Mixed scale joint graphical lasso

EUGEN PIRCALABELU, GERDA CLAESKENS*

KU Leuven, ORSTAT and Leuven Statistics Research Center, Naamsestraat 69, 3000 Leuven, Belgium
gerda.claeskens@kuleuven.be

LOURENS J. WALDORP

University of Amsterdam, Department of Psychological Methods, Weesperplein 4,1018 Amsterdam, The Netherlands

SUMMARY

We have developed a method for estimating brain networks from fMRI datasets that have not all been measured using the same set of brain regions. Some of the coarse scale regions have been split in smaller subregions. The proposed penalized estimation procedure selects undirected graphical models with similar structures that combine information from several subjects and several coarseness scales. Both within-scale edges and between-scale edges that identify possible connections between a large region and its subregions are estimated.

Keywords: Fused and group lasso; Joint graphical lasso; l_1 penalization; Mixed scale data; Sparsistency.

1. INTRODUCTION

This work is motivated by the need to jointly estimate multiple undirected graphs from data that do not all have the same measurement scale. The example used here originates from resting state functional magnetic resonance imaging (rsfMRI) where brain activity measurements at voxel level were recorded for two subjects. The rsfMRI images were segmented first in 68 atlas-based regions of interest (ROIs, i.e., sets of voxels that form nonoverlapping parts of the brain), corresponding to a coarse measurement scale with anatomically large brain regions. Based on the procedure of Hagmann *and others* (2008), several of the large ROIs were further split in two or more smaller regions, resulting in a finer scale with 114 ROIs. For each subject, for all ROIs at each scale, $n = 240$ volumes of the blood oxygenation level dependent (BOLD) signal have been obtained. The researcher has knowledge about which ROIs have been split, and a similar functional behavior for a region and its subregions is expected. The setup bears some similarity to the multiresolution framework of Choi and Willsky (2007) and Choi *and others* (2010). However, by nature of the splitting process, the measurement for the coarse region cannot be algebraically reconstructed from the measurements of its subregions.

The objectives are twofold. First, we wish to jointly estimate sparse undirected graphs that show cerebral pathways between ROIs, where each ROI is associated with a node in the graph, using all the available data at coarse and finer scales. Analyzing data available at only one scale would make the implicit assumption that the chosen scale is in a sense "best", and it has been shown that the definition of the ROIs

---

*To whom correspondence should be addressed.

has a high impact on the estimated graph (see, e.g., Zalesky *and others*, 2010; Schmittmann *and others*, 2015). Our procedure avoids such a selection. To overcome the differences in the dimensionality between measurement scales, our procedure introduces two algebraic operators that make use of the knowledge about which finer scale regions are subregions of a coarse scale region. Second, dependencies exist between a coarse scale node and its corresponding finer scale nodes due to the experimental design. This we refer to as "splits". Our procedure is directed at estimating such splits. These splits are of interest because if the subregions are all connected, then it seems reasonable to assume that these subregions act as one, and so the region one scale up is good enough for modeling brain connectivity. In this sense, estimating different scales simultaneously with our proposed method indicates a data-driven mixture of scales to use for modeling connectivity.

The graphs are estimated by enforcing sparsity via ℓ_1-based penalization. For estimating undirected graphs, see Yuan and Lin (2007), Banerjee *and others* (2008), Friedman *and others* (2008), Bickel and Levina (2008a,b), Cai *and others* (2010), and Leng and Tang (2012) among others. See also Bühlmann and Van De Geer (2011) and Hastie *and others* (2015) for regularization based on ℓ_1-penalties and generalizations. Joint estimation of multiple sparse graphs has been studied by Guo *and others* (2011), Danaher *and others* (2014), Gaskins and Daniels (2013), Zhao *and others* (2014), and Mohan *and others* (2014). Fan *and others* (2016) give an overview.

We estimate graphs that show interactions between the regions both "within" (to reveal brain pathways between ROIs) and "between" coarseness scales (to reveal dependencies between coarser and finer nodes). The method can accommodate data from more than one subject, more than two measurement scales and an unequal numbers of splits for different ROIs. Important for the method to work is the information on which regions have been split, how many splits they have at each scale, and whether the measurements for coarser ROIs are not deterministically obtained from the partitioned ones. We constrain the graphs for the different scales to be "similar" to each other by using the "fused" or "group" graphical lasso penalties as in Danaher *and others* (2014). The current setting is different in that first their method requires to have the same measurement scale to combine graphs, and second we allow for connections between the split and unsplit ROIs.

2. Notation

We associate each random variable with a node in an undirected graph $G(E, V)$, where $V = \{1, \ldots, p\}$ represents the set of nodes and E is the set of undirected edges $i - j$ between a pair of nodes (i, j). A superscript denotes the measurement scale. Let $X^{(k)}$ be the random vector of variables that correspond to the ROIs for scale k, with length $q^{(k)}$. When $k = 1$, $X^{(1)}$ collects all variables at the coarse scale. All other scales are splits of the coarsest scale. We call scale $k > 1$ finer than scale 1, or conversely scale 1 is coarser than scale k.

Define the vector X and its concentration matrix Θ, both partitioned according to the lengths $q^{(k)}$, $k = 1, \ldots, K$, where the submatrices Θ^{kk} are the inverse covariance matrices corresponding to each of the K coarseness scales, while the matrices $\Theta^{kk'}$ with $k \neq k'$ are the across-scale inverse covariance matrices between scales k and k'. If an element $\theta_{i,j} \neq 0$, then an edge links nodes i and j in the graph $G(E, V)$, and there is thus a one-to-one correspondence between Θ and $G(E, V)$,

$$X = (X^{(1)^\mathsf{T}}, X^{(2)^\mathsf{T}}, \ldots, X^{(K)^\mathsf{T}})^\mathsf{T} \sim N(\mathbf{0}, \Sigma), \ \Sigma^{-1} \equiv \Theta \equiv \begin{pmatrix} \Theta^{11} & \ldots & \Theta^{1K} \\ \vdots & \ddots & \vdots \\ \Theta^{K1} & \ldots & \Theta^{KK} \end{pmatrix}.$$

We assume that Σ is invertible, and hence components at different scales cannot be algebraically related. We estimate the matrices Θ^{kk} in such a way that they are sparse and "similar" to each other, since they

Fig. 1. An illustrative example involving ROIs 4 and 28 at two coarseness scales. Region 28 is split, first in two, then in three subregions, whereas region 4 is unsplit at both scales. Panel (a) shows edges at the coarse scale, (b) shows edges at the fine scale, and (c) shows edges between coarse regions and finer subregions. Note that the shape and size of the drawn regions do not correspond to the actual irregular shape of the anatomical regions in the data example. This figure serves as an illustration of the splitting process only.

differ only in the measurement scale. The nonzero elements in Θ^{kk} are interpreted as the "within" scale edges. The non-zero elements in the off-diagonal matrices $\Theta^{kk'}$ ($k' \neq k$) are interpreted as the "between" scale edges. The between scale graphs are constrained to be sparse as we only allow for the estimation of edges corresponding to splits of the same regions across the scales, that is, coarse region i is only allowed to connect to its subregions and not to subregions of another coarse region. Since we know exactly which regions are split, the sparsity pattern of the off-diagonal matrices is partially known, while this is unknown for the matrices on the diagonal. Figure 1 shows an example using ROIs 4 and 28 at two scales. At both scales, ROI 4 is unsplit, whereas ROI 28 is split first in two and next in three subregions. The matrix Θ^{11} (within edges at the coarse scale) is represented in (a); (b) depicts Θ^{22}. The between scale edges, or splits, contained in Θ^{12} are in (c). The goal is to estimate all these edges.

Denote by $k \in \{1, \ldots, K\}$ the coarseness scale, by $h \in \{1, \ldots, H\}$ a coarse ROI, and by $l \in \{1, \ldots, L\}$ a partition of a coarse ROI. The couple $\{k, h(l)\}$ denotes partition l of region h at scale k. In Figure 1, $\{2, 28(2)\}$ refers to the second partition of ROI 28 at the second scale. There are K scales and H regions at the coarsest scale. The number of partitions within a scale is denoted by L_k and can vary (e.g. at scale $k = 3$, we can have more splits of the region h than at scale $k = 2$). The number of partitions can vary for each region (e.g., h' can have more splits than h). If a coarse ROI has not been split, the region is used unchanged in the other scales.

The dimensions of each submatrix Θ^{kk} are dictated by the number of splits encountered at scale k and can all be different. To induce similarity between scales, we couple a region that is split into, say, three regions at a finer scale to the original single region. This coupling requires that we have the same dimension in both scales (three in this case). To tackle this dimensionality problem, we introduce in Section 5 the "expand" operator, with the objective of transforming all submatrices to a common dimension that is most informative.

The set $\mathcal{W} = \{\theta_{i,j}^{kk}, k = 1, \ldots, K; i, j = 1, \ldots, q^{(k)}\}$ contains all elements of the block matrices on the main diagonal. If an element of $\mathcal{W}$ is nonzero, an undirected edge is present between different regions *within* a scale. In the example of Figures 1(a) and (b), $\mathcal{W}$ corresponds to the parameters related to edges using nodes from a fixed scale. *Within* a scale, a node may connect to all other nodes within the same scale. *Between* scales, coarser nodes can connect only to their own finer partitions. The set $\mathcal{B} = \{\theta_{i,j}^{kk'} | k \neq k', i$ corresponds to $h(l)$ and j corresponds to $h(l'), \forall h$ and $l, l' = 1, \ldots, L_k$ and $k, k' = 1, \ldots, K\}$ contains all entries of off-diagonal submatrices. In Figure 1(c), $\mathcal{B}$ corresponds to parameters related to edges

between the partitions of an ROI across scales. This reflects interest in parameters that (if non-zero) correspond to an undirected edge *between* different splits (l and l') of a region (h) across different scales (k and k').

3. ESTIMATION METHOD

Since the number of nodes at each measurement scale is not the same due to the splitting, the vector lengths $q^{(k)}$, $k = 1, \ldots, K$ are also different. As a first step, we bring all components $\boldsymbol{X}^{(k)}$ to the same length by applying the "expand" operator. When a certain coarse region has been split in, say, three subregions in its finest measurement scale, the expand operator repeats the measurement for that region three times. As a result, all expanded vectors have the same length. The mathematical details about this operator are given in Section 5.

Let $\boldsymbol{Y} = \mathrm{Ex}(\boldsymbol{X}, D)$ be the expanded vector based on the measurements $\boldsymbol{X}$ and the knowledge of the region splits D. For a sample $\boldsymbol{Y}_1, \ldots, \boldsymbol{Y}_n$, denote by $\boldsymbol{S}_{\mathrm{Ex}}$ the empirical variance matrix and let $\boldsymbol{\Sigma}_{\mathrm{Ex}}$ be the true variance matrix of $\boldsymbol{Y}$. We minimize the following objective function

$$q(\boldsymbol{\Psi}) = \left\{ \mathrm{trace}(\boldsymbol{S}_{\mathrm{Ex}} \boldsymbol{\Psi}) - \log \det \boldsymbol{\Psi} + \sum_{i \neq j} p_{\lambda_{n1}, \lambda_{n2}}(|\psi_{i,j}^{11}|, \ldots, |\psi_{i,j}^{KK}|) \right\}, \tag{3.1}$$

over symmetric positive definite matrices $\boldsymbol{\Psi}$, having the same dimension as $\boldsymbol{\Sigma}_{\mathrm{Ex}}$ and acting as a pseudo-concentration matrix since $\boldsymbol{\Sigma}_{\mathrm{Ex}}$ is not invertible due to singularities induced by the expand operator, hence $\boldsymbol{\Psi}^{-1} \neq \boldsymbol{\Sigma}_{\mathrm{Ex}}$. $\boldsymbol{\Psi}$ contains "partial parameters" that are averaged by the reduce operator to obtain the parameters of interest.

The penalty function $p_{\lambda_{n1}, \lambda_{n2}}$ is a convex real valued function depending on two regularization values ($\lambda_{n1}, \lambda_{n2}$) that forces small entries to be shrunk to zero through λ_{n1}, while enforcing similarity between subgraphs from different coarseness scales through λ_{n2}. For notational simplicity, we use $p_{\lambda_{n1}, \lambda_{n2}}(|\psi_{i,j}^{11}|, \ldots, |\psi_{i,j}^{KK}|)$ to denote $p_{\lambda_{n1}, \lambda_{n2}}(|\psi_{i,j}^{11}|, \ldots, |\psi_{i,j}^{1K}|, |\psi_{i,j}^{22}|, \ldots, |\psi_{i,j}^{2K}|, \ldots, |\psi_{i,j}^{KK}|)$. To ensure similarity between the concentration submatrices $\boldsymbol{\Theta}^{kk}$, we use a "fused" (FGL) or a "group" (GGL) graphical lasso penalty as in Danaher *and others* (2014). For the fused lasso penalty, see Tibshirani *and others* (2005), Höfling *and others* (2010), Liu *and others* (2010), and Yang *and others* (2015). The group lasso penalty has been introduced by Yuan and Lin (2006) as a form of shrinkage that allows certain groups of parameters to be jointly estimated as zero or nonzero values. The two penalties in our context take the form,

$$p_{\lambda_{n1}, \lambda_{n2}}^{FGL}(|\psi_{i,j}^{11}|, \ldots, |\psi_{i,j}^{KK}|) = \lambda_{n1} \sum_k \sum_{k'} |\psi_{i,j}^{kk'}| + \lambda_{n2} \sum_k \sum_{k' > k} |\psi_{i,j}^{kk} - \psi_{i,j}^{k'k'}|,$$

$$p_{\lambda_{n1}, \lambda_{n2}}^{GGL}(|\psi_{i,j}^{11}|, \ldots, |\psi_{i,j}^{KK}|) = \lambda_{n1} \sum_k \sum_{k'} |\psi_{i,j}^{kk'}| + \lambda_{n2} \left\{ \sum_k (\psi_{i,j}^{kk})^2 \right\}^{1/2},$$

where $\{\psi_{i,j}^{11}, \ldots, \psi_{i,j}^{KK}\}$ are the elements in the expanded matrix that correspond to the unexpanded elements $\{\theta_{i,j}^{11}, \ldots, \theta_{i,j}^{KK}\}$ from the set $\mathcal{W} \cup \mathcal{B}$. If there is an ROI that has been split, the cardinality of $\{\psi_{i,j}^{11}, \ldots, \psi_{i,j}^{KK}\}$ is larger than that of $\{\theta_{i,j}^{11}, \ldots, \theta_{i,j}^{KK}\}$. For both penalties, the first term regularizes all allowed edges (both within and between scales), while the second part, related to λ_{n2}, regularizes the similarity of the edges between the scales. All entries of $\boldsymbol{\Psi}$ for which the corresponding entries in $\boldsymbol{\Theta}$ are not in $\mathcal{W} \cup \mathcal{B}$ are defined as 0 due to considering only the connections between ROIs and their split versions as enforced by the design.

FGL penalizes the differences between matrix entries, thus making them more similar to each other, while GGL encourages that nonzero entries of the concentration matrices occur at the same places. By setting entries nonzero at the same positions across the coarseness scales, the group penalty also enforces the within group matrices to be similar to each other. When the ordering of the scales is important, the fused penalty is more appropriate. When the group formed by the coarse regions and all its splits is of interest, the group penalty might be more appropriate. Both penalties encourage shared sparsity patterns across the different scales, and as a consequence, the matrices or rather their graph representations are close together. They have both been applied, so far, only for cases with the same coarseness, whereas we extend these ideas to different scales of coarseness. Using Danaher *and others* (2014), one cannot simultaneously estimate several sparse graphs that allow between-coarseness scale edges. Moreover, their technique is constructed for the case where the number of regions for both scales is the same, which for our problem is not the case, as we make a clear distinction between coarser and finer regions.

By optimizing (3.1), we obtain an estimated matrix which has a larger dimension than desired. To bring all estimated matrices back to their original dimensions, we apply the "reduce" operator on $\hat{\mathbf{\Psi}}$ (see Section 5 for details). Roughly, if a coarse scale region is expanded in three subregions, the estimates for the three regions are combined in a (weighted) average, a single estimate, for the coarse scale region. We define the resulting reduced matrix as the mixed scale joint graphical lasso (msJGL) estimate of the concentration matrix, since it jointly estimates the between- and within-scale edges across all coarseness scales.

4. ALGORITHM FOR THE MIXED SCALE JOINT GRAPHICAL LASSO

The proposed algorithm, based on the ADMM algorithm presented in Boyd *and others* (2011), allows one part of the $\mathbf{\Psi}$ matrix to be estimated using the FGL or GGL penalty, while the other part of the matrix is estimated using a regular ℓ_1-penalty. All off-diagonal elements of $\mathbf{\Psi}$ corresponding to the set $\mathcal{B}$ receive only the ℓ_1-penalty, while all elements of the submatrices on the main diagonal of $\mathbf{\Psi}$, or equivalently the elements in the set $\mathcal{W}$, receive both the ℓ_1-penalty and either the group or fused penalty. The structural 0s in the off-diagonal matrices are obtained by taking a large enough penalty such that all entries corresponding to *unallowed* edges in $G(E, V)$ are set to zero. Let $\mathbf{I}$ and $\mathbf{0}$ be the identity and the null matrix of the appropriate dimension. The algorithm is described as follows.

Step 1: Apply the "expand" operator to construct $\mathbf{S}_{\mathrm{Ex}}$ and set $\hat{\mathbf{\Psi}} = \mathbf{I}$, $\mathbf{U} = \mathbf{0}$, $\mathbf{Z} = \mathbf{0}$.

Step 2: Update $\hat{\mathbf{\Psi}} = \arg\min_{\mathbf{\Psi}}(\mathrm{trace}(\mathbf{S}_{\mathrm{Ex}}\mathbf{\Psi}) - \log\det\mathbf{\Psi} + \frac{\rho}{2}||\mathbf{\Psi} - \mathbf{Z} + \mathbf{U}||_F^2)$, where $||\cdot||_F$ is the Frobenious norm. The solution is obtained in closed form. Compute the eigen decomposition $\rho(\mathbf{Z} - \mathbf{U}) - \mathbf{S}_{\mathrm{Ex}} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^\mathsf{T}$, where $\rho > 0$ is a fixed scalar and form a diagonal matrix with $\tilde{X}_{ii} = \{\mathrm{eig}_i + (\mathrm{eig}_i^2 + 4\rho)^{1/2}\}/(2\rho)$, where eig_i is the i-th eigenvalue. Update $\hat{\mathbf{\Psi}} = \mathbf{Q}\tilde{X}\mathbf{Q}^\mathsf{T}$.

Step 3: Split $\hat{\mathbf{\Psi}}$, $\mathbf{U}$, $\mathbf{Z}$ into block matrices according to the number of scales.

Step 4: Update the block submatrices $\mathbf{Z}^{kk}$ with $k = 1, \ldots, K$, for which the FGL or GGL penalty is used, $\mathbf{Z}^{kk} = \arg\min_{\mathbf{Z}^{kk}} \sum_{k=1}^{K} \frac{\rho}{2}||\hat{\mathbf{\Psi}}^{kk} - \mathbf{Z}^{kk} + \mathbf{U}^{kk}||_F^2 + \sum_{i \neq j} p_{\lambda_{n1}, \lambda_{n2}}(|z_{i,j}^{11}|, \ldots, |z_{i,j}^{KK}|)$. The $\mathbf{Z}^{kk}$ can take different values for the FGL and GGL penalties, but in both cases this amounts to applying the soft-thresholding operator on a linear combination of matrices.

Step 5: Update $\mathbf{Z}^{kk'}$ with $k \neq k'$ for which the ℓ_1-penalty is used as $\mathbf{Z}^{kk'} = \mathrm{Soft}_{\lambda_{n1}/\rho}(\mathbf{\Psi}^{kk'} + \mathbf{U}^{kk'})$, where $\mathrm{Soft}_{\lambda_{n1}/\rho}$ is the soft-threshold operator using λ_{n1}/ρ as thresholding value.

Step 6:   Update $U = \hat{\Psi} + U - Z$ and repeat steps 2–6 until convergence.

Step 7:   "Reduce" $\hat{\Psi}$ to obtain $\hat{\Theta}_{\text{Red}}$, the msJGL estimated concentration matrix.

5. THE EXPAND AND REDUCE OPERATORS

We denote by $\boldsymbol{a} \circ \boldsymbol{b}$ the elementwise and by $\boldsymbol{a} \otimes \boldsymbol{b}$ the Kronecker product between the vectors $\boldsymbol{a}$ and $\boldsymbol{b}$. By $\mathbf{1}_q$ we denote a vector of length q with elements equal to 1. The vector $\boldsymbol{s}^{(k)}$ records for each coarsest region the number of splits at scale k. Let $d_h; h = \{1, \ldots, H\}$ be the maximal number of edges that can be set between splits of a region h from different measurement scales, hereby accounting for the hierarchical structure by only counting edges between consecutive scales. Let $\boldsymbol{X}^{(k)}_{(h)}$ be the subvector of splits of region h for scale k. Further, let $v_h^{(k)} = d_h/\#\boldsymbol{X}^{(k)}_{(h)}$, where the symbol # denotes the number of elements in the vector.

DEFINITION 1   Let $\boldsymbol{X} = (\boldsymbol{X}^{(1)\mathsf{T}}, \boldsymbol{X}^{(2)\mathsf{T}}, \ldots, \boldsymbol{X}^{(K)\mathsf{T}})^{\mathsf{T}}$. Define a design vector $\boldsymbol{D} = \boldsymbol{s}^{(1)} \circ \boldsymbol{s}^{(2)} \circ \ldots \circ \boldsymbol{s}^{(K)} \equiv (d_1, d_2, \ldots, d_H)^{\mathsf{T}}$ that records the product of the number of splits across all coarseness scales. The expand operator produces the vector $\text{Ex}(\boldsymbol{X}, \boldsymbol{D})$ of length $\sum_{h=1}^{H} d_h K$, which is defined as $\text{Ex}(\boldsymbol{X}, \boldsymbol{D}) = \left(\text{Ex}(\boldsymbol{X}^{(1)}, \boldsymbol{D}), \text{Ex}(\boldsymbol{X}^{(2)}, \boldsymbol{D}), \ldots, \text{Ex}(\boldsymbol{X}^{(K)}, \boldsymbol{D})\right)^{\mathsf{T}}$, where $\text{Ex}(\boldsymbol{X}^{(k)}, \boldsymbol{D}) = \left((\boldsymbol{X}^{(k)}_{(1)} \otimes \mathbf{1}_{v_1^{(k)}})^{\mathsf{T}}, (\boldsymbol{X}^{(k)}_{(2)} \otimes \mathbf{1}_{v_2^{(k)}})^{\mathsf{T}}, \ldots, (\boldsymbol{X}^{(k)}_{(h)} \otimes \mathbf{1}_{v_h^{(k)}})^{\mathsf{T}}\right), \ k = 1, \ldots, K$.

This operator ensures that all submatrices $\boldsymbol{\Psi}^{kk'}$ have the same dimension, which, in turn, allows to enforce similarity between coarse and finer scales and to connect splits across scales.

DEFINITION 2   Let $\boldsymbol{\Psi}$ be an expanded matrix and $\boldsymbol{P}$ a projection matrix of dimension $\sum_{k=1}^{K} \sum_{h=1}^{H} s_h^{(k)} \times \sum_{h=1}^{H} d_h K$. The number of rows corresponds to the length of the unexpanded vector $\boldsymbol{X}$, while the number of columns corresponds to the length of the expanded vector $\text{Ex}(\boldsymbol{X}, \boldsymbol{D})$. Let $\boldsymbol{\Theta}_{\text{Red}} = \text{Red}(\boldsymbol{\Psi}) = \boldsymbol{P}\boldsymbol{\Psi}\boldsymbol{P}^{\mathsf{T}}$ be the reduced matrix. $\boldsymbol{P}$ is constructed by placing $1/v_h^{(k)}$ on the row corresponding to the unexpanded region h from scale k and *all* the columns for the expanded region.

Using $1/v_h^{(k)}$ in the matrix $\boldsymbol{P}$ assigns equal weights to all splits of a region. Weighted averages where splits get different weights can be obtained by using other values. The matrix $\boldsymbol{\Theta}_{\text{Red}}$ has the same dimension and structure as $\boldsymbol{\Theta}$. Only the entries that have been expanded in $\boldsymbol{\Psi}$ are reduced, as all other entries remain unchanged.

The reduce operator is not rank preserving but has the following properties: (i) $\boldsymbol{\Theta}_{\text{Red}}$ contains a 0 on position (i, j) if all entries in $\boldsymbol{\Psi}$ pertaining to the couple (i, j) are also 0 and (ii) if $\boldsymbol{\Psi}$ is symmetric and full rank, then also $\boldsymbol{\Theta}_{\text{Red}}$ is a symmetric full rank matrix. Property (i) is similar to the "OR" rule of Meinshausen and Bühlmann (2006). If the entries in the expanded matrix are nonzero and do not cancel each other, then the reduced entry will also be nonzero. Property (ii) holds because the reduce operator is equivalent to applying elementary operations on the rows and columns of $\boldsymbol{\Psi}$, which can be organized such that $\boldsymbol{\Theta}_{\text{Red}}$ becomes the upper-left submatrix and the off-diagonal blocks are each other's transpose. Following Proposition 16.2 from Gallier (2011, pp. 435), if the matrix $\boldsymbol{\Psi}$ is positive definite, also $\boldsymbol{\Theta}_{\text{Red}}$ is positive definite, implying full rank. The symmetry of $\boldsymbol{\Theta}_{\text{Red}}$ follows that of $\boldsymbol{\Psi}$.

6. THEORETICAL PROPERTIES

Due to the eigenvalue decomposition of the expanded matrix, the complexity of the algorithm in Section 4 is of the order $O\left((\sum_{h=1}^{H} d_h K)^3\right)$. Danaher *and others* (2014) and Witten *and others* (2011) investigated

improvements in computational speed when the concentration matrix is block diagonal. In this case, one can apply within each scale the FGL, GGL, or graphical lasso (GL) on only a smaller subset of variables. A similar argument cannot be made for the msJGL, due to the design of the problem: The off-diagonal elements provide dependence information between larger and smaller partitions of the same anatomical ROI.

The objective function (3.1) is closed, proper, and convex, and we assume that its Lagrangian has a saddle point, then the ADMM algorithm is guaranteed to converge (Boyd *and others*, 2011).

Let $S_n = \{(i,j)|\psi_{i,j}^{kk',0} \neq 0; k, k' = 1, \ldots, K\}$, where the true expanded matrix is $\mathbf{\Psi}^0 = (\psi_{i,j}^{kk',0})$, and let $s_n = \#S_n - p_n$ denote the number of off-diagonal nonzero elements in S_n. Let $\mathbf{\Phi}^0 \equiv (\mathbf{\Psi}^0)^{-1}$ be the inverse of the pseudo-concentration matrix $\mathbf{\Psi}^0$. We stress that $\mathbf{\Phi}^0$ is not a proper covariance matrix but a pseudo-covariance matrix. Assume that (a) there exist constants τ_1 and τ_2 such that $0 < \tau_1 < \mathrm{eig_{min}}(\mathbf{\Phi}^0) < \mathrm{eig_{max}}(\mathbf{\Phi}^0) < \tau_2 < \infty$; (b) the sequences $a_n = \max_{i,j \in S_n}\max_{k,k'=1,\ldots,K} \left(|(\partial p_{\lambda_{n1},\lambda_{n2}}(|\psi_{i,j}^{11,0}|, \ldots, |\psi_{i,j}^{KK,0}|))/(\partial|\psi_{i,j}^{kk'}|)|\right) = O\left(\{(p_n/s_{n+1} + 1)\log p_n/n\}^{1/2}\right)$ and $b_n = \max_{i,j \in S_n}\max_{k,k',k'',k'''=1,\ldots,K} \left(|(\partial^2 p_{\lambda_{n1},\lambda_{n2}}(|\psi_{i,j}^{11,0}|, \ldots, |\psi_{i,j}^{KK,0}|))/(\partial|\psi_{i,j}^{kk'}|\partial|\psi_{i,j}^{k''k'''}|)|\right) = o(1)$, and (c) for any nonrandom matrices $\mathbf{A}$, $\mathbf{B}$ for which the operator norm $\|\mathbf{A}\| = O(1)$, $\|\mathbf{B}\| = O(1)$, the quantity $\max_{i,j} |(\mathbf{A}(\mathbf{S}_{\mathrm{Ex}} - \mathbf{\Phi}^0)\mathbf{B})_{i,j}| = O_p((\log p_n/n)^{1/2})$.

Condition (a) guarantees that the eigenvalues of $\mathbf{\Phi}^0$ are well behaved. The sequence a_n in condition (b) is connected to the bias of estimating nonzero entries and represents the maximal value over indices (i,j) of any of the K components in the partial derivative vector. And condition (c) mimics the result of Lemma 2 from Lam and Fan (2009) for the extended matrices.

PROPOSITION 1 *Under conditions (a)–(c) if (i) $n^{-1}\log p_n = O(\lambda_{n1}^2)$ and (ii) $(p_n + s_n)n^{-1}(\log p_n)^k = O(1)$ for some $k > 1$, there exists a minimizer $\hat{\mathbf{\Psi}}$ such that $||\hat{\mathbf{\Psi}} - \mathbf{\Psi}^0||_F^2 = O\{(p_n + s_n)n^{-1}\log p_n\}$.*

PROPOSITION 2 *Under conditions (a)–(c) and conditions of Proposition 1 for any local minimizer that satisfies $||\hat{\mathbf{\Psi}} - \mathbf{\Psi}^0||_F^2 = O_P\{(p_n + s_{n1})n^{-1}\log p_n\}$ and $||\hat{\mathbf{\Psi}} - \mathbf{\Psi}^0||^2 = O_P(\eta_n)$ for a sequence $\eta_n \to 0$, if the sequence $\{\sqrt{n^{-1}\log(p_n)} + \sqrt{\eta_n} + \lambda_{n2}\frac{\theta_{i,j}^{kk}}{\sqrt{\sum_{k=1}^K(\theta_{i,j}^{kk})^2}}I(k = k')\} = O(\lambda_{n1})$ for the GGL penalty or if $\{\sqrt{n^{-1}\log(p_n)} + \sqrt{\eta_n} + \lambda_{n2}\sum_{k''>k}\mathrm{sgn}(\theta_{i,j}^{kk} - \theta_{i,j}^{k''k''})I(k = k')\} = O(\lambda_{n1})$ for the FGL penalty, then with probability tending to 1, $\hat{\psi}_{ij}^{kk'} = 0$ for all $(i,j) \in S^c$.*

Proof. The proofs follow from Theorems 1 and 2 from Lam and Fan (2009) which follow the lines of Rothman *and others* (2008) and Bickel and Levina (2008a,b). We show that (i) $P(\inf_{U \in \mathcal{A}} q(\mathbf{\Psi}^0 + \mathbf{\Delta}_U) > q(\mathbf{\Psi}^0)) \to 1$ which implies that there exists a minimizer in the set $\{\mathbf{\Psi}^0 + \mathbf{\Delta}_U : ||\mathbf{\Delta}_U||_F^2 \leq C_1^2\alpha_n^2 + C_2^2\beta_n^2\}$ such that $||\hat{\mathbf{\Psi}} - \mathbf{\Psi}^0||_F^2 = O_p(\alpha_n + \beta_n)$ and (ii) that the sign of $\frac{\partial q(\mathbf{\Psi})}{\partial \psi_{ij}^{kk'}}$ when $i, j \in S^c$, depends only on $\mathrm{sgn}(\psi_{ij}^{kk'})$ which implies that $\forall k, k'$ $\hat{\psi}_{ij}^{kk'} = 0$ with $i, j \in S_n^c$ with probability tending to 1.

For (i) it can be shown that $q(\mathbf{\Psi}) - q(\mathbf{\Psi}^0)$ can be decomposed as $I_1 + I_2 + I_3$, a sum which is asymptotically positive. For (ii) we can show that $\left|\frac{\psi_{i,j}^{kk}}{\sqrt{\sum_{k=1}^K(\psi_{i,j}^{kk})^2}}\right| \leq 1$ and $\left|\sum_{k'';\, k''>k}\mathrm{sgn}(\psi_{i,j}^{kk} - \psi_{i,j}^{k''k''})\right| \leq K - 1$ (since it is a finite sum, as K is fixed, of 1s or -1s). If $\psi_{i,j}^{kk'}$ lies in a small neighborhood of 0 (excluding the value 0), as long as the conditions of Proposition 2 hold in the case of the GGL or FGL penalty, then the sign of the derivative will depend on the sign of $\psi_{i,j}^{kk'}$ only. This is because we can choose the λ_{n1} sequence to be large such that it dominates the remaining terms. This sparsistency rate implies that for entries that do not get the GGL/FGL penalty, the rate is $\sqrt{n^{-1}\log(p_n)} + \sqrt{\eta_n}$, which is the same as for GL. For entries that receive the GGL/FGL penalty, the rate is worse, due to the extra term associated with λ_{n2}. □

7. SIMULATION STUDY

We have evaluated the performance of msJGL against the performance of GL on seven measures: (i) TPR = (#estimated edges that are true edges)/(#true edges), this is the true positive rate, larger is better; (ii) FPR = (#estimated edges that are NOT true edges)/(#edges that are NOT present in the true graph), this is the false positive rate, smaller is better; (iii) FDR = (#estimated edges that are NOT true edges)/ (#estimated edges), this is the false discovery rate, smaller is better; (iv) SI = 1 − (#estimated edges)/(#possible edges), this is the sparsity index; (v) SHD = #edge additions/deletions such that estimated graph = true graph (structural Hamming distance, smaller is better); (vi) $F_1 = 2PR/(P + R)$ where $P =$ (#estimated edges that are true edges)/(#estimated edges); $R =$ TPR (F_1 score, Jardine and van Rijsbergen, 1971, larger is better); and (vii) FL= $||\hat{\Theta}_{Red} − \Theta||_F = \sqrt{\sum_{i=1}^{p} \sum_{j=1}^{p} |\hat{\theta}_{ij_{Red}} − \theta_{ij}|^2}$, where p represents the number of columns of the matrix Θ (Frobenius loss, smaller is better).

Data corresponding to two different coarseness scales have been used for the simulation study. For the second scale, a graph with 300 or 600 nodes was generated. For the first scale, a graph with 1/3rd of this number of nodes was generated, where each node was obtained from concatenating the "finer" nodes from the second scale. The number of regions that were combined to obtain a coarser region varied randomly, meaning that some nodes were obtained by merging more regions than other nodes. Specifically, we have first sampled with replacement 300 indices from the set $\{1, \ldots, 100\}$ or 600 indices from the set $\{1, \ldots, 200\}$, making sure that each integer from the set appeared at least once. The number of times an index appeared is how many finer nodes are combined to obtain the coarser node. For both scales, the graph structure generated was either "random", "hub", "cluster", "banded" or "scale-free" (see Figure 2).

From the composed graph, we have defined a Σ^{-1} matrix as follows. The adjaceny matrix of the graph (that contained only 0 or 1 values, where 1 on row i and column j denotes the presence of an edge between nodes i and j) was multiplied with the value 3.9, and then an eigenvalue decomposition was performed. The diagonal values of the matrix were replaced by the absolute value of the minimal eigenvalue, to which the value 0.4 was added to ensure positive definiteness. With n either 100, 300, or 3000, data were generated from a normal distribution of mean vector 0 and with Σ as covariance matrix. Note that as in the real example from Section 8 both graphs are available for all the samples. The number of simulation runs was set at 500. Both msJGL and GL use the ADMM algorithm in the optimization process.


Fig. 2. Simulated data. Schematic representation of concentration matrices corresponding to four graphs used in the simulation study. For two scales, a random (a), hub (b), cluster (c), or scale-free (d) graph has been generated. Black dots represent an edge between nodes or equivalently a nonzero element in the Θ matrix. In each graph, the "smaller" bulk of points (top left) represents the graph for the first scale and the "larger" bulk of points (bottom right) represents the graph for the second scale. The coarser regions of scale 1 have been split in finer regions on which the second scale has been measured. The splits are denoted by the two "lines" above/below the diagonal.
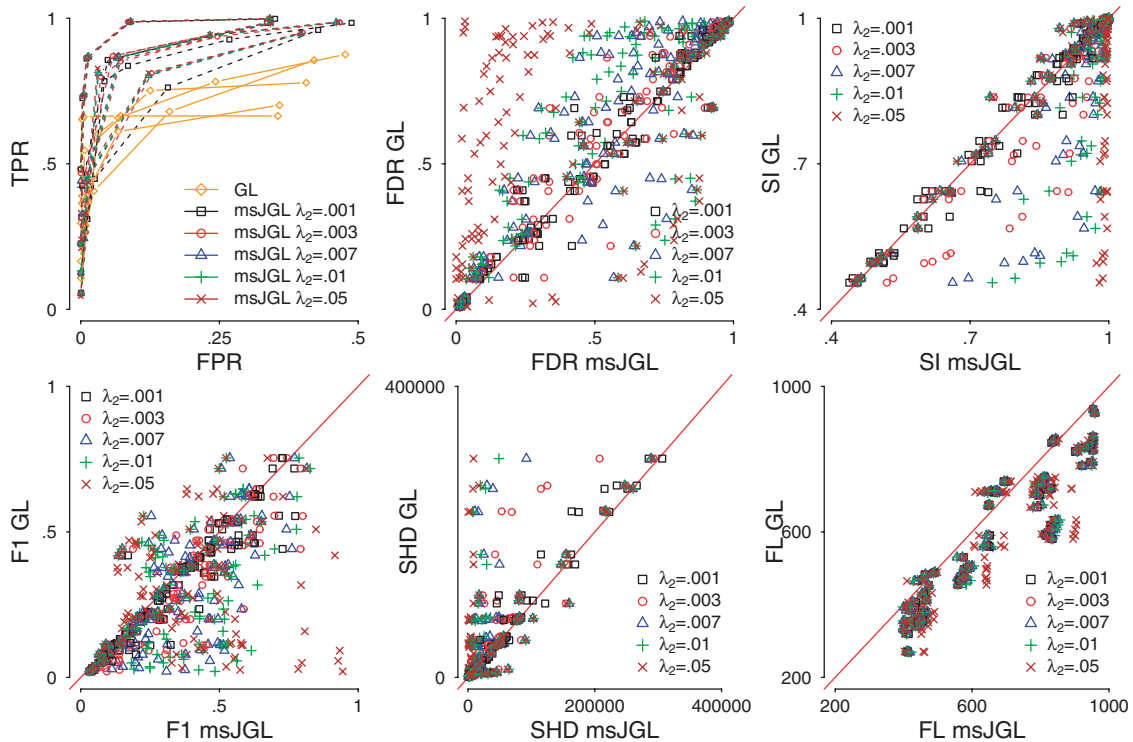
Fig. 3. Simulated data. The panels present the TP rate against FP rate, FD rate, Sparsity index (top row) and F_1 score, Hamming distance and Frobenius loss (bottom row) for msJGL and GL.

In the GL case, we estimate from the data two separate graphs (corresponding to each scale) using a separate GL for each of them. Note that there is no involvement of the reduce/expand operators when using GL and note too that the estimated graphs based on GL do not take into account any desire of obtaining graphs that are similar to each other, nor do they account that some nodes in the larger graph are actually obtained from splitting nodes in the coarser graph; as such GL is insensitive to dependencies induced by splits.

Figure 3 presents the obtained results. Each symbol represents the average over all simulation runs within one simulation setting, where a certain p, n, λ_{n1}, λ_{n2} and penalty was used. For this study, both λ_{n1} and λ_{n2} take a value in the fixed set {0.001, 0.003, 0.007, 0.01, 0.05} and once a λ_{n1} is selected, then it is used for both msJGL and GL. For the FDR, SHD, and FL plots, a value above the diagonal indicates a more favorable position of msJGL against GL. For the F_1 and SI plots, a value below the diagonal indicates a more favorable position of msJGL against GL.

The results indicate that generally the msJGL for a fixed FP rate provided larger TP rates. For ease of exposition, Figure 3 (top left panel) shows the TPR and FPR performance for the case where the number of nodes was 800 (200 coarse scale nodes and 600 finer scale nodes), the sample size was 300, and the FGL penalty was used. For each of the 5 types of graphs, a curve is presented. A similar behavior was observed for the other settings and when using the group penalty. In a large majority of cases, the SHD, F_1, and FDR measures were either comparable or better for msJGL. With respect to the Frobenius loss, the results indicate that the performance is either comparable to that of the GL or slightly worse due to the constraint of the similarity of submatrices, which forces entries in the concentration matrix to have similar

values. In general, increasing the λ_{n2} regularization parameter has as a direct effect a slight improvement in the performance, increasing at the same time the sparsity of the msJGL graph as this increases the regularization imposed on the estimated concentration matrix. While λ_{n1} involves shrinking single entries in the concentration matrix, λ_{n2} influences groups of parameters.

8. rsfMRI example

For two subjects that were instructed to stay alert and focus on a white fixation cross, rsfMRI data were acquired. The obtained images at scale 1 were segmented into 68 atlas-based ROIs, while for the second scale 114 atlas-based ROIs have been used (Desikan *and others*, 2006). The 68 ROIs correspond to a coarser scale of measurement, which implies that the brain regions under investigation were anatomically larger. In the second step, some of the large regions were further split into several smaller regions, resulting in a total of 114 ROIs for which the cerebral activity has been measured. For all ROIs, we obtained $n = 240$ volumes of the BOLD signal. See Schmittmann *and others* (2015) for more information regarding the acquisition of the data.

We want to investigate (i) if there are links present in the estimated networks due to the splitting of regions between scales that connect coarser regions with their smaller counterparts and (ii) which links are present between regions within a given coarseness scale. To answer these questions, we have applied the msJGL algorithm to the data from both subjects using the FGL and GGL penalties on a grid of regularization parameters $(\lambda_{n1}, \lambda_{n2})$. The final regularization parameters have been selected to convey sufficient information about brain pathways without having the graphs too cluttered nor having them too sparse. A cross-validation scheme or an information criterion-based selection could also be used to select an appropriate regularization value.

Figure 4 shows that both the FGL and GGL estimated between-scale edges ("split" dependencies) as most of the entries are nonzero. This suggests that conditional independencies between the coarser and finer scale splits are not supported by the data.

Figure 5 presents both the common and the unique edges pertaining to the estimated graphs for both subjects at each scale (upper row) and the estimated splits across the scales (bottom row). The graphs appear stable across subjects, which is seen by the percentage of common edges within scales (97.9% for


Fig. 4. fMRI data. Schematic representation of estimated concentration matrices with the fused (left panel) and group (right panel) msJGL procedure using $(\lambda_{n1}, \lambda_{n2}) = (.4, .1)$. The black dots represent an edge between nodes or equivalently a nonzero element in the estimated concentration matrix. The square "bulks" of points denote the within coarseness scale edges, while the "lines" above/below the diagonal denote the between scale edges. Each panel represents two subjects with two scales (coarse and fine).

Fig. 5. fMRI data. Fused msJGL within scale (top row) and between scale (bottom row) graphs. Both the coarse scale (top left) and the fine scale (top right) are presented. Full line edges are common edges for both subjects, while dashed and dot-dashed edges are estimated for one subject but not for the other. The values of the regularization parameters are $(\lambda_{n1}, \lambda_{n2}) = (.4, .1)$.

the coarse scale and 98.2% for the fine scale) and between scales (93.7%). Our method explicitly estimates the splits and provides information on how the different scales are functionally related. Here we can see that most differences between the subjects within and between scales are in the parietal and prefrontal areas. This is in line with poorer reliability in these brain areas as found in Mueller *and others* (2015).

It is striking that the coarser regions do *not* connect with some of the finer subregions. The coarser regions connect at most to one or two finer subregions, indicating that some of the coarser regions are formed by grouping together several heterogeneous finer subregions. This would indicate that the function of one or more subregions is different. This is likely to be found in rather arbitrary parcellations (Zalesky *and others*, 2010). This is the case for ROIs 23, 29, and 30. The most extreme case is that of the coarse ROI 29, which for one subject connects to only one split, although the left and right hemispheres contained each four splits, suggesting that this coarser region is a conglomerate of regions that exhibit different cerebral activity. On the other hand, regions 8, 9, 10, 12, 13, 14, 16, 30, and 32 from the left hemisphere,

as well as regions 7, 8, 9, 10, 12, 16, 31, and 32 from the right hemisphere seem to be homogeneous regions, as links are present between the coarser regions and all of their finer splits.

The impact of using either a group or fused penalty can be important for the estimated structure and sparsity of the networks, but there is an agreement with respect to which regions are deemed important. Regions 10, 24, 25, 26, 28, 29, 31, and 32 are highly connected (see Figure 5) with the rest of the ROIs for both scales and both penalties.

The analysis was performed on a standard laptop. As a test, we let the algorithm perform 1000 iterations, which took around 2 minutes when applied to one subject and about 86 minutes for two subjects. In both cases, two coarseness scales have been used.

9. Discussion

We have developed a new method of jointly estimating graphs where the nodes come from mixed coarseness scales. The approach is motivated by an fMRI dataset where the brain image has been "partitioned" in various ROIs in an incremental manner. The cerebral activity has been measured first for 68 ROIs and then for 114 ROIs, where the latter finer ROIs were created by splitting the coarser ROIs. Using the proposed method, we were able to identify certain brain regions which exhibit either a homogeneous or heterogeneous cerebral activity pattern. The method has direct applicability beyond fMRI data and in other areas where data on different scales are observed and where the joint estimation of graphs that resemble each other is desired.

Having multiple coarseness scales sets the identification of an optimal coarseness scale of the data as an open problem, for which a scientist should perform the analysis. Different scales lead to some qualitative differences in the conclusions, but one would hope that the decision on the scales is invariant. Such questions related to the selection of an optimal coarseness scale are the subject of ongoing research. The proposed method avoids selecting one such optimal scale and produces interpretable results at all available scales.

Acknowledgments

Funding

References

Banerjee, O., El Ghaoui, L. and 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.

Bickel, P. J. and Levina, E. (2008a). Covariance regularization by thresholding. *The Annals of Statistics* **36**, 2577–2604.

Bickel, P. J. and Levina, E. (2008b). Regularized estimation of large covariance matrices. *The Annals of Statistics* **36**, 199–227.

Boyd, S., Parikh, N., Chu, E., Peleato, B. and Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning* **3**, 1–122.

BÜHLMANN, P. AND VAN DE GEER, S. (2011). *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Berlin Heidelberg: Springer-Verlag.

CAI, T. T., ZHANG, C.-H. AND ZHOU, H. H. (2010). Optimal rates of convergence for covariance matrix estimation. *The Annals of Statistics.* **38**, 2118–2144.

CHOI, M.J., CHANDRASEKARAN, V. AND WILLSKY, A.S. (2010). Gaussian multiresolution models: exploiting sparse Markov and covariance structure. *IEEE Transactions on Signal Processing* **58**, 1012–1024.

CHOI, M.J AND WILLSKY, A.S. (2007). Multiscale Gaussian graphical models and algorithms for large-scale inference. In: *IEEE/SP 14th Workshop on Statistical Signal Processing, 2007*. pp. 229–233.

DANAHER, P., WANG, P. AND WITTEN, D. (2014). The joint graphical lasso for inverse covariance estimation across multiple classes. *Journal of the Royal Statistical Society, Series B,* **76**, 373–397.

DESIKAN, R. S., SÉGONNE, F., FISCHL, B., QUINN, B. T., DICKERSON, B. C., BLACKER, D., BUCKNER, R. L., DALE, A. M., MAGUIRE, R. P., HYMAN, B. T., ALBERT, M. S. AND *others*. (2006). An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage* **31**, 968–980.

FAN, J., LIAO, Y. AND LIU, H. (2016). An overview on the estimation of large covariance and precision matrices. *The Econometrics Journal* **19**, C1–C32.

FRIEDMAN, J., HASTIE, T. AND TIBSHIRANI, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9**, 432–441.

GALLIER, J. (2011). *Geometric Methods and Applications: For Computer Science and Engineering*. New York: Springer.

GASKINS, J. T. AND DANIELS, M. J. (2013). A nonparametric prior for simultaneous covariance estimation. *Biometrika* **100**, 125–138.

GUO, J., LEVINA, E., MICHAILIDIS, G. AND ZHU, J. (2011). Joint estimation of multiple graphical models. *Biometrika* **98**, 1–15.

HAGMANN, P., CAMMOUN, L., GIGANDET, X., MEULI, R., HONEY, C. J., WEDEEN, V. J. AND SPORNS, O. (2008). Mapping the structural core of human cerebral cortex. *PLoS Biology* **6**, e159.

HASTIE, T., TIBSHIRANI, R. AND WAINWRIGHT, M. (2015). *Statistical Learning with Sparsity: The Lasso and Generalizations*, Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Boca Raton: Taylor & Francis.

HÖFLING, H., BINDER, H. AND SCHUMACHER, M. (2010). A coordinate-wise optimization algorithm for the fused lasso. *Technical Report*, arXiv:1011.6409.

JARDINE, N. AND VAN RIJSBERGEN, C.J. (1971). The use of hierarchic clustering in information retrieval. *Information Storage and Retrieval* **7**, 217–240.

LAM, C. AND FAN, J. (2009). Sparsistency and rates of convergence in large covariance matrix estimation. *The Annals of Statistics* **37**, 4254–4278.

LENG, C. AND TANG, C. Y. (2012). Sparse matrix graphical models. *Journal of the American Statistical Association* **107**, 1187–1200.

LIU, J., YUAN, L. AND YE, J. (2010). An efficient algorithm for a class of fused lasso problems. In: Rao, B., Krishnapuram, B., Tomkins, A. and Qiang, Y. (editors). Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. ACM: New York. pp. 323–332.

MEINSHAUSEN, N. AND BÜHLMANN, P. (2006). High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics* **34**, 1436–1462.

MOHAN, K., LONDON, P., FAZEL, M., WITTEN, D. AND LEE, S. (2014). Node-based learning of multiple Gaussian graphical models. *Journal of Machine Learning Research* **15**, 445–488.

MUELLER, S., WANG, D., FOX, M. D., PAN, R., LU, J., LI, K., SUN, W., BUCKNER, R. L. AND LIU, H. (2015). Reliability correction for functional connectivity: Theory and implementation. *Human brain mapping* **36**, 4664–4680.

ROTHMAN, A. J., BICKEL, P. J., LEVINA, E. AND ZHU, J. (2008). Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics* **2**, 494–515.

SCHMITTMANN, V. D., BORSBOOM, D., JAHFARI, S., SAVI, A. AND WALDORP, L. J. (2015). Making large-scale networks from fMRI data. *PLoS One* **10**, e0129074.

TIBSHIRANI, R., SAUNDERS, M., ROSSET, S., ZHU, J. AND KNIGHT, K. (2005). Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society, Series B,* **67**, 91–108.

WITTEN, D. M., FRIEDMAN, J. H. AND SIMON, N. (2011). New insights and faster computations for the graphical lasso. *Journal of Computational and Graphical Statistics* **20**, 892–900.

YANG, S., LU, Z., SHEN, X., WONKA, P. AND YE, J. (2015). Fused multiple graphical lasso. *SIAM Journal on Optimization* **25**, 916–943.

YUAN, M. AND LIN, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society, Series B,* **68**, 49–67.

YUAN, M. AND LIN, Y. (2007). Model selection and estimation in the Gaussian graphical model. *Biometrika* **94**, 19–35.

ZALESKY, A., FORNITO, A., HARDING, I.H., COCCHI, L., YÜCEL, M., PANTELIS, C. AND BULLMORE, E. T. (2010). Whole-brain anatomical networks: does the choice of nodes matter? *Neuroimage* **50**, 970–983.

ZHAO, S. D., CAI, T. T. AND LI, H. (2014). Direct estimation of differential networks. *Biometrika* **101**, 253–268.