



On the application of Gaussian graphical models to paired data problems

Saverio Ranciati¹ · Alberto Roverato²

Received: 14 February 2024 / Accepted: 25 September 2024 / Published online: 30 October 2024
© The Author(s) 2024

Abstract

Gaussian graphical models are nowadays commonly applied to the comparison of groups sharing the same variables, by jointly learning their independence structures. We consider the case where there are exactly two dependent groups and the association structure is represented by a family of coloured Gaussian graphical models suited to deal with paired data problems. To learn the two dependent graphs, together with their across-graph association structure, we implement a fused graphical lasso penalty. We carry out a comprehensive analysis of this approach, with special attention to the role played by some relevant submodel classes. In this way, we provide a broad set of tools for the application of Gaussian graphical models to paired data problems. These include results useful for the specification of penalty values in order to obtain a path of lasso solutions and an ADMM algorithm that solves the fused graphical lasso optimization problem. Finally, we carry out a simulation study to compare our method with the traditional graphical lasso, and present an application of our method to cancer genomics where it is of interest to compare cancer cells with a control sample from histologically normal tissues adjacent to the tumor. All the methods described in this article are implemented in the R package `pdglasso` available at <https://github.com/savranciati/pdglasso>.

Keywords ADMM algorithm · Coloured Gaussian graphical model · Conditional independence · Fused lasso penalty · Graphical lasso · Symmetry

1 Introduction

Graphical models are powerful tools for expressing the relationships between variables. In Gaussian graphical models (GGMs) the dependency structure is obtained by associating an undirected graph to the concentration matrix, that is the inverse of the covariance matrix. The graph has one vertex for every variable and every missing edge implies that the corresponding entry of the concentration matrix is equal to zero; see Lauritzen (1996).

In recent years, there has been a great deal of interest in the joint learning of multiple networks, where the obser-

vations come from two or more groups sharing the same variables. For instance, Danaher et al. (2014) considered the case of gene expression measurements collected from the cancer tissue of a sample of patients and from the normal tissue of a control sample. Hence, the association structure of each group is represented by a network and methods for the joint learning have been developed to deal with the fact that networks are expected to share similar patterns while retaining individual features. In this framework, the literature has mostly focused on the case where the groups are independent so that every network is a distinct unit, disconnected from the other networks; see Tsai et al. (2022) for a review. Thus, specific methods that deal with the across network association are required in the case where the groups are dependent. Xie et al. (2016) considered the case of gene expression data obtained from multiple tissues from the same individual and modelled the cross-graph dependence between groups by means of a latent vector representing systemic variation manifesting simultaneously in all groups; see also Zhang et al. (2022). Roverato and Nguyen (2022, 2024) focused on paired data problems where there are exactly two

✉ Saverio Ranciati
saverio.ranciati2@unibo.it

Alberto Roverato
alberto.roverato@unipd.it

¹ Department of Statistical Sciences, University of Bologna, Bologna, Italy

² Department of Statistical Sciences, University of Padova, Padova, Italy

dependent groups. Paired data are commonly originated from experimental designs where each subject is measured twice under two different conditions, or time points, as well as in matched pairs designs. For instance, in cancer genomics it is common practice to take control samples from histologically normal tissues adjacent to the tumor (Aran et al. 2017). Roverato and Nguyen (2022) approached the problem by considering a family of coloured GGMs (Højsgaard and Lauritzen 2008), named pdRCON models, in which similarities between groups are represented by symmetries which can involve both the graph structure and the concentration values, in the form of equality constraints. One of the appealing features of this approach is that the resulting model has a graph for each of the two groups with the cross-graph dependence explicitly represented by the edges joining one group with the other, which may themselves present symmetries; we refer to Sect. 2 for details. In a related framework, Ranciati et al. (2021) considered the task of learning brain networks from fMRI data, and approached the problem by means of a fused graphical lasso procedure, named the symmetric graphical lasso, designed with the specific aim to identify symmetries between the left and the right hemisphere.

In this paper, we introduce the *graphical lasso for paired data* (pdglasso) that extends the symmetric graphical lasso to deal with the wider family of pdRCON models, and then provide a set of methods which are meant to establish an extensive collection of tools for the immediate application of GGMs to the analysis of paired data. We provide an alternating directions method of multiplier (ADMM) algorithm that solves the pdglasso optimization problem and, furthermore, the R package (R Core Team 2023) `pdglasso` that implements the ADMM algorithm as well as a function for the computation of maximum likelihood estimates, and other utility functions to deal efficiently with the objects resulting from the analysis. For details about the package we refer to Section ‘Data availability’ of this paper. We analyse the role played by relevant submodel classes characterized by a fully symmetric structure. More specifically, full symmetry can be a property of the overall structure, or it may be confined to some specific components, such as the diagonal entries of the concentration matrix, the inside-group structure or the across-group structure. We provide results on the values of the penalty terms required to obtain either a diagonal, a block-diagonal or a fully symmetric solution, which are useful in the specification of the grid of penalty values required to obtain a path of pdglasso solutions. We then carry out a simulation study that shows that our pdglasso method has comparable performances to glasso when there are no symmetries, and improves on the latter when symmetries are present. Furthermore, we present an application to the analysis of gene expression data in cancer genomics.

The rest of this paper is organized as follows. Section 2 provides the background on coloured GGMs and on graph-

ical models for paired data, as required for this paper. The pdglasso problem is introduced in Sect. 3, whereas Sect. 4 deals with the practical application of the method, including the maximum theoretical values of the penalty parameters. Section 5 is devoted to some relevant families of submodels. The ADMM algorithm for the optimization of the penalized likelihood can be found in Sect. 6. The simulation study and the application to gene expression data are described in Sect. 7 and 8, respectively. Finally, Sect. 9 contains a brief discussion. Proofs are deferred to Appendix B, whereas the remaining three appendices contain a comprehensive presentation of fully symmetric models, some illustrative examples, and some additional material concerning the application to gene expression data, respectively.

2 Coloured Gaussian graphical models for paired data

Let Y_V be a continuous random vector indexed by a finite set $V = \{1, \dots, p\}$. We denote by $\Sigma = \{\sigma_{ij}\}_{i,j \in V}$ and $\Sigma^{-1} = \Theta = \{\theta_{ij}\}_{i,j \in V}$ the covariance and the concentration matrix of Y_V , respectively. Both Σ and Θ belong to the set \mathcal{S}_p^+ of (symmetric) $p \times p$ positive definite matrices and we assume, without loss of generality, that Y_V has zero mean vector. An undirected graph with vertex set V is a pair $\mathcal{G} = (V, E)$ where E is an edge set that is a set of pairs of distinct vertices and, with a slight abuse of notation, we will sometimes write $\{i, j\} \in \mathcal{G}$, in place of $\{i, j\} \in E$, to mean that the edge $\{i, j\}$ belongs to E . We say that the concentration matrix Θ is adapted to a graph $\mathcal{G} = (V, E)$ if every missing edge of \mathcal{G} corresponds to a zero entry in Θ ; formally, $\{i, j\} \notin E$ implies $\theta_{ij} = 0$ for every $i, j \in V$ with $i \neq j$. A Gaussian graphical model (GGM) with graph \mathcal{G} is the family of multivariate normal distributions for Y_V whose concentration matrix is adapted to \mathcal{G} . These models are also known with the name of covariance selection models or concentration graph models (Lauritzen 1996).

In paired data problems, the set V is naturally partitioned into a *Left* and a *Right* block, $V = L \cup R$ with $|L| = |R| = q = p/2$ and every variable in Y_L has a corresponding variable in Y_R . We set $i' = i + q$ for every $i \in L$ and, without loss of generality, we index the variables so that Y_i corresponds to $Y_{i'}$ for every $i \in L$. In this way, it also holds that $L = \{1, \dots, q\}$ and $R = \{q + 1, \dots, p\}$. Accordingly, the concentration matrix Θ can be naturally partitioned into four blocks,

$$\Theta = \begin{pmatrix} \Theta_{LL} & \Theta_{LR} \\ \Theta_{RL} & \Theta_{RR} \end{pmatrix}. \quad (1)$$

The subgraph of \mathcal{G} induced by a subset $A \subseteq V$ is denoted by $\mathcal{G}_A = (A, E_A)$, where $\{i, j\} \in E_A$ if and only if both

$i, j \in A$ and $\{i, j\} \in E$. If Θ is adapted to \mathcal{G} then it holds that Θ_{LL} and Θ_{RR} are adapted to \mathcal{G}_L and \mathcal{G}_R , respectively. Thus, \mathcal{G}_L and \mathcal{G}_R are the group specific graphs and the interest is for similarities involving the independence structures of the two groups, that we call *structural symmetries*. We distinguish between two types of structural symmetries. When for a pair $i, j \in L$, with $i \neq j$ the edges $\{i, j\}$ and $\{i', j'\}$ are either both present or both missing in \mathcal{G}_L and \mathcal{G}_R , respectively, we say that there is an *inside-block* structural symmetry. On the other hand, Y_L and Y_R are not expected to be independent, and symmetries can also appear in the cross-group association structure, thereby involving the edges connecting the vertices in L with the vertices in R . Hence, we say that an *across-block* structural symmetry is present if the edges $\{i, j'\}$ and $\{i', j\}$ are either both present or both missing in \mathcal{G} .

In a GGM every entry of the concentration matrix can be naturally associated with either a vertex or an edge of the graph and, thus, structural symmetries in the graph imply that the corresponding entries in the concentration matrix are both either zero or non-zero. Hence, a stronger class of models could be defined that specifies equality constraints between such entries. Equality constraints encode stronger similarities and result in more parsimonious models. From this viewpoint, every structural symmetry due to a pair of missing edges also corresponds to a *parametric symmetry* because if, for example, both $\{i, j\} \notin \mathcal{G}_L$ and $\{i', j'\} \notin \mathcal{G}_R$, then also the associated parameters have the same value, $\theta_{ij} = \theta_{i'j'} = 0$. This idea can be extended to symmetric pairs of non-missing edges, and we say that there is an *inside-block parametric symmetry* if for a pair $i, j \in L$, with $i \neq j$, it holds that $\theta_{ij} = \theta_{i'j'}$. On the other hand, *across-block parametric symmetries* are also possible in the case where, for some $i, j \in L$, it holds that $\theta_{ij'} = \theta_{i'j}$, with $i \neq j$. Finally, *parametric symmetries* can also be satisfied by the diagonal entries of Θ and we say that there is a *vertex parametric symmetry* if for $i \in L$ it holds that $\theta_{ii} = \theta_{i'i'}$. Appendix C gives some examples of pdRCON models with a detailed description of the different types of symmetries they may represent.

Gaussian graphical models with additional restrictions on the parameter space were introduced in the seminal paper by Højsgaard and Lauritzen (2008) with the name of coloured GGMs. The family of coloured GGMs characterized by equality Restrictions on CONcentration values are known as RCON models, whereas the subfamily of RCON models comprising the equality constraints suited for paired data problems, as described above, was introduced by Roverato and Nguyen (2022) with the name of RCON models for paired data (pdRCON models); see also Ranciaci et al. (2021) for an application of RCON models for paired data that only involves inside-block and vertex symmetries.

Coloured GGMs are typically represented by coloured graphs where vertices and edges depicted in black identify unconstrained parameters, whereas other colours are used to

identify subsets of parameters which are constrained to having the same value. In this way, a different colour is required for every distinct equality constraint. We note, however, that this representation is redundant in the case of pdRCON models, because equality constraints may only involve specific pairs of parameters. Indeed, in order to identify the model from the graph it is sufficient to be able to distinguish whether a parameter is constrained or not. It follows that a pdRCON model can be unambiguously identified also if all the non-black vertices and edges are depicted by the same colour. This makes the graphical representation more readable for large graphs and, thus, here we will follow this rule and simply refer to coloured vertices and edges in contrast to uncoloured (i.e. black) ones. In Appendix C, in order to make graphs readable also in black and white printing, we distinguish between uncoloured and coloured edges by using thin and thick lines, respectively.

3 The graphical lasso for paired data

For a sample $y_V^{(1)}, \dots, y_V^{(n)}$ of i.i.d. observations of $Y_V \sim N(0, \Sigma)$, the maximum likelihood estimator of Θ is the value that maximises the log-likelihood function

$$l(\Theta) = \log \det(\Theta) - \text{tr}(S\Theta) \tag{2}$$

over \mathcal{S}_p^+ where $S = \{s_{ij}\}_{i,j \in V}$ is the matrix $S = n^{-1} \sum_{i=1}^n y_V^{(i)} (y_V^{(i)})^\top$. In GGMs the interest is for the zero pattern of Θ and Yuan and Lin (2007) proposed an estimator of Θ obtained from the minimization of the penalized log-likelihood

$$\mathcal{L}_{\lambda_1}(\Theta) = -l(\Theta) + \mathcal{P}_{\lambda_1}(\Theta) \tag{3}$$

with $\mathcal{P}_{\lambda_1}(\Theta) = \lambda_1 \|\Theta\|_1$, where $\|\cdot\|_1$ denotes the ℓ_1 -norm, that is the sum of the absolute values of the entries of the matrix, and λ_1 is a nonnegative regularization parameter. The minimization of (3) over \mathcal{S}_p^+ is known as the *graphical lasso* problem and a number of algorithms have been proposed for its solution; see, among others, Friedman et al. (2008) and Boyd et al. (2011, Section 6.5). Unlike the maximum likelihood estimator, for $\lambda_1 > 0$ the solution to the graphical lasso exists also when $p > n$, in which case S is singular. The high popularity of the graphical lasso is due to the fact that it simultaneously performs estimation and model selection within the family of GGMs. Indeed, due to the geometry of the ℓ_1 -penalty, some of the off-diagonal entries of the concentration matrix are shrunk to exactly zero, with a non-decreasing level of sparsity as λ_1 increases.

In order to simultaneously perform estimation and model selection within the family of pdRCON models we introduce

an additional penalty term

$$Q_{\lambda_2}(\Theta) = \lambda_2 (\|\Theta_{LL} - \Theta_{RR}\|_1 + \|\Theta_{LR} - \Theta_{RL}\|_1) \quad (4)$$

to (3) so as to obtain,

$$\mathcal{L}_{\lambda_1, \lambda_2}(\Theta) = -l(\Theta) + \mathcal{P}_{\lambda_1}(\Theta) + Q_{\lambda_2}(\Theta), \quad (5)$$

and estimate Θ by the *paired data graphical lasso* (pdglasso) estimator $\hat{\Theta} = \arg \min_{\Theta} \mathcal{L}_{\lambda_1, \lambda_2}(\Theta)$, where minimization is taken over \mathcal{S}_p^+ .

Note that $Q_{\lambda_2}(\Theta)$ in (4) is a fused type lasso penalty (Tibshirani et al. 2005; Hoefling 2010) with regularization parameter $\lambda_2 \geq 0$. For $\lambda_2 = 0$ the pdglasso estimator coincides with the graphical lasso estimator whereas positive values of λ_2 encourage parametric symmetries. More specifically, the term $\lambda_2 \|\Theta_{LL} - \Theta_{RR}\|_1 = \lambda_2 \sum_{i \in L} |\theta_{ii} - \theta_{i'i'}| + \lambda_2 \sum_{i, j \in L; i \neq j} |\theta_{ij} - \theta_{i'j'}|$ encourages both vertex parametric symmetries and inside-block parametric symmetries, whereas $\lambda_2 \|\Theta_{LR} - \Theta_{RL}\|_1 = \lambda_2 \sum_{i, j \in L} |\theta_{ij'} - \theta_{i'j}|$ encourages across-block parametric symmetries; we remark that, in the latter sum, $|\theta_{ij'} - \theta_{i'j}| = 0$ whenever $i = j$ because Θ is a symmetric matrix. Hence, like the graphical lasso, the pdglasso is a sparse estimate of Θ when λ_1 is large, and has many parametric symmetries when λ_2 is large.

The idea of using a fused penalty to induce equality constraints between concentration parameters is not new, and it was first introduced by Danaher et al. (2014) in the context of joint learning of multiple GGMs for independent samples. Subsequently, Ranciati et al. (2021) considered the case of two possibly non independent groups and the pdglasso problem in (5) extends the work of Ranciati et al. (2021) also including cross-group symmetries.

4 Selection of penalty parameters and application issues

This section deals with the practical application of the method. Specifically, we provide the maximum theoretical values of the penalty parameters, give details on the computation of the eBIC criterion we use for the selection of the model and, finally, discuss the role played by the unit of measurement of the variables.

The application of the pdglasso requires the initial definition of a grid of penalty parameter values to obtain a path of pdglasso solutions. To this aim, it is useful to identify the values of λ_1 and λ_2 that return a diagonal and a fully symmetric solution, respectively, to be set as maximum values of the grid. Formally, we say that a pdRCON model is fully symmetric if and only if both $\Theta_{LL} = \Theta_{RR}$ and $\Theta_{LR} = \Theta_{RL}$. We remark that the equality $\Theta_{LL} = \Theta_{RR}$ implies (i) full vertex parametric symmetry in the sense that $\theta_{ii} = \theta_{i'i'}$ for every

$i \in L$ and (ii) both structural and parametric inside-block symmetry because $\theta_{ij} = \theta_{i'j'}$ for every $i, j \in L$ with $i \neq j$ also implies that $\theta_{ij} = 0$ if and only if $\theta_{i'j'} = 0$. Similarly, the equality $\Theta_{LR} = \Theta_{RL}$ implies both structural and parametric across-block symmetry. Hence, we have the following,

Theorem 4.1 *A sufficient condition for the solution $\hat{\Theta}$ to the pdglasso to be fully symmetric is that $\lambda_2 \geq \lambda_2^{sym}$, where $\lambda_2^{sym} = \max \{|s_{ij} - s_{i'j'}|/2, |s_{i'j} - s_{ij'}|/2; i, j \in L\}$.*

Proof See Appendix B.1. □

The pdglasso problem coincides with the graphical lasso in the case where λ_2 is set to zero, and it is a well-known result that if $\lambda_1 \geq \max_{i, j \in V; i \neq j} |s_{ij}|$ then the solution to the graphical lasso is a diagonal matrix so that the selected graph is fully disconnected; see among others Mazumder and Hastie (2012, Section 2.1). The following proposition shows that this is still true in the pdglasso problem for any $\lambda_2 \geq 0$.

Proposition 4.2 *A sufficient condition for the solution $\hat{\Theta}$ to the pdglasso to be a diagonal matrix is that $\lambda_1 \geq \lambda_1^{diag}$, where $\lambda_1^{diag} = \max \{|s_{ij}|; i, j \in V \text{ with } i \neq j\}$.*

Proof See Appendix B.2. □

For the sake of completeness, we also consider the case where the solution $\hat{\Theta}$ to the pdglasso is block diagonal with all the entries of $\hat{\Theta}_{LR}$ equal to zero. This is of interest because $\hat{\Theta}_{LR} = O_q$ implies that there are no across-block edges so that the two groups are independent. For the case where $\lambda_2 = 0$, i.e. in the graphical lasso, this problem was considered by Witten et al. (2011) and Mazumder and Hastie (2012) where it is proved that such block diagonal structure is obtained whenever $\lambda_1 \geq \max_{i, j \in L} |s_{ij'}|$. Here, we show that the latter result is still true in the pdglasso problem for any $\lambda_2 \geq 0$.

Proposition 4.3 *A sufficient condition for the solution $\hat{\Theta}$ to the pdglasso to be block diagonal with blocks $\hat{\Theta}_{LL}$ and $\hat{\Theta}_{RR}$ is that $\lambda_1 \geq |s_{ij}|$ for every $i \in L$ and $j \in R$.*

Proof See Appendix B.3. □

In the applications considered in this paper, the above results on the penalty parameters are used as follows. Firstly, we apply the pdglasso procedure with $\lambda_2 = 0$ to a sequence of m values of λ_1 equally spaced, on the log-scale, between $\frac{\lambda_1^{diag}}{m}$ and λ_1^{diag} , and then identify the optimal value of λ_1 , among the m considered. Next, with the latter value of λ_1 , we apply again the pdglasso to a sequence of m values of λ_2 equally spaced, on the log-scale, between $\frac{\lambda_2^{sym}}{m}$ and λ_2^{sym} . The optimal solution is then chosen among the latter $m + 1$ solutions identified, including that with $\lambda_2 = 0$. As a goodness of fit measure for a model M we use the extended Bayesian Information Criterion (eBIC) of Foygel and Drton (2010),

$$eBIC = -n l(\hat{\Theta}^{mle}) + \log(n) d + 4 d \gamma \log(p), \quad (6)$$

where we either set $\gamma = 0$, thereby reverting to the classical BIC, or $\gamma = 0.5$, as suggested by Foygel and Drton (2010), when a higher level of sparsity is required. Furthermore, d is the number of parameters, computed by subtracting from the number of parameters of the saturated model, that is $p(p + 1)/2$, both the number of zero constraints and the number of equality constraints characterizing the relevant pdRCON model. It is also worth recalling that $\hat{\Theta}^{\text{mle}}$ in (6) is the (unpenalized) Maximum Likelihood Estimate (MLE) of Θ (Foygel and Drton 2010), that we compute by means of the same algorithm implemented to solve the pdglasso problem; see Sect. 6 for details. This model selection procedure is implemented in the function `pdRCON.fit` of the `pdglasso` package.

The application considered in Sect. 8 concerns the identification of gene networks from gene expression data, whereas Ranciati et al. (2021) considered the identification of brain networks from fMRI data. These are both relevant areas of application where variables are measured on the same unit. We remark that this is an important issue because lasso methods are not invariant to scalar multiplication of the variables (Hastie et al. 2015, p. 9). Thus, when the variable units are not the same, it is common practice to standardized the data before applying the traditional graphical lasso (Carter et al. 2024). However, as noticed by Højsgaard and Lauritzen (2008, Section 3.4), RCON models are not invariant under rescaling, in the sense that standardization will not preserve the original structure of colour classes. This means that, although the application of the pdglasso to the sample correlation matrix may still represent a useful way to choose parsimonious GGMs, care needs to be taken in the interpretation of the resulting symmetries. Roverato and Nguyen (2024) considered the implementation of greedy search procedures on the space of pdRCON models, for which the standardization of the variables is not needed. However, they are computationally demanding and, thus, their use is limited to problems of small dimension. Furthermore, unlike graphical lasso methods, their application in the high-dimensional setting, with p larger than n , is not straightforward.

5 pdRCON submodel classes

When fitting a pdRCON model to a set of data it can be useful to restrict the analysis to one of some relevant submodel classes. This may be motivated by a number of different reasons, including the interpretability of the selected model or the need to keep the model dimension low, for example when n is small relative to p . The function `admm.pdglasso` of the `pdglasso` package implements a more flexible version of the fused penalty (4), given by

$$Q_{\lambda_2}(\Theta) = \lambda_2^{(V)} \|\text{diag}(\Theta_{LL}) - \text{diag}(\Theta_{RR})\|_1 + \lambda_2^{(I)} \|\Theta_{LL}^* - \Theta_{RR}^*\|_1 + \lambda_2^{(A)} \|\Theta_{LR} - \Theta_{RL}\|_1 \tag{7}$$

where $\Theta_{LL}^* = \Theta_{LL} - \text{diag}(\Theta_{LL})$ and, similarly for Θ_{RR}^* ; note that $\text{diag}(\Theta)$ refers to a diagonal matrix having the same diagonal as Θ . In this way, there is one regularization parameter associated with every type of parametric symmetry and, more specifically, $\lambda_2^{(V)}$ is associated with vertex symmetry, $\lambda_2^{(I)}$ with inside-block symmetry and $\lambda_2^{(A)}$ with across-block symmetry. For a given value of $\lambda_2 > 0$, each of the three $\lambda_2^{(\cdot)}$ parameters can take any of the values in the set $\{0, \lambda_2, \text{Inf}\}$. For instance, the user can impose (i) no constraints involving vertex symmetries by setting $\lambda_2^{(V)} = 0$, (ii) the amount of vertex symmetry regularisation implied by $\lambda_2^{(V)} = \lambda_2$ and (iii) full vertex symmetry by setting $\lambda_2^{(V)} = \text{Inf}$. The same can be done independently for $\lambda_2^{(I)}$ and $\lambda_2^{(A)}$ thereby allowing the user to select a model within one of $|\{0, \lambda_2, \text{Inf}\}|^3 = 9$ different pdRCON submodel classes, and in the following we look at some of these submodels in more detail; see also Appendix C for some examples with $|L| = |R| = 3$.

The restrictions used to define an RCON model are linear in the concentration matrix so that the log-likelihood function is concave (Højsgaard et al. 2007; Højsgaard and Lauritzen 2008; Gehrmann 2011) and, thus, the pdglasso in (5) is a convex optimization problem. It is therefore computationally convenient to define symmetries in terms of equality of concentration values, although the interpretation of such constraints may not be straightforward. Højsgaard and Lauritzen (2008, Section 8) remarked that the comparison of concentration values is meaningful only when variables are measured on comparable scales, and we note that this is a condition that is naturally satisfied in our framework because, firstly, for every $i \in L$ the corresponding variables Y_i and $Y_{i'}$ measure a common feature and, secondly, pdRCON models are defined in such a way that comparisons only involve comparable concentration values, associated with pairs of corresponding variables. Hence, for instance, the equality $\theta_{ii} = \theta_{i'i'}$, for $i \in L$, is equivalent to $\sigma_{ii|V \setminus \{i\}}^2 = \sigma_{i'i'|V \setminus \{i'\}}^2$, that is, that the corresponding variables Y_i and $Y_{i'}$ have the same partial variance; recall that $\sigma_{ii|V \setminus \{i\}}^2 = \theta_{ii}^{-1}$ (see, e.g. Lauritzen 1996, Section 5.1.3). Furthermore, if for instance, for $j \in L$ with $j \neq i$ it holds that both $\theta_{ii} = \theta_{i'i'}$ and $\theta_{jj} = \theta_{j'j'}$, then the constraint $\theta_{ij} = \theta_{i'j'}$ implies the equality of the two partial correlations $\rho_{ij|V \setminus \{i,j\}} = \rho_{i'j'|V \setminus \{i',j'\}}$; see Appendix A for details. It is therefore of interest to consider the subfamily of pdRCON models satisfying full vertex symmetry, that is $\theta_{ii} = \theta_{i'i'}$ for every $i \in L$, which are easy to interpret because every equality constraint between off-diagonal concentrations implies the equality of the corresponding partial correlations. In order to restrict the analysis within the family of pdRCON models satisfying full vertex symmetry is

sufficient to set $\lambda_2^{(V)} = \text{Inf}$, but we remark that the equality of partial variances is a strong assumption that should be properly verified.

It may be the case that the substantive research hypothesis underlying the analysis only concerns the comparison of the association structure of the first group with that of the second. More formally, the focus may be on the comparison of Θ_{LL} and Θ_{RR} , which are adapted to \mathcal{G}_L and \mathcal{G}_R , respectively, whereas the cross-group association Θ_{LR} is regarded as a nuisance parameter. In this case, across-block symmetries are of no interest and if the sample size is large it may make sense to set $\lambda_2^{(A)} = 0$. On the other hand, for smaller sample sizes it is also possible to set $\lambda_2^{(A)} = \text{Inf}$ so as to impose full across-block symmetry with the aim to reduce model dimensionality.

Finally, the relevant research question may require the identification of a model that is fully symmetric, in the sense that there are no differences between the two groups. A fully symmetric model can be obtained by setting $\lambda_2^{(V)} = \lambda_2^{(I)} = \lambda_2^{(A)} = \text{Inf}$ and may be used, for instance, as a benchmark for the comparison with an arbitrarily selected model. Interestingly, fully symmetric models belong to the family of RCON models satisfying permutation symmetry, and we refer to Appendix A for a more detailed account on the properties of this submodel class.

6 Implementation via ADMM algorithm

Following Danaher et al. (2014) and Ranciati et al. (2021) we solve the pdglasso optimization problem (5) using an alternating directions method of multipliers (ADMM) algorithm; see Boyd et al. (2011, Section 6.5). This is obtained by splitting the procedure into two nested optimization problems both solved by a specific ADMM algorithm, and where the inner ADMM algorithm is written in a form that makes use of the results for the fused lasso signal approximator given in Friedman et al. (2007). Our ADMM algorithm is implemented in the function `admm.pdglasso` of the `pdglasso` package. The latter extends the algorithm of Ranciati et al. (2021) to include across-block symmetries but, in fact, it has been designed to solve the more general problem involving the penalty (7), and thus it makes it possible for the user to optimize (5) either over the entire family of pdRCON models or over one of the subfamilies of pdRCON models described in Sect. 5. Furthermore, the function `pdRCON.mle` of the same package exploits the ADMM algorithm for the computation of the MLEs. This is achieved by allowing different values of the penalty parameters for the different zero and equality constraints, and then setting the penalty term to a sufficiently large value for the constraints characterizing the model and to zero otherwise.

In the following, we provide a detailed description of the algorithm, and refer to Boyd et al. (2011) for its convergence properties. The pdglasso solution is obtained by looping, over iterations $l = 1, 2, \dots$, the following three steps (see Boyd et al. 2011, equations (3.5) to (3.6)):

- (1) $\Theta^{l+1} := \arg \min_{\Theta} \left(-\log \det(\Theta) + \text{tr}(S\Theta) + \frac{\rho_1}{2} \left\| \Theta - Z^l + U^l \right\|_F^2 \right);$
- (2) $Z^{l+1} := \arg \min_Z \left(\mathcal{P}_{\lambda_1}(Z) + \mathcal{Q}_{\lambda_2}(Z) + \frac{\rho_1}{2} \left\| \Theta^{l+1} - Z + U^l \right\|_F^2 \right);$
- (3) $U^{l+1} := U^l + \Theta^{l+1} - Z^{l+1},$

where $\|\cdot\|_F$ denotes the Frobenius norm, $\rho_1 > 0$ is the step size and U and Z are $p \times p$ matrices initialized with all entries equal to zero. The solution to step (1) can be obtained in analytic form as detailed in Boyd et al. (2011, Section 3.1.1); see also Ranciati et al. (2021).

More specific to our implementation is the solution of step (2). Consider a matrix Q whose columns and rows are indexed by $V = L \cup R$, and let $\text{vech}(\cdot)$ and $\text{vd}(\cdot)$ denote the half-vectorization and the diagonal extraction operator, respectively. Hence, we define the vector $v(Q)$ as,

$$v(Q)^T = \begin{bmatrix} \text{vd}(Q_{LL})^T & \text{vd}(Q_{RR})^T & \text{vech}(Q_{LL})^T & \text{vech}(Q_{RR})^T \\ & & \text{vech}(Q_{LR})^T & \text{vech}(Q_{RL})^T & \text{vd}(Q_{LR})^T \end{bmatrix}^T,$$

and then we set $z^l = v(Z^l)$ and $b^l = v(\Theta^l) + v(U^l)$. All equality constraints are encoded in a matrix F , made up of the following three row blocks: $[I_q, -I_q, O_{q,4s+q}]$, $[O_{s,2q}, I_s, -I_s, O_{s,2s+q}]$ and $[O_{s,2q+2s}, I_s, -I_s, O_{qq}]$, where I and O are the identity and zero matrix, respectively, and $s = q(q - 1)/2$. Thus, step (2) can be stated as the following self-standing optimization problem,

$$\arg \min_z \left(\frac{1}{2} \|z - b\|_2^2 + \lambda'_1 \|z\|_1 + \|F(\lambda'_2 \circ z)\|_1 \right),$$

where \circ denotes the element-wise product, $\lambda'_1 = \lambda_1/\rho_1$, $\lambda'^T_2 = \begin{bmatrix} \lambda_2^{(V)}/\rho_1 \mathbf{1}_{2q}^T & \lambda_2^{(I)}/\rho_1 \mathbf{1}_{2s}^T & \lambda_2^{(A)}/\rho_1 \mathbf{1}_{2s+q}^T \end{bmatrix}$, and $\mathbf{1}_q$ the unit vector of length q .

Due to Friedman et al. (2007, Lemma A.1), we can firstly solve

$$\arg \min_z \left(\frac{1}{2} \|z - b\|_2^2 + \|F(\lambda'_2 \circ z)\|_1 \right),$$

with $\lambda'_1 = 0$, and then obtain the solution for the case $\lambda'_1 > 0$ via a soft-thresholding operation. This is a generalized lasso

problem (Tibshirani and Taylor 2011), which we solve with the following inner ADMM procedure (Boyd et al. 2011, Section 6.4.1)

- (i) $z^{m+1} := (I + \rho_2 F^T F)^{-1} \{b + \rho_2 F^T (v^m - t^m)\};$
- (ii) $v^{m+1} := \mathcal{S}_{\lambda'_2/\rho_2} (F z^{m+1} + t^m);$
- (iii) $t^{m+1} := t^m + F z^{m+1} - v^{m+1}.$

The vectors v and t are initialized with zero entries, the scalar $\rho_2 > 0$ refers to the step size of the inner ADMM and $\mathcal{S}_{\lambda'_2/\rho_2}(\cdot)$ represents the soft-thresholding operator, such that any input less than λ'_2/ρ_2 (in absolute value) is set to zero, otherwise it gets shrunk by that threshold. The stopping rule of the algorithm is based on the primal and dual residual (Boyd et al. 2011, Section 3.3.1), computed at each iteration, which are compared to a given numerical precision threshold.

Finally, we remark that, in the computer implementation of the above ADMM algorithm, efficiency has been achieved by making the step sizes ρ_1 and ρ_2 adaptive and, furthermore, both by properly reducing the dimension of the matrix F according to the submodel class of interest, and by encoding the results of the products involving the matrix F in the form of less expensive vectorized computations, explicitly based on the (sparse) structure of such matrix. In a call to the function `pdRCON.fit` on a grid of 20 values for each penalty parameter, the ADMM algorithm is run 80 times, that is 40 times on a range of different penalty values and 40 times for the computation the MLEs for the eBIC. On a recent hardware (CPU 1.8GHz Intel Core i7 quad-core, memory 16 GB 2133 MHz LPDDR3), the individual cost for an ADMM execution is approximately 1 s. These computations refer to execution times for the simulation study described in Sect. 7.

7 Simulation study

pdRCON models are GGMs with additional equality constraints, which we call parametric symmetries. It follows that a GGM is a pdRCON model with no parametric symmetry, and it is therefore of interest to investigate the behaviour of the pdglasso method when the data are generated from a GGM and, similarly, the behaviour of the glasso method when parametric symmetries are indeed present.

In the following, we compare the behaviour of glasso and pdglasso by means of simulated data, and consider three different scenarios: (i) GGMs with no parametric symmetry, (ii) pdRCON models where 50% of the parameters are involved in a parametric symmetry and the remaining 50% are unconstrained and, (iii) fully symmetric pdRCON models. For each scenario we set $p = 100$ and randomly generated 10 models with graph density equal to 0.20, where

this density is computed as the number of present edges over the number of edges of the complete graph. A Gaussian distribution, identified by its covariance matrix, was obtained for every model and from each of the resulting 30 distributions we randomly generated 7 datasets of size $n = \{100, 150, 200, 300, 500, 1000, 1500\}$. To each of the 210 resulting datasets, we applied the procedure described in Sect. 4 with grid length $m = 20$ and $\gamma = 0$. All the computations were carried out by using the R package `pdglasso`, and we now describe the procedure used to randomly generate the covariance matrices characterizing the Gaussian distributions used in the simulations. The procedure implemented in the function `GGM.simulate` to obtain a covariance matrix $\Sigma^{(*)}$ relative to a GGM is as follows. Firstly, a $p \times p$ positive definite matrix S^* is randomly generated from a Wishart distribution with matrix parameter equal to the identity matrix. Next, $(S^*)^{-1}$ is computed and an undirected graph \mathcal{G} , with the required density, is obtained in such a way that the edges of \mathcal{G} correspond to the largest absolute values of the off-diagonal entries of $(S^*)^{-1}$. Finally, a positive definite matrix $\Sigma^{(*)}$, that identifies a distribution from the GGM represented by \mathcal{G} , is obtained by applying to S^* the procedure for the computation of the MLEs. The same idea is exploited by the function `pdRCON.simulate` to add to $\Sigma^{(*)}$ the required amount of symmetries so as to obtain a covariance matrix from a pdRCON model. For information about the relevant material to reproduce the simulations see the Section ‘Data availability’ of this paper.

For every scenario and sample size, the average behaviour across the 10 selected models of pdglasso was compared with that of glasso. Because glasso is not designed to identify parametric symmetries, we compared the two methods with respect to their performance in learning the graph structure and the inverse covariance matrix Θ .

The performance of the procedures in structure learning was measured through the Positive Predicted Value (PPV) and the True Positive Rate (TPR). These are given in Fig. 1 where also a measure of parsimony is provided, by the number of parameters of the selected models. The panels composing the first column of Fig. 1 show that, for the no symmetry case where the true model is a GGM, the pdglasso method selected less parsimonious models, thereby resulting in a worse performance, compared to glasso, with respect to PPV but an improvement in terms of TPR; these differences between the two methods are more noticeable for larger sample sizes. Plots in the second column show that in the 50% symmetry case the two methods performed similarly. For the full symmetry case, third column, pdglasso selected more parsimonious models, with a better performance in terms of PPV and producing TPR values comparable to those of glasso.

As overall measures of performance in structure learning we considered the harmonic mean of PPV and TPR, called the

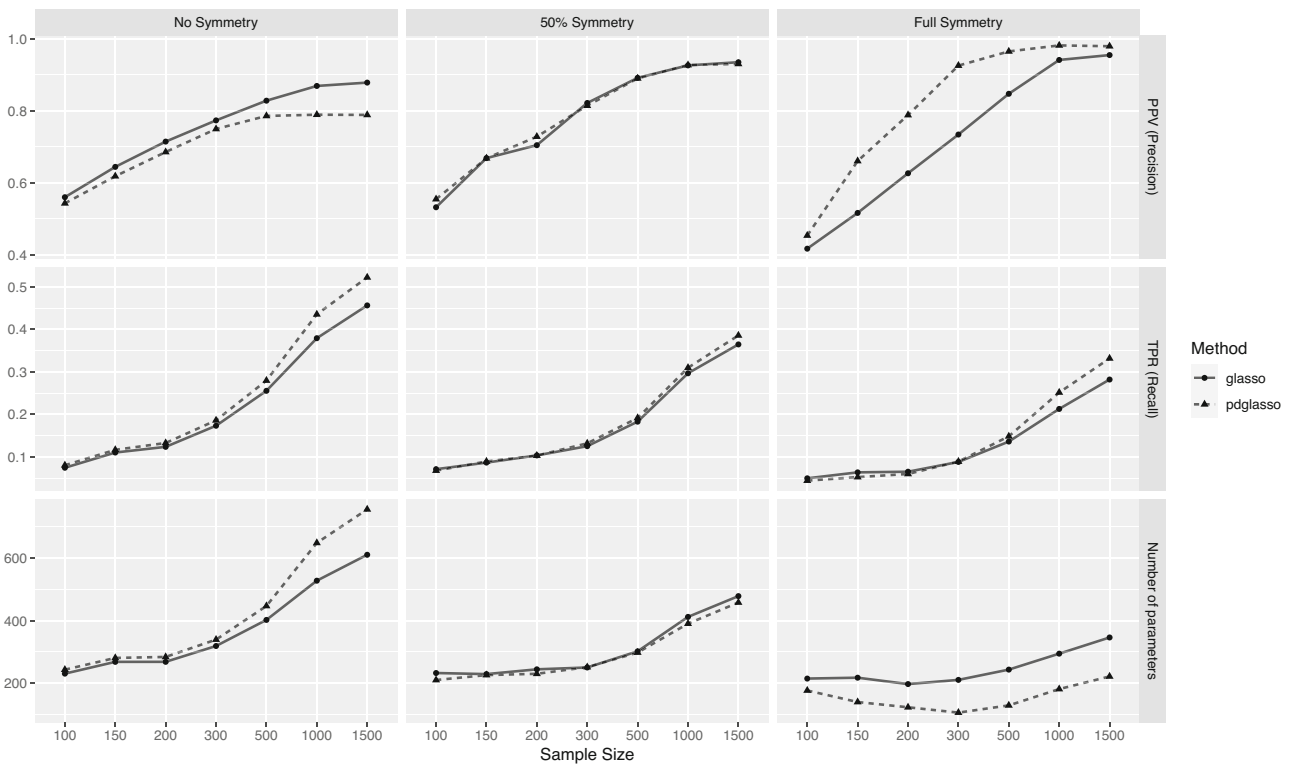


Fig. 1 Comparison of the pdglasso and the glasso methods for the three scenarios: “no symmetry”, “50% symmetry” and “full symmetry”. Panels in the first row give the Positive Predicted Value (PPV) also called precision, those in the second row the True Positive Rate (TPR) also

called recall, whereas the third row gives the number of parameters. Every value is the average of the corresponding quantities computed on the 10 selected models for every scenario and sample size

F1 score, and the Matthews correlation coefficient (MCC). Furthermore, we assessed the performance with respect to the estimate of the inverse covariance by means of the Frobenius norm loss and of the entropy loss (Dey and Srinivasan 1985). More specifically the top panels of Fig. 2 compare the two method by means of the differences of the F1 scores and the MCC values, respectively, with positive values in favour of pdglasso whereas the bottom panels give the relative differences of the Frobenius norm loss and of the entropy loss, with negative values in favor of pdglasso. Figure 2 shows that, although there are some cases where glasso performs moderately better than pdglasso especially for small sample sizes and in the no symmetry case, there is a clear evidence that, as expected, the pdglasso method provided better results in the models presenting parametric symmetry. Additionally, in the case where the data come from a GGM the pdglasso method provided results comparable with those of glasso.

8 Analysis of breast cancer gene expression data

We illustrate the use of our proposed method in a gene expression paired data problem concerning breast cancer. The samples refer to $n = 114$ individuals with both tumor and healthy adjacent tissue information, hence the paired data nature of the problem. A curated set of $q = 89$ genes of the Hedgehog Pathway, known for their involvement in breast cancer (Song et al. 2014; Kubo et al. 2004), is extracted from MSigDB Collections (Subramanian et al. 2005). For each individual, we consider gene-level transcription estimates, that is $\log_2(Y + 1)$ transformed normalized counts, across the same set of selected genes in both tissues ($p = 89 \times 2 = 178$), which lead to the 114×178 final data matrix for the analysis.

We focus on two pdRCON submodel classes. Firstly, we consider a pdRCON model (labelled as fV) where $\lambda_2^{(V)} = \text{Inf}$ and $\lambda_2^{(I)} = \lambda_2^{(A)} = \lambda_2$, so as to force full vertex symmetry, and penalize with magnitude λ_2 the remaining blocks. The second model considered (labelled as fVIA) is a fully symmetric model, which is obtained by setting $\lambda_2^{(V)} = \lambda_2^{(I)} = \lambda_2^{(A)} = \text{Inf}$. More details on these submodel classes can be found in Sect. 5 and Appendix A. For

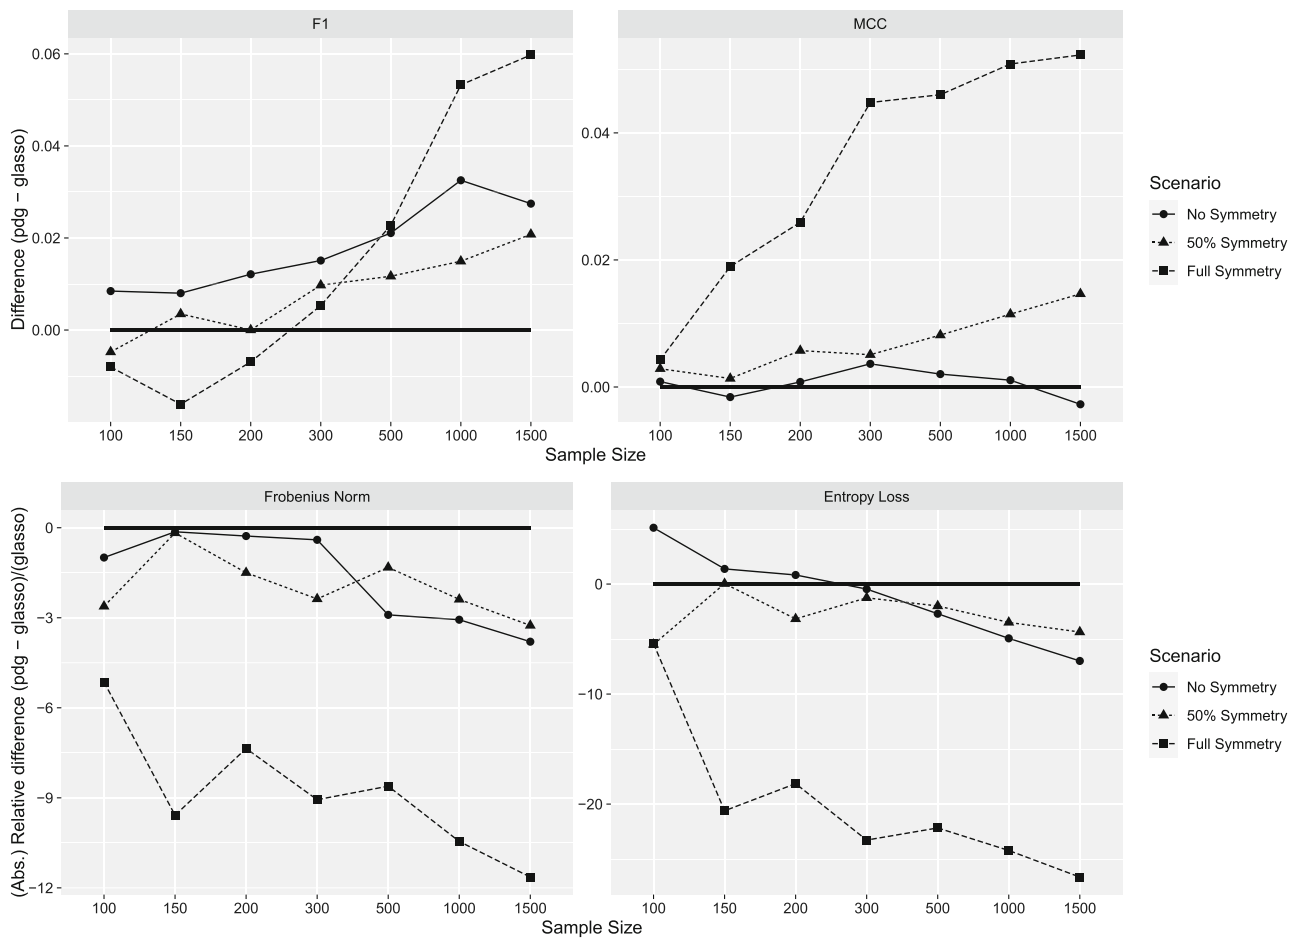


Fig. 2 Comparison of the pdglasso and the glasso methods for the three scenarios: “no symmetry”, “50% symmetry” and “full symmetry”. Top panels: differences of the F1 score values and of the Matthews correlation coefficient (MCC) values, with positive value denoting a better performance of pdglasso. Bottom panels: relative differences of

the Frobenius norm loss values for the inverse covariance and of the entropy loss values, with negative value denoting a better performance of pdglasso. Every value is the average of the corresponding quantities computed on the 10 selected models for every scenario and sample size

the selection of the optimal penalty values we apply the procedure described in Sect. 4, with grid length $m = 20$, and where the eBIC is applied with $\gamma = 0.5$, because domain experts are interest in recovering a sparse graph with the most relevant connections. The selected models, denoted as mod_{fV} and mod_{fVIA} , respectively, are visually depicted in Figs. 7 and 8 of Appendix D. They are very sparse and, for this reason, in the following we will only focus on the symmetries, either structural and parametric, that involve edges which are present in the graphs. Accordingly, to ease the reading, we will simply talk of symmetries without recalling that we are only considering present edges.

An overview of the structure of the models involved in this analysis is provided in Table 1. The graph associated to model mod_{fV} has an overall density lower than 1%, with inside-block density larger than the across-block one. In particular, the graph has only 7 edges connecting genes across

the two different groups (tumor vs healthy), whereas the two group specific graphs are similar in terms of density, with 56 edges between variables in the tumor group and 54 in the healthy group. The most connected gene, both inside and across blocks, is the one named LRP2. More specifically, in the tumor group 16 out of the 56 edges (28.6%) connect LRP2 to other inside-block genes, whereas in the healthy group 37 out of the 54 edges (68.5%) involve LRP2. Furthermore, most of the across-block edges, i.e. 5 out of the 7 (71.4%), involve LRP2 and, finally, 12 of the 24 coloured edge involve LRP2, that is 50% of the identified parametric symmetries concern this gene.

We turn now to the fully symmetric model mod_{fVIA} and note that, although in this case the identified graph has a larger number of edges, with a density that is almost three times greater than that of the graph associated with mod_{fV} , the number of parameters of mod_{fVIA} is only 1.3 times

Table 1 Summary statistics on the structure of the models identified by the application of pdglasso: mod_{fVIA} (last column); mod_{fV} and its three submodels considered (second to fifth column); sub1_{fV} : across-block edges set to zero; sub2_{fV} : inside-block structural asymmetries removed; sub3_{fV} : inside-block structural symmetries converted to parametric symmetries. All models have full vertex symmetry

Summary	Reference model and submodels				
	mod_{fV}	sub1_{fV}	sub2_{fV}	sub3_{fV}	$fVIA$
Total edges	117	110	49	117	300
Graph density	0.74 %	0.70 %	0.31 %	0.74 %	1.90 %
<i>inside block</i>					
Total edges	110	110	42	110	286
Structural (non-zero) symmetric edges	18	18	18	0	0
Parametric (non-zero) symmetric edges	24	24	24	42	286
<i>across block</i>					
Total edges	7	0	7	7	14
Structural (non-zero) symmetric edges	0	0	0	0	0
Parametric (non-zero) symmetric edges	0	0	0	0	10

larger than that of mod_{fV} , which shows how the identification of parametric symmetries can provide an effective way to increase parsimony in the selection of a GGM.

We also compare models mod_{fV} with three of its submodels: (i) sub1_{fV} , where across-block edges are removed; (ii) sub2_{fV} , where inside-block structural asymmetries are removed; (iii) sub3_{fV} , where all inside-block structural symmetries are turned into parametric symmetries. The three submodels are not nested within one another but they all are with respect to mod_{fV} .

We first consider model sub1_{fV} . Unlike the case of independent samples, in paired data problems a key issue is represented by the cross-graph structure that encodes the across-group dependence. It is interesting to note that there are potentially $q^2 = 7921$ across-graph edges, but the selected model mod_{fV} shows a very sparse across-graph association structure with as few as 7 edges, most of which involving the gene LRP2. It is therefore of interest to quantify the significance of the across-graph structure by comparing the model mod_{fV} with its submodel sub1_{fV} obtained by assuming that the group specific subgraphs are disconnected, that is $\Theta_{LR} = O_q$. This comparison is carried out by means of the Likelihood Ratio Test (LRT) at a significance level of $\alpha = 0.05$, given in Table 2, that shows that the model sub1_{fV} seems not to provide an adequate fit to the data, and thus we can conclude that the across-groups structure cannot be ignored, and should thus be retained.

In the comparison of the two groups, differences are represented by asymmetries, and this motivates the comparison of the selected model mod_{fV} with its submodels sub2_{fV} and sub3_{fV} . As shown but juxtaposing the second and fourth column of Table 1, among the 110 inside-block edges of the reference model 68 are asymmetric edges and thus not involved in any structural or parametric symmetry and model sub2_{fV} is obtained by removing such 68 edges. Hence, comparing mod_{fV} with sub2_{fV} amounts to test the presence of structural differences between the group specific subgraphs.

Finally, comparing mod_{fV} with sub3_{fV} amounts to test if the model can be further simplified by assuming that every structural symmetry is also parametric. As shown in Table 2 the empirical evidence is in favor of the hypotheses that structural asymmetries are present and that not all structural symmetries are also parametric.

9 Conclusions

We have considered the problem of joint learning of GGMs for paired data, in an approach that implements a fused graphical lasso penalty to identify a model within a suited family of coloured graphical models. We have addressed a number of issues with the aim of providing the results and tools required for an immediate and knowledgeable application of the method. The latter include an ADMM algorithm for the optimization of the penalized likelihood, some results required when computing a path of lasso solutions, the description of the features of some model subfamilies of specific interest and, finally, an R package where all the methods have been implemented.

We deem that one appealing feature of our approach is the interpretation of the selected model. When the number of variables is large, it is not straightforward to visualize and interpret a GGM, and this task is even more challenging when the analysis involves two or more, possibly dependent, networks. The application of pdRCON models is restricted to the case where two groups are considered, but it provides a transparent representation of the across group association. Furthermore, as shown in the application of Sect. 8, one can meaningfully summarize the relationship between groups by focusing on the amount of both structural and parametric symmetry/asymmetry. The scope of the analysis may also justify the restriction to specific submodels such as the model with full vertex symmetry of the fully symmetric model.

Table 2 Summary of Likelihood Ratio Tests (LRTs) performed on each submodel against the reference mod_{EV}

H_0 : submodel is preferable	mod_{EV}	sub1_{EV}	sub2_{EV}	sub3_{EV}
Number of parameters	194	187	126	185
Deviance	12242.22	12285.79	14218.92	12393.96
LRT value	–	43.56	1976.69	151.73
Degrees of freedom (df)	–	7	68	9
Critical $\chi_{df,\alpha}$ quantile at $\alpha = 0.95$	–	14.07	88.25	16.92

Appendix

A Fully symmetric pdRCON models

One of the main motivations for the analysis of paired data is the need to identify similarities, i.e. symmetries, and differences between the two groups. Thus, in this framework, a central role is played by the subclass of pdRCON models that show no differences between the two groups. It is of theoretical interest to understand the properties of this model class and, furthermore, in applications it may be useful to estimate a model within this class because it could be regarded as a benchmark for the comparison with an arbitrarily selected model, and the quantification of the amount of asymmetry it shows. We call these models *fully symmetric* because, as shown below, their parametric symmetries are not restricted to the equalities of concentration coefficients but extend to partial correlation coefficients, variances, covariances and correlation coefficients. Theorem 4.1 provides a result concerning the value of λ_2 such that the solution to the *pdglasso* is a fully symmetric model, whereas an instance of a fully symmetric pdRCON model is given in Example C.4 of Appendix C.

As far as the interpretation of the model is of concern, a useful property of fully symmetric pdRCON models is that every inside and across-block parametric equality constraint, for instance $\theta_{ij} = \theta_{i'j'}$, also implies that the corresponding partial correlations have the same values, i.e. $\rho_{ij|V \setminus \{i,j\}} = \rho_{i'j'|V \setminus \{i',j'\}}$. This follows immediately from the fact that the partial correlation between Y_i and Y_j given the remaining variables can be computed from concentrations as $\rho_{ij|V \setminus \{i,j\}} = -\theta_{ij} / \sqrt{\theta_{ii}\theta_{jj}}$ (see, e.g. Lauritzen 1996, Section 5.1.3) and in a fully symmetric model it also holds that both $\theta_{ii} = \theta_{i'i'}$ and $\theta_{jj} = \theta_{j'j'}$.

An equivalent way to define fully symmetric pdRCON models is by requiring that the random vector $(Y_L, Y_R)^\top$ has the same distribution as that of $(Y_R, Y_L)^\top$; that is Y_L and Y_R are exchangeable. This shows that fully symmetric pdRCON models belong to the family of coloured GGMS which satisfy permutation symmetry, as defined in Højsgaard and Lauritzen (2008, Section 5); see also Gehrman (2011) and Graczyk et al. (2022). More specifically, if we denote by J the permutation matrix that exchanges Y_L and

Y_R , i.e. $J \times (Y_L, Y_R)^\top = (Y_R, Y_L)^\top$, then we have,

$$J = \begin{pmatrix} O_q & I_q \\ I_q & O_q \end{pmatrix} \text{ and thus } J\Theta J = \begin{pmatrix} \Theta_{RR} & \Theta_{RL} \\ \Theta_{LR} & \Theta_{LL} \end{pmatrix}, \tag{A.1}$$

where I_q and O_q are the $q \times q$ identity and zero matrices, respectively. Recall that permutation matrices are orthogonal and, furthermore, J is symmetric so that $J = J^\top = J^{-1}$. We can thus say that a pdRCON model is fully symmetric if and only if $J\Theta J = \Theta$, in words, if and only if Θ is *invariant under the action of J* . It is straightforward to see that the latter equality is equivalent to $(J\Theta J)^{-1} = \Theta^{-1}$ and thus to $J\Sigma J = \Sigma$, and this shows that in a fully symmetric model Σ is also invariant under the action of J . This implies that parametric symmetry holds also for to the entries of the variance matrix, that is $\Sigma_{LL} = \Sigma_{RR}$ and $\Sigma_{LR} = \Sigma_{RL}$ and it straightforward to see that also the correlation matrix P satisfies the same property, i.e. both $P_{LL} = P_{RR}$ and $P_{LR} = P_{RL}$.

Finally, we notice that for the family of fully symmetric models, the MLE of Θ can be obtained by instead considering the matrix $\bar{S} = (S + JSJ)/2$ and finding the MLE of the corresponding GGM without symmetry restrictions (Højsgaard and Lauritzen 2008, Section 5.2). This MLE can be calculated explicitly when \mathcal{G} is decomposable, or alternatively a standard algorithm, such as the iterative proportional scaling, can be used (Lauritzen 1996, page 146).

B Proofs

B.1 Proof of Theorem 4.1

For a positive definite matrix Θ we consider J in (A.1) and set the matrix $\bar{\Theta} = \frac{1}{2}(\Theta + J\Theta J)$. Note that $J^{-1} = J$ so that $J\bar{\Theta}J = \frac{1}{2}(J\Theta J + JJ\Theta JJ) = \frac{1}{2}(J\Theta J + \Theta) = \bar{\Theta}$, that is, $\bar{\Theta}$ is fully symmetric. It is also worth remarking that $\bar{\Theta}$ is positive definite by construction.

We now show that if λ_2 is chosen to satisfy the inequality of Theorem 4.1 then for any positive definite matrix Θ it holds that $\mathcal{L}_{\lambda_1, \lambda_2}(\bar{\Theta}) \leq \mathcal{L}_{\lambda_1, \lambda_2}(\Theta)$ so that the matrix which minimises $\mathcal{L}_{\lambda_1, \lambda_2}(\cdot)$ in (5) must be fully symmetric. Concretely, we show that $\mathcal{L}_{\lambda_1, \lambda_2}(\bar{\Theta}) - \mathcal{L}_{\lambda_1, \lambda_2}(\Theta) \leq 0$ and, to this aim, we analyse, in turn, three components of this difference.

We first show that $\{-\log \det(\bar{\Theta})\} - \{-\log \det(\Theta)\} \leq 0$, which is equivalent to $\log \det(\Theta) \leq \log \det(\bar{\Theta})$ and therefore to $\det(\Theta) \leq \det(\bar{\Theta})$. Because both Θ and $J\Theta J$ are positive definite the following inequality holds (see, e.g., Lütkepohl 1996, p. 55 eqn (14))

$$\begin{aligned} \det(\bar{\Theta})^{1/p} &= \det\left(\frac{1}{2}\Theta + \frac{1}{2}J\Theta J\right)^{1/p} \\ &\geq \det\left(\frac{1}{2}\Theta\right)^{1/p} + \det\left(\frac{1}{2}J\Theta J\right)^{1/p} \\ &= 2 \det\left(\frac{1}{2}\Theta\right)^{1/p} = \det(\Theta)^{1/p} \end{aligned}$$

where we have used the facts that $\det(\Theta) = \det(J\Theta J)$, because J is a permutation matrix, and that $\det\left(\frac{1}{2}\Theta\right) = \frac{1}{2^p} \det(\Theta)$. We have thus shown that $\det(\bar{\Theta})^{1/p} \geq \det(\Theta)^{1/p}$ and this immediately implies that $\det(\bar{\Theta}) \geq \det(\Theta)$ as required.

We now show that $+\lambda_1 \|\bar{\Theta}\|_1 - \lambda_1 \|\Theta\|_1 \leq 0$, i.e. that $\|\bar{\Theta}\|_1 - \|\Theta\|_1 \leq 0$. We first notice that $\|\Theta\|_1$ can be computed on the block components (1) of Θ as follows,

$$\|\Theta\|_1 = \|\Theta_{LL}\|_1 + \|\Theta_{RR}\|_1 + \|\Theta_{LR}\|_1 + \|\Theta_{RL}\|_1.$$

Similarly,

$$\begin{aligned} \|\bar{\Theta}\|_1 &= \|\bar{\Theta}_{LL}\|_1 + \|\bar{\Theta}_{RR}\|_1 + \|\bar{\Theta}_{LR}\|_1 + \|\bar{\Theta}_{RL}\|_1 \\ &= \frac{1}{2}(\|\Theta_{LL} + \Theta_{RR}\|_1 + \|\Theta_{RR} + \Theta_{LL}\|_1 \\ &\quad + \|\Theta_{LR} + \Theta_{RL}\|_1 + \|\Theta_{RL} + \Theta_{LR}\|_1) \\ &= \|\Theta_{LL} + \Theta_{RR}\|_1 + \|\Theta_{LR} + \Theta_{RL}\|_1. \end{aligned}$$

Hence,

$$\begin{aligned} \|\bar{\Theta}\|_1 - \|\Theta\|_1 &= \|\Theta_{LL} + \Theta_{RR}\|_1 - (\|\Theta_{LL}\|_1 + \|\Theta_{RR}\|_1) \\ &\quad + \|\Theta_{LR} + \Theta_{RL}\|_1 - (\|\Theta_{LR}\|_1 + \|\Theta_{RL}\|_1), \end{aligned}$$

and it follows that $\|\bar{\Theta}\|_1 - \|\Theta\|_1 \leq 0$ because for any pair of real numbers a and b it holds that $|a + b| \leq |a| + |b|$.

In order to complete the proof we have now to show that

$$\begin{aligned} \text{tr}(S\bar{\Theta}) - \text{tr}(S\Theta) - \lambda_2 (\|\Theta_{LL} - \Theta_{RR}\|_1 + \|\Theta_{LR} - \Theta_{RL}\|_1) \\ \leq 0. \end{aligned} \tag{B.1}$$

Note that the term $\lambda_2 (\|\bar{\Theta}_{LL} - \bar{\Theta}_{RR}\|_1 + \|\bar{\Theta}_{LR} - \bar{\Theta}_{RL}\|_1)$ does not appear in (B.1) because it is equal to zero by construction. We first consider the difference

$$\begin{aligned} \text{tr}(S\bar{\Theta}) - \text{tr}(S\Theta) &= \text{tr}(S\bar{\Theta} - S\Theta) = \text{tr}\{S(\bar{\Theta} - \Theta)\} \\ &= \text{tr}\left\{S\left(\frac{1}{2}\Theta + \frac{1}{2}J\Theta J - \Theta\right)\right\} \end{aligned}$$

$$\begin{aligned} &= \text{tr}\left\{\frac{1}{2}S(J\Theta J - \Theta)\right\} \\ &= \frac{1}{2} \sum_{i=1}^p \sum_{j=1}^p s_{ij}(J\Theta J - \Theta)_{ij}, \end{aligned} \tag{B.2}$$

where (B.2) follows from the fact that both S and $(J\Theta J - \Theta)$ are symmetric matrices. Hence, if we write (B.2) by distinguishing the terms corresponding to each of the four blocks of $(J\Theta J - \Theta)$ as in (1) we obtain,

$$\begin{aligned} \text{tr}(S\bar{\Theta}) - \text{tr}(S\Theta) &= \frac{1}{2} \sum_{i=1}^q \sum_{j=1}^q \{s_{ij}(\theta_{i'j'} - \theta_{ij}) + s_{i'j'}(\theta_{ij} - \theta_{i'j'}) \\ &\quad + s_{ij'}(\theta_{ij} - \theta_{i'j'}) + s_{i'j}(\theta_{ij'} - \theta_{i'j})\} \\ &= \frac{1}{2} \sum_{i=1}^q \sum_{j=1}^q \{(s_{i'j'} - s_{ij})(\theta_{ij} - \theta_{i'j'}) + (s_{i'j} - s_{ij'})(\theta_{ij'} - \theta_{i'j})\} \end{aligned} \tag{B.3}$$

$$\leq \frac{1}{2} \sum_{i=1}^q \sum_{j=1}^q (|s_{i'j'} - s_{ij}||\theta_{ij} - \theta_{i'j'}| + |s_{i'j} - s_{ij'}||\theta_{ij'} - \theta_{i'j}|). \tag{B.4}$$

Furthermore,

$$\begin{aligned} \lambda_2 (\|\Theta_{LL} - \Theta_{RR}\|_1 + \|\Theta_{LR} - \Theta_{RL}\|_1) &= \sum_{i=1}^q \sum_{j=1}^q \lambda_2 (|\theta_{ij} - \theta_{i'j'}| + |\theta_{ij'} - \theta_{i'j}|). \end{aligned} \tag{B.5}$$

Hence, from (B.3), (B.5) and then (B.4) we obtain

$$\begin{aligned} &= \text{tr}(S\Theta) - \text{tr}(S\Theta) - \lambda_2 (\|\Theta_{LL} - \Theta_{RR}\|_1 + \|\Theta_{LR} - \Theta_{RL}\|_1) \\ &= \frac{1}{2} \sum_{i=1}^q \sum_{j=1}^q \{(s_{i'j'} - s_{ij})(\theta_{ij} - \theta_{i'j'}) + (s_{i'j} - s_{ij'})(\theta_{ij'} - \theta_{i'j})\} \\ &\quad - \sum_{i=1}^q \sum_{j=1}^q \lambda_2 (|\theta_{ij} - \theta_{i'j'}| + |\theta_{ij'} - \theta_{i'j}|) \\ &\leq \frac{1}{2} \sum_{i=1}^q \sum_{j=1}^q (|s_{i'j'} - s_{ij}||\theta_{ij} - \theta_{i'j'}| + |s_{i'j} - s_{ij'}||\theta_{ij'} - \theta_{i'j}|) \\ &\quad - \sum_{i=1}^q \sum_{j=1}^q \lambda_2 (|\theta_{ij} - \theta_{i'j'}| + |\theta_{ij'} - \theta_{i'j}|) \\ &= \sum_{i=1}^q \sum_{j=1}^q \{(|s_{i'j'} - s_{ij}|/2 - \lambda_2) |\theta_{ij} - \theta_{i'j'}| \\ &\quad + (|s_{i'j} - s_{ij'}|/2 - \lambda_2) |\theta_{ij'} - \theta_{i'j}|\}, \end{aligned}$$

that establishes (B.1) because we have assumed that λ_2 satisfies the inequality of Theorem 4.1 so that both $(|s_{i'j'} - s_{ij}|/2 - \lambda_2)$ and $(|s_{i'j} - s_{ij'}|/2 - \lambda_2)$ are smaller or equal to zero for every $i, j \in L$.

B.2 Proof of Proposition 4.2

For a positive definite matrix Θ we denote by $\bar{\Theta}$ the diagonal matrix with the same diagonal as Θ , that is $\bar{\Theta} = \text{diag}(\Theta)$, and note that $\bar{\Theta}$ is positive definite by construction. Hence, we now show that if $\lambda_1 \geq |s_{ij}|$ for every $i \neq j \in V$ then for any positive definite matrix Θ it holds that $\mathcal{L}_{\lambda_1, \lambda_2}(\bar{\Theta}) \leq \mathcal{L}_{\lambda_1, \lambda_2}(\Theta)$ so that the matrix which minimises $\mathcal{L}_{\lambda_1, \lambda_2}(\cdot)$ in (5) must be diagonal. Concretely, we show that $\mathcal{L}_{\lambda_1, \lambda_2}(\bar{\Theta}) - \mathcal{L}_{\lambda_1, \lambda_2}(\Theta) \leq 0$ and, to this aim, we analyse, in turn, three components of this difference.

We first show that $\{-\log \det(\bar{\Theta})\} - \{-\log \det(\Theta)\} \leq 0$. The latter is equivalent to $\log \det(\Theta) \leq \log \det(\bar{\Theta})$, and therefore to $\det(\Theta) \leq \det(\bar{\Theta})$, which follows immediately from Hadamard’s inequality; see, for instance, Lütkepohl (1996, p. 54 eqn (3)).

We now show that $\text{tr}(S\bar{\Theta}) - \text{tr}(S\Theta) + \lambda_1 \|\bar{\Theta}\|_1 - \lambda_1 \|\Theta\|_1 \leq 0$. Firstly, we notice that

$$\begin{aligned} \text{tr}(S\bar{\Theta}) - \text{tr}(S\Theta) &= \text{tr}(S\bar{\Theta} - S\Theta) = \text{tr}\{S(\bar{\Theta} - \Theta)\} \\ &= -\sum_{i=1}^p \sum_{\substack{j=1 \\ j \neq i}}^p s_{ij} \theta_{ij} \leq \sum_{i=1}^p \sum_{\substack{j=1 \\ j \neq i}}^p |s_{ij}| |\theta_{ij}|, \end{aligned}$$

where the third equality follows from the fact that both S and $(\bar{\Theta} - \Theta)$ are symmetric matrices and, more specifically, that $(\bar{\Theta} - \Theta)$ is the matrix obtained by setting to zero the diagonal entries of $-\Theta$. Similarly,

$$\lambda_1 \|\bar{\Theta}\|_1 - \lambda_1 \|\Theta\|_1 = -\sum_{i=1}^p \sum_{\substack{j=1 \\ j \neq i}}^p \lambda_1 |\theta_{ij}|;$$

so that,

$$\begin{aligned} \text{tr}(S\bar{\Theta}) - \text{tr}(S\Theta) + \lambda_1 \|\bar{\Theta}\|_1 - \lambda_1 \|\Theta\|_1 &= -\sum_{i=1}^p \sum_{\substack{j=1 \\ j \neq i}}^p s_{ij} \theta_{ij} - \sum_{i=1}^p \sum_{\substack{j=1 \\ j \neq i}}^p \lambda_1 |\theta_{ij}| \\ &\leq \sum_{i=1}^p \sum_{\substack{j=1 \\ j \neq i}}^p |s_{ij}| |\theta_{ij}| - \sum_{i=1}^p \sum_{\substack{j=1 \\ j \neq i}}^p \lambda_1 |\theta_{ij}| \\ &= \sum_{i=1}^p \sum_{\substack{j=1 \\ j \neq i}}^p (|s_{ij}| - \lambda_1) |\theta_{ij}| \leq 0, \end{aligned}$$

where the latter inequality holds true because, by assumption, $(|s_{ij}| - \lambda_1) \leq 0$ for every $i \neq j \in V$.

Finally, we show that the difference

$$\lambda_2 (\|\bar{\Theta}_{LL} - \bar{\Theta}_{RR}\|_1 + \|\bar{\Theta}_{LR} - \bar{\Theta}_{RL}\|_1)$$

$$- \lambda_2 (\|\bar{\Theta}_{LL} - \Theta_{RR}\|_1 + \|\bar{\Theta}_{LR} - \Theta_{RL}\|_1) \tag{B.6}$$

is smaller or equal to zero. We first note that $\|\bar{\Theta}_{LR} - \bar{\Theta}_{RL}\|_1 = 0$ because $\bar{\Theta}$ is diagonal and, next, that $\|\bar{\Theta}_{LL} - \bar{\Theta}_{RR}\|_1 - \|\Theta_{LL} - \Theta_{RR}\|_1 = -\sum_{i=1}^p \sum_{\substack{j=1 \\ j \neq i}}^p |\theta_{ij} - \theta_{i'j'}|$. Hence, (B.6) can be written as

$$-\sum_{i=1}^q \sum_{\substack{j=1 \\ j \neq i}}^q \lambda_2 |\theta_{ij} - \theta_{i'j'}| - \sum_{i=1}^q \sum_{j=1}^q \lambda_2 |\theta_{ij'} - \theta_{i'j}|$$

that is trivially non-positive because $\lambda_2 \geq 0$, and this completes the proof.

B.3 Proof of Proposition 4.3

For a positive definite matrix Θ we let $\bar{\Theta}$ be the block diagonal matrix with blocks Θ_{LL} and Θ_{RR} . We remark that $\bar{\Theta}$ is positive definite by construction, and that the condition $|s_{ij}| \leq \lambda_1$ for every $i \in L$ and $j \in R$ can be written as $|s_{ij'}| \leq \lambda_1$ for every $i, j \in L$. Hence, we now show that if $\lambda_1 \geq |s_{ij'}|$ for every $i, j \in L$ then for any positive definite matrix Θ it holds that $\mathcal{L}_{\lambda_1, \lambda_2}(\bar{\Theta}) \leq \mathcal{L}_{\lambda_1, \lambda_2}(\Theta)$ so that the matrix which minimises $\mathcal{L}_{\lambda_1, \lambda_2}(\cdot)$ in (5) must be block diagonal with $\bar{\Theta}_{LR}$ equal to zero. Concretely, we show that $\mathcal{L}_{\lambda_1, \lambda_2}(\bar{\Theta}) - \mathcal{L}_{\lambda_1, \lambda_2}(\Theta) \leq 0$ and, to this aim, we analyse, in turn, three components of this difference.

We first show that $\{-\log \det(\bar{\Theta})\} - \{-\log \det(\Theta)\} \leq 0$. The latter is equivalent to $\log \det(\Theta) \leq \log \det(\bar{\Theta})$, and therefore to $\det(\Theta) \leq \det(\bar{\Theta})$, which follows immediately from Fisher’s inequality; see, for instance, Lütkepohl (1996, p. 54 eqn (5)). We now show that $\text{tr}(S\bar{\Theta}) - \text{tr}(S\Theta) + \lambda_1 \|\bar{\Theta}\|_1 - \lambda_1 \|\Theta\|_1 \leq 0$. Firstly, we notice that

$$\begin{aligned} \text{tr}(S\bar{\Theta}) - \text{tr}(S\Theta) &= \text{tr}(S\bar{\Theta} - S\Theta) = \text{tr}\{S(\bar{\Theta} - \Theta)\} \\ &= -2 \sum_{i=1}^q \sum_{j=1}^q s_{ij'} \theta_{ij'} \leq 2 \sum_{i=1}^q \sum_{j=1}^q |s_{ij'}| |\theta_{ij'}|, \end{aligned}$$

where the third equality follows from the fact that both S and $(\bar{\Theta} - \Theta)$ are symmetric matrices and, more specifically, that both $(\bar{\Theta} - \Theta)_{LL}$ and $(\bar{\Theta} - \Theta)_{RR}$ are equal to zero. Similarly,

$$\lambda_1 \|\bar{\Theta}\|_1 - \lambda_1 \|\Theta\|_1 = -2 \sum_{i=1}^q \sum_{j=1}^q \lambda_1 |\theta_{ij'}|.$$

Hence,

$$\begin{aligned} \text{tr}(S\bar{\Theta}) - \text{tr}(S\Theta) + \lambda_1 \|\bar{\Theta}\|_1 - \lambda_1 \|\Theta\|_1 &= -2 \sum_{i=1}^q \sum_{j=1}^q s_{ij'} \theta_{ij'} - 2 \sum_{i=1}^q \sum_{j=1}^q \lambda_1 |\theta_{ij'}| \end{aligned}$$

$$\begin{aligned} &\leq 2 \sum_{i=1}^q \sum_{j=1}^q |s_{ij'}| |\theta_{ij'}| - 2 \sum_{i=1}^q \sum_{j=1}^q \lambda_1 |\theta_{ij'}| \\ &= 2 \sum_{i=1}^q \sum_{j=1}^q (|s_{ij'}| - \lambda_1) |\theta_{ij'}| \leq 0 \end{aligned}$$

where the latter inequality holds true because, by assumption, $(|s_{ij'}| - \lambda_1) \leq 0$ for every $i, j \in L$.

In order to complete the proof we have now to show that the difference

$$\begin{aligned} &\lambda_2 (\|\bar{\Theta}_{LL} - \bar{\Theta}_{RR}\|_1 + \|\bar{\Theta}_{LR} - \bar{\Theta}_{RL}\|_1) \\ &\quad - \lambda_2 (\|\Theta_{LL} - \Theta_{RR}\|_1 + \|\Theta_{LR} - \Theta_{RL}\|_1) \end{aligned} \tag{B.7}$$

Fig. 3 Coloured graph and concentration matrix of the pdRCON model in Example C.1

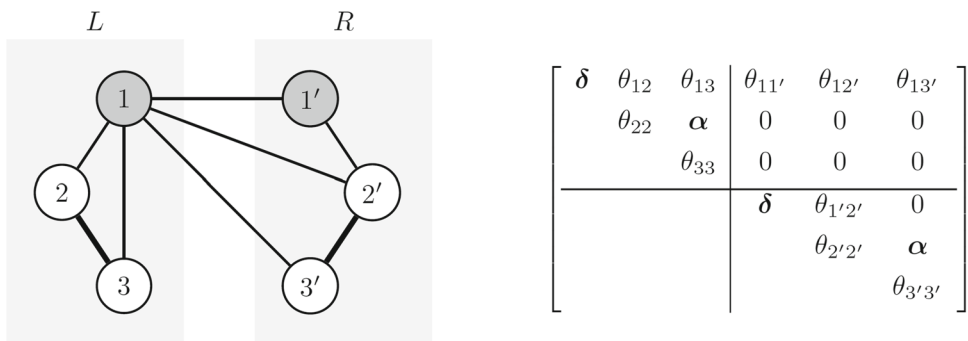


Fig. 4 Coloured graph and concentration matrix of the pdRCON model in Example C.2




Fig. 5 Coloured graph and concentration matrix of the pdRCON model in Example C.3





Fig. 6 Coloured graph and concentration matrix of the pdRCON model in Example C.4



is smaller or equal to zero. This can be shown by noting that $\|\tilde{\Theta}_{LR} - \tilde{\Theta}_{RL}\|_1 = 0$ because $\tilde{\Theta}$ is block diagonal, and that, by construction, $\|\tilde{\Theta}_{LL} - \tilde{\Theta}_{RR}\|_1 = \|\Theta_{LL} - \Theta_{RR}\|_1$ so that $\|\tilde{\Theta}_{LL} - \tilde{\Theta}_{RR}\|_1 - \|\Theta_{LL} - \Theta_{RR}\|_1 = 0$. Hence (B.7) simplifies to $-\lambda_2 \|\Theta_{LR} - \Theta_{RL}\|_1$ that is trivially non-positive because $\lambda_2 \geq 0$.

C Examples of pdRCON models

We present here some examples of pdRCON models with a detailed description of the different types of symmetry they include and of the relevant equality constraints. All the models considered involve $p = 6$ variables so that $L = \{1, 2, 3\}$ and $R = \{1', 2', 3'\}$ and, for each of them, we give both the coloured graph and the concentration matrix. Recall that the coloured graph representing a pdRCON model may contain two types of vertices and edges, namely, coloured and uncoloured. In order to make our graphs readable also in black and white printing, the colour white is used to denote uncoloured vertices whereas coloured vertices are in gray. On the other hand, uncoloured and coloured edges are represented by thin and thick black lines, respectively. Shaded areas are used to highlight the subgraphs \mathcal{G}_L and \mathcal{G}_R relative to the two groups. Finally, in our representation of the concentration matrices, only the upper-triangular part is given and the entries involved in equality constraints are in bold.

Example C.1 (Fig. 3) In the pdRCON model of Fig. 3 the edges $\{2, 3\} \in \mathcal{G}_L$ and $\{2', 3'\} \in \mathcal{G}_R$ are both present and such that $\theta_{23} = \theta_{2'3'} = \alpha$; that is, they form an inside-block parametric symmetry. On the other hand the edges $\{1, 2\} \in \mathcal{G}_L$ and $\{1', 2'\} \in \mathcal{G}_R$ are both present, thereby forming an inside-block structural symmetry, but not a parametric symmetry because the corresponding concentrations θ_{12} and $\theta_{1'2'}$ are not constrained to be equal. In this model there is also one vertex symmetry, encoded by the equality constraints $\theta_{11} = \theta_{1'1'} = \delta$, whereas the only existing across-block symmetries are those that involve missing edges, and thus zero concentrations.

Example C.2 (Fig. 4) In the pdRCON model of Fig. 4 the edges $\{1, 2'\}$ and $\{2, 1'\}$ are both present and form the inside-block parametric symmetry with $\theta_{12'} = \theta_{21'} = \gamma$. On the other hand, the edges $\{2, 3'\}$ and $\{3, 2'\}$ are both present,

thereby forming an across-block structural symmetry, but not a parametric symmetry because the corresponding concentrations $\theta_{23'}$ and $\theta_{32'}$ are not constrained to be equal. In this model, there are neither vertex symmetries nor inside-block symmetries. Indeed, the inner structure of the two groups is very different because \mathcal{G}_L is a fully disconnected whereas \mathcal{G}_R is complete.

Example C.3 (Fig. 5) The pdRCON model of Fig. 5 shows a large amount of symmetry. More specifically, there is (i) full vertex symmetry with $\theta_{11} = \theta_{1'1'} = \delta_1, \theta_{22} = \theta_{2'2'} = \delta_2$ and $\theta_{33} = \theta_{3'3'} = \delta_3$; (ii) full across-block parametric symmetry with $\theta_{12'} = \theta_{21'} = \gamma_1$ and $\theta_{23'} = \theta_{3'2} = \gamma_2$ and, finally, (iii) there is full inside-block structural symmetry. However, there are no inside-block parametric symmetries involving present edges, and therefore this is not a fully symmetric model.

Example C.4 (Fig. 6) The structure of the graph in Fig. 6 is the same as that considered in Example C.3, but in this case there are additional equality constraints so that full parametric symmetry is achieved. It may be worth noting that, although this is a fully symmetric model, thin lines are used to depict the edges $\{1, 1'\}$ and $\{2, 2'\}$. Indeed, these edges are associated with the diagonal entries of Θ_{LR} and, in pdRCON models, these parameters are not considered for possible equality constraints.

D Plots of the two selected models for the breast cancer analysis of Sect. 8

Visual depiction of the two graphs associated to: (i) pdRCON model with forced full vertex symmetry and penalization on inside- and across-blocks (mod_{fV} , Fig. 7); (ii) fully symmetric pdRCON model (mod_{fVIA} , Fig. 8). In both plots, the character coding is the following: *empty circle*, for structural non-zero symmetries (unicode character U+25CB, white circle); *full circle*, for parametric non-zero symmetries (unicode character U+25CF, black circle); *diagonal slash* for asymmetric edges (unicode character U+2571, box drawings light diagonal upper right to lower left). Labels for columns and rows end with the string “_T” if the genes are measured on tumor tissues and “_H” if measured on healthy tissues. Lower triangular portion is not shown due to the symmetric nature of the matrix.




Fig. 8 Depiction of the graph associated to model mod_{FVIA} (fully symmetric model), selected according to eBIC ($\gamma = 0.5$), and fit on the breast cancer data of Sect. 8; for row and column labels, “_T” denotes

tumor tissue and “_H” healthy tissue; character coding: *empty circle*, structural non-zero symmetries; *full circle*, parametric non-zero symmetries; *diagonal slash*, asymmetric edges

Acknowledgements We would like to thank two anonymous referees for their valuable comments, and Alberto Tonolo, Ugo Ala and Sara Ferrini for useful discussion and the support provided in pre-processing the cancer genomics data.

Author Contributions S.R. and A.R. equally contributed to the idea behind the work, and the production of the article itself. This also applies to the prepared material within the manuscript (Figures, Tables). Supplementary material also follows the same contribution statement.

Funding Open access funding provided by Alma Mater Studiorum - Università di Bologna within the CRUI-CARE Agreement. Financial support was provided by the European Union - Next Generation EU, Mission 4 Component 2 - CUP C53D23002580006 (MUR-PRIN grant 2022SMNNKY)

Data availability All implemented scripts and functions used throughout the paper (and more) are available at the following GitHub repository: <https://github.com/savranciati/pdglasso>. For reproducibility reasons, both original data (downloaded from <https://xenabrowser.net/>) and sample covariance matrices are available as part of the R package `pdglasso`, together with their metadata information. Scripts and datasets pertaining to the simulation study are available online as Supplementary Material.

Declarations

Conflict of interest The authors declare no competing interest

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

- Aran, D., Camarda, R., Odegaard, J., Paik, H., Oskotsky, B., Krings, G., Goga, A., Sirota, M., Butte, A.J.: Comprehensive analysis of normal adjacent to tumor transcriptomes. *Nat. Commun.* **8**(1), 1–14 (2017)
- Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* **3**(1), 1–122 (2011)
- Carter, J.S., Rossell, D., Smith, J.Q.: Partial correlation graphical LASSO. *Scand. J. Stat.* **51**(1), 32–63 (2024)
- Danaher, P., Wang, P., Witten, D.M.: The joint graphical lasso for inverse covariance estimation across multiple classes. *J. Royal Stat. Soc.: Series B (Statistical Methodology)* **76**(2), 373–397 (2014)
- Dey, D.K., Srinivasan, C.: Estimation of a covariance matrix under Stein's loss. *Ann. Stat.* **13**(4), 1581–1591 (1985)
- Foygel, R., Drton, M.: Extended Bayesian information criteria for Gaussian graphical models. *Adv. Neural Inf. Process Syst.* **23**, 604–612 (2010)
- Friedman, J., Hastie, T., Höfling, H., Tibshirani, R.: Pathwise coordinate optimization. *Annals Appl. Stat.* **1**(2), 302–332 (2007)
- Friedman, J., Hastie, T., Tibshirani, R.: Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9**(3), 432–441 (2008)
- Gehrmann, H.: Lattices of graphical Gaussian models with symmetries. *Symmetry* **3**(3), 653–679 (2011)
- Graczyk, P., Ishi, H., Kołodziejek, B., Massam, H.: Model selection in the space of Gaussian models invariant by symmetry. *Ann. Stat.* **50**(3), 1747–1774 (2022)
- Hastie, T., Tibshirani, R., Wainwright, M.: *Statistical Learning with Sparsity: The Lasso and Generalizations*. CRC Press (2015)
- Hoefling, H.: A path algorithm for the fused lasso signal approximator. *J. Comput. Graph. Stat.* **19**(4), 984–1006 (2010)
- Højsgaard, S., Lauritzen, S.L.: Graphical Gaussian models with edge and vertex symmetries. *J. Royal Stat. Soc: Series B (Statistical Methodology)* **70**(5), 1005–1027 (2008)
- Højsgaard, S., Lauritzen, S.L., et al.: Inference in graphical Gaussian models with edge and vertex symmetries with the `gRc` package for R. *J. Stat. Softw.* **23**(6), 1–26 (2007)
- Kubo, M., Nakamura, M., Tasaki, A., Yamanaka, N., Nakashima, H., Nomura, M., Kuroki, S., Katano, M.: Hedgehog signaling pathway is a new therapeutic target for patients with breast cancer. *Can. Res.* **64**(17), 6071–6074 (2004)
- Lauritzen, S.L.: *Graphical models*. Oxford University Press (1996)
- Lütkepohl, H.: *Handbook of Matrices*. John Wiley and Sons, New York (1996)
- Mazumder, R., Hastie, T.: Exact covariance thresholding into connected components for large-scale graphical lasso. *J. Mach. Learn. Res.* **13**(27), 781–794 (2012)
- R Core Team: *R: A Lang. Environ. Stat. Computing*. R Foundation for Statistical Computing, Vienna, Austria (2023)
- Ranciati, S., Roverato, A., Luati, A.: Fused graphical lasso for brain networks with symmetries. *J. Roy. Stat. Soc.: Ser. C (Appl. Stat.)* **70**(5), 1299–1322 (2021)
- Roverato, A., Nguyen, D. N.: Model inclusion lattice of coloured Gaussian graphical models for paired data. In A. Salmerón and R. Rumi (Eds.), *Proceedings of the 11th International Conference on Probabilistic Graphical Models*, Volume 186 of *Proceedings of Machine Learning Research*, pp. 133–144. PMLR (2022)
- Roverato, A., Nguyen, D.N.: Exploration of the search space of Gaussian graphical models for paired data. *J. Mach. Learn. Res.* **25**(92), 1–41 (2024)
- Song, D., Cui, M., Zhao, G., Fan, Z., Nolan, K., Yang, Y., Lee, P., Ye, F., Zhang, D.Y.: Pathway-based analysis of breast cancer. *American journal of translational research* **6**(3), 302 (2014)
- Subramanian, A., Tamayo, P., Mootha, V.K., Mukherjee, S., Ebert, B.L., Gillette, M.A., Paulovich, A., Pomeroy, S.L., Golub, T.R., Lander, E.S., Mesirov, J.P.: Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proc. Natl. Acad. Sci.* **102**(43), 15545–15550 (2005)
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., Knight, K.: Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **67**(1), 91–108 (2005)
- Tibshirani, R.J., Taylor, J.: The solution path of the generalized lasso. *Ann. Stat.* **39**(3), 1335–1371 (2011)
- Tsai, K., Koyejo, O., Kolar, M.: Joint Gaussian graphical model estimation: A survey. *Wiley Interdisciplinary Reviews: Computational Statistics* **14**(6), e1582 (2022)
- Witten, D.M., Friedman, J.H., Simon, N.: New insights and faster computations for the graphical lasso. *J. Comput. Graph. Stat.* **20**(4), 892–900 (2011)
- Xie, Y., Liu, Y., Valdar, W.: Joint estimation of multiple dependent Gaussian graphical models with applications to mouse genomics. *Biometrika* **103**(3), 493–511 (2016)
- Yuan, M., Lin, Y.: Model selection and estimation in the Gaussian graphical model. *Biometrika* **94**(1), 19–35 (2007)

Zhang, H., Huang, X., Arshad, H.: Comparing dependent undirected Gaussian networks. *Bayesian Anal.* **1**, 1–26 (2022)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.