

ORIGINAL ARTICLE

Fused graphical lasso for brain networks with symmetries

Saverio Ranciati¹  | Alberto Roverato² | Alessandra Luati¹ 

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

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

Correspondence

Saverio Ranciati, Department of Statistical Sciences, University of Bologna, Bologna, Italy.

Email: saverio.ranciati2@unibo.it

Funding information

Air Force Office of Scientific Research,
Grant/Award Number: FA9550-17-1-0039

Abstract

Neuroimaging is the growing area of neuroscience devoted to produce data with the goal of capturing processes and dynamics of the human brain. We consider the problem of inferring the brain connectivity network from time-dependent functional magnetic resonance imaging (fMRI) scans. To this aim we propose the symmetric graphical lasso, a penalized likelihood method with a fused type penalty function that takes into explicit account the natural symmetrical structure of the brain. Symmetric graphical lasso allows one to learn simultaneously both the network structure and a set of symmetries across the two hemispheres. We implement an alternating directions method of multipliers algorithm to solve the corresponding convex optimization problem. Furthermore, we apply our methods to estimate the brain networks of two subjects, one healthy and one affected by mental disorder, and to compare them with respect to their symmetric structure. The method applies once the temporal dependence characterizing fMRI data have been accounted for and we compare the impact on the analysis of different detrending techniques on the estimated brain networks. Although we focus on brain networks, symmetric graphical lasso is a tool which can be more generally applied to learn multiple networks in a context of dependent samples.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* published by John Wiley & Sons Ltd on behalf of Royal Statistical Society.

KEYWORDS

ADMM algorithm, fMRI data, graphical model with symmetries, time series, undirected graphical models

1 | INTRODUCTION

A brain network is a model of a nervous system represented as a set of nodes, also called vertices, interconnected by a set of edges; see Bullmore and Bassett (2011) for a review on the use of brain graphs for modelling the human brain connectome. Within the domain of human brain mapping, great interest has been posed in the estimation of brain networks from functional magnetic resonance imaging (fMRI) data (Smith et al., 2011).

Functional MRI is a non-invasive technique for collecting data on brain activity, with a good resolution in terms of space and time. Essentially, fMRI measures the increase in the oxygenation level at some specific brain region, as long as an increase in blood flow occurs, due to some brain activity. The latent signal in the observed fMRI data is referred to as the blood oxygenation level-dependent (BOLD) signal. The BOLD signal arises from the interplay of blood flow, blood volume and blood oxygenation in response to changes in neural activity. Under an active state, the local concentration of oxygenated haemoglobin increases, with a corresponding increase in the homogeneity of magnetic susceptibility, which, in turn, results in an increase in MRI signal. Typically, active states are observed in task-based experiments in response to an exogenous event. In recent years, the attention has been concentrating towards resting state fMRI (R-fMRI) data, collected on subjects at rest and in the absence of any external stimulus, as the key to understand the neuronal organization of the brain through the investigation of the spatial and temporal structure of spontaneous neural activity. Smith et al. (2009) carried out an analysis to assess how functional networks at rest match the ones detected under activation tasks. The authors conclude that the resting brain functional dynamics are fully utilizing the set of functional networks exhibited by the brain over the range of its possible tasks. In the review paper by Biswal et al. (2010), R-fMRI is described as the candidate approach capable of addressing the core challenge in neuroimage, that is, the development of common paradigms for interrogating the functional systems in the brain, without the constraints of a priori hypotheses.

The construction of a network from fMRI data requires first the identification of a set of functional vertices, such as spatial regions of interest (ROIs), and then the analysis of connectivity patterns across ROIs. It is also relevant that the human brain has a natural symmetric structure. More specifically, it is made up of two hemispheres such that for every spatial ROI on the left hemisphere there is a homologous ROI on the right hemisphere. Accordingly, in the brain network one can identify pairs of homologous vertices and edges, and R-fMRI studies have suggested a highly symmetric connectivity; see Section 2 for additional details.

In this paper, we address the problem of estimating the brain network from R-fMRI data by keeping symmetries into explicit account. Special attention is also posed in the temporal dependence characterizing fMRI data and the impact of alternative detrending approaches on the estimated brain network. We remark that we consider undirected graphs, that is, graphs where the edges are of symmetric type. However, following Højsgaard and Lauritzen (2008), throughout this paper the term ‘symmetry’ refers to similarities across the two hemispheres with respect to their network structure and equality in parameter values.

Undirected graphical models (Lauritzen, 1996) are widely applied in network modelling from fMRI data (see, among others, Marrelec et al., 2006; Smith et al., 2011; Zhu & Cribben, 2018). In this framework, the network structure follows from the sparsity pattern of the inverse covariance matrix

of ROI values, and a popular approach to estimate sparse undirected graphical models is the graphical lasso technique (Banerjee et al., 2008; Friedman et al., 2008). This method is based on the optimization of a penalized log-likelihood function, where the role of the penalty term is to encourage sparsity in the network. One drawback of the graphical lasso, in the present context, is the fact that it ignores the symmetric structure of the brain. For this reason, we propose a fused graphical lasso method based on the optimization of a penalized log-likelihood where the penalty function is obtained by the sum of two distinct terms. Like the graphical lasso, the first term encourages sparsity in the solution. On the other hand, the second one is a fused type penalty (Tibshirani et al., 2005) that encourages symmetry by penalizing differences between the left and right hemispheres. As detailed more formally in Section 3, symmetries are implemented in the form of equality constraints between entries of the inverse covariance matrix of ROI values. This leads to a convex optimization problem and we provide an alternating direction method of multiplier (ADMM) algorithm for its solution.

The method applies once the temporal dependence characterizing fMRI data have been accounted for and the BOLD signal has been extracted. We assume a simple decomposition of pre-processed R-fMRI data into an unobserved signal plus noise. The underlying hypotheses on the two latent variables are related to the evolution of the components in time and determine the method adopted for their estimation. As the dynamics of fMRI time series are controversial, we shall assess the impact of detrending on the estimated ROI connectivity network using three methods, representative of different approaches to trend estimation, based on different assumptions and including possible misspecification. In particular, we shall specify a linear Gaussian multivariate parametric model, a non-linear observation-driven model for unobserved components and possibly heavy tailed data, and a non-parametric local polynomial regression method. Details are deferred to Section 4.2.

We carried out an extensive analysis and provide an illustration using R-fMRI data from two representative subjects who have similar characteristics in terms of age and handedness, although one of the two is healthy while the other has been diagnosed with some mental disorder. We may anticipate that the results show a lack of brain asymmetry in the latter individual.

In summary, the novel contribution of this paper is twofold. First, we introduce a fused graphical lasso approach to estimate sparse undirected graphical models with a specific symmetric structure and provide an ADMM algorithm for its solution. An implementation of the latter, written in the R language (R Core Team, 2020), can be found at <https://github.com/savranciati/sgl>. Second, we compare the impact on the analysis of different detrending techniques, thereby providing insight into the robustness of the estimated network on such a preliminary step.

1.1 | Related works and possible applications

The novel contribution of this paper pertains the research area usually referred to as *joint learning of multiple graphical models*. In this framework, the observations come from two or more groups where each group shares the same variables and some of the dependence structure. Accordingly, every group is associated with a network and it is expected that some edges are common across all groups and other edges are unique to each group. More specifically, the literature has focused on the case where the groups correspond to independent experimental conditions so that every network is a distinct unit, disconnected from the other networks; see, among others, Danaher et al. (2014); Yang et al. (2015). Examples of possible applications include genetic networks (see Danaher et al., 2014) and brain networks from neuroimaging data (see Yang et al., 2015). In the former, the groups are normal tissue from healthy subjects and one or more different types of cancer tissues whereas in the latter groups correspond to normal subjects and subjects with different degrees of cognitive impairment.

Our work is motivated by brain network inference where the groups are given by the two hemispheres. The independence assumption does not hold in this case because every observation from the left hemisphere is paired with an observation from the right hemisphere. Thus, existing methods for joint learning of graphical models no longer apply, and one should distinguish between *independent samples* and *paired data*, the latter being a largely unexplored area of research. Note also that, unlike the case of independent samples, with paired data the networks associated with different groups are not expected to be disconnected from each other. Although we focus on brain networks, our approach applies to other settings, and another relevant example comes from cancer genomics where control samples are often obtained from histologically normal tissues adjacent to the tumour (NAT), so that every observation from a cancer tissue is paired with an observation from a normal tissue; see for example, Aran et al. (2017).

A central role in the theory of joint learning of multiple graphical models for independent samples is played by the group and the fused graphical lasso (Dahner et al., 2014). The penalty term of the group graphical lasso encourages both sparsity and a similar structure of the networks. On the other hand, the fused graphical lasso encourages sparsity and, at the same time, the parameters of the model to be identical across groups. In this way, the groups are encouraged to have both similar network structure and identical parameter values, thereby typically resulting in more parsimonious models. Procedures for applying both the group and the fused graphical lasso are not available for paired data and the symmetric graphical lasso introduced in this paper contributes to fill this gap.

Interestingly, the application of fused graphical lasso for paired data results in a model that belongs to the family of Gaussian graphical models with edge and vertex symmetries, shortly RCON models, introduced by Højsgaard and Lauritzen (2008). Note that, as pointed out in that paper, symmetry restrictions in the multivariate Gaussian distribution have a long history and RCON models can be identified as a special case within this framework. For recent applications of these models see Gao and Massam (2015); Vinciotti et al. (2016); Massam et al. (2018). Although the theory of estimation and testing for RCON models is well-established, a procedure that performs model selection within the family of RCON models is not available, with the relevant exception of the procedures introduced by Gehrmann (2011) and by Li et al. (2020), within the frequentist and the Bayesian approach, respectively, which are of theoretical interest but whose computational complexity restricts their application to low-dimensional settings. More specifically, the problem of model selection for RCON model is discussed in Gehrmann (2011) where it is shown that the number of RCON models grows super-exponentially in the number of variables. For this reason, Gehrmann (2011) suggested that lasso procedures with fused type penalties might represent a useful alternative to traditional model selection approaches. The symmetric graphical lasso does not constitute a general solution to this problem but it represents, to the best of our knowledge, the first instance of a lasso procedure specifically designed for RCON models.

A further contribution of the paper is concerned with the prewhitening of fMRI data that, measured at each region, are characterized by temporal dependence. There is a longstanding debate on the dynamic properties of fMRI data and parametric models have been employed along with fully non-parametric methods. Autoregressive errors have been considered, see for example, Worsley et al. (2002), Lindquist (2008) and Zhu and Cribben (2018), as well as fractional noise error processes, as in Bullmore et al. (2003) or change point methods, see Aston and Kirsch (2012). Semiparametric methods and high pass filters are also applied to fMRI data, see Zhang and Yu (2008) and Schmal et al. (2017), who use the Hodrick–Prescott filter as in St. John and Doyle (2015). Lund (2006) deduced that no commonly accepted model for noise in fMRI exists and that regressors may whiten the noise as well as non-parametric smoothing methods. In a Bayesian setting, the relevance of prewhitening has been investigated by Kundu and Risk (2020), who model the temporal covariance under an

inverse-Wishart prior. The discussion in the latter paper leads to the overall conclusion that prewhitening is a crucial step yet not fully solved when modelling fMRI data. Against this background, the paper provides a contribution by assessing the impact of different signal extraction methods, applied to the same dataset, to the whitening of the original, temporally correlated, series.

2 | PROBLEM AND DATA DESCRIPTION

Structural symmetry of human brain is concerned with anatomical or physiological similarities between the left and the right hemisphere. Otherwise, functional asymmetry is referred to activity-related differences, in a similar way in which left and right hands operate differently, although being anatomically symmetric. As it is related to behavioural differences, functional asymmetry, also known as lateralization, is usually detected with respect to some specific tasks, the most relevant being connected to language organization and handedness. Non-invasive methods for exploring the brain organization with respect to lateralization are electroencephalography, positron emission tomology and fMRI, the latter being the most used in research, which generally display bilateral activations that contrast with the asymmetric effect of lateralization.

So far, R-fMRI studies have suggested a highly symmetric connectivity. Indeed, along with the recognition of the relevance of analysing the brain at rest, the focus has moved from detecting functional asymmetries to detecting structural symmetries. In some sense, the two methods are complementary, but clearly task-based analyses tend to evidence asymmetries whereas resting state analyses are designed to shed light on symmetric structures. In a recent R-fMRI analysis, Raemaekers et al. (2018) focus on differences between hemispheres that are reflected in asymmetric functional connectivity in resting state subjects and recognize that any asymmetries are prone to be relatively minute. They also observe that a direct quantification of the extent of the hemispheric symmetry is missing.

The fused graphical lasso procedure introduced in this paper provides a methodological contribution for analysing functional symmetries between the left and right hemisphere of the brain. We apply our method to a multimodal imaging dataset which comes from a pilot study of the Enhanced Nathan Kline Institute-Rockland Sample project. This project aims at providing a large cross-sectional sample of publicly shared multimodal neuroimaging data and psychological information to support and motivate researchers in the relevant scientific goal of understanding the mechanisms underlying the complex brain system. A detailed description of the project, scopes and technical aspects can be found at http://fcon_1000.projects.nitrc.org/indi/enhanced/. The pilot NKI1 study comprises multimodal imaging data and subject-specific covariates for $n = 24$ subjects. Detailed information can be found at http://fcon_1000.projects.nitrc.org/indi/CoRR/html/nki_1.html.

For each subject several information are collected as personal covariates, such as anxiety diagnosis, age, gender and handedness. The fMRI time series are recorded at $p = 70$ spatial ROI, clustering anatomically close and functionally similar voxels. The way ROIs are defined can depend on the scope of the analysis or on the design of the experiments and it has implications for fMRI analysis, see the discussion by Poldrack (2007). In our case, ROIs are pre-defined according to the Desikan atlas, see Desikan et al. (2006). For any region, additional information on 3D spatial locations, hemisphere and lobe membership is available. As we have ROI-specific information, we apply a ROI analysis approach, based on the given anatomical parcellation. An alternative approach is to conduct a whole-brain voxel-wise analysis, on a finer scale, but such an approach is computationally expensive, sensitive to noise, and often difficult to interpret. The optimal means of combining voxels into functionally distinct ROIs remains to be determined. The issue of parcellation is largely discussed in Craddock

et al. (2012), where the authors develop a spatially constrained spectral clustering approach for group clustering of the whole resting state fMRI data into functionally and spatially coherent regions.

As far as dynamic functional activity is concerned, the dataset we are focusing on is composed of time series data collected for each of the 24 subjects in an imaging session. This imaging technology monitors brain functional activity in different regions via dynamic changes in blood flow creating a low-frequency BOLD signal when the subject is not performing an explicit task during the imaging session. In the present NKI1 study, the subjects are simply asked to stay awake with eyes open. Focusing on subject i and on scan k , where $i = 1: 24$ and $k = 1: 2$, we have 70×404 matrix whose rows contain the dynamic activity data of the brain regions, collected at $T = 404$ equally spaced times (time lag is 1400 ms).

The data are provided by Greg Kiar and Eric Bridgeford from NeuroData at Johns Hopkins University, who graciously pre-processed the raw DTI and R-fMRI imaging data available at http://fcon_1000.projects.nitrc.org/indi/CoRR/html/nki_1.html, using the pipelines ndmg and C-PAC.

3 | OVERVIEW ON THE METHODOLOGICAL FRAMEWORK

Let $X^{(t)}$ be a p dimensional time series vector collecting the fMRI series observed on each single subject over $p = 70$ regions, $t = 1, \dots, T$ where $T = 404$, the length of each time series. We assume the general signal plus noise decomposition for $X^{(t)}$,

$$X^{(t)} = M^{(t)} + Y^{(t)} \quad (1)$$

where $M^{(t)}$ is the vector collecting the BOLD signal and $Y^{(t)}$ is the idiosyncratic noise component. Our input data for the analysis of the ROI network association structure will be the estimate of $Y^{(t)}$, obtained by contrast as the residual vector once the BOLD signal $M^{(t)}$ is extracted (see Section 4.2). More specifically, we denote by $V = \{1, \dots, p\}$ the set indexing the $p = 70$ brain regions and by $Y_V = (Y_1, \dots, Y_{70})^\top$ the zero mean residual vector, where we have dropped the time index as we assume that the time series dynamics are fully captured by the time-varying BOLD signal (Section 4.2).

We assume $Y_V \sim N_p(0, \Sigma)$ and consider the ROI connectivity network obtained from the application of the theory of undirected graphical models (Lauritzen, 1996). In this framework, the network structure follows from the sparsity pattern of the concentration matrix $\Theta = \Sigma^{-1}$. More specifically, if the entry θ_{ij} of Θ , with $i \neq j$, is such that $\theta_{ij} \neq 0$ then the brain regions indexed by i and j are connected by an edge in the network. Conversely, for every missing edge in the network the corresponding entry of Θ is equal to zero. Concentrations can be interpreted by exploiting their connection with partial correlation and regression coefficients, because for every pair $i, j \in V$ with $i \neq j$ it holds that (see Cox & Wermuth, 1996, Section 3.2),

$$\rho_{ij|V \setminus \{i,j\}} = -\frac{\theta_{ij}}{\sqrt{\theta_{ii}\theta_{jj}}}, \quad \beta_{i \leftarrow j|V \setminus \{i,j\}} = -\frac{\theta_{ij}}{\theta_{ii}} \quad \text{and} \quad \sigma_{i|V \setminus \{i\}}^2 = \theta_{ii}^{-1}, \quad (2)$$

where $\rho_{ij|V \setminus \{i,j\}}$ is the partial correlation between Y_i and Y_j given the remaining components $Y_{V \setminus \{i,j\}}$ whereas $\beta_{i \leftarrow j|V \setminus \{i,j\}}$ is the regression coefficient of Y_i on Y_j given $Y_{V \setminus \{i,j\}}$ and $\sigma_{i|V \setminus \{i\}}^2$ is the partial variance of Y_i given $Y_{V \setminus \{i\}}$. Hence, if the brain regions indexed by i and j are not connected by an edge it holds that $\theta_{ij} = 0$ and this is equivalent to $\rho_{ij|V \setminus \{i,j\}} = 0$ but also to $\beta_{i \leftarrow j|V \setminus \{i,j\}} = 0$ and to $\beta_{j \leftarrow i|V \setminus \{i,j\}} = 0$. Furthermore, in this case, Y_i and Y_j are conditionally independent given $Y_{V \setminus \{i,j\}}$.

Every region in the left hemisphere has a homologous region in the right hemisphere so that the vector Y_V can be naturally partitioned into two subvectors. More formally, we set $q=p/2$ and let the sets $L = \{1, \dots, q\}$ and $R = \{q+1, \dots, p\}$ index the subvectors Y_L and Y_R associated with the left and right hemispheres, respectively, so that the region relative to Y_i of Y_L is homologous to the region relative to Y_{i+q} of Y_R ; furthermore, to shorten the notation, we set $i' = i + q$ for every $i \in L$. Accordingly, the concentration matrix Θ can be partitioned as

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

We investigate the presence of symmetries in the ROI association network that take the form of identities between concentrations in Θ_{LL} with the corresponding concentrations in Θ_{RR} . This is motivated by the interpretation of such equality restrictions that, by (2), allows one to identify equality relationships involving partial correlation and regression coefficients. Specifically:

1. Equalities involving the diagonal entries imply equality in partial covariances, that is $\theta_{ii} = \theta_{i'i'}$ implies $\sigma_{i|V \setminus \{i\}}^2 = \sigma_{i'|V \setminus \{i'\}}^2$.
2. If in addition to the equality $\theta_{ii} = \theta_{i'i'}$ in (i) it also holds that $\theta_{ij} = \theta_{i'j'}$ then we have $\beta_{i \leftarrow j|V \setminus \{i,j\}} = \beta_{i' \leftarrow j'|V \setminus \{i',j'\}}$ so that the contribution of Y_j to the prediction of Y_i is identical to the contribution of $Y_{j'}$ to the prediction of $Y_{i'}$.
3. If in addition to the equalities $\theta_{ii} = \theta_{i'i'}$ and $\theta_{ij} = \theta_{i'j'}$ in (ii) it also holds that $\theta_{ji} = \theta_{j'i'}$ then the partial correlation between Y_i and Y_j is identical to that between $Y_{i'}$ and $Y_{j'}$; formally $\rho_{ij|V \setminus \{i,j\}} = \rho_{i'j'|V \setminus \{i',j'\}}$. It is also worth remarking that in this case it follows from (ii) that both $\beta_{i \leftarrow j|V \setminus \{i,j\}} = \beta_{i' \leftarrow j'|V \setminus \{i',j'\}}$ and $\beta_{j \leftarrow i|V \setminus \{i,j\}} = \beta_{j' \leftarrow i'|V \setminus \{i',j'\}}$.

4 | BACKGROUND

4.1 | Graphical models, graphical lasso and symmetries

We represent the ROI connectivity network by means of an undirected graph $\mathcal{G} = (V, E)$ where the vertex set V indexes the brain regions and $E \subset V \times V$ is a set of edges, which are unordered pairs of vertices. Let Y_V be a multivariate normal random vector with zero mean vector, variance and covariance matrix $\Sigma = \{\sigma_{ij}\}_{i,j \in V}$ and concentration matrix $\Sigma^{-1} = \Theta = \{\theta_{ij}\}_{i,j \in V}$. The concentration graph model (Cox & Wermuth, 1996) with graph $\mathcal{G} = (V, E)$ is the family of multivariate normal distributions with $\Theta \in \mathcal{S}^+(V)$, the set of (symmetric) positive definite matrices which have zero elements $\theta_{ij} = 0$ whenever $\{i, j\} \notin E$. The latter model has also been called a covariance selection model (Dempster, 1972) and a graphical Gaussian model (Whittaker, 1990); we refer the reader to Lauritzen (1996) for details and discussion.

Let $S = n^{-1} \sum_{i=1}^n y_V^{(i)} (y_V^{(i)})^\top$ be the matrix of sums of squares and products for a sample $y_V^{(1)}, \dots, y_V^{(n)}$ of n i.i.d. observations of Y_V . The maximum likelihood estimator (MLE) $\hat{\Theta}^{\text{mle}}$ in the concentration graph model with graph \mathcal{G} maximizes the log-likelihood

$$l(\Theta) = \log \det(\Theta) - \text{tr}(S\Theta), \quad (3)$$

subject to $\Theta \in \mathcal{S}^+(\mathcal{G})$; see Lauritzen (1996) Section 5.2) for details. On the other hand, the structure of a concentration graph can be estimated from data by determining the zero entries of the concentration matrix. We refer the reader to Drton and Maathuis (2017) for a comprehensive review on structure learning for graphical models.

In recent years, much interest has focused on the estimation of concentration graph models through the use of ℓ_1 (lasso) regularization. More specifically, Yuan and Lin (2007), Banerjee et al. (2008) and Friedman et al. (2008) proposed the graphical lasso estimator

$$\hat{\Theta}^{\text{gl}} = \arg \min_{\Theta} \{ -\log \det(\Theta) + \text{tr}(S\Theta) + \lambda \|\Theta\|_1 \} \quad (4)$$

where minimization is over the set \mathcal{S}^+ of $p \times p$ positive definite matrices, $\lambda \geq 0$ and the ℓ_1 -norm $\|\Theta\|_1$ is the sum of the absolute values of the elements of Θ . The graphical lasso adds to the log-likelihood function from (3) a ℓ_1 -penalty pushing the solutions to be sparse, in the sense that due to the geometry of the ℓ_1 -penalty, typically some of the off-diagonal entries of the correlation matrix are shrunk to exactly zero. The term λ is the regularization parameter that controls the amount of shrinkage applied to the elements of Θ , and therefore controlling the sparsity of the solution. Thus, graphical lasso is an effective procedure that conducts model selection and estimation simultaneously. Finally, we remark that for $\lambda > 0$ the minimum in (4) is achieved uniquely because the objective is strictly convex, and this holds true also in high-dimensional settings where $p > n$.

Højsgaard and Lauritzen (2008) investigated the properties of subfamilies of concentration graph models, named RCON models, obtained by imposing additional equality restrictions between specified entries of the concentration matrix. RCON models are commonly referred to as coloured graphical models because equality constraints can be represented by colouring of edges and vertices of the concentration graph \mathcal{G} . Edges of the same colour correspond to off-diagonal entries of Θ with identical values, and similarly for vertices with respect to diagonal entries. The model is thus identified by the structure of \mathcal{G} together with a collection of colour classes. Højsgaard and Lauritzen (2008) showed that, as well as concentration graph models, RCON models are regular exponential families and provided an algorithm for the computation of the MLE of Θ , implemented in the R package `gRc` (Højsgaard & Lauritzen, 2007).

4.2 | Time series analysis

To assess the impact of detrending on our procedure, we consider three different specifications for the latent components in Equation (1). Each one is representative of a wide class of methods for signal extraction and is based on different assumptions on the latent components and their dependence relation. In particular, we specify a Gaussian vector autoregressive model (Section 4.2.1), a univariate Student-t score-driven model (Section 4.2.2) and a local polynomial regression filter, that is the Henderson filter (Section 4.2.3). In the univariate case, we shall denote the elements of the vectors $X^{(t)}, M^{(t)}, Y^{(t)}$ as $x^{(t)}, \mu^{(t)}, y^{(t)}$ respectively.

4.2.1 | Vector autoregressive models

In the class of multivariate linear models, we consider a first-order vector autoregressive process, VAR(1), see Tunnicliffe-Wilson et al. (2015), where

$$M^{(t)} = \Phi X^{(t-1)}$$

and $Y^{(t)}$ is assumed to be multivariate normal with zero mean, covariance matrix Σ and uncorrelated with X_s for $s < t$. The coefficient matrix $\Phi \in \mathbb{R}^{p \times p}$ is required to have eigenvalues that are in modulus smaller than one and it is usually estimated by least squares. Under distributional assumptions on $Y^{(t)}$ maximum likelihood estimation can be carried out and for VAR processes of higher order, the latter can be selected by means of information criteria.

4.2.2 | Score-driven models

Among non-linear models for unobserved components, we focus on the class of score-driven models, recently introduced by Creal et al. (2013) and Harvey (2013) as flexible observation-driven models for time-varying parameters that characterize a given conditional distribution. Specifically, we consider the first-order dynamic conditional score (DCS) model for the location discussed by Harvey and Luati (2014), where each $x^{(t)}$ is assumed to be conditionally distributed as a Student-t random variable with ν degrees of freedom, $x^{(t)} | \mathcal{F}_{t-1} \sim t_\nu(\mu^{(t)}, \sigma^2)$, with the filtration \mathcal{F}_s representing the information set up to time s . The signal $\mu^{(t)}$ is estimated based on an autoregressive mechanism,

$$\hat{\mu}^{(t)} = \omega + \phi \hat{\mu}^{(t-1)} + \kappa \hat{u}^{(t-1)}$$

where $\hat{u}^{(t)}$ is a realization of a martingale difference sequence, that is, $E(u^{(t)} | \mathcal{F}_{t-1}) = 0$, proportional to the score of the conditional likelihood of the time-varying location, that is, $u^{(t)} \propto (\partial/\partial \mu^{(t)}) \mathcal{L}(\mu^{(t)} | \mathcal{F}_{t-1})$, $|\phi| < 1$ and $\hat{\mu}^{(0)}$ is set equal to a fixed value. In this framework, the dynamic BOLD signal is updated by a filter that is robust with respect to extreme values (Calvet et al., 2015). The robustness comes from the properties of the martingale difference sequence $u^{(t)}$: if the data arise from a heavy tail distribution, then the score $\hat{u}^{(t)}$ is less sensitive to extreme values than the score of a Gaussian distribution or than the innovation error $\hat{v}^{(t)} = x^{(t)} - \hat{\mu}^{(t)}$. An important property of the proposed specification is that it encompasses the Gaussian case, in that the score of the Student-t converges to that of the Gaussian distribution when the degrees of freedom tend to infinity. In practice, if a score-driven model is specified when the underlying dataset is in fact Gaussian, a very high value for the degrees of freedom is estimated and a Gaussian model is eventually fitted with the time-varying parameter updated through the Kalman filter. The static parameters, $\omega, \nu, \phi, \kappa, \sigma$, are consistently estimated by maximum likelihood and asymptotic standard errors can be derived (see Harvey, 2013; Harvey & Luati, 2014).

It is important to remark that by applying this method, we are taking into account the possibility that the distribution of the input vector, $Y^{(t)}$, is misspecified, as it is allowed to come from an heavy tailed, rather than Gaussian, distribution.

4.2.3 | Local polynomial regression

Filters that arise from fitting a local polynomial have a well-established tradition in time series analysis and signal extraction, see Cleveland and Loader (1996). With no parametric assumptions on the error term, the signal is approximated locally by a polynomial of degree d , so that in the neighbourhood of time t , for $t = h+1, \dots, n-h$ one has, for $j = 0, \pm 1, \dots, \pm h$, $\mu^{(t+j)} = \beta_0 + \beta_1 j + \beta_2 j^2 + \dots + \beta_d j^d$. Using this design, the estimate of the trend at time t is simply given by the intercept, $\hat{\mu}^{(t)} = \hat{\beta}_0$. Provided that $2h \geq d$, the $d+1$ unknown coefficients β_k , $k = 0, \dots, d$, can be estimated by the method of weighted least squares (see Proietti & Luati, 2007) which eventually produce the trend estimate at time t as the result of a weighted average,

$$\hat{\mu}^{(t)} = \sum_{j=-h}^h w_j x^{(t+j)}.$$

The Henderson filter (Henderson, 1916) arises as the weighted least squares estimator of a local cubic trend, that is, $d=3$, at time t using $2h+1$ consecutive observations. Henderson (1916) addressed the problem of defining a set of weights that maximize the smoothness of the estimated trend, in the sense that the variance of its third differences is minimum. He showed that up to a factor of proportionality, the resulting weights are the following $w_j \propto [(h+1)^2 - j^2][(h+2)^2 - j^2][(h+3)^2 - j^2]$.

Note that, with local polynomial regression methods, $2h$ trend estimates are missing, corresponding to the first and last h time points. Even if the latter are not relevant in the present paper, the reader is referred to Proietti and Luati (2008) for estimation of the signal at the boundaries by asymmetric filters.

5 | THE SYMMETRIC GRAPHICAL LASSO

5.1 | The penalized log-likelihood

In order to encourage both sparsity in the graph structure and similarity across the two brain hemispheres, we introduce a specific fused-type penalty (Hoefling, 2010; Tibshirani & Taylor, 2011) especially designed to encourage the equality between the concentration values of the relevant sub-graphs. Hence, we propose the following estimator of Θ , which we name the symmetric graphical lasso estimator,

$$\hat{\Theta}^{\text{sgl}} = \arg \min_{\Theta} \{ -\log \det(\Theta) + \text{tr}(S\Theta) + \lambda_1 \|\Theta\|_1 + \lambda_2 \|\Theta_{LL} - \Theta_{RR}\|_1 \}, \quad (5)$$

where $\lambda_1, \lambda_2 \geq 0$ are regularization parameters that control the amount of shrinkage. Equation (5) is obtained by adding to the (minus) log-likelihood in Equation (3) a convex penalty function obtained as the sum of two ℓ_1 -penalties, that is, the penalty $\lambda_1 \|\Theta\|_1$ that, like the graphical lasso, for large values of λ_1 encourages sparsity in $\hat{\Theta}^{\text{sgl}}$, and the penalty $\lambda_2 \|\Theta_{LL} - \Theta_{RR}\|_1$ that, for large values of λ_2 encourages the elements of $\hat{\Theta}_{LL}^{\text{sgl}}$ to be identical to the corresponding elements of $\hat{\Theta}_{RR}^{\text{sgl}}$ (Danaher et al., 2014; Tibshirani et al., 2005). Recall that, as described in Section 3, such equality constraints may, in turn, imply the equality of other quantities of interest, such as regression coefficients and partial correlation coefficients.

One of the appealing features of the lasso is that it typically performs model selection and estimation simultaneously. From this perspective, it is worth remarking that the symmetric graphical lasso performs model selection and estimation within the class of RCON models. More precisely, it is suited to identify colour classes of the form $\{\theta_{ij}, \theta_{i'j'}\}$ corresponding to $\hat{\theta}_{ij}^{\text{sgl}} = \hat{\theta}_{i'j'}^{\text{sgl}}$, which are of natural interest in the analysis of brain networks.

5.2 | An algorithm for the symmetric graphical lasso problem

In order to solve Equation (5) we use an ADMM algorithm. A comprehensive exposition of ADMM algorithm can be found in Boyd et al. (2011) whereas we refer to Danaher et al. (2014) and Tan et al. (2014), and references therein, for applications of ADMM to related problems. ADMM is an attractive algorithm for this problem because it allows us to split the optimization procedure into two

nested, less involved, convex optimization problems. These can be both solved using suitable ADMM algorithms.

First, we note that the optimization problem in Equation (5) is equivalent to minimize with respect to Θ and Z the quantity

$$-\log \det(\Theta) + \text{tr}(S\Theta) + \lambda_1 \|Z\|_1 + \lambda_2 \|Z_{LL} - Z_{RR}\|_1, \quad (6)$$

where Θ and Z are restricted to belong to \mathcal{S}^+ and subject to the linear constraint $Z=\Theta$. We remark that Z_{LL} and Z_{RR} in Equation (6) are the relevant diagonal submatrices of Z . Hence, the scaled form of the augmented Lagrangian can be written as (Boyd et al., 2011, Section 3.1.1),

$$\begin{aligned} L_{\rho_1}(\Theta, Z, U) = & -\log \det(\Theta) + \text{tr}(S\Theta) + \lambda_1 \|Z\|_1 + \lambda_2 \|Z_{LL} - Z_{RR}\|_1 + \\ & + \frac{\rho_1}{2} \|\Theta - Z + U\|_F^2 - \frac{\rho_1}{2} \|U\|_F^2, \end{aligned} \quad (7)$$

where U is the scaled dual variable and the symbol $\|\cdot\|_F$ denotes the Frobenius norm, that is, the square root of the sum of the squared entries of its argument. The ADMM algorithm for the optimization of (7) uses the augmented Lagrangian parameter $\rho_1 > 0$ as ‘step size’ and, when the algorithm is at convergence, due to the constraint $Z = \Theta$, the last two terms of Equation (7) cancel out and one obtains the solution to (5).

ADMM iterates three fundamental steps in order to minimize (7) (see Boyd et al., 2011, Equations (3.5) to (3.6)). More formally, we initialize Z^1 and U^1 equal to the zero matrix and for $l = 1, 2, 3, \dots$ the updates for the quantities (Θ, Z, U) are obtained as (see also Boyd et al., 2011, Section 6.6):

1. $\Theta^{l+1} := \arg \min_{\Theta} \left(-\log \det(\Theta) + \text{tr}(S\Theta) + \frac{\rho_1}{2} \|\Theta - Z^l + U^l\|_F^2 \right)$;
2. $Z^{l+1} := \arg \min_Z \left(\lambda_1 \|Z\|_1 + \lambda_2 \|Z_{LL} - Z_{RR}\|_1 + \frac{\rho_1}{2} \|\Theta^{l+1} - Z + U^l\|_F^2 \right)$;
3. $U^{l+1} := U^l + \Theta^{l+1} - Z^{l+1}$.

The implementation of step (3) is straightforward and in the following we describe steps (1) and (2) in detail.

Step (1) has an analytical solution, with computational complexity given by performing an eigendecomposition of a $p \times p$ matrix. More specifically, if QDQ^T is the eigendecomposition of $\rho_1(Z^l - U^l) - S$ then the solution is given by $\Theta^{l+1} := Q\tilde{D}Q^T$ where \tilde{D} is the diagonal matrix with i th diagonal entry $(d_{ii} + \sqrt{d_{ii}^2 + 4\rho_1})/(2\rho_1)$ and d_{ii} is the i th diagonal entry of D . Note that the diagonal entries of \tilde{D} are always positive because $\rho_1 > 0$ and therefore $\Theta^{l+1} \in \mathcal{S}^+$, as required. Finally, we remark that this step of ADMM coincides with the corresponding step of ADMM for graphical lasso and the reader can see Boyd et al. (2011) Section 6.6) for further details.

We turn now to step (2) of the algorithm. For a matrix Q with rows and columns indexed by $V = L \cup R$ we let $v(Q)$ be the vector defined as

$$v(Q)^T = [\text{vech}(Q_{LL})^T \quad \text{vech}(Q_{RR})^T \quad \text{vec}(Q_{LR})^T],$$

where $\text{vec}(\cdot)$ and $\text{vech}(\cdot)$ are the vectorization and half-vectorization operators respectively. Hence, we set

$$z = v(Z), \quad b^l = v(\Theta^l) + v(U^l) \quad \text{and} \quad F = \begin{bmatrix} I - I O \end{bmatrix}$$

where I is the identity matrix of dimension $q(q+1)/2$ and O is the $q(q+1)/2 \times q^2$ zero matrix. We can thus write the second step of the main ADMM algorithm in the form,

$$\arg \min_z \left(\frac{1}{2} \|z - b\|_2^2 + \lambda'_1 \|z\|_1 + \lambda'_2 \|Fz\|_1 \right), \quad (8)$$

where $\|\cdot\|_2$ is the Euclidean norm, $\lambda'_1 = \lambda_1/\rho_1$, $\lambda'_2 = \lambda_2/\rho_1$ and, to simplify the notation, we have omitted the superscript from b . Equation (8) shows that the optimization problem in the second step of ADMM is a special variant of the classical fused lasso called the fused lasso signal approximator. This allows us to exploit known results for this class of problems. More specifically, Friedman et al. (2007) Lemma A.1) showed that if a solution of (8) for $\lambda'_1 = 0$ and $\lambda'_2 > 0$ is known, then the solution for $\lambda'_1 > 0$ can be easily obtain in closed form through a soft-thresholding operation. Hence, we can focus on the solution of

$$\arg \min_z \left(\frac{1}{2} \|z - b\|_2^2 + \lambda'_2 \|Fz\|_1 \right), \quad (9)$$

that is a generalized lasso problem (Tibshirani & Taylor, 2011), and an ADMM algorithm for its solution can be found in Boyd et al. (2011) Section 6.4.1). Concretely, the ADMM algorithm iterates until convergence through the following steps:

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

In step (i), I is an identity matrix of appropriate dimension and $\rho_2 > 0$ is the ‘step size’ for the inner ADMM. In step (ii), $\mathcal{S}_\kappa(\cdot)$ is the soft-thresholding operator (see Boyd et al., 2011, Section 4.4.3). The vectors v and t mimic the role of Z and U of the outer ADMM and can be initialized to the zero vector. If we denote by $z_{[\lambda'_1=0, \lambda'_2]}$ the optimal solution at convergence of (9), then we can apply Friedman et al. (2007) Lemma A.1) and adjust $z_{[\lambda'_1=0, \lambda'_2]}$ element-wise to obtain the optimal solution of (8) for a given $\lambda'_1 > 0$ as $z_{[\lambda'_1, \lambda'_2]} = \mathcal{S}_{\lambda_1/\rho_1}(z_{[\lambda'_1=0, \lambda'_2]})$. The update Z^{l+1} for step (2) of the outer ADMM algorithm is thus the symmetric matrix such that $v(Z^{l+1}) = z_{[\lambda'_1, \lambda'_2]}$.

Finally, as stopping rule we use a tolerance check on the total relative change of the current estimate of the solution. In particular, if $\frac{\|\Theta^m - \Theta^{m-1}\|}{\|\Theta^{m-1}\|}$ is lower than the chosen tolerance, the algorithm is stopped and assumed to be at convergence.

6 | SIMULATION STUDY

We carry out a simulation study which aims to assess the performance of symmetric graphical lasso in a framework that mimics the structure of the R-fMRI data in Section 2. For this reason, we apply our procedure to simulated datasets sampled from normally distributed random vectors Y_V of $|V| = p = 70$ variables with $V = L \cup R$, as in Section 3. We consider two scenarios, denoted by A and B, that differ in their edge and symmetry degrees, with scenario A being sparser than B. The experiment is designed as follows. First, we randomly generate two undirected graphs, \mathcal{G}^A and \mathcal{G}^B , with edge degrees, computed as the ratio of the number of edges of the graph over the number of edges of the complete graph, $p(p-1)/2$, equal to $d^A = 23.1\%$ and $d^B = 31.6\%$ respectively. Next, for each scenario, we

randomly generate four positive definite concentration matrices $\Theta_{\text{true},i}^A$ and $\Theta_{\text{true},i}^B$, for $i = 1, \dots, 4$ with zero pattern corresponding to the missing edges of \mathcal{G}^A and \mathcal{G}^B respectively. The concentration matrices are constructed in order for a given proportion of randomly selected pairs of homologous concentrations across the two brain hemispheres to have the same value. More specifically, we focus on present edges, with nonzero concentration values, and the proportion of pairs of symmetric nonzero concentration is $d_{\text{sym}}^A = 10.8\%$ for scenario A and $d_{\text{sym}}^B = 30.1\%$ for scenario B. The eight generated concentration matrices characterize eight normal distributions with zero mean vector, and from each of these distributions we extract nine i.i.d. samples of size $n = 400$, so as to resemble the sample size of the data in Section 2.

In the penalized likelihood framework, a controversial question is how to choose the regularization parameter and several methods have been proposed in the literature. This issue is beyond the scope of this paper and, in order to avoid that the performance of symmetric graphical lasso is confounded by the choice of a selection method, we follow an ‘oracle’ procedure, described in the following. In each of the 72 generated datasets, we apply the graphical lasso and choose the value of λ_1 that produces a graph with an edge density equal to the density of the graph used to simulate the data. Next, conditional on the selected value of λ_1 , we apply the symmetric graphical lasso for 10 different logarithmically spaced values of λ_2 .

For every selected model, we consider some well-established measures to assess the performance in recovering the graph structure. The same quantities are then adapted to assess the performance in recovering the symmetric structure. Specifically, we compute the edge positive-predicted value (ePPV), also called precision, as the ratio between the number of true edges (eTP) and the number of edges (#edges) in the selected graph, and the symmetry positive-predicted value (sPPV) as the ratio between the number of true symmetric nonzero concentrations (sTP) and the number of nonzero symmetric concentrations (#symm) in the estimated concentration matrix. Furthermore, we compute the edge true-positive rate (eTPR) as the ratio between eTP and the number of edges (eP) in the true graph and the symmetry true-positive rate (sTPR) as the ratio between sTP and the number symmetric nonzero concentrations (sP) in the true concentration matrix. Similarly, we compute the edge true-negative rate (eTNR) and the symmetry true-negative rate (sTNR). In this way, we consider the quantities

$$\text{ePPV} = \frac{\text{eTP}}{\#\text{edges}}, \quad \text{eTPR} = \frac{\text{eTP}}{\text{eP}} \quad \text{and} \quad \text{eTNR} = \frac{\text{eTN}}{\text{eN}}, \quad (10)$$

which we use to assess how much the symmetric graphical lasso procedure recovers the graph structure, and the quantities

$$\text{sPPV} = \frac{\text{sTP}}{\#\text{symm}}, \quad \text{sTPR} = \frac{\text{sTP}}{\text{sP}} \quad \text{and} \quad \text{sTNR} = \frac{\text{sTN}}{\text{sN}}, \quad (11)$$

used to assess the ability of the symmetric graphical lasso to identify symmetries. It is worth remarking that symmetric graphical lasso tends to encourage equality between both zero and nonzero concentrations but we evaluate its performance only with respect to nonzero concentrations whose identification is of greater interest in applied contexts.

Table 1 summarizes the behaviour of the symmetric graphical lasso for a dataset in the scenario A. More specifically, we report the performance measures for the model selected by the graphical lasso and each of the 10 models obtained from the 10 values of λ_2 in the application of the symmetric graphical lasso. As shown in the first two lines of Table 1, the results for the graphical lasso and

TABLE 1 Performance measures in Equations (10) and (11) from the application of graphical lasso ($g\ell$) and symmetric graphical lasso ($sg\ell$) to one of the datasets in the simulated scenario A. All values, except #edges and #symm, are reported in percentages and the line in bold corresponds to the model for which $sTPR+sTNR$ is maximal

Method	Graph structure				Symmetric nonzero concentrations				
	ePPV	eTPR	eTNR	#edges	sPPV	sTPR	sTNR	sTPR+sTNR	#symm
$g\ell$	56.09	56.29	86.82	558	–	–	–	–	–
	56.01	56.12	86.82	557	0.00	0.00	100.00	100.00	0
	55.94	55.94	86.82	556	0.00	0.00	100.00	100.00	0
$sg\ell$ (increasing values of λ_2)	56.39	55.58	87.14	548	0.00	0.00	99.65	99.65	2
	57.89	55.40	87.95	532	25.00	1.67	98.77	100.44	8
	58.98	54.32	88.70	512	41.67	8.33	97.37	105.70	20
	60.16	54.86	89.13	507	42.86	20.00	95.26	115.26	39
	56.45	50.36	88.38	496	37.31	41.67	89.82	131.49	83
	53.51	48.02	87.52	499	28.87	46.67	82.28	128.95	129
	52.98	48.02	87.25	504	28.00	46.67	81.23	127.90	135
	52.98	48.02	87.25	504	28.00	46.67	81.23	127.90	135

the symmetric graphical lasso with the lowest value of λ_2 are virtually identical, which is expected given the equivalence between our proposed approach and graphical lasso when $\lambda_2 = 0$. Increasing values of λ_2 tend to increase the sparsity of the selected graph, which can be explained by the fact that symmetric graphical lasso encourages symmetries also for zero concentrations. As a consequence, increasing values of λ_2 tend to correspond to a moderate decrease in the values of ePPV and eTPR; on the other hand, eTNR tends to increase as λ_2 increases.

To choose the value of λ_2 we adopt the ‘oracle’ criterion that selects the model corresponding to the highest sum $sTPR+sTNR$, highlighted in bold in Table 1. We note that this corresponds to the sparsest graph, with 496 present edges. Focusing on the symmetric concentrations, we see, as expected, an increase in $sTPR$ for increasing values of λ_2 , with a steady decrease in terms of $sTNR$, with an appreciable value of 89.92 for the selected model. If we piece together both the considerations, we can see from Table 1 that the price in terms of eTPR and eTNR paid by using symmetric graphical lasso instead of the graphical lasso is worth the additional information we gain by recovering the symmetric structure of the two blocks of the concentration matrix, and the associated graph.

We apply this procedure to the 72 generated datasets thereby obtaining 72 models identified by graphical lasso and 72 models identified by symmetric graphical lasso. The result of these analyses is summarized in Table 2 and Figure 1. These show that, as far as the structure of the graph is concerned, the graphs obtained from symmetric graphical lasso have smaller values of eTPR and ePPV with respect to graphical lasso, and higher values of eTNR. However, the reduction in eTPR and ePPV is moderate in the sparser scenario A, about 10% for eTPR and 2.5% for ePPV, and, in fact, quite small in the denser scenario B, about 5% for eTPR and 1.5% for ePPV. Nonetheless, these moderate reductions in eTPR and ePPV are compensated by an increase in eTNR and, most importantly, by a satisfying performance in terms of recovery of symmetries. Interestingly, the symmetric graphical lasso seems to have a very similar behaviour in the two scenarios as far as $sTPR$ and $sTNR$ are concerned, but scenario B shows higher values of sPPV.

TABLE 2 Performance measures in Equations (10) and (11) from the application of graphical lasso (gl) and symmetric graphical lasso (sgl) to the 8×9 simulated datasets. Data are reported as mean (and standard deviation) computed across the nine replicated datasets for each of the eight environments. Values are reported in percentages for ePPV, eTPR, eTNR and density; average number of edges, average number of symmetric concentrations and their standard deviations are rounded

Environment		Graph structure					Symmetric nonzero concentrations				
Scenario	Code	Method	ePPV	eTPR	eTNR	#edges	Density	sPPV	sTPR	sTNR	#symm
A	A.1	gl	57.8 (1.3)	57.9 (1.3)	87.3 (0.4)	560 (1)	23.2 (0.1)	-	-	-	-
		sgl	54.8 (1.9)	47.4 (1.7)	88.24 (0.6)	483 (9)	20.0 (0.4)	30.6 (4.8)	49.5 (5.7)	87.9 (2.7)	99 (18)
	A.2	gl	57.6 (1.1)	57.6 (1.2)	87.6 (0.3)	547 (1)	22.7 (0.1)	-	-	-	-
		sgl	54.7 (2.1)	48.2 (2.4)	88.3 (0.6)	482 (12)	20. (0.5)	28.1 (5.9)	44.4 (4.4)	87.4 (3.6)	98 (21)
A	A.3	gl	57.6 (1.5)	57.6 (1.4)	87.3 (0.5)	556 (1)	23.0 (0.1)	-	-	-	-
		sgl	55.2 (2.5)	47.46 (2.4)	88.4 (1.0)	479 (23)	19.8 (1.0)	25.5 (4.2)	41.3 (6.2)	86.8 (3.3)	100 (21)
	A.4	gl	55.8 (1.2)	57.8 (1.3)	86.7 (0.4)	559 (1)	23.1 (0.1)	-	-	-	-
		sgl	54.0 (1.3)	47.7 (1.4)	87.8 (0.6)	494 (14)	20.5 (0.6)	29.5 (5.3)	43.9 (4.5)	88.6 (2.7)	92 (17)
B	B.1	gl	57.4 (0.6)	57.4 (0.6)	80.4 (0.3)	763 (2)	31.6 (0.1)	-	-	-	-
		sgl	56.0 (1.3)	52.3 (1.2)	81.1 (0.9)	711 (19)	29.5 (0.8)	54.1 (5.4)	40.9 (3.2)	85.9 (3.8)	137 (22)
	B.2	gl	57.8 (1.0)	57.8 (1.0)	81.1 (0.5)	748 (1)	31.0 (0.1)	-	-	-	-
		sgl	56.0 (1.6)	51.9 (1.8)	81.7 (0.7)	694 (12)	28.7 (0.5)	54.4 (2.4)	40.4 (3.6)	86.6 (1.9)	133 (14)
B	B.3	gl	58.1 (0.9)	58.2 (0.9)	80.8 (0.4)	758 (1)	31.4 (0.1)	-	-	-	-
		sgl	56.8 (1.1)	52.6 (0.7)	81.7 (0.8)	702 (15)	29.1 (0.6)	53.7 (5.5)	41.4 (3.3)	85.4 (4.0)	140 (23)
B	B.4	gl	58.2 (1.5)	58.2 (1.5)	81.3 (0.7)	745 (1)	30.9 (0.1)	-	-	-	-
		sgl	57.0 (1.7)	53.2 (1.9)	82.1 (0.8)	695 (14)	28.8 (0.6)	52.7 (3.7)	44.5 (4.5)	84.0 (2.7)	152 (17)




FIGURE 1 Difference between a given performance measure (top panel: ePPV; middle panel: eTPR; bottom panel: eTNR) computed on the model selected by symmetric graphical lasso and the same measure computed on the model selected by graphical lasso, for each of the 8×9 datasets. Every boxplot summarizes the nine datasets of the corresponding environment

7 | ANALYSIS OF R-fMRI DATA

The symmetric graphical lasso is applied to the data described in Section 2. The analysis is focused on two subjects, indexed as subject 18 and subject 22, who show homogeneous characteristics in terms of age and handedness but different diagnosis status. According to the available explanatory variables, subject 18 is 46 years old, right-handed, and healthy, whereas subject 22 is 42 years old, right-handed as well, but had a current and recurring diagnosis of drug abuse and mental disorders at the time of the fMRI scan recording.

We first account for the temporal dependence, with the tools and methodologies discussed in Section 4.2. In particular, for each subject, we obtain a matrix of residuals of dimension $n \times p$ from a VAR(1) model, a first-order score-driven model, and a Henderson filter with $h = 6$ (a 13-term weighted average). We then apply the symmetric graphical lasso to the residuals. As criteria for choosing the optimal value of λ_1 and λ_2 we use the Bayesian information criterion (BIC) and the extended BIC (eBIC; Foygel & Drton, 2010),

$$\text{eBIC} = -2l(\hat{\Theta}^{\text{mle}}) + \log(n)d + 4d\gamma \log(p),$$

where $l(\hat{\Theta}^{\text{mle}})$ and d denote the maximized log-likelihood function and the number of free parameters of the relevant model respectively. The eBIC depends on the parameter $\gamma \in [0; 1]$ that controls how much the criterion prefers simpler models. The limit case $\gamma = 0$ corresponds to the classical BIC. As suggested in Foygel and Drton (2010) we set $\gamma = 0.5$. For the computation of the maximum likelihood estimates within the family of RCON models, we use the `gRc` package for R (Højsgaard & Lauritzen, 2007). As a joint grid search over λ_1 and λ_2 could be computationally prohibitive (see also Danaher et al., 2014), we fix first λ_2 to a low value—close to zero—while performing a dense grid search over λ_1 . After selecting the best value of λ_1^* , a conditional sweep on a grid of 20 equally spaced values on a logarithmic scale for λ_2 can be performed to select the final pair of optimal values $(\lambda_1^*, \lambda_2^*)$.

The empirical results of graphical lasso (`gl`) and symmetric graphical lasso (`sgl`) fit on the residuals estimates for the two subjects are reported in Table 3 (vector autoregressive model, VAR), Table 4 (score-driven model, DCS) and Table 5 (Henderson filter, H13), according to the different

TABLE 3 Empirical results for graphical lasso (`gl`) symmetric graphical lasso (`sgl`) fit on residuals from a VAR(1) model on two subjects. The last three columns provide a description of the symmetric structure, that is the number of pairs of symmetric: (i) edges, (ii) off-diagonal nonzero concentrations and (iii) diagonal concentrations

Subject	Criterion	Method	#edges	Density	Pairs of symmetric		
					edges	nonzero concentrations	
					off-diagonal	diagonal	
18	BIC	gl	876	36.27%	120	–	–
		sgl	910	37.68%	171	87	6
	eBIC $\gamma=0.5$	gl	366	15.16%	48	–	–
		sgl	373	15.45%	92	89	31
22	BIC	gl	879	36.40%	128	–	–
		sgl	891	36.89%	202	160	14
	eBIC $\gamma=0.5$	gl	624	25.84%	91	–	–
		sgl	649	26.87%	158	149	27

TABLE 4 Empirical results for graphical lasso (gl) symmetric graphical lasso (sgl) fit on residuals from a score-driven model on two subjects. The last three columns provide a description of the symmetric structure, that is the number of pairs of symmetric: (i) edges, (ii) off-diagonal nonzero concentrations and (iii) diagonal concentrations

Subject	Criterion	Method	#edges	Density	Pairs of symmetric		
					edges	nonzero concentrations	
						off-diagonal	diagonal
18	BIC	gl	815	33.75%	115	–	–
		sgl	826	34.20%	195	175	22
	eBIC $\gamma=0.5$	gl	513	21.24%	72	–	–
		sgl	524	21.70%	123	120	27
22	BIC	gl	894	37.02%	121	–	–
		sgl	885	36.65%	167	110	8
	eBIC $\gamma=0.5$	gl	625	25.88%	91	–	–
		sgl	640	26.50%	149	134	14

TABLE 5 Empirical results for graphical lasso (gl) symmetric graphical lasso (sgl) fit on residuals from a local polynomial regression with $p = 3$ —Henderson filter weights—and $h = 6$ on two subjects. The last three columns provide a description of the symmetric structure, that is the number of pairs of symmetric: (i) edges, (ii) off-diagonal nonzero concentrations and (iii) diagonal concentrations

Subject	Criterion	Method	#edges	Density	Pairs of symmetric		
					edges	concentrations	
						off-diagonal	diagonal
18	BIC	gl	940	38.82%	133	–	–
		sgl	937	38.80%	188	135	15
	eBIC $\gamma=0.5$	gl	285	11.80%	33	–	–
		sgl	282	11.68%	42	21	10
22	BIC	gl	913	37.81%	119	–	–
		sgl	870	36.02%	173	123	10
	eBIC $\gamma=0.5$	gl	644	26.67%	81	–	–
		sgl	645	26.71%	156	145	26

filtering techniques. For sake of comparison, we report the results obtained by using both eBIC and its limit value BIC. However, we focus on the models obtained from the minimization of eBIC that are more parsimonious than the corresponding models selected from BIC, in particular for the symmetric graphical lasso. A more direct visual representation of the results detailed in Tables 3–5 is summarized in Figures 2 and 3, which provide a graphical representation of the brain symmetry structure. Specifically, the edges of the graphs encode symmetric nonzero off-diagonal concentrations whereas shaded vertices denote symmetric diagonal concentrations. While Figure 2 summarizes the results for the two subjects across the three filtering methods, Figure 3 reports the symmetries which turn out to be common to the three methods for subject 18 (left) and 22 (right). Moreover, in every Figure we




FIGURE 2 Graphical representation of the models in Tables 3, 4 and 5 obtained from the application of symmetric graphical lasso with eBIC. Edges encode symmetric nonzero concentrations whereas shaded vertices represent symmetric diagonal concentrations. From left to right: VAR(1); score-driven model; local polynomial regression—Henderson filter with $p = 3$ and $h = 6$. From top to bottom: subject 18; subject 22

omit non-symmetric edges from the visualization, in order to highlight the novelty aspect of the analysis and also to facilitate the reader with graphs that would otherwise be too densely plotted to be appreciable; nevertheless, the overall edge density for each result is reported in the summary Tables 3–5.

The first evident result is that, regardless of the filtering method, subject 22 shows a denser and more symmetric graph than subject 18. Also, the three models selected by eBIC for subject 22 have both similar densities and similar amount of symmetric edges and nonzero concentrations; compare, for instance, the 649 edges of subject 22 and the 373 edges of subject 18 in Table 3, related to VAR estimation. On the other hand, filtering has an impact on subject 18, as it is evident that the three methods induce a different density (from 11.68% of H13 in Table 5, to 21.70% of DCS in Table 4) and a different amount of symmetries in the resulting graph (from the 42 pairs of symmetric edges of H13 to the 123 of DCS). These considerations are ever more pronounced in Figure 2, which shows how the graph associated with local polynomial regression is the one exhibiting the lowest number of symmetric concentrations, both diagonal (shaded nodes) and off-diagonal (black solid lines). The second lowest number of symmetric concentrations is observed for the graph identified from the residuals of a VAR(1) model, whereas the residuals from the score-driven model bring to a graph with the highest number of symmetric off-diagonal concentrations. This is in line with the idea that a robust detrending method leaves more information in the residuals; on the other hand, methods that tend to overfit the data, such as a high-degree local polynomial regression, may allocate most of the data dependence structure in the signal component, rather than in the noise. This can be quantitatively

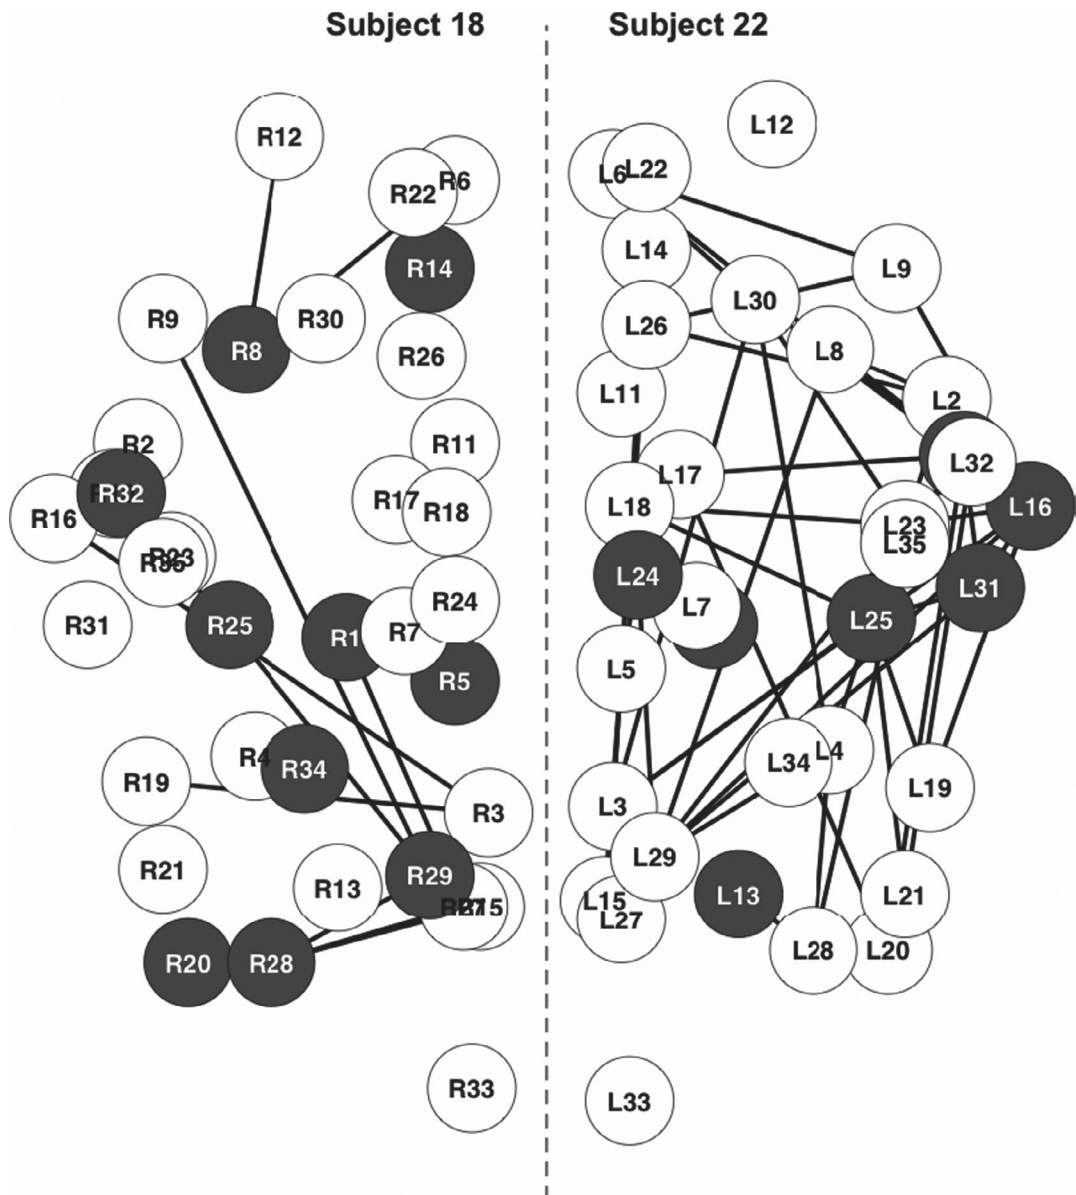


FIGURE 3 Common graph structure across the three different filter techniques obtained from the intersection of the graphs in Figure 2, for subject 18 (left of vertical dashed line) and subject 22 (right of vertical dashed line)

assessed by comparing the number of pairs of symmetrical concentrations, both off-diagonal and diagonal, reported as the last two columns of Tables 3–5.

Using different filtering techniques also allows one to extract the common information on the symmetric structure across the three filtering methods, by retaining the shared graph of the concentrations for each subject, that is, the graph resulting from the intersection of the three graphs in Figure 2. After this marginalization, and in order to give a better insight, we juxtapose them in Figure 3. In this side-by-side comparison, subject 22 exhibits a higher number of off-diagonal concentrations, and it is worth noting both subject seems to share approximately the same number of core symmetric diagonal values.

In conclusion, we may envisage two main empirical findings emerging from the analysis. The first evidence is concerned with the fact that the subject with a diagnosis of disorder shows a more symmetric brain structure than the healthy one. This is in line with several studies that are in favour of an evidence of lack of asymmetry in schizophrenic patients, see Sun et al. (2015). Nevertheless, the literature is quite controversial on this theme, see the review paper by Stephane et al. (2001) Section 3.1.4) and our results just refer to a pair of subjects. Second, the impact of the detrending method appears to be stronger in a subject who presents a less symmetric brain structure while it seems to be irrelevant in the subject who shows a more defined symmetric pattern. In any case, the combination of different filtering methods may shed light on the core symmetries that characterize the brain of different subjects.

8 | DISCUSSION

The procedures introduced in this paper have been developed with a focus on the identification of brain networks from R-fMRI data. In this respect, we have first considered the problem of removing the temporal dependence from data and then we have designed our methods and algorithms to suit the natural partition of the brain into hemispheres. Nevertheless, symmetric graphical lasso identifies a model within the class of RCON models and, as such, it has a potentially wider range of applications. To the best of our knowledge, the symmetric graphical lasso proposed in this contribution is the first instance where the lasso procedure is used to perform model selection within the class of coloured graphical models. In this way, we are able to identify symmetries characterized by equality constraints in the entries of the concentration matrix. Methods not explicitly designed to identify symmetric nonzero concentrations, such as the graphical lasso, can still identify symmetries in the graph structure but only in terms of edges being present or absent. Tables 3–5 compare the graphical lasso and the symmetric graphical lasso in their ability to identify symmetric edges, and show that our method encourages structural symmetries without affecting the graph sparsity.

Our proposed approach has an associated computational complexity comparable with that of methods using penalties of similar type, such as a conventional fused or group lasso (Danaher et al., 2014). The computational effort associated with using the ADMM algorithm to solve this class of convex quadratic programming optimizations is admittedly the eigendecomposition of a $p \times p$ matrix (Tibshirani et al., 2005). In the application considered in this manuscript, the dimensionality of the problem is bounded by the atlases used in resting state fMRI studies: usually, these atlases identify a number of ROIs close to $p = 70$, as in the data analysed in Section 5, which makes estimation times of our procedure not a concern. If indeed other applications are considered, such as cancer genomic where the number of variables is in the order of thousands, then exploring some potential pre-screening procedures would be a very interesting avenue to pursue, maybe adapting those mentioned in Danaher et al. (2014) and Yang et al. (2015) to our symmetric penalty.

Future research directions involve the specification of a convex penalty function that allows a more flexible specification of colour classes, which could affect other sub-components of the main concentration matrix, as well as consider different types of constraints.

SOFTWARE

The code implementing the ADMM algorithm described in this paper is available at the following GitHub repository: <https://github.com/savranciati/sgl>.

ACKNOWLEDGEMENTS

The authors thank Søren Højsgaard for the support provided with `gRc` package, Veronica Vinciotti for useful discussions, as well as both referees and associated editor for providing feedbacks that improved the quality of the paper. The authors AR and SR were supported by the Air Force Office of Scientific Research under award number FA9550-17-1-0039. We also thank Antonio Canale, Daniele Durante, Lucia Paci and Bruno Scarpa for introducing us to the challenging dataset analysed in the paper. These data were provided by Greg Kiar and Eric Bridgeford from NeuroData at Johns Hopkins University, who graciously pre-processed the raw DTI and R-fMRI imaging data available at http://fcon_1000.projects.nitrc.org/indi/CoRR/html/nki_1.html, using the pipelines ndmg and C-PAC.

ORCID

Saverio Ranciati  <http://orcid.org/0000-0001-7880-9465>

Alessandra Luati  <http://orcid.org/0000-0001-6407-9385>

REFERENCES

- Aran, D., Camarda, R., Odegaard, J., Paik, H., Oskotsky, B., Krings, G. et al. (2017) Comprehensive analysis of normal adjacent to tumor transcriptomes. *Nature communications*, 8, 1–14.
- Aston, J. & Kirsch, C. (2012) Evaluating stationarity via change-point alternatives with applications to fMRI data. *Annals of Applied Statistics*, 6, 1906–1948.
- Banerjee, O., Ghaoui, L.E. & d’Aspremont, A. (2008) Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *Journal of Machine Learning Research*, 9, 485–516.
- Biswal, B.B., Mennes, M., Zuo, X.-N., Gohel, S., Kelly, C., Smith, S.M. et al. (2010) Toward discovery science of human brain function. *PNAS*, 107, 4734–4739.
- Boyd, S., Parikh, N., Chu, E., Peleato, B. & Eckstein, J. (2011) Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3, 1–122.
- Bullmore, E.T. & Bassett, D.S. (2011) Brain graphs: graphical models of the human brain connectome. *Annual review of clinical psychology*, 7, 113–140.
- Bullmore, E., Fadili, J., Breakspear, M., Salvador, R., Suckling, J. & Brammer, M. (2003) Wavelets and statistical analysis of functional magnetic resonance images of the human brain. *Statistical methods in medical research*, 12, 375–399.
- Calvet, L.E., Czellar, V. & Ronchetti, E. (2015) Robust filtering. *Journal of the American Statistical Association*, 110, 1591–1606. Available from: <https://doi.org/10.1080/01621459.2014.983520>.
- Cleveland, W. & Loader, C. (1996) *Models for dependent time series*. Berlin: Springer.
- Cox, D.R. & Wermuth, N. (1996) *Multivariate dependencies: models, analysis and interpretation*. London: Chapman and Hall.
- Craddock, R.C., James, G.A., Holtzheimer, P.E., Hu, X.P. & Mayberg, H.S. (2012) A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Human Brain Mapping*, 33, 1914–1928.
- Creal, D., Koopman, S.J. & Lucas, A. (2013) Generalized autoregressive score models with applications. *Journal of Applied Econometrics*, 28, 777–795.
- Danaher, P., Wang, P. & Witten, D.M. (2014) The joint graphical lasso for inverse covariance estimation across multiple classes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76, 373–397.
- Dempster, A.P. (1972) Covariance selection. *Biometrics*, 28, 157–175.
- Desikan, R.S., Ségonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D. et al. (2006) An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage*, 31, 968–980. Available from: <https://www.sciencedirect.com/science/article/pii/S1053811906000437>.
- Drton, M. & Maathuis, M.H. (2017) Structure learning in graphical modeling. *Annual Review of Statistics and Its Application*, 4, 365–393.
- Foygel, R. & Drton, M. (2010) Extended bayesian information criteria for gaussian graphical models. In *Advances in neural information processing systems*, pp. 604–612.
- Friedman, J., Hastie, T., Höfling, H. & Tibshirani, R. (2007) Pathwise coordinate optimization. *The Annals of Applied Statistics*, 1, 302–332.

- Friedman, J., Hastie, T. & Tibshirani, R. (2008) Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9, 432–441.
- Gao, X. & Massam, H. (2015) Estimation of symmetry-constrained Gaussian graphical models: application to clustered dense networks. *Journal of Computational and Graphical Statistics*, 24, 909–929.
- Gehrmann, H. (2011) Lattices of graphical Gaussian models with symmetries. *Symmetry*, 3, 653–679.
- Højsgaard, S. & Lauritzen, S.L. (2007) Inference in graphical gaussian models with edge and vertex symmetries with the GRC package for R. *Journal of Statistical Software*, 23, 1–26.
- Højsgaard, S. & Lauritzen, S. (2008) Graphical Gaussian models with edge and vertex symmetries. *Journal of the Royal Statistical Society: Series B*, 70, 1005–1027. Available from: <https://rss.onlinelibrary.wiley.com/doi/10.1111/j.1467-9868.2008.00666.x>.
- Harvey, A.C. (2013) *Dynamic models for volatility and heavy tails: with applications to financial and economic time series, vol. 52*. Cambridge: Cambridge University Press.
- Harvey, A. & Luati, A. (2014) Filtering with heavy tails. *Journal of the American Statistical Association*, 109, 1112–1122.
- Henderson, R. (1916) Note on graduation by adjusted average. *Transactions of the Actuarial Society of America*, 17, 43–48.
- Hoefling, H. (2010) A path algorithm for the fused lasso signal approximator. *Journal of Computational and Graphical Statistics*, 19, 984–1006.
- Kundu, S. & Risk, B.B. (2020) Scalable bayesian matrix normal graphical models for brain functional networks. *Biometrics*, 1–12.
- Lauritzen, S.L. (1996) *Graphical Models*. Oxford: Oxford University Press.
- Li, Q., Gao, X. & Massam, H. (2020) Bayesian model selection approach for coloured graphical gaussian models. *Journal of Statistical Computation and Simulation*, 90, 2631–2654. Available from: <https://doi.org/10.1080/00949655.2020.1784175>.
- Lindquist, M.A. (2008) The statistical analysis of fMRI data. *Statistical Science*, 23, 439–464.
- Lund, T.E. (2006) Non-white noise in fMRI: Does modelling have an impact? *Neuroimage*, 29, 1639–1651.
- Marrelec, G., Krainik, A., Duffau, H., Péligrini-Issac, M., Lehericy, S., Doyon, J. et al. (2006) Partial correlation for functional brain interactivity investigation in functional MRI. *Neuroimage*, 32, 228–237.
- Massam, H., Li, Q. & Gao, X. (2018) Bayesian precision and covariance matrix estimation for graphical Gaussian models with edge and vertex symmetries. *Biometrika*, 105, 371–388. Available from: <https://doi.org/10.1093/biomet/asx084>.
- Poldrack, R.A. (2007) Region of interest analysis for fMRI. *Social Cognitive and Affective Neuroscience*, 2, 67–70.
- Proietti, T. & Luati, A. (2007) Least square regression: graduation and filters. In: Boumans, M. (Ed.) *Measurement in Economics: A Handbook*. chap. 16. Amsterdam: Academic Press, Elsevier, pp. 377–411.
- Proietti, T. & Luati, A. (2008) Real time estimation in local polynomial regression, with application to trend-cycle analysis. *Annals of Applied Statistics*, 2, 1523–1553.
- R Core Team (2020) *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. Available from: <https://www.R-project.org/>.
- Raemaekers, M., Schellekens, W., Petridou, N. & Ramsey, N. (2018) Knowing left from right: asymmetric functional connectivity during resting state. *Brain Structure and Function*, 223, 1909–1922.
- Schmal, C., Myung, J., Herzel, H. & Bordyugov, G. (2017) Moran's I quantifies spatio-temporal pattern formation in neural imaging data. *Bioinformatics*, 33, 3072–3079.
- Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., et al. (2009) Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences*, 106, 13040–13045. Available from: <https://www.pnas.org/content/106/31/13040>.
- Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., Beckmann, C.F., Nichols, T.E. et al. (2011) Network modelling methods for fMRI. *Neuroimage*, 54, 875–891.
- St. John, P. & Doyle, F. (2015) Quantifying stochastic noise in cultured circadian reporter cells. *PLOS Computational Biology*, 11, 1–17.
- Stephane, M., Barton, S. & Boutros, N. (2001) Auditory verbal hallucinations and dysfunction of the neural substrates of speech. *Schizophrenia Research*, 50, 61–78.
- Sun, Y., Chen, Y., Collinson, S.L., Bezerianos, A. & Sim, K. (2015) Reduced hemispheric asymmetry of brain anatomical networks is linked to schizophrenia: a connectome study. *Cerebral Cortex*, 27, 602–615.

- Tan, K.M., London, P., Mohan, K., Lee, S.-I., Fazel, M. & Witten, D. (2014) Learning graphical models with hubs. *Journal of Machine Learning Research*, 15, 3297–3331. Available from: <http://jmlr.org/papers/v15/tan14b.html>.
- Tibshirani, R.J. & Taylor, J. (2011) The solution path of the generalized lasso. *The Annals of Statistics*, 39, 1335–1371.
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J. & Knight, K. (2005) Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67, 91–108.
- Tunncliffe-Wilson, G., Reale, M. & Haywood, J. (2015) *Models for Dependent Time Series*. London: Chapman & Hall.
- Vinciotti, V., Augugliaro, L., Abbruzzo, A. & Wit, E.C. (2016) Model selection for factorial Gaussian graphical models with an application to dynamic regulatory networks. *Statistical Applications in Genetics and Molecular Biology*, 15, 193–212.
- Whittaker, J. (1990) *Graphical Models in Applied Multivariate Analysis*. Chichester: John Wiley & Sons.
- Worsley, K.J., Liao, C., Aston, J., Petre, V., Duncan, G., Morales, F. & Evans, A. (2002) A general statistical analysis for fMRI data. *Neuroimage*, 15, 1–15.
- Yang, S., Lu, Z., Shen, X., Wonka, P. & Ye, J. (2015) Fused multiple graphical lasso. *SIAM Journal on Optimization*, 25, 916–943.
- Yuan, M. & Lin, Y. (2007) Model selection and estimation in the Gaussian graphical model. *Biometrika*, 94, 19–35.
- Zhang, C. & Yu, T. (2008) Semiparametric detection of significant activation for brain fMRI. *Annals of Statistics*, 38, 1693–1725.
- Zhu, Y. & Cribben, I. (2018) Sparse graphical models for functional connectivity networks: best methods and the auto-correlation issue. *Brain connectivity*, 8, 139–165.

How to cite this article: Ranciati S, Roverato A, Luati A. Fused graphical lasso for brain networks with symmetries. *J R Stat Soc Series C*. 2021;00:1–24. <https://doi.org/10.1111/rssc.12514>