

LEARNING POPULATION AND SUBJECT-SPECIFIC BRAIN CONNECTIVITY NETWORKS VIA MIXED NEIGHBORHOOD SELECTION

BY RICARDO PIO MONTI^{*}, CHRISTOFOROS ANAGNOSTOPOULOS^{*} AND GIOVANNI MONTANA^{*,†}

Imperial College London^{} and King's College London[†]*

In neuroimaging data analysis, Gaussian graphical models are often used to model statistical dependencies across spatially remote brain regions known as functional connectivity. Typically, data is collected across a cohort of subjects and the scientific objectives consist of estimating population and subject-specific connectivity networks. A third objective that is often overlooked involves quantifying inter-subject variability and thus identifying regions or sub-networks that demonstrate heterogeneity across subjects. Such information is crucial to thoroughly understand the human connectome. We propose *Mixed Neighborhood Selection* to simultaneously address the three aforementioned objectives. By recasting covariance selection as a neighborhood selection problem we are able to efficiently learn the topology of each node. We introduce an additional mixed effect component to neighborhood selection to simultaneously estimate a graphical model for the population of subjects as well as for each individual subject. The proposed method is validated empirically through a series of simulations and applied to resting state data for healthy subjects taken from the ABIDE consortium.

1. Introduction. In this work we consider the estimation of multiple related Gaussian graphical models (GGMs) in the context of high-dimensional data. This topic has received considerable attention in recent years, however, an aspect which has been overlooked is that of understanding uncertainty across estimated GGMs. In many applications this is fundamental, yet it has received limited attention. To address this issue, we propose a novel model through which to estimate multiple GGMs which directly allows for variability to be identified on an edge-by-edge basis.

This work is motivated by the study of neuroimaging data where GGMs are often used to model statistical dependencies across spatially remote brain regions known as functional connectivity [Friston, 2011]. While traditional neuroimaging studies focused on the roles of specific brain regions, there has recently been a significant shift towards understanding the connectivity across regions [Smith, 2012]. This shift has been catalyzed by recent advances in the understanding of brain connectivity and its intimate relationship with diverse aspects of neuroscience; ranging from cognitive ability [Greicius et al., 2003, Gusnard and Raichle, 2001] to neurodegenerative and psychiatric diseases [Fox and Greicius, 2010]. As a result, the study of functional connectivity is of paramount

Keywords and phrases: functional connectivity, neuroimaging, graphical models, inter-subject variability

importance and provides a fundamental tool through which to understand the organization of the human brain [Van Den Heuvel and Pol, 2010]. It follows that quantifying heterogeneity across multiple subjects in order to relate such variability to physiological or genetic traits is an important problem in neuroscience [Dubois and Adolphs, 2016].

One of the hallmarks of neuroimaging data is its reproducible nature. Observed patterns in connectivity have been shown to demonstrate reproducible properties across subjects [Damoiseaux et al., 2006, Zuo et al., 2010]. The reproducible nature of functional connectivity networks motivates the need for novel methodologies with two overriding objectives. First, there is a need to exploit the presence of shared connectivity structure in order to yield more accurate network estimates for each subject. Second, there is also a critical need to understand and quantify inter-subject variability in the context of functional connectivity [Kelly et al., 2012, Mueller et al., 2013]. By quantifying variability across a cohort of subjects, such methods are able to untangle the characteristics which define a population from subject-specific idiosyncrasies. Such methods therefore open the door to a more intimate understanding of the properties of brain networks [Fallani et al., 2014].

Within the neuroimaging literature there are two main avenues of research which aim to address the aforementioned challenges. The first involves learning a separate GGM for each subject. While methods such as the Graphical Lasso [Friedman, Hastie and Tibshirani, 2008] are often employed to address the high-dimensional nature of the data, more sophisticated techniques are able to exploit the reproducible nature of connectivity via the introduction of novel regularization schemes [Varoquaux et al., 2010, Danaher, Wang and Witten, 2014]. Such methods propose to jointly estimate networks across subjects under some constraints over edges. In this manner, the edge structure of each subject is informed by the estimated structure of all remaining subjects.

The second approach is to learn a single GGM that is representative of the entire population of brain networks. Such a strategy is able to alleviate issues caused by the high-dimensional nature of the data by combining observations across subjects (albeit in a potentially naïve manner). However, the question of understanding variability across the population is often sidelined [Fallani et al., 2014].

The objective of this work is to reconcile the two popular approaches presented above, thus allowing for accurate network estimation at subject-specific and population levels while also quantifying variability present across a cohort. The proposed method, named Mixed Neighborhood Selection (MNS), is based on the neighborhood selection method introduced by Meinshausen and Bühlmann [2006]. By recasting covariance selection as a series of linear regression problems, neighborhood selection methods are able to learn the local network topology of each region. MNS extends neighborhood selection by incorporating an additional random effect component. This corresponds to learning a novel model for covariance structure across a cohort of subjects. In the proposed model the conditional independence structure for each subject is decomposed as the union of a population covariance structure together with subject-specific idiosyncrasies. Decomposing the edge structure in this manner serves to directly model inter-subject variability and provides a much richer model of functional connectivity. In particular, the proposed




FIG 1. *Toy motivating example to illustrate the capabilities of the proposed method. Networks were simulated with $p = 5$ nodes and with $n = 8$ observations per subject for $N = 4$ subjects. The networks for two of the subjects is shown in (a) while the networks for the remaining two is shown in (b). Solid and dashed edges indicate positive and negative partial correlations respectively. A significant proportion of the edges are shared across subjects with a single variable edge. The results for our proposed method are shown in (c): thin black lines indicate edges shared by the entire population while thick gray edges indicate highly variable edges. Estimated edge coefficients for edges A and B are shown as obtained by the MNS algorithm as well as by applying the Graphical Lasso to each dataset independently in (d): Dashed lines indicate the estimated population edge value while the solid line is the estimated probability density function of that edge based on the random effects. Triangular points indicate edge values as estimated by the Graphical Lasso while circular points indicate subject-specific MNS estimates.*

method is able to partition edges according to their reproducibility across the cohort. In doing so, MNS provides an additional layer of information which can be exploited to further understand functional connectivity. Moreover, by effectively differentiating between reproducible edges present across the entire cohort and highly variable edges, the proposed method is able to share information across subjects in a discriminative manner, leading to more reliable network estimates.

To illustrate the capabilities of the proposed method we present a brief motivating example, shown in Figure [1]. We consider a scenario where the population consists of four individuals whose functional connectivity networks share a common structure but also demonstrate some variability. In particular, one edge varies across subjects such that two subjects exhibit the functional connectivity shown in Figure [1a] and the remaining two Figure [1b]; the edge in question (edge A) is shown to vary from positive to negative across groups. In such a scenario, it is of scientific interest both to uncover the correct functional connectivity networks as well as to correctly identify edges which are variable within the population. MNS is designed to achieve both these goals. The results are shown in Figure [1c] where the black lines indicate edges shared across the entire population. The thick gray edges indicate *random effect* edges that demonstrate high variability. Figure [1d] shows the estimated edge coefficients for two edges of interest when estimated using the proposed method and the Graphical Lasso. Figure [1d] demonstrates that the proposed method is able to correctly recover covariance

structure as well as discriminate edges according to their reproducibility over the cohort. This is in contrast to traditional methods such as the Graphical Lasso¹. As a result, it follows that identifying variable edges is challenging even in low dimensions.

The proposed method is detailed in Section 2. We present an extensive simulation study in Section 3. The proposed method is applied to resting-state fMRI data from the ABIDE consortium in Section 4. We conclude with a discussion in Section 5.

2. Methods. To set notation, we assume we have access to fMRI time series across a cohort of N subjects. For the i th subject, it is assumed we observe an n -dimensional fMRI time series across p fixed regions of interest. We write $V = \{1, \dots, p\}$ to denote the set of regions or nodes and refer to the dataset for the i th subject by $X^{(i)} \in \mathbb{R}^{n \times p}$. Further, we write $X_v^{(i)} \in \mathbb{R}^{n \times 1}$ to denote the time-series for any node $v \in V$. Similarly, we let $X_{\setminus v}^{(i)} \in \mathbb{R}^{n \times (p-1)}$ denote the times-series across all remaining nodes.

Throughout this work it is assumed that the data of each subject follows a stationary multivariate Gaussian distribution. Since our primary interest is the estimation of functional connectivity networks, summarized in the inverse covariance matrix, we assume without loss of generality that each $X^{(i)}$ corresponds to n samples from a multivariate Gaussian distribution with zero mean and covariance given by $\Sigma^{(i)}$.

2.1. Modeling connectivity through GGMs. Functional connectivity can be measured through a myriad of techniques, the simplest of which involves estimating the correlation between the time series of any pair of nodes [Smith, 2012]. An alternative approach taken in this work is the use of partial correlations, where pairwise correlations between nodes are studied once the effects of all other nodes have been removed [Varoquaux and Craddock, 2013]. Such methods have been employed extensively within the neuroimaging literature [Smith et al., 2011].

Under the assumption of Gaussianity, estimating functional connectivity networks based on partial correlations is equivalent to learning the conditional independence structure for each subject. The conditional independence structure can be succinctly represented as a graphical model, $G^{(i)} = (V, E^{(i)})$, where the edge set, $E^{(i)}$, encodes conditional dependencies across a fixed set of nodes, V . Formally, the edge set summarizes the non-zero entries in the precision matrix, thus:

$$(2.1) \quad E^{(i)} = \text{supp} \left(\left(\Sigma^{(i)} \right)^{-1} \right) = \left\{ (j, k) : \left(\Sigma^{(i)} \right)^{-1}_{j,k} \neq 0 \right\}.$$

The resulting edges are taken to be indicative of functional relationships between spatially remote regions of the brain, allowing the estimated graphical model to be interpreted as a functional connectivity network.

Learning conditional independence graphs is a challenging problem that is further exacerbated by the high-dimensional nature of fMRI data. As a result, regularization is often introduced [Smith et al., 2011, Varoquaux and Craddock, 2013]. In particular,

¹The Graphical Lasso was run independently for each subject. The regularization parameter for each subject was selected using cross-validation.

the use of neighborhood selection methods introduced by [Meinshausen and Bühlmann \[2006\]](#) have been widely adopted within the neuroimaging community [[Lee et al., 2011](#), [Chung et al., 2015](#), [Belilovsky, Varoquaux and Blaschko, 2016](#)]. As these methods will form the backbone for the proposed method, we formally discuss neighborhood selection below.

2.1.1. Neighborhood selection. The intuition behind neighborhood selection stems from the fact that we may learn the conditional independence structure across all nodes by iteratively learning the conditional independence structure of each node. The latter is referred to as the neighborhood for each node $v \in V$. We write $\hat{n}e^{(i)}(v)$ to denote the estimated neighborhood of node v at the i th subject.

[Meinshausen and Bühlmann \[2006\]](#) propose to learn the neighborhood of each node $v \in V$ by considering the optimal prediction of $X_v^{(i)}$ given the time series of the remaining nodes. This results in the formulation of the following linear model for node v :

$$(2.2) \quad X_v^{(i)} = X_{\setminus v}^{(i)} \beta^{(i),v} + \epsilon^{(i),v}$$

where $\epsilon^{(i),v} \sim \mathcal{N}(0, \sigma^2 I)$ is white noise. In such a regression model it follows that nodes that are not in the neighborhood of v will be omitted from the set of optimal predictors. Thus neighborhood selection can be reformulated as subset selection in a linear model. The latter problem has received considerable attention, one notable solution being that of the Lasso [[Tibshirani, 1996](#)]. Briefly, the Lasso imposes a constraint on the ℓ_1 norm of the regression coefficients leading to parsimonious solutions whilst remaining convex.

The neighborhood selection approach described in [Meinshausen and Bühlmann \[2006\]](#) proceeds by solving the following convex optimization problem for each node v :

$$(2.3) \quad \hat{\beta}^{(i),v} = \underset{\beta^{(i),v} \in \mathbb{R}^{p-1}}{\operatorname{argmin}} \left\{ \frac{1}{2} \|X_v^{(i)} - X_{\setminus v}^{(i)} \beta^{(i),v}\|_2^2 + \lambda \|\beta^{(i),v}\|_1 \right\}.$$

Due to the parsimony property of the Lasso, some elements of $\hat{\beta}^{(i),v} \in \mathbb{R}^{p-1}$ will be shrunk to zero, effectively removing these nodes from the optimal prediction set. An estimate for the neighborhood of v is subsequently defined as:

$$(2.4) \quad \hat{n}e^{(i)}(v) = \left\{ u \in V \setminus \{v\} : \hat{\beta}_u^{(i),v} \neq 0 \right\}.$$

That is to say, the neighborhood of v is the set of all nodes included in the Lasso solution. Given an estimate of the neighborhood of all nodes, the edge structure for a graphical model can then be obtained using one of the following rules:

$$(2.5) \quad E_{OR}^{(i)} = \{(v, u) : u \in \hat{\beta}^{(i),v} \text{ or } v \in \hat{\beta}^{(i),u}\} \text{ or } E_{AND}^{(i)} = \{(v, u) : u \in \hat{\beta}^{(i),v} \text{ and } v \in \hat{\beta}^{(i),u}\}.$$

Throughout this work the AND rule was employed. This decision was based on the conservative nature of this rule, where edges are only reported if they are present in the neighborhoods of both of the relevant nodes.

2.2. *Mixed Neighborhood Selection.* In this section we formally detail the proposed methodology. We begin by describing a novel model for the covariance structure across a cohort of subjects in Section 2.2.1. The associated estimation framework and algorithm are detailed in Sections 2.2.2 and 2.2.3 respectively. Finally, parameter tuning is discussed in Section 2.2.4.

2.2.1. *A novel covariance model.* We propose to model the covariance structure for each subject as the union of a shared covariance structure together with subject-specific idiosyncrasies. The proposed model is based on the assumption that there exists a shared covariance structure which manifests itself across all subjects together with subject-specific deviations from this structure. The latter allows our model to accommodate inter-subject variability which cannot be ignored. As a result, we model the conditional independence structure of each subject as the union of the support of a sparse population network and a subject-specific network. Formally, the support for each subject's conditional independence structure, originally defined in equation (2.1), is modeled as:

$$(2.6) \quad E^{(i)} = E^{pop} \cup \tilde{E}^{(i)}$$

Here we interpret E^{pop} as the population edges which encode the conditional independence structure shared across the entire population. Under the assumption of Gaussianity, it follows that E^{pop} is associated with a population precision matrix, $\Theta^{pop} \in \mathbb{R}^{p \times p}$. From the perspective of covariance structure, E^{pop} encodes the maximal conditional dependence structure shared across all subjects. On the other hand, it is $\tilde{E}^{(i)}$ which encodes subject-specific deviations from the population covariance structure. We define $\tilde{E} = \bigcup_{i=1}^N \tilde{E}^{(i)}$ as the set of edges demonstrating variability across the entire population of N subjects. This variability may either be attributed to the nature of the edge (i.e., positive or negative partial correlations as in the motivating example described in Figure [1]) or partial presence of the edge (i.e., the edge is only present in some subjects).

The objective of the proposed method therefore corresponds to accurately identifying both E^{pop} and $\tilde{E}^{(i)}$. Given E^{pop} and $\tilde{E}^{(i)}$, one can infer $E^{(i)}$ and \tilde{E} . However, by focusing on E^{pop} and $\tilde{E}^{(i)}$, as opposed to directly considering subject-specific edges, a far richer description of functional architecture is obtained. In the case of the motivating example presented in Section 1, $\tilde{E} = \tilde{E}^{(i)} = \{A\}$ while the remaining edges are captured in E^{pop} . From the perspective of neuroimaging, partitioning edges in this manner is fundamental to further understanding the functional architecture of the brain [Kanai and Rees, 2011].

It is useful to note that the model described in equation (2.6) generalizes two typical approaches in the study of functional connectivity. The traditional method of estimating a single population network, Θ^{pop} , by concatenating data across all subjects is equivalent to the assumption that $\tilde{E} = \emptyset$. Such an approach is burdened by the sizable assumption that observations across all subjects share an identical conditional independence structure. Conversely, the approach of estimating a functional connectivity network for each subject independently corresponds to the assumption that $E^{pop} = \emptyset$. In such a scenario, there is no advantage to be gained by sharing information across subjects. Typically, we would expect the true underlying network structure across subjects to lie somewhere along the spectrum between these two extremes; thus justifying the proposed model.

2.2.2. *Estimation framework.* The covariance model described in Section 2.2.1 provides a rich framework through which to understand connectivity across a cohort of subjects. In order to learn the associated parameters, we look to extend neighborhood selection. As a result, we consider learning the neighborhood of node $v \in V$ over a cohort of N subjects by studying the following linear mixed effect model:

$$(2.7) \quad X_v^{(i)} = X_{\setminus v}^{(i)} \beta^v + X_{\setminus v}^{(i)} \tilde{b}^{(i),v} + \epsilon^{(i),v} \quad \text{for } i = 1, \dots, N.$$

Recall that $X_v^{(i)}$ denotes the time series at node v for subject i . The model described in equation (2.7) directly extends traditional neighborhood selection model by introducing random effect terms, $\tilde{b}^{(i),v}$, for each subject. We note that β^v corresponds to the shared population neighborhood.

The random effects are assumed to follow a multivariate Gaussian distribution, $\tilde{b}^{(i),v} \sim \mathcal{N}(0, \Phi^v)$, independently of $\epsilon^{(i),v}$. The choice of covariance structure for random effects is crucial to both estimating the model as well as to its interpretability. While it is possible to motivate many choices for $\Phi^v \in \mathbb{R}^{p-1 \times p-1}$, in this work we limit ourselves to the scenario where $\Phi^v = \sigma^2 \text{diag}(\sigma^{v2})$. Here $\sigma^v \in \mathbb{R}^{p-1}$ is a vector describing the standard deviation of the neighborhood of v across the cohort of N subjects. A large value of σ_u^v would be indicative of heterogeneity in the edge between nodes v and u .

For any node $v \in V$, the model described in equations (2.7) is easily interpretable. The population, or fixed effects, neighborhood is captured in β^v . These are the effects that are shared across the entire cohort of subjects and correspond to the set of edges in E^{pop} . Meanwhile, the random effects are able to capture subject-specific deviations from the population neighborhood and can thereby be employed to obtain a network for each subject. Formally, the random effects captured in σ^v correspond to the set of highly variable edges, \tilde{E} . Finally, we are also able to obtain estimates of $\tilde{b}^{(i),v}$, which can be employed to obtain subject-specific networks. These values correspond to the subject-specific idiosyncrasies, $\tilde{E}^{(i)}$.

2.2.3. *Estimation algorithm.* The model described in Section 2.2.2 contains the following parameters, $\phi^v = (\beta^v, \sigma^v, \sigma^2) \in \mathbb{R}^{2(p-1)+1}$, which must be estimated for each node $v \in V$. Given ϕ^v we can subsequently obtain the best linear unbiased predictions (BLUPs) for each of the random effects, $\tilde{b}^{(i),v}$, across subjects [Pinheiro and Bates, 2000]. In this work ϕ^v is estimated in a maximum likelihood framework, where the negative log-likelihood for node v is proportional to:

$$(2.8) \quad \mathcal{L}(\phi^v) = \sum_{i=1}^N \frac{1}{2} \log \det V_v^{(i)} + \frac{1}{2} \left(X_v^{(i)} - X_{\setminus v}^{(i)} \beta^v \right)^T V_v^{(i)-1} \left(X_v^{(i)} - X_{\setminus v}^{(i)} \beta^v \right),$$

where we define $V_v^{(i)}$ to be the variance structure for node v at subject i :

$$(2.9) \quad V_v^{(i)} = \sigma^2 \left(X_{\setminus v}^{(i)} \text{diag}(\sigma^{v2}) \left(X_{\setminus v}^{(i)} \right)^T + I \right).$$

where we write I to denote the identity matrix.

In order to simplify future discussion we re-parameterize the random effects component of the mixed effect model, described in equation (2.7), as follows:

$$(2.10) \quad \tilde{b}^{(i),v} = \text{diag}(\sigma^v) b^{(i),v},$$

where $b^{(i),v} \sim \mathcal{N}(0, \sigma^2 I)$.

In this work random effects are treated as latent variables and an EM algorithm is employed [McLachlan and Krishnan, 2007]. Fitting linear mixed effects models in this manner is a popular approach first posited by Dempster, Laird and Rubin [1977] and for which many efficient algorithms have been proposed [Meng and Van Dyk, 1998]. In the context of this work, such an approach will prove beneficial when regularization constraints are introduced. Assuming the random effects, $b^{(i),v}$, are observed the complete data log-likelihood is proportional to:

$$(2.11) \quad \mathcal{L}_c(\phi^v) = \sum_{i=1}^N \frac{n+p}{2} \log \sigma^2 + \frac{1}{2\sigma^2} \left(\left\| X_v^{(i)} - X_{\setminus v}^{(i)} \beta^v - X_{\setminus v}^{(i)} \text{diag}(\sigma^v) b^{(i),v} \right\|_2^2 + b^{(i),v T} b^{(i),v} \right).$$

Regularization is introduced for two reasons. First, sparse solutions remain feasible when only a reduced number of observations or subjects are available. Second, parsimonious solutions remain easily interpretable even in the presence of many nodes. As a result, we impose an ℓ_1 penalty on both the fixed as well as random effects. In terms of the random effects we penalize the variance terms, σ^v . Should a variance be shrunk to zero, the resulting random effect is effectively removed from the model. The introduction of sparsity inducing penalties yields the following penalized complete-data log-likelihood:

$$(2.12) \quad \mathcal{L}_c^{\lambda_1, \lambda_2}(\phi^v) = \mathcal{L}_c(\phi^v) + \lambda_1 \|\beta^v\|_1 + \lambda_2 \|\sigma^v\|_1,$$

where λ_1 and λ_2 are regularization parameters. Sparsity at the population level is enforced by λ_1 , while λ_2 encourages sparsity in the random effects by shrinking the standard deviation terms, σ^v .

The proposed EM algorithm involves iteratively computing the conditional expectation of latent variables, $Q(\phi; \phi^v)$, in our case the random effects, and minimizing the expected conditional log-likelihood with respect to parameters ϕ^v . The expectation step (E-step) can be computed in closed form as follows:

$$(2.13) \quad b^{(i),v} = \left(\text{diag}(\sigma^v) X_{\setminus v}^{(i) T} X_{\setminus v}^{(i)} \text{diag}(\sigma^v) + I \right)^{-1} X_{\setminus v}^{(i) T} \text{diag}(\sigma^v) \left(X_v^{(i)} - X_{\setminus v}^{(i)} \beta^v \right)$$

independently for each subject $i = 1, \dots, N$. It is clear from equation (2.13) that if σ_u^v is shrunk to zero then the u th entry of $b^{(i),v}$ will also be zero for all subjects.

In the minimization step (M-step) the latent variables, $b^{(i),v}$, are assumed to be observed. We therefore learn (β^v, σ^v) by solving the following convex problem:

$$(2.14) \quad (\beta^v, \sigma^v) = \underset{(\beta^v \in \mathbb{R}^p, \sigma^v \in \mathbb{R}_+^p)}{\text{argmin}} \left\{ \left\| X_v^{(i)} - X_{\setminus v}^{(i)} \beta^v - X_{\setminus v}^{(i)} \text{diag}(b^{(i),v}) \sigma^v \right\|_2^2 + \lambda_1 \|\beta^v\|_1 + \lambda_2 \|\sigma^v\|_1 \right\}.$$

We note that equation (2.14) is a Lasso problem with distinct regularization parameters applied to the fixed and random effects components respectively. A vast range of efficient

algorithms can be employed to solve equation (2.14). In this work gradient descent algorithms [Friedman et al., 2007] were employed. The motivation behind this choice was that due to the iterative nature of the EM algorithm employed, a Lasso problem must be solved at each iteration. It follows that while solutions from one iteration to the next will typically not be identical they will be similar. As a result, computational gains may be obtained by using past solutions as good initializations for the Lasso problem at each iteration. Gradient descent algorithms are particularly well-suited for such tasks. The MNS procedure is described in Algorithm 1.

Algorithm 1: Mixed Neighborhood Selection

Input: Data across N subjects, $\{X^{(i)} : i = 1, \dots, N\}$, regularization parameters, λ_1, λ_2 .

```

1 begin
2   for  $v \in \{1, \dots, V\}$  do
3     Define initial estimates:  $\beta^v = \mathbf{0}$ ,  $\sigma^v = \mathbf{1}$ ,  $\sigma = 1$  and  $b^{(i),v} = \mathbf{0}$ 
4     while not converged do
5       Update  $(\beta^v, \sigma^v)$  by solving equation (2.14) // M-step
6       Estimate latent variables using equation (2.13) for  $i \in \{1, \dots, N\}$  // E-step
7     Store  $\beta^v$ ,  $\sigma^v$  and  $\{b^{(i),v}\}_{i=1}^N$ 
8      $E^{pop} = \{(u, v) : \beta_u^v \neq 0 \text{ and } \beta_v^u \neq 0\}$ 
9      $\tilde{E} = \{(u, v) : \sigma_u^v \neq 0 \text{ and } \sigma_v^u \neq 0\}$ 
10     $\tilde{E}^{(i)} = \{(u, v) : b_u^{(i),v} \neq 0 \text{ and } b_v^{(i),u} \neq 0\}$ 
11 return  $E^{pop}$ ,  $\tilde{E}$  and  $\tilde{E}^{(i)}$  for  $i = 1, \dots, N$ 

```

2.2.4. Parameter tuning. The proposed method requires the tuning of two regularization parameters which govern the nature of the estimated population and subject-specific networks respectively. Large values of λ_1 will lead to sparse networks at the population level. Conversely, selecting large λ_2 will penalize the variance of the random effects leading to sparse subject-specific contributions to covariance structure.

Moreover, in the class of models considered in this work each covariate may contribute to the fixed as well as random effect structure. Such a parameterization may lead to problems regarding the interpretability of estimated models. For example, over-penalizing the fixed effects may lead to over-estimation of the random effect variances as compensation [Schelldorfer, Bühlmann and Van de Geer, 2011]. The choice of regularization parameters is therefore a delicate issue which must be handled with care.

While information theoretic methods such as the Bayesian Information Criterion may be employed for the purpose of tuning regularization parameters, in this work we employ cross-validation. We note that such an approach is frequently employed within neuroimaging applications [Varoquaux et al., 2010, Varoquaux and Craddock, 2013]. Formally, the data across all subjects is divided into K folds. For each fold, the data from the remaining $K - 1$ folds is employed to fit the penalized linear mixed model described in Section 2.2.2. The resulting model is then used to predict the unseen data and the mean square error is noted. This procedure is repeated over all nodes and across

all subjects, with the parameters minimizing total mean square error selected.

3. Simulation Study. In this section we evaluate the performance of the proposed method using simulated data that is representative of functional imaging data. We assess the empirical performance of the MNS algorithm in three distinct settings which correspond to correctly reporting the edge structure of the population, subject-specific and highly variable network edges respectively. The first task corresponds to correctly recovering E^{pop} while the second requires learning subject-specific edge structure, $E^{(i)}$, defined in equation (2.1). Finally, the task of recovering variable edges is equivalent to learning the set of variable edges, \tilde{E} .

3.1. Network simulation. In order to perform such a study we require a method through which to simulate population and subject-specific networks. While numerous algorithms have been proposed to generate random individual networks, there has been limited work on algorithms to simulate multiple clustered networks. Notably, there is no documented method through which to generate networks from a cohort of related subjects that demonstrate the characteristics observed in real fMRI data; namely a shared core structure which is reproducible across all subjects together with significant inter-subject variability in the remaining edges [Bullmore and Sporns, 2009].

To address this issue we propose a novel method of simulating networks. The proposed algorithm is motivated by an exploratory data analysis of resting state fMRI data. We briefly outline the proposed algorithm in this section with further details provided in Supplement B.

The underlying idea behind the proposed network simulation method is that key properties observed in fMRI data should be present. As such, the proposed method consists of a set of population edges E^{pop} which are sampled according to the preferential attachment model of Barabási and Albert [1999]. These edges constitute the core, reproducible connectivity structure which will be present across all subjects. Thereafter, a set of variable edges, \tilde{E} , is selected uniformly at random across all edges. For each subject, edges in \tilde{E} , are included in the subject-specific network, $\tilde{E}^{(i)}$, with some fixed probability τ , yielding clustered networks where there is a clear shared structure together with diverse subject-specific idiosyncrasies.

The proposed method was employed to simulate synthetic data for a cohort of $N = 10$ subjects. The number of nodes was fixed at $p = 50$. For each subject, data consisted of n samples from a multivariate Gaussian with zero mean and covariance specified by $\tilde{E}^{(i)}$. Data was simulated with a varying number of observations per subject, $n \in \{50, 100, 200\}$.

3.2. Alternative models. Throughout this simulation the performance of the MNS algorithm was benchmarked against the current state of the art in each of the three settings described above. In the case of estimating the population network, the *Graphical Lasso* (Glasso) [Friedman, Hastie and Tibshirani, 2008] was employed. Such an approach has been used extensively in the neuroimaging community to learn functional connectivity networks across populations [Smith et al., 2011]. An approach based on resampling

and randomization was also employed. This approach, which we refer to as the *Stability* approach, is outlined in Supplement A. We note that while this approach is inspired by the recently proposed R^3 method of Narayan, Allen and Tomson [2015], the objective here is different.

The problem of estimating subject-specific functional connectivity networks has received considerable attention. In this simulation study we compare the performance of the proposed method with the two penalized likelihood methods presented in Varoquaux et al. [2010] and Danaher, Wang and Witten [2014]. Each of these methods can be seen as a special case of the Joint Graphical Lasso (JGL) framework proposed by Danaher, Wang and Witten [2014], as a result we refer to each as the the JGL-Group or the JGL-Fused algorithms respectively. The Glasso algorithm is also employed in this context.

As far as we are aware there are no alternative methods available which address the problem of recovering highly variable edges. The aforementioned *Stability* approach was therefore employed as a benchmark.

3.3. Performance measures. Throughout this simulation the task of recovering covariance structure is treated as a binary classification task. Thus performance is measured according to the proportion of edges which are correctly reported as being either present or absent. To compare performance across various algorithms we employ receiver operating characteristic (ROC) curves, which illustrate the performance of a binary classifier by plotting the true positive rate against false positive rate across a range of regularization parameters [Krzanowski and Hand, 2009].

The use of ROC curves requires a single, sparsity-inducing parameter to be varied across a range of possible values. In the case of the MNS algorithm both the population and subject-specific parameters can affect sparsity. As a result, we look to reparameterize the MNS penalty as follows:

$$(3.1) \quad \lambda_1 = \alpha \lambda$$

$$(3.2) \quad \lambda_2 = \sqrt{2}(1 - \alpha)\lambda$$

where α controls the ratio of sparsity between the population and subject-specific contributions and λ the overall sparsity. Thus α is fixed allowing λ to vary. While no such adjustments are needed in the case of the JGL-Fused algorithm, we follow the same parameterization described in equations (3.1) and (3.2) in the case of the JGL-Group algorithm².

3.4. Simulation results. In this section we present the results to the simulation study described above. We begin by first considering performance in the context of recovering the set of variable edges in Section 3.4.1. This problem is fundamental within many neuroscientific applications, however, it has received limited attention to date [Kelly et al., 2012]. Results for the more frequently studied problems of recovering population and subject covariance structure are presented in Section 3.4.2 and 3.4.3 respectively.

²Note this same parameterization was employed in Danaher, Wang and Witten [2014]

Throughout this simulation the MNS algorithm was run with $\alpha = 0.25$ while sparsity parameter λ varied as described in equations (3.1) and (3.2). The same parameterization was employed for the JGL-Group algorithm with $\alpha = 0.15$ selected. In the case of the JGL-Fused algorithm, $\lambda_2 = 0.2$ was employed. Finally, the *Stability* algorithm was run with $B = 10,000$ bootstrap iterations per subject and $c = 0.25$.




FIG 2. Simulation results for all five algorithms across all tasks. Recovery of variable edges is considered in the top panel, population network recovery is shown in the middle panel and finally the bottom panel shows subject-specific network recovery. This simulation was performed with $p = 50$ nodes and $n \in \{50, 100, 200\}$ observations.

3.4.1. *Variable network recovery.* Understanding variability in covariance structure across a cohort of subjects is a fundamental problem in neuroscience. In particular, understanding whether this variation can be attributed to phenotypic characteristics or other sources of noise is crucial in further understanding the human connectome.

The results shown in the top panel of Figure [2] demonstrate that the proposed MNS algorithm is able to accurately identify edges which demonstrate variability across a cohort of subjects, in contrast to the *Stability* method. Briefly, the *Stability* method (described in Supplement A) treats the presence or absence of edges at a subject level as a Bernoulli random variable. A hierarchical random effects model is then proposed to

model the presence or absence of an edge across all subjects. The resulting estimate of the edge variability is then employed to discriminate between variable and non-variable edges. The *Stability* method therefore corresponds to a two-step procedure where variability is only studied after networks have been estimated for subjects independently. Conversely, the proposed method simultaneously learns subject-specific, population and variable networks, resulting in significant improvements in performance. Further results are given in Table [1] where the true positive rate (TPR) and false positive rate (FPR) are reported for selected regularization parameters.

3.4.2. Population network recovery. Obtaining an accurate understanding of a population level covariance structure is a challenging problem due to the high inter-subject variability. As mentioned previously, it is imperative to differentiate between subject-specific idiosyncrasies and behavior which is reproducible across the entire cohort. A popular approach often taken in neuroimaging studies is to estimate a single network using data from all subjects, thus effectively concatenating all data. Such an approach corresponds to the sizable assumption that $\tilde{E} = \emptyset$. This approach is included in this simulation together with the aforementioned *Stability* approach.

Results are shown in the middle panel of Figure [2]. It is interesting to note that for small sample sizes (i.e., $n = 50$ or $n = 100$) the *Stability* approach is out-performed by the Glasso. As mentioned in Section 3.4.1, we attribute this drop in performance to the two-step design of the *Stability* method where information is only shared across subjects *after* networks have been estimated. It is only when the number of observations increases that reliable estimates of uncertainty can be obtained. Conversely, the difference in performance between the Glasso algorithm and the MNS algorithm is due to the presence of heterogeneous edges, implying $\tilde{E} \neq \emptyset$. Thus, by providing a more sophisticated model for inter-subject variability, the MNS algorithm is able to obtain more reliable population network estimates.

3.4.3. Subject-specific network recovery. Finally, we consider the recovery of subject-specific networks. This problem has received considerable attention in recent years and a range of methods have been proposed. The underlying theme in these methods revolves around effectively sharing information across subjects. In the case of the methods proposed by Varoquaux et al. [2010] and Danaher, Wang and Witten [2014], information is shared across subjects via the introduction of regularization penalties over the edge structure. In this manner, the covariance structure of an individual subject is informed by the estimated covariance structure across all remaining subjects. However, a short-coming of the aforementioned methods is that regularization is applied in an indiscriminate manner. By enforcing either a group or fused lasso penalty on all entries of precision matrices, such methods effectively encourage information to be shared homogeneously across all edges. We envisage a scenario where edges may be ordered according to their variability. For example, within the context of functional connectivity networks there is evidence to suggest that variability in connectivity is directly modulated by factors such as the distance between regions [Power et al., 2012]. The proposed MNS algorithm is able to address precisely this issue. By discriminating between subject-specific and

Algorithm	n	Population		Subject		Variance	
		TPR	FPR	TPR	FPR	TPR	FPR
MNS	50	0.76	0.12	0.75	0.33	0.54	0.06
	100	0.77	0.11	0.80	0.32	0.70	0.03
	200	0.75	0.11	0.82	0.30	0.79	0.02
Glasso	50	0.69	0.27	0.88	0.83	NA	
	100	0.70	0.27	0.83	0.66		
	200	0.68	0.27	0.85	0.58		
Stability	50	0.56	0.20	NA		0.54	0.24
	100	0.59	0.20			0.64	0.18
	200	0.78	0.35			0.71	0.15
JGL Group	50	NA		0.86	0.71	NA	
	100			0.83	0.62		
	200			0.82	0.57		
JGL Fused	50	NA		0.78	0.51	NA	
	100			0.79	0.51		
	200			0.79	0.50		

TABLE 1

Performance of all five algorithms. The true positive rate (TPR) and false positive rate (FPR) is reported for each of the three tasks: recovering population, subject and variance networks.

population edges, it is able to effectively vary the how extensively information is shared across subjects on an edge-by-edge basis. As a result, the MNS algorithm is able to more reliably recover subject-specific covariance structure.

3.5. Further experiments. One of the assumptions of the proposed method is that data follow a multivariate Gaussian distribution. While this assumption is commonplace in the analysis of fMRI data [Lindquist, 2008], we also consider the performance of the MNS algorithm in the context of non-Gaussian data. To study the robustness of the MNS algorithm, the simulation study presented above was repeated with data generated according to a multivariate t -distribution. Detailed results are reported in Supplement C. The results indicate that in comparison to alternative methods, the MNS algorithm is robust in the presence of non-Gaussian data. We attribute this behavior to the fact that the covariance model underlying the MNS algorithm explicitly models heterogeneity over subjects, thus allowing the MNS algorithm to better tolerate contaminated data.

Furthermore, simulating networks as described in Section 3.1 is one of many possible methods which could be employed. To provide a thorough and fair comparison an additional simulation was also performed where networks were simulated as described in Danaher, Wang and Witten [2014]. This simulation was proposed with the objective of providing empirical evidence regarding how accurately subject-specific networks could be reported. It is therefore not well suited for examining how reliably the population or variance networks can be reported. The results are presented in Supplement C.

4. Application. In this section the proposed MNS algorithm is applied to resting-state fMRI data from the ABIDE consortium [Di Martino et al., 2014]. While the ABIDE dataset contains data corresponding to healthy subjects and Autism Spectrum Disorder (ASD) subjects, we chose only to study healthy controls here as the focus of this work

consisted in fully understanding uncertainty across a single population of subjects. The decision to study the ABIDE dataset in this manner was motivated by the fact that it is an open-source dataset which has been previously studied in the context of functional connectivity. Data from the University of Utah School of Medicine (USM) site was considered here, a choice motivated by results suggesting the USM site contained high-quality data [Nielsen et al., 2013]. The data therefore consisted of 43 healthy subjects with ages ranging from 8 to 40 years old.

4.1. *Data acquisition and processing.* Data was downloaded from the Autism Brain Imaging Data Exchange (ABIDE) [Di Martino et al., 2014]. Data were preprocessed via a CPAC³ pipeline from the ABIDE repository. Preprocessing involved slice time correction, motion correction and intensity normalization followed by regression of motion parameters. Linear and quadratic trends were removed from frequency drifts. Mean time-courses were then extracted from 116 regions defined by the Automated Anatomical Labeling (AAL) atlas, yielding 200 observations over 116 nodes for each subject.

4.2. *Results.* The MNS algorithm requires the specification of two regularization parameters, each of which controls the population and subject-specific topology of each node. As discussed in Section 2.2.4, parameters were selected on the basis of a 10-fold cross-validation framework.

One of the advantages of the proposed MNS algorithm is that it is able to simultaneously estimate both a population network, corresponding to reproducible edges which are present across the entire cohort of subjects, as well as a network quantifying variability on an edge-by-edge basis. The latter network is able to succinctly summarize variability across a cohort of subjects. Finally, the MNS algorithm also yields estimates of subject-specific connectivity networks, allowing connectivity to be studied in three distinct yet complimentary approaches which we discuss below.

The top panel of Figure [3] shows the estimated population network, indicating the edges which were identified as being consistently present across the entire cohort. The network has an estimated edge density of around 10%, indicating that approximately 90% of edges are not present across all subjects. We also note there is strong inter-hemispheric coupling as would be expected in resting-state connectivity. More importantly, the bottom panel of Figure [3] shows the estimated variability network, corresponding to the collection of edges that were identified as demonstrating variability across the cohort of subjects. In the case of the variability network, the edge thickness is proportional to the estimated variance of the random effect. We note that there is strong inter-hemispheric variability, in particular between the left and frontal gyrus as well as between the left and right postcentral gyrus. There also appears to be a region of variability centered around the cerebellum. We note that the aforementioned regions are in brain areas with relatively high susceptibility to artifact and sensitivity to changes in brain shape, such as the medial prefrontal cortex [Nielsen et al., 2013].

One of the strengths of the MNS algorithm is that this variability can be further

³see <http://fcp-indi.github.com> for further details




FIG 3. *Estimated population network (top) and variable edge network (bottom). Edge thickness is proportional to the magnitude of the partial correlation across edges (or variance in the case of variance network). The left panel displays the coronal view, showing the strong inter-hemispheric connectivity. The Saggital view is shown in the middle panel. We note there is strong connectivity centered around the frontal cortex (located to the right) as well as high variability in the cerebellum (located to the bottom left). Finally, the axial view is provided in the right panel where strong inter-hemispheric connectivity is again visible across the frontal gyrus (located towards the top) as well as the postcentral gyrus (located midway down the brain). Further details regarding the brain regions shown is provided in Supplement D.*

studied to obtain a deeper understanding regarding the characteristics that define differences in connectivity over a cohort. In Figure [4] the variability of two edges is studied in detail. The edges correspond to inter-hemispheric connectivity between the left and right frontal gyrus and postcentral gyrus respectively. The histograms capture the distribution of estimated edges across the cohort of subjects. As these edges are estimated to be variable across subjects, the proposed method learns a distinct partial correlation for each subject. The color of histograms visualizes the mean age of subjects within each bin, indicating that bilateral connectivity across the frontal gyrus and postcentral gyrus increases significantly with age⁴. At a higher level, these results are consistent with previous literature which suggests that connectivity increases across distant brain regions during development [Fair et al., 2009] and serve to highlight the maturation of a dual-control system within brain networks [Fair et al., 2007].

Finally, the MNS algorithm also provides estimates of subject-specific functional connectivity networks. As a result, the proposed method can be used to study connectivity on a subject-by-subject basis. Here we study various properties associated with the estimated functional connectivity network for each subject, in particular we look to study potential changes in connectivity that are associated with the age of subjects.

It has been suggested that the structure of functional connectivity networks in children

⁴Significant at the $\alpha = 1\%$ level using Spearman's rank correlation coefficient.




FIG 4. The two histograms show the estimated partial correlations across two edges highlighted on the left. Each of the histograms shows the distribution over estimated partial correlations between the left and right frontal gyrus and postcentral gyrus respectively. The color of bins is indicative of mean age of subjects within that bin.

is driven by anatomical proximity, with a high connectivity between spatially adjacent regions, while the corresponding structure in adults reflects the integration of remote brain regions. To study this hypothesis the average distance between functionally connected brain regions was estimated on a subject-by-subject basis (i.e., using the subject-specific estimates of functional connectivity). The left panel of Figure [5] shows the average distance between functionally connected brain regions as a function of the subjects age. We note there is a significant positive correlation at the $\alpha = 1\%$ level using Spearman's rank-order correlation, placing the results in line with other results in the literature [Fair et al., 2007, 2009].

To obtain a more detailed understanding of changes occurring in the functional connectivity two further network statistics are studied; the mean betweenness centrality of nodes and the transitivity of estimated networks. The betweenness centrality of a node is a measure of its centrality or importance within a network [Rubinov and Sporns, 2010] and is defined as the fraction of all shortest paths passing through a node. Nodes with high betweenness centrality are seen to be bridge connections across many nodes, thereby making their presence in a network important. The mean betweenness centrality across all nodes can be interpreted as a measure of the efficiency in a network. On the other hand, transitivity is a measure of network segregation which quantifies the presence of clusters in the network. In the context of functional connectivity networks, high transitivity suggests an organization of statistical dependencies indicative of segregated neural processing [Rubinov and Sporns, 2010].

The middle and right panels of Figure [5] shows the mean betweenness centrality and the transitivity of estimated networks as a function of age. These results would indicate an increase in network segregation and specialization during development, a finding that is consistent with previous literature [Fair et al., 2007, 2009].

5. Discussion. We have considered the task of estimating multiple related GGMs. In particular, we have focused on three closely related challenges: recovering population




FIG 5. Multiple network properties are plotted as a function of the subject ages. Estimated p -values, obtained using Spearman’s rank-order measure of correlation, are shown in bottom right corner of each plot. **Left:** the anatomical distance between functionally connected regions. **Middle:** the mean betweenness centrality of nodes. **Right:** the transitivity (i.e., clustering coefficient).

and subject-specific covariance structure as well as identifying the set of edges demonstrating heterogeneity across networks. The latter is fundamental in the context of many applications, yet it has received limited attention. The proposed methodology is able to simultaneously address all three aforementioned challenges by considering a novel model for covariance structure across a cohort of subjects. Formally, the proposed model looks to decompose covariance structure as the union of population effects, which are reproducible across subjects, with subject-specific idiosyncrasies.

The underlying covariance model results in several important benefits, principal among which is the ability of the MNS algorithm to accurately identify heterogeneous edges. As a result, the MNS algorithm is able to borrow information across subjects in a discriminative manner. This is in contrast to many of the current methodologies which share information in an indiscriminate fashion (e.g., via the use of regularization penalties whose parameterization is fixed across edges).

The capabilities of the proposed MNS algorithm have been demonstrated using both simulated as well as resting-state data taken from the ABIDE consortium [Di Martino et al., 2014]. Throughout the simulation study, care was taken to ensure that the underlying covariance structure closely resembled the frequently reported properties of fMRI data as well as to consider the robustness of the proposed algorithm.

The MNS algorithm requires the specification of two regularization parameters, λ_1 and λ_2 , each of which has a natural interpretation. The first parameter controls the sparsity in the population node topologies while the second controls the sparsity of the subject-specific edges. We employ a cross-validation to tune both parameters, as is frequently the case in the context neuroimaging data analysis [Varoquaux and Craddock, 2013]. The MNS algorithm, together with network simulation methods described in this work have been implemented as an R package named MNS, which can be downloaded from the Comprehensive R Archive Network (CRAN).

In conclusion, the MNS algorithm provides a novel methodology through which to understand variability across multiple related GGMs. Furthermore, by providing a refined model for the covariance structure, the proposed method is also able to accurately recover both population and subject-specific functional connectivity networks.

Acknowledgements. We thank Ai Chung, Dean Bodenham, Georg Hahn and Romy Lorenz for helpful discussions. We also thank the Editor, the Associate Editor and two referees for their valuable comments which helped significantly improve the manuscript.

References.

- BARABÁSI, A.-L. and ALBERT, R. (1999). Emergence of scaling in random networks. *Science* **286** 509–512.
- BELILOVSKY, E., VAROQUAUX, G. and BLASCHKO, M. (2016). Testing for differences in Gaussian Graphical Models: Applications to brain connectivity. In *Neural Information Processing Systems* 595–603.
- BULLMORE, E. and SPORNS, O. (2009). Complex brain networks: Graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience* **10** 186–198.
- CHUNG, M., HANSON, J., YE, J., DAVIDSON, R. and POLLAK, S. (2015). Persistent homology in sparse regression and its application to brain morphometry. *IEEE Transactions on Medical Imaging* **34** 1928–1939.
- DAMOISEAUX, J., ROMBOUTS, S., BARKHOF, F., SCHELTENS, P., STAM, C., SMITH, S. and BECKMANN, C. (2006). Consistent resting-state networks across healthy subjects. *Proceedings of the National Academy of Sciences* **103** 13848–13853.
- 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.
- DEMPSTER, A., LAIRD, N. and RUBIN, D. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B* **39** 1–38.
- DI MARTINO, A., YAN, C., LI, Q., DENIO, E., CASTELLANOS, F., ALAERTS, K., ANDERSON, J., ASSAF, M., BOOKHEIMER, S. and DAPRETTO, M. (2014). The Autism Brain Imaging Data Exchange: Towards a large-scale evaluation of the intrinsic brain architecture in Autism. *Molecular Psychiatry* **19** 659–667.
- DUBOIS, J. and ADOLPHS, R. (2016). Building a science of individual differences from fMRI. *Trends in Cognitive Sciences* **20** 425–443.
- FAIR, D., DOSENBACH, N., CHURCH, J., COHEN, A., BRAHMBHATT, S., MIEZIN, F., BARCH, D., RAICHLE, M., PETERSEN, S. and SCHLAGGAR, B. (2007). Development of distinct control networks through segregation and integration. *Proceedings of the National Academy of Sciences* **104** 13507–13512.
- FAIR, D., COHEN, A., POWER, J., DOSENBACH, N., CHURCH, J., MIEZIN, F., SCHLAGGAR, B. and PETERSEN, S. (2009). Functional brain networks develop from a local to distributed organization. *PLoS Computational Biology* **5** e1000381.
- FALLANI, F., RICHIARDI, J., CHAVEZ, M. and ACHARD, S. (2014). Graph analysis of functional brain networks: Practical issues in translational neuroscience. *Philosophical Transactions of the Royal Society B: Biological Sciences* **369** 1–17.
- FOX, M. and GREICIUS, M. (2010). Clinical applications of resting state functional connectivity. *Frontiers in Systems Neuroscience* **4** 1–13.
- FRIEDMAN, J., HASTIE, T. and TIBSHIRANI, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9** 432–441.
- FRIEDMAN, J., HASTIE, T., HÖFLING, H. and TIBSHIRANI, R. (2007). Pathwise coordinate optimization. *The Annals of Applied Statistics* **1** 302–332.
- FRISTON, K. (2011). Functional and effective connectivity: A review. *Brain connectivity* **1** 13–36.
- GREICIUS, M., KRASNOW, B., REISS, A. and MENON, V. (2003). Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences* **100** 253–258.

- GUSNARD, D. and RAICHLER, M. (2001). Searching for a baseline: Functional imaging and the resting human brain. *Nature Reviews Neuroscience* **2** 685–694.
- KANAI, R. and REES, G. (2011). The structural basis of inter-individual differences in human behaviour and cognition. *Nature Reviews Neuroscience* **12** 231–242.
- KELLY, C., BISWAL, B., CRADDOCK, C., CASTELLANOS, X. and MILHAM, M. (2012). Characterizing variation in the functional connectome: Promise and pitfalls. *Trends in Cognitive Sciences* **16** 181–188.
- KRZANOWSKI, W. and HAND, D. (2009). *ROC curves for continuous data*. CRC Press.
- LEE, H., LEE, D., KANG, H., KIM, B. and CHUNG, M. (2011). Sparse brain network recovery under compressed sensing. *IEEE Transactions on Medical Imaging* **30** 1154–1165.
- LINDQUIST, M. (2008). The statistical analysis of fMRI data. *Statistical Science* **23** 439–464.
- MCLACHLAN, G. and KRISHNAN, T. (2007). *The EM algorithm and extensions* **382**. John Wiley & Sons.
- MEINSHAUSEN, N. and BÜHLMANN, P. (2006). High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics* 1436–1462.
- MENG, X. and VAN DYK, D. (1998). Fast EM-type implementations for mixed effects models. *Journal of the Royal Statistical Society: Series B* **60** 559–578.
- MUELLER, S., WANG, D., FOX, M., YEO, T., SEPULCRE, J., SABUNCU, M., SHAFEE, R., LU, J. and LIU, H. (2013). Individual variability in functional connectivity architecture of the human brain. *Neuron* **77** 586–595.
- NARAYAN, M., ALLEN, G. and TOMSON, S. (2015). Two sample inference for populations of graphical models with applications to functional connectivity. *arXiv preprint arXiv:1502.03853*.
- NIELSEN, J., ZIELINSKI, B., FLETCHER, T., ALEXANDER, A., LANGE, N., BIGLER, E., LAINHART, J. and ANDERSON, J. (2013). Multisite functional connectivity MRI classification of Autism: ABIDE results. *Frontiers in Human Neuroscience* **7** 72–83.
- PINHEIRO, J. and BATES, D. (2000). *Mixed-effects models in S and S-PLUS*. Springer Science & Business Media.
- POWER, J., BARNES, K., SNYDER, A., SCHLAGGAR, B. and PETERSEN, S. (2012). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *NeuroImage* **59** 2142–2154.
- RUBINOV, M. and SPORNS, O. (2010). Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage* **52** 1059–1069.
- SCHELLDORFER, J., BÜHLMANN, P. and VAN DE GEER, S. (2011). Estimation for High-Dimensional Linear Mixed-Effects Models Using l1-Penalization. *Scandinavian Journal of Statistics* **38** 197–214.
- SMITH, S. (2012). The future of fMRI connectivity. *NeuroImage* **62** 1257–1266.
- SMITH, S., MILLER, K., SALIMI-KHORSHIDI, G., WEBSTER, M., BECKMANN, C., NICHOLS, T., RAMSEY, J. and WOOLRICH, M. (2011). Network modelling methods for fMRI. *NeuroImage* **54** 875–891.
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)* 267–288.
- VAN DEN HEUVEL, M. and POL, H. (2010). Exploring the brain network: A review on resting-state fMRI functional connectivity. *European Neuropsychopharmacology* **20** 519–534.
- VAROQUAUX, G. and CRADDOCK, C. (2013). Learning and comparing functional connectomes across subjects. *NeuroImage* **80** 405–415.
- VAROQUAUX, G., GRAMFORT, A., POLINE, J. and THIRION, B. (2010). Brain covariance selection: Better individual functional connectivity models using population prior. In *Neural Information Processing Systems* 2334–2342.
- ZUO, X., DI MARTINO, A., KELLY, C., SHEHZAD, Z., GEE, D., KLEIN, D., CASTELLANOS, X., BISWAL, B. and MILHAM, M. (2010). The oscillating brain: Complex and reliable. *NeuroImage* **49** 1432–1445.

R. P. MONTI, C. ANAGNOSTOPOULOS
 DEPARTMENT OF MATHEMATICS,
 IMPERIAL COLLEGE LONDON
 SW7 2AZ, LONDON, UK
 E-MAIL: ricardo.monti08@imperial.ac.uk
canagnos@imperial.ac.uk

G. MONTANA
 DEPARTMENT OF BIOMEDICAL ENGINEERING,
 KING'S COLLEGE LONDON,
 ST THOMAS' HOSPITAL, LONDON SE1 7EH, UK
 E-MAIL: giovanni.montana@kcl.ac.uk