics and Computing*, 20(3), 343–356.
- Lyu, Y., Xue, L., Zhang, F., Koch, H., Saba, L., Kechris, K., & Li, Q. (2018). Condition-adaptive fused graphical lasso (CFGL): An adaptive procedure for inferring condition-specific gene co-expression network. *PLoS computational Biology*, 14(9), e1006436.
- Maugis, C., Celeux, G., & Martin-Magniette, M.-L. (2009a). Variable selection for clustering with Gaussian mixture models. *Biometrics*, 65(3), 701–709.
- Maugis, C., Celeux, G., & Martin-Magniette, M.-L. (2009b). Variable selection in model-based clustering: A general variable role modeling. *Computational Statistics & Data Analysis*, 53(11), 3872–3882.
- McLachlan, G.J., & Peel, D. (1998). Robust cluster analysis via mixtures of multivariate t-distributions. In *Joint IAPR international workshops on statistical techniques in pattern recognition (SPR) and structural and syntactic pattern recognition (SSPR)* (pp. 658–666). Springer.
- McLachlan, G.J., Peel, D., & Bean, R. (2003). Modelling high-dimensional data by mixtures of factor analyzers. *Computational Statistics & Data Analysis*, 41(3-4), 379–388.
- McNicholas, P.D. (2016). Model-based clustering. *Journal of Classification*, 33(3), 331–373.
- McNicholas, P.D., ElSherbiny, A., McDaid, A.F., & Murphy, T.B. (2019). pgmm: Parsimonious gaussian mixture models. R package version 1.2.4. <https://CRAN.R-project.org/package=pgmm>

- McNicholas, P. D., & Murphy, T. B. (2008). Parsimonious gaussian mixture models. *Statistics and Computing*, 18(3), 285–296.
- Meinshausen, N., Bühlmann, P., & et al. (2006). High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, 34(3), 1436–1462.
- Meyer, D., Dimitriadou, E., Hornik, K., Weingessel, A., & Leisch, F. (2020). e1071: Misc Functions of the Department of Statistics, Probability Theory Group (Formerly: E1071), TU Wien. R package version 1.7-4.
- Mohan, K., London, P., Fazel, M., Witten, D., & Lee, S. (2014). Node-based learning of multiple gaussian graphical models. *Journal of Machine Learning Research*, 15(1), 445–488.
- Pan, W., & Shen, X. (2007). Penalized model-based clustering with application to variable selection. *Journal of Machine Learning Research*, 8, 1145–1164.
- Polson, N.G., & Scott, J.G. (2010). Shrink globally, act locally: Sparse bayesian regularization and prediction. *Bayesian Statistics*, 9(501-538), 105.
- Pourahmadi, M. (2013). *High-dimensional covariance estimation wiley series in probability and statistics*. New York: Wiley.
- R Core Team. (2022). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Raftery, A.E., & Dean, N. (2006). Variable selection for model-based clustering. *Journal of the American Statistical Association*, 101(473), 168–178.
- Russell, N., Murphy, T.B., & Raftery, A.E. (2015). Bayesian model averaging in model-based clustering and density estimation. arXiv:1506.09035.
- Scheinberg, K., Ma, S., and Goldfarb, D. (2010). Sparse inverse covariance selection via alternating linearization methods. In Proceedings of the 23rd International Conference on Neural Information Processing Systems.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2), 461–464.
- Scrucca, L., Fop, M., Murphy, T.B., & Raftery, A.E. (2016). mclust 5: Clustering, classification and density estimation using gaussian finite mixture models. *The R Journal*, 8(1), 289–317.
- Scrucca, L., & Raftery, A.E. (2015). Improved initialisation of model-based clustering using Gaussian hierarchical partitions. *Advances in Data Analysis and Classification*, 9(4), 447–460.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267–288.
- Vrbik, I., & McNicholas, P.D. (2014). Parsimonious skew mixture models for model-based clustering and classification. *Computational Statistics & Data Analysis*, 71, 196–210.
- Wang, H. (2012). Bayesian graphical lasso models and efficient posterior computation. *Bayesian Analysis*, 7(4), 867–886.
- Wang, H. (2014). Coordinate descent algorithm for covariance graphical lasso. *Statistics and Computing*, 24(4), 521–529.
- Wei, Y., & McNicholas, P.D. (2015). Mixture model averaging for clustering. *Advances in Data Analysis and Classification*, 9(2), 197–217.
- Whittaker, J. (1990). *Graphical models in applied multivariate statistics*. New York: Wiley.
- Witten, D.M., Friedman, J.H., & Simon, N. (2011). New insights and faster computations for the graphical lasso. *Journal of Computational and Graphical Statistics*, 20(4), 892–900.
- Xie, B., Pan, W., & Shen, X. (2008). Penalized model-based clustering with cluster-specific diagonal covariance matrices and grouped variables. *Electronic Journal of Statistics*, 2, 168.
- Zhao, T., Liu, H., Roeder, K., Lafferty, J., & Wasserman, L. (2012). The huge package for high-dimensional undirected graph estimation in R. *Journal of Machine Learning Research*, 13(1), 1059–1062.
- Zhou, H., Pan, W., & Shen, X. (2009). Penalized model-based clustering with unconstrained covariance matrices. *Electronic Journal of Statistics*, 3, 1473–1496.
- Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American Statistical Association*, 101(476), 1418–1429.
- Zou, H., Hastie, T., & Tibshirani, R. (2007). On the “degrees of freedom” of the lasso. *The Annals of Statistics*, 35(5), 2173–2192.