

A Novel Sparse Group Gaussian Graphical Model for Functional Connectivity Estimation

Bernard Ng^{1,2}, Gaël Varoquaux¹, Jean Baptiste Poline¹, and Bertrand Thirion¹

¹ Parietal Team, Neurospin, INRIA Saclay, France

² FIND Lab, Stanford University, United States

bernardying@gmail.com

Abstract. The estimation of intra-subject functional connectivity is greatly complicated by the small sample size and complex noise structure in functional magnetic resonance imaging (fMRI) data. Pooling samples across subjects improves the conditioning of the estimation, but loses subject-specific connectivity information. In this paper, we propose a new sparse group Gaussian graphical model (SGGGM) that facilitates joint estimation of intra-subject and group-level connectivity. This is achieved by casting functional connectivity estimation as a regularized consensus optimization problem, in which information across subjects is aggregated in learning group-level connectivity and group information is propagated back in estimating intra-subject connectivity. On synthetic data, we show that incorporating group information using SGGGM significantly enhances intra-subject connectivity estimation over existing techniques. More accurate group-level connectivity is also obtained. On real data from a cohort of 60 subjects, we show that integrating intra-subject connectivity estimated with SGGGM significantly improves brain activation detection over connectivity priors derived from other graphical modeling approaches.

Keywords: brain connectivity, fMRI, Gaussian graphical model, regularized consensus optimization, sparse inverse covariance estimation.

1 Introduction

Accumulating evidence suggests that a prominent effect of neurological diseases is abnormal alterations in functional connectivity [1, 2], which is fundamental to brain function. In recent years, functional magnetic resonance imaging (fMRI) has become the primary means for investigating this integrative property of the brain. Particularly relevant to clinical applications is the discovery of synchronized, ongoing brain activity even when a person is at rest [3]. The sets of brain areas that show such synchronized activity during resting state (RS) are broadly conceptualized as networks. These RS networks have been shown to display high resemblance to those evoked by task [4], thus demonstrating the presence of important structures in RS connectivity patterns that reflect brain organization [3]. This finding has significant implications for studying diseased populations, since patients often have trouble performing certain tasks, but have much less difficulties lying at rest in the scanner.

Inference of functional connectivity from RS-fMRI data is commonly performed by computing the Pearson's correlation between the observations of different brain areas. Due to the small sample size and the complex noise structure in RS-fMRI data, reliable estimation of intra-subject functional connectivity is extremely challenging [1, 5]. Also, Pearson's correlation cannot distinguish whether two brain areas are directly connected or indirectly connected through a third area [1]. This distinction is crucial if one is interested in the underlying connection structure of the brain [6]. To improve correlation estimation given limited noisy samples, a number of techniques based on l_2 and l_1 regularization has been proposed. l_2 regularization amounts to adding a scaled identity matrix to the sample covariance matrix for reducing estimation errors and improving the conditioning, which in turn enables stable matrix inversion [7]. The implication of being able to stably invert a covariance matrix is that elements of the inverse covariance matrix are proportional to partial correlations, which accounts for indirect influences [6]. The drawback of l_2 regularization is that it is limited to uniform shrinkage of the off diagonal elements [5]. This limitation can be mitigated using a l_1 regularization approach [5], commonly referred to as sparse Gaussian graphical model (SGGM), in which sparsity is imposed on the inverse covariance estimates to learn the partial correlation structure in a data-driven fashion. Imposing sparsity also helps reduce estimation errors. In the context of functional connectivity estimation, enforcing sparsity conforms to past findings that the connection structure of the brain is, in fact, sparse [1].

In settings where the number of parameters is greater than the number of samples, which is typical for single-subject RS-fMRI data, accurate correlation estimation is far from trivial even with l_2 and l_1 regularization. To this end, we have previously proposed extending SGGM by exploiting anatomical connectivity [8] as well as commonalities across subjects [5]. Alternatively, a more widely-used approach is to pool data across subjects [9] by either averaging their correlation matrices or concatenating their RS-fMRI time courses to increase the number of samples. Although this approach sacrifices subject-specific information, it generates group-level correlation estimates, which are useful for population comparisons.

The optimal way for deriving a representative group correlation matrix, while accounting for inter-subject variability, is still an open question. Approaches based on Bayesian networks that enable learning of conditional independence structure, i.e. brain connection structure, common across subjects have been proposed, but these approaches do not scale well with increasing number of brain areas [6]. Recently, a probabilistic model that enables integration of fMRI and diffusion MRI (dMRI) has been put forth [9], but this model assumes each brain connection is independent, which complicates the separation of direct connections from indirect connections [1].

In this paper, we propose a sparse group Gaussian graphical model (SGGGM) that permits joint estimation of intra-subject and group-level functional connectivity. We cast the estimation as a regularized consensus optimization problem [10], in which each subject's data is modeled as a GGM with intra-subject inverse covariance matrices tied across subjects by a sparse latent group inverse covariance matrix. Commonalities across subjects are thus exploited in handling noise and the problem of limited samples in each subject's data. The general idea of aggregating information across

subjects is akin to SGGM with group Lasso [5]. The difference is that SGGGM encourages intra-subject inverse covariance estimates to be *similar* to a sparse latent group inverse covariance estimate but does not impose all intra-subject inverse covariance estimates to have exactly the same sparsity pattern, i.e. the same connection structure. We show on simulated data with similar sample-to-parameter ratio as in real fMRI experiments that SGGGM significantly outperforms state-of-the-art covariance estimation techniques at both intra-subject and group level. Moreover, we illustrate on real data from a cohort of 60 subjects that incorporating intra-subject connectivity estimated using SGGGM significantly increases sensitivity in brain activation detection over connectivity computed with other widely-used methods.

2 Methods

The goal of this work is to address the challenge of reliable inverse covariance estimation under the setting where the number of samples is much less than the number of parameters and the samples are contaminated by strong noise. The state-of-the-art method for estimating a well-conditioned sparse inverse covariance matrix under this setting is SGGM, which we summarize in Section 2.1. We then describe our proposed model, SGGGM, for coalescing information across subjects in jointly estimating intra-subject and group-level inverse covariance matrices in Section 2.2. A quantitative scheme for validation is discussed in Section 2.3.

2.1 Sparse Gaussian Graphical Model

Given a $d \times d$ sample covariance matrix, $\mathbf{S}$, in which the samples are drawn from a multivariate Gaussian distribution, estimating a well-conditioned sparse invariance covariance matrix, $\hat{\mathbf{\Lambda}}$, can be formulated as the following optimization problem [11]:

$$\min_{\mathbf{\Lambda} > 0} tr(\mathbf{S}\mathbf{\Lambda}) - \log det(\mathbf{\Lambda}) + \lambda \|\mathbf{\Lambda}\|_1, \quad (1)$$

in which we search over the space of $d \times d$ positive definite matrices, $\mathbf{\Lambda} > 0$, to minimize the negative log-likelihood of a multivariate Gaussian distribution, $l(\mathbf{\Lambda}) = tr(\mathbf{S}\mathbf{\Lambda}) - \log det(\mathbf{\Lambda})$, while promoting a sparse estimate, $\hat{\mathbf{\Lambda}}$, by minimizing the l_1 -norm of the off diagonal elements, which we denote as $\|\mathbf{\Lambda}\|_1$. The level of sparsity is governed by λ , which can be selected in a data-driven manner through cross-validation (Section 2.2). In the context of functional connectivity estimation, $\mathbf{S}$ corresponds to the correlation matrix estimated from the RS-fMRI time courses of d brain areas of a given subject. (1) can be efficiently solved using algorithms, such as QUadratic Inverse Covariance (QUIC) [11] and Alternating Direction Method of Multipliers (ADMM) [10]. The latter algorithm is described in the next section.

2.2 Sparse Group Gaussian Graphical Model

Given $N \times d \times d$ sample covariance matrices, $\mathbf{S}^s$, where s is the subject index, we post the joint estimation of intra-subject inverse covariance, $\hat{\mathbf{\Lambda}}^s$, and group-level inverse covariance, $\hat{\mathbf{\Lambda}}^G$, as a regularized consensus optimization problem [10].

Regularized Consensus Optimization. For the case in which the solution is restricted to reside in the space of positive definite matrices, the regularized consensus optimization problem written in the unconstrained form is given by [10]:

$$\min_{\mathbf{\Lambda}^s > 0, \mathbf{\Lambda}^G > 0} J(\mathbf{\Lambda}^G) + \sum_{s=1}^N l_s(\mathbf{\Lambda}^s) + \frac{\rho}{2} \|\mathbf{\Lambda}^s - \mathbf{\Lambda}^G\|_F^2. \quad (2)$$

The intuition behind (2) is that each $\mathbf{\Lambda}^s$ should explain its own observations as encouraged through the individual loss, $l_s(\mathbf{\Lambda}^s)$, but to handle noise and the issue of limited samples in each subject's data, commonalities across subjects are exploited by penalizing deviations of $\mathbf{\Lambda}^s$ from $\mathbf{\Lambda}^G$, as imposed by minimizing $\|\mathbf{\Lambda}^s - \mathbf{\Lambda}^G\|_F^2$ where $\|\cdot\|_F^2$ denotes the Frobenius norm. The degree of this penalty is governed by ρ . For learning inverse covariance, we set $l_s(\mathbf{\Lambda}^s)$ to $\text{tr}(\mathbf{S}^s \mathbf{\Lambda}^s) - \log \det(\mathbf{\Lambda}^s)$ as in the SGGM formulation and $J(\mathbf{\Lambda}^G)$ to $\lambda \|\mathbf{\Lambda}^G\|_1$ for enforcing sparsity. We use ADMM to solve (2).

ADMM. Given initial $\mathbf{\Lambda}^s(0)$, $\mathbf{\Lambda}^G(0)$, and $\mathbf{U}^s(0)$, the overall idea of ADMM is to alternately update these matrices by minimizing the augmented Lagrangian [10]:

$$L_\rho(\mathbf{\Lambda}^s, \mathbf{\Lambda}^G, \mathbf{U}^s) = J(\mathbf{\Lambda}^G) + \sum_{s=1}^N \left(l_s(\mathbf{\Lambda}^s) + \mathbf{U}^s (\mathbf{\Lambda}^s - \mathbf{\Lambda}^G) + \frac{\rho}{2} \|\mathbf{\Lambda}^s - \mathbf{\Lambda}^G\|_F^2 \right), \quad (3)$$

where $\mathbf{U}^s$ is a $d \times d$ Lagrangian multiplier matrix. With $\mathbf{\Lambda}^G$ and $\mathbf{U}^s$ fixed, we find the update for $\mathbf{\Lambda}^s$ by minimizing $L_\rho(\mathbf{\Lambda}^s, \mathbf{\Lambda}^G(k), \mathbf{U}^s(k))$ over $\mathbf{\Lambda}^s$:

$$\mathbf{\Lambda}^s(k+1) = \arg \min_{\mathbf{\Lambda}^s > 0} l_s(\mathbf{\Lambda}^s) + \frac{\rho}{2} \|\mathbf{\Lambda}^s - \mathbf{\Lambda}^G(k) + \mathbf{U}^s(k)\|_F^2, \quad (4)$$

where $l_s(\mathbf{\Lambda}^s) = \text{tr}(\mathbf{S}^s \mathbf{\Lambda}^s) - \log \det(\mathbf{\Lambda}^s)$ for inverse covariance estimation and k denotes the iteration number. For this choice of $l_s(\mathbf{\Lambda}^s)$, there is an analytic solution for $\mathbf{\Lambda}^s(k+1)$ [10], computed by taking the derivative of (4) and setting that to zero, which gives:

$$\rho \mathbf{\Lambda}^s - (\mathbf{\Lambda}^s)^{-1} = \rho(\mathbf{\Lambda}^G(k) - \mathbf{U}^s(k)) - \mathbf{S}^s. \quad (5)$$

Since all matrices in (5) are symmetric, the right hand side can be decomposed into $\mathbf{Q}\mathbf{\Gamma}\mathbf{Q}^T$, where $\mathbf{Q}$ is the eigenmatrix of $\rho(\mathbf{\Lambda}^G(k) - \mathbf{U}^s(k)) - \mathbf{S}^s$, and $\mathbf{\Gamma}$ is a diagonal matrix containing the eigenvalues, γ_i . Multiplying both side by $\mathbf{Q}^T$ and $\mathbf{Q}$, results in:

$$\rho\tilde{\mathbf{\Lambda}}^s - (\tilde{\mathbf{\Lambda}}^s)^{-1} = \mathbf{\Gamma}, \quad (6)$$

where $\tilde{\mathbf{\Lambda}}^s = \mathbf{Q}^T \mathbf{\Lambda}^s \mathbf{Q}$. An elegant way for solving (6) is to assume $\tilde{\mathbf{\Lambda}}^s$ is a diagonal matrix, which amounts to finding $\tilde{\mathbf{\Lambda}}^s_{ii}$ such that $\rho \tilde{\mathbf{\Lambda}}^s_{ii} - 1/\tilde{\mathbf{\Lambda}}^s_{ii} = \gamma_i$, solution of which is given by the quadratic formula. $\mathbf{\Lambda}^s = \mathbf{Q} \tilde{\mathbf{\Lambda}}^s \mathbf{Q}^T$ is thus a solution of (4).

After computing $\mathbf{\Lambda}^s(k+1)$ for all subjects, we fix $\mathbf{\Lambda}^s$ and $\mathbf{U}^s$, and find the update for $\mathbf{\Lambda}^G$ by minimizing $L_\rho(\mathbf{\Lambda}^s(k+1), \mathbf{\Lambda}^G, \mathbf{U}^s(k))$ over $\mathbf{\Lambda}^G$:

$$\mathbf{\Lambda}^G(k+1) = \arg \min_{\mathbf{\Lambda}^G > 0} J(\mathbf{\Lambda}^G) + \sum_{i=1}^N \frac{\rho}{2} \|\mathbf{\Lambda}^s(k+1) - \mathbf{\Lambda}^G + \mathbf{U}^s(k)\|_F^2. \quad (7)$$

where $J(\mathbf{\Lambda}^G) = \lambda \|\mathbf{\Lambda}^G\|_1$ to impose sparsity. Since (7) is exactly the proximal operator of $J(\mathbf{\Lambda}^G)$, $\mathbf{\Lambda}^G(k+1)$ can be found by element-wise soft thresholding [10]:

$$\mathbf{\Lambda}^G_{ij}(k+1) = \left(\bar{\mathbf{\Lambda}}_{ij}(k+1) + \frac{\bar{\mathbf{U}}_{ij}(k)}{\rho} - \frac{\lambda}{N\rho} \right)_+ - \left(-\bar{\mathbf{\Lambda}}_{ij}(k+1) - \frac{\bar{\mathbf{U}}_{ij}(k)}{\rho} - \frac{\lambda}{N\rho} \right)_+, \quad (8)$$

where $(a)_+ = a$ if $a > 0$ and $(a)_+ = 0$ if $a \leq 0$. $\bar{\mathbf{\Lambda}}_{ij}$ and $\bar{\mathbf{U}}_{ij}$ denote the subject average of $\mathbf{\Lambda}^s_{ij}$ and $\mathbf{U}^s_{ij}$, respectively.

Finally, we fix $\mathbf{\Lambda}^s$ and $\mathbf{\Lambda}^G$, and minimize $L_\rho(\mathbf{\Lambda}^s(k+1), \mathbf{\Lambda}^G(k+1), \mathbf{U}^s)$ over $\mathbf{U}^s$ to find an update for $\mathbf{U}^s$. The solution is given by:

$$\mathbf{U}^s(k+1) = \mathbf{U}^s(k) + \rho(\mathbf{\Lambda}^s(k+1) - \mathbf{\Lambda}^G(k+1)). \quad (9)$$

(4), (7), and (9) are repeated until convergence. We highlight that (4) is separable across subjects. Thus, the estimation of $\mathbf{\Lambda}^s$ can be distributed over multiple processors, in which the computational cost of each $\mathbf{\Lambda}^s$ estimation is approximately that of an eigenvalue decomposition, i.e. $O(d^3)$. For initialization, we set $\mathbf{U}^s(0)$ to $\mathbf{0}_{d \times d}$ and $\mathbf{\Lambda}^G(0)$ to the subject average inverse covariance, with each subject's inverse covariance estimated using oracle approximating shrinkage (OAS), which is an l_2 regularization technique with a closed-form solution for the optimal shrinkage parameter [7]. We note that for $N = 1$, (1) and (2) are equivalent, thus the above procedures can be directly applied for pure single-subject sparse inverse covariance estimation.

Convergence Criteria. We adopt the convergence criteria in [10], which are based on the differences between the intra-subject inverse covariance estimates and their means, $\|\mathbf{R}(k)\|_F < \varepsilon_{\text{pri}}$, and the differences between the mean inverse covariance of two consecutive iterations, $\|\mathbf{Q}(k)\|_F < \varepsilon_{\text{dual}}$, as summarized below:

$$\|\mathbf{R}(k)\|_F^2 = \sum_{s=1}^N \|\mathbf{\Lambda}^s(k) - \bar{\mathbf{\Lambda}}(k)\|_F^2, \quad \|\mathbf{Q}(k)\|_F^2 = N\rho^2 \|\bar{\mathbf{\Lambda}}(k) - \bar{\mathbf{\Lambda}}(k-1)\|_F^2, \quad (10)$$

$$\varepsilon_{\text{pri}} = d\varepsilon_{\text{abs}} + \varepsilon_{\text{rel}} \max\{\|\bar{\mathbf{\Lambda}}(k)\|_F, \|\mathbf{\Lambda}^G(k)\|_F\}, \quad (11)$$

$$\varepsilon_{\text{dual}} = d\varepsilon_{\text{abs}} + \varepsilon_{\text{rel}} \|\rho\bar{\mathbf{U}}(k)\|_F. \quad (12)$$

Both ε_{abs} and ε_{rel} are set to 10^{-4} in this work. We note that one would normally prefer finding solutions with ε_{abs} and ε_{rel} being as small as computationally practical. However, in the present context of functional connectivity estimation, we anticipate inter-subject differences, thus setting ε_{abs} and ε_{rel} too low could falsely force all subjects to have overly-similar intra-subject connectivity estimates.

Parameter Selection. SGGGM requires setting two parameters: λ and ρ . Since the estimated connection structure highly depends on λ , we employ cross validation in combination with a refined grid search strategy similar to the approach we have taken in [8] to find the optimal λ . We proceed by first selecting a range of λ and temporally splitting the data into K folds. We then estimate $\mathbf{\Lambda}^G$ on $K-1$ training folds and compute the log data likelihood of the left-out test fold given $\mathbf{\Lambda}_{\text{train}}^G$ for each λ , i.e. $\Sigma_s(\log\det(\mathbf{\Lambda}_{\text{train}}^G) - \text{tr}(\mathbf{S}_{\text{test}}^s \mathbf{\Lambda}_{\text{train}}^G))$. We define the optimal λ as the one that gives the largest average log data likelihood across folds. We then search within a refined range around the optimal λ and repeat the procedure. The initial λ range is set as $[\lambda_{\text{max}}/100, \lambda_{\text{max}}]$, where $\lambda_{\text{max}} = \max_i |\bar{\mathbf{S}}_{ij}|$, $i \neq j$ is the value beyond which off-diagonal elements of $\mathbf{\Lambda}^G$ are shrunk to 0. $\bar{\mathbf{S}}_{ij}$ denotes the mean sample covariance across subjects. For each refinement level, we evenly distribute 5 grid points on a logarithm scale to cover the λ range. K is set to 3, and we refine the search 3 times within the interval defined by the two grid points adjacent to the optimal λ found at the previous refinement level. For ρ , we fix it to 1 to ensure convergence [10] as well as to reduce the computation load arising from nesting the λ search into an internal cross-validation loop. We defer investigation of efficient means for joint selection of λ and ρ to future work.

2.3 Validation

We base our validation on increased sensitivity in group activation detection by using our recently proposed model that permits incorporation of connectivity into the estimation of activation effects [12]. The rationale is that connectivity estimates that

better reflect the intrinsic wiring of the brain would presumably result in higher detection. By using this validation scheme, *quantitative* comparison of different connectivity estimation techniques is facilitated. We summarized below our model in [12]:

$$\mathbf{Y}^s \sim N(\mathbf{A}^s \mathbf{X}^s, \mathbf{V}_1^s) = \frac{1}{|2\pi\mathbf{V}_1^s|^{\frac{n}{2}}} \exp\left(\frac{-tr((\mathbf{Y}^s - \mathbf{A}^s \mathbf{X}^s)^T \mathbf{V}_1^{s^{-1}} (\mathbf{Y}^s - \mathbf{A}^s \mathbf{X}^s))}{2}\right) \quad (13)$$

$$\mathbf{A}^s \sim MN(0, \mathbf{V}_2^s, \alpha \mathbf{X}^s \mathbf{X}^{sT}) = \frac{|\alpha \mathbf{X}^s \mathbf{X}^{sT}|^{\frac{d}{2}}}{|2\pi\mathbf{V}_2^s|^{\frac{m}{2}}} \exp\left(-\frac{\alpha}{2} tr(\mathbf{X}^{sT} \mathbf{A}^{sT} \mathbf{V}_2^{s^{-1}} \mathbf{A}^s \mathbf{X}^s)\right), \quad (14)$$

where $\mathbf{Y}^s$ is a $d \times n$ matrix containing task fMRI time courses of d brain areas of subject s . $\mathbf{X}^s$ is a $m \times n$ matrix with m regressors along the rows for modeling the expected task-evoked responses [13]. $\mathbf{A}^s$ is a $d \times m$ activation effect matrix to be estimated. $\mathbf{V}_1^s$ and $\mathbf{V}_2^s$ are $d \times d$ covariance matrices of $\mathbf{Y}^s$ and $\mathbf{A}^s$, respectively. $\mathbf{X}^s \mathbf{X}^{sT}$ models the correlations between the m experimental conditions. $MN(0, \mathbf{V}_2^s, \alpha \mathbf{X}^s \mathbf{X}^{sT})$ denotes the matrix normal distribution, which serves as a conjugate prior of (13). The influence of this prior on $\mathbf{A}$ is governed by α , which can be optimized by maximizing model evidence [12]. We assume $\mathbf{V}_1^s = \mathbf{I}_{d \times d}$ as conventionally done, and $\mathbf{V}_2^s$ is where we inject our RS connectivity estimates. We note that all parameters in the model can be estimated in a data-driven manner without any manual interventions [12], thus providing an objective framework for evaluating different connectivity estimation methods.

3 Materials

3.1 Synthetic Data

For validation, we created 100 synthetic datasets with a sample-to-parameter ratio similar to that in a typical fMRI connectivity study. Specifically, each dataset comprised 10 subjects with 500 regional time courses of 186 time samples, and was generated as follows. First, we created a random 500×500 sparse positive definite matrix, Λ^G , with 20% of the elements being non-zero to approximately match the sparsity level observed when we applied SGGM to the real data. Λ^G corresponds to the group inverse covariance matrix, representative of the 10 subjects. We then randomly drew 10 positive definite matrices, Λ^s , $s = 1$ to 10, from $Wishart(\Lambda^G, \nu)$, where ν is the degrees of freedom set to ten times the number of parameters, i.e. $10 \times 500 \times 499/2$, to ensure that Λ^s is well-conditioned for matrix inversion. Λ^s corresponds to the inverse covariance matrix of each subject s , which was used to generate regional time courses by drawing samples from $N(0, \Sigma^s)$, where Σ^s is the matrix inverse of Λ^s .

3.2 Real Data

60 healthy subjects were recruited and scanned at multiple centers. Each subject was asked to perform 10 language, computation, and sensorimotor tasks similar to those in [14], as fMRI data were acquired over a duration of ~ 5 min. ~ 7 min of RS-fMRI data were also collected. Scanning was performed using 3T scanners from multiple manufacturers with TR = 2200 ms, TE = 30 ms, and flip angle = 75° . The task fMRI data were corrected for slice timing and head motions, temporally detrended, and spatially normalized using the SPM8 software. The RS-fMRI data were similarly preprocessed except we applied a band-pass filter with cutoff frequencies at 0.01 and 0.1 Hz. White matter and cerebrospinal fluid confounds were regressed out from the gray matter voxel time courses. To create a finer brain parcellation than provided by standard brain atlases (< 150 regions), we divided the brain into 500 parcels by concatenating RS-fMRI time courses across subjects and applying Ward clustering [15]. We then averaged the voxel time courses within each parcel to generate parcel time courses. These time courses were normalized by subtracting the mean and dividing by the standard deviation to account for scanner variability across imaging centers.

4 Results and Discussion

To investigate the gain of jointly estimating intra-subject and group-level connectivity under a unified model, we compared SGGGM against a number of state-of-the-art techniques on both synthetic and real data. Contrasted methods at the intra-subject level include SGGM [1] and OAS [7]. We excluded Pearson's correlation, since it does not provide an invertible correlation matrix when the number of parameters is greater than the number of samples, which is the case for our real data and quite generally in fMRI-based functional connectivity analysis. At the group level, we examined different ways of computing the mean from the SGGM and OAS intra-subject connectivity estimates as well as SGGM and OAS applied to concatenated observations across subjects. Specifically, we considered the conventionally-used Euclidean mean, $\Sigma_s \mathbf{\Lambda}^s / N$, and the Log-Euclidean mean [16], $\expm(\Sigma_s \logm(\mathbf{\Lambda}^s) / N)$, where $\expm(\cdot)$ and $\logm(\cdot)$ denote matrix exponential and matrix logarithm, respectively. The reason for using Log-Euclidean mean is that $\mathbf{\Lambda}^s$ lives on the space of positive definite matrices, which is not a vector space. Applying $\logm(\cdot)$ preserves the spectral characteristics of the input matrices, which avoids the undesirable "swelling effect" observed with Euclidean mean [16]. We note that when estimating group connectivity by concatenating observations across subjects, it is important to normalize the subjects' observations by subtracting the mean and dividing by the standard deviation to reduce inter-subject variability.

4.1 Synthetic Data Results

To assess the accuracy of the contrasted inverse covariance estimation methods, we computed the affine invariant distance [16] between the estimates and ground truth:

$$d(\hat{\Lambda}, \Lambda_{\text{gnd}}) = \| \log m(\Lambda_{\text{gnd}}^{-1/2} \hat{\Lambda} \Lambda_{\text{gnd}}^{-1/2}) \|_F^2, \quad (15)$$

where $\hat{\Lambda}$ denotes the estimated inverse covariance of a given method and Λ_{gnd} is the ground truth inverse covariance. The results are summarized in Fig. 1. At intra-subject level, SGGGM significantly outperformed SGGM and OAS, which demonstrates the benefits of incorporating group information into intra-subject inverse covariance estimation when the subjects belong to the same population. At the group level, using SGGGM also resulted in significantly more accurate inverse covariance estimates. Interestingly, using concatenated observations performed better than Euclidean and Log-Euclidean means. We speculate the reason is that the increased number of samples by concatenating observations improves the conditioning of the estimation, whereas using subject means amounts to averaging poorly estimated intra-subject inverse covariance matrices. The higher accuracy achieved with SGGM over OAS for the case where observations are concatenated is also likely due to a similar reason, in which the increased sample size enables SGGM to more accurately learn the support, i.e. the sparsity pattern, of the ground truth inverse covariance matrix.




Fig. 1. Synthetic data results. Subscripts “_E”, “_{LE}”, and “_C” denote estimation with Euclidean mean, Log-Euclidean mean, and concatenated observations. Using SGGGM resulted in significantly more accurate inverse covariance estimates than the contrasted methods.

4.2 Real Data Results

Validation on real data is greatly complicated by the lack of ground truth. To safely base our validation on increased sensitivity in brain activation detection, we employed the max- t permutation test [17] to enforce strict control on false positive rate. With detection sensitivity in group activation being the validation criterion, using a common group connectivity prior or even injecting group information into intra-subject connectivity estimation could bias the results. To remove this bias, we first randomly permuted the rows and columns of a given estimated connectivity matrix, $\hat{\Lambda}$, and computed the percentage of parcels found to be activated. This percentage indicates the amount of detections that could be obtained given the same set of estimated

connectivity values but with the connection structure randomized. This process was repeated 50 times with the maximum difference in the percentages of detected parcels between the permuted cases and that obtained with ordinary least square (OLS), i.e. standard univariate analysis without any prior [13], subtracted from the original percentage of detected parcels found with $\hat{\Lambda}$. This procedure was applied for all methods that used group information in the connectivity estimation. To test the generality of SGGGM, we examined the mean detection rate averaged over the 10 experimental conditions and 21 contrasts of interest between these conditions.

Quantitative results obtained by incorporating intra-subject connectivity priors into task activation detection are shown in Fig. 2(a). OLS is also plotted to serve as a baseline. Using SGGGM significantly increased detection sensitivity over SGGM and OAS for a typical p -value range of 0 to 0.05. Significance is declared based on a permutation described in [8]. Our results thus indicate that there are commonalities across subjects that can be exploited to improve intra-subject connectivity estimation. Denoting group-level connectivity estimation using SGGM and OAS with Euclidean mean, Log-Euclidean mean, and concatenated observations as SGGM_E , OAS_E , SGGM_{LE} , OAS_{LE} , SGGM_C , and OAS_C in Fig. 2(b), SGGGM was found to significantly outperform SGGM_E , OAS_E , and SGGM_{LE} . SGGGM also resulted in slightly more detections than OAS_{LE} and similar performance compared to SGGM_C and OAS_C . Overall, our results show that SGGGM provides relevant priors for activation detection at both intra-subject and group level.




Fig. 2. Real data results. Percentage of parcels with significant activation averaged across contrasts v.s. p -value thresholds displayed. Subscripts “ E ”, “ LE ”, and “ C ” denote group connectivity estimation with Euclidean mean, Log-Euclidean mean, and concatenated observations.

Qualitative results obtained by incorporating intra-subject connectivity priors are shown in Fig. 3. We did not include the group connectivity prior results due to space limitation. Areas in red are parcels detected with SGGGM only. The other coloured areas are parcels detected by SGGGM as well as one or more of the contrasted methods (OLS, OAS, and SGGM). For the task in which subjects performed calculations following auditory instructions, only SGGGM detected the right auditory thalamus (Fig. 3(a)), which is responsible for relaying information to the auditory cortex. For the case where the calculations to be performed was visually presented, SGGGM detected a much wider extent of the dorsal anterior cingulate cortex (Fig. 3(b)), which

is involved with executive processing and attention. Moreover, for the task in which subjects listened to sentences, only SGGGM detected the right Broca's area (Fig. 3(c)), which is associated with processing of subtle features in speech, e.g. prosody.




Fig. 3. Activation maps with intra-subject connectivity priors. Red = detected by SGGGM only. Other colours indicate areas detected by SGGGM and one or more of the contrasted methods.

5 Conclusions

We proposed a new sparse graphical model for joint estimation of intra-subject and group-level functional connectivity. We showed on synthetic data that SGGGM provides significantly more accurate inverse covariance estimates than state-of-the-art techniques at both the intra-subject and group level. We also demonstrated on real data that incorporating intra-subject connectivity priors learned from SGGGM results in higher sensitivity in activation detection compared to SGGM and OAS. Based on our results, a couple of insights can be drawn. First, sample size appears to be the limiting factor to functional connectivity estimation, as suggested by the improvements obtained by concatenating observations in contrast to using subject means. Thus, concatenating data across subjects, while appropriately accounting for inter-subject variability, could be a compelling strategy if only group connectivity is of interest. Moreover, since more sensitive activation detection can be achieved by incorporating a group connectivity prior, one can envision building a functional connectivity atlas by applying SGGGM to datasets from a large cohort of subjects and using the resulting connectivity estimate as a generic prior for future activation studies.

Acknowledgements. This work was supported by the Genim ANR-10-BLAN-0128 grant, NSERC, and the Berkeley-INRIA-Stanford grant. The data were acquired within the IMAGEN project. Jean Baptiste Poline was partly funded by the IMAGEN project, which receives funding from the E.U. Community's FP6, LSHM-CT-2007-037286. This manuscript reflects only the authors' views and the Community is not liable for any use that may be made of the information contained therein.

References

1. Huang, S., Li, J., Sun, L., Ye, J., Fleisher, A., Wu, T., Chen, K., Reiman, E.: Learning Brain Connectivity of Alzheimer's Disease by Sparse Inverse Covariance Estimation. *Neuroimage* 50, 935–949 (2010)

2. Delbeuck, X., Van der Linden, M., Collette, F.: Alzheimer's Disease as a Disconnection Syndrome? *Neuropsychol. Rev.* 13, 79–92 (2003)
3. Fox, M.D., Raichle, M.E.: Spontaneous Fluctuations in Brain Activity Observed with Functional Magnetic Resonance Imaging. *Nat. Rev. Neurosci.* 8, 700–711 (2007)
4. Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R., Beckmann, C.F.: Correspondence of the Brain's Functional Architecture During Activation and Rest. *Proc. Natl. Acad. Sci.* 106, 13040–13045 (2009)
5. Varoquaux, G., Gramfort, A., Poline, J.B., Thirion, B.: Brain Covariance Selection: Better Individual Functional Connectivity Models Using Population Prior. In: *Advances in Neural Information Processing Systems*, vol. 23, pp. 2334–2342 (2010)
6. Smith, S.: The Future of fMRI Connectivity. *NeuroImage* 62, 1257–1266 (2012)
7. Chen, Y., Wiesel, A., Eldar, Y.C., Hero, A.O.: Shrinkage Algorithms for MMSE Covariance Estimation. *IEEE Trans. Sig. Proc.* 58, 5016–5029 (2010)
8. Ng, B., Varoquaux, G., Poline, J.-B., Thirion, B.: A Novel Sparse Graphical Approach for Multimodal Brain Connectivity Inference. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) *MICCAI 2012, Part I. LNCS*, vol. 7510, pp. 707–714. Springer, Heidelberg (2012)
9. Venkataraman, A., Rathi, Y., Kubicki, M., Westin, C.F., Golland, P.: Joint Modeling of Anatomical and Functional Connectivity for Population Studies. *IEEE Trans. Med. Imaging* 31, 164–182 (2012)
10. Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J.: Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Found. Trend Mach. Learn.* 3, 1–122 (2010)
11. Hsieh, C.J., Sustik, M.A., Dhillon, I.S., Ravikumar, P.: Sparse Invers Covariance Matrix Estimation Using Quadratic Approximation. In: *Advances in Neural Information Processing Systems*, vol. 24, pp. 2330–2338 (2011)
12. Ng, B., Abugharbieh, R., Varoquaux, G., Poline, J.B., Thirion, B.: Connectivity-Informed fMRI Activation Detection. In: Fichtinger, G., Martel, A., Peters, T. (eds.) *MICCAI 2011, Part II. LNCS*, vol. 6892, pp. 285–292. Springer, Heidelberg (2011)
13. Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.B., Frith, C.D., Frackowiak, R.S.J.: Statistical Parametric Maps in Functional Imaging: A General Linear Approach. *Hum. Brain Mapp.* 2, 189–210 (1995)
14. Pinel, P., Thirion, B., Meriaux, S., Jobet, A., Serres, J., Le Bihan, D., Poline, J.B., Dehaene, S.: Fast Reproducible Identification and Large-scale Databasing of Individual Functional Cognitive Networks. *BioMed. Central Neurosci.* 8, 91 (2007)
15. Michel, V., Gramfort, A., Varoquaux, G., Eger, E., Keribin, C., Thirion, B.: A Supervised Clustering Approach for fMRI-based Inference of Brain States. *Patt. Recog.* 45, 2041–2049 (2012)
16. Arsigny, V., Fillard, P., Pennec, X., Ayache, N.: Fast and Simple Calculus on Tensors in the Log-Euclidean Framework. In: Duncan, J., Gerig, G. (eds.) *MICCAI 2005, Part I. LNCS*, vol. 3749, pp. 115–122. Springer, Heidelberg (2005)
17. Nichols, T., Hayasaka, S.: Controlling the Familywise Error Rate in Functional Neuroimaging: a Comparative Review. *Stat. Methods Med. Research* 12, 419–446 (2003)