

Deriving a Multi-subject Functional-Connectivity Atlas to Inform Connectome Estimation

Ronald Phlypo^{1,2}, Bertrand Thirion^{1,2}, and Gaël Varoquaux^{1,2}

¹ Parietal Team, Inria Saclay-Île-de-France, Saclay, France

{ronald.phlypo,gael.varoquaux,bertrand.thirion}@inria.fr

² CEA, DSV, I²BM, Neurospin bât 145, 91191 Gif-Sur-Yvette, France

Abstract. The estimation of functional connectivity structure from functional neuroimaging data is an important step toward understanding the mechanisms of various brain diseases and building relevant biomarkers. Yet, such inferences have to deal with the low signal-to-noise ratio and the paucity of the data. With at our disposal a steadily growing volume of publicly available neuroimaging data, it is however possible to improve the estimation procedures involved in connectome mapping. In this work, we propose a novel learning scheme for functional connectivity based on sparse Gaussian graphical models that aims at minimizing the bias induced by the regularization used in the estimation, by carefully separating the estimation of the model support from the coefficients. Moreover, our strategy makes it possible to include new data with a limited computational cost. We illustrate the physiological relevance of the learned prior, that can be identified as a functional connectivity atlas, based on an experiment on 46 subjects of the Human Connectome Dataset.

Keywords: functional connectivity, sparse Gaussian graphical models.

1 Functional Connectivity and Rest fMRI

Functional connectivity (FC) is a simple measure of the interactions between brain regions. First introduced for electro-physiological recordings, such as spike-train recordings or electro-encephalography, it is now commonly used in spatially-resolved neuro-imaging modalities despite their poor temporal resolution at the scale of a second or more. Reported first as an alternative tool for positron emission tomography analysis [1], it has become a prominent tool in functional magnetic resonance imaging (fMRI) analysis. Functional connectivity has initiated a paradigm shift in fMRI studies, since it makes possible to scan subjects without engaging them in a controlled task. This is especially important when dealing with non-cooperative subjects, such as children or patients with neuropathologic disorders [2, among others]. In addition, it creates an opportunity to study the brain as a whole, through its global interactions, rather than through local effects.

Although FC is an established framework in neuroimaging, there is still no consensus on how to measure it from the data [3]. The connectivity pattern can be seen as a graph with its nodes (vertices) and connections (edges). Nodes are associated with elementary brain regions, between which connection strengths are then estimated, and the graph structure as a whole is referred to as a functional connectome. Nodes may either be predefined regions of interest (ROIs) or voxels [4], or may consist of distributed structures estimated from the data [5,6]. Once the nodes have been defined, edge strengths are inferred either through Pearson’s correlation coefficients [4], partial correlation coefficients [7], or mutual information [8]. Characteristic network properties are then derived from these graph structures based on graph theoretic measures such as small-worldness, modularity, *etc* [9].

The central challenge in functional connectome estimation is that it is ill-posed, due to the small number of temporal samples in typical datasets. Regularization, *e.g.* with a sparsity assumption in the connections, is critical to recover connections from the data [3], yet at the expense of a bias on the estimated connectome. Ng *et al* [10] have shown that introducing a non-uniform connectivity prior derived from anatomical connectivity greatly improves estimation. However, the matching between functional and anatomical connectivity is often very imperfect [11]. For this reason, it is important to construct connectivity priors from functional connectivity information itself.

Here, we use a group of subjects to build such a prior in the form of an atlas that characterizes the stability of brain connections. Our estimation method is based on the precision matrix (inverse of correlation matrix, related to the partial correlation coefficients), since it has proven sensitive to correctly detect existing connections in a graph structure [3]. In contrast to existing stability studies, we do not focus on global properties of the graph, as in [9,12]. Rather, we aim at quantifying the probability of a connection being present between two pre-defined regions by means of a group study, and use this as a prior in the estimation of single subject FC. To alleviate the bias of sparsity-inducing penalization, we use a two-stage estimation procedure to obtain edge strengths.

Here, we focus on the qualitative properties of the model, which sketches a new generation of FC atlases, and leave quantitative validation for future work.

2 Data Used in the Study

In this contribution we use resting state data as provided by the minimal pre-processed data [13] of the human connectome project (HCP) [14] (release Q2). This comprises 46 subjects. All data were recorded in 3T MRI scanners with a repetition time of 720 ms. Data have been appropriately masked and motion corrected using the estimated motion parameters provided by HCP. In addition, we corrected for confounds using high variance confound regression [15] based on five principal components of 2% highest variance voxel time series of whole brain image, and using five principal components of ventricle and white matter voxels (using standard masks provided in FSL, the mask of the latter eroded by 1 voxel to avoid overlap with gray matter).

3 Learning the Functional Connectivity Model

3.1 Graph Representation of Functional Connectivity

Given a set of p brain regions, the goal is to estimate an undirected graphical model $\mathcal{G}(V, E, \Omega)$ as a connectome representation. The vertex set $V = \{v_i, i \in [[p]]\}$ represents weighted averages of time series with respect to each of the p given ROIs (here a smooth positive function over a bounded subset of \mathbb{R}^3). Weights reflect ‘probabilities’ of a voxel to belong to the given ROI. In this work, we have chosen the Harvard-Oxford lateralised probabilistic cortical atlas ($p = 96$) [16] (see, e.g., [12] for ROI abbreviations and acronyms used in this contribution).

The edge set $E = \{e_{ij}, (i, j) \in [[p]]^2\}$ contains e_{ij} if and only if the ROIs represented by v_i and v_j are functionally connected, given all other $v_k, k \notin \{i, j\}$. Edge weights $\Omega = \{\omega(e_{ij}), e_{ij} \in E\}$ represent signed connectivity strengths. Negative strengths correspond to a phase inversion between the time series of interest. When modeling FC, it is often assumed that the graph is sparse, which corresponds to a cardinality of E being (largely) inferior to $p(p - 1)/2$. This hypothesis is justified based on evidence provided by structural connectivity measurements [17].

3.2 Estimation of the Graph Model Parameters

Estimating a sparse graphical-model is an NP-hard problem, and we choose graphical lasso [18,19] as an appropriate convex relaxation of our problem. We represent the graphical model through its adjacency matrix $\Theta \stackrel{\text{def.}}{=} (\theta_{ij})_{i,j}$, which under a Gaussian hypothesis corresponds to the rescaled precision matrix (inverse of the covariance matrix). Hence, $\theta_{ij} = \theta_{ji} = 0$ if and only if the two regions represented by v_i and v_j are independent, conditional on all other regions [20]. Graphical lasso combines the maximum likelihood estimate with a convex sparsity inducing penalty with an objective function given by¹

$$\hat{\Theta}_E = \arg \min_{\Theta \succ \mathbf{0}} \varphi_E(\Theta) \stackrel{\text{def.}}{=} -\log \det(\Theta) + \text{trace}(\Theta \mathbf{S}) + \|\Lambda \circ \Theta\|_1 \quad (1)$$

where \mathbf{S} is the sample correlation matrix of the time-series represented by the nodes in V . We minimize $\varphi_E(\Theta)$ of Eq. (1) over the space of symmetric positive definite matrices using the algorithm of [21], which guarantees $\hat{\Theta}_E \succ \mathbf{0}$ at any step of the algorithm, as such providing a procedure more robust against ill-conditioning of \mathbf{S} than those originally proposed in [18,19].

The estimate $\hat{\Theta}_E$ is a biased estimate of Θ , and we only use it to infer the edge set E . We thus break our estimation scheme down into three steps: (i) estimation

¹ The original implementation of the penalization in graph lasso reads $\lambda \|\Theta\|_1$, but since no penalization of the diagonal terms is required, we may write $\lambda \left\| (\mathbf{1}\mathbf{1}^\top - \text{Id}_p) \circ \Theta \right\|_1$, where $\mathbf{1}$ is a vector of ones in \mathbb{R}^p and $\cdot \circ \cdot$ denotes the element-wise or Hadamard product. We use here an extension of this model, replacing $\lambda (\mathbf{1}\mathbf{1}^\top - \text{Id}_p)$ by a symmetric matrix $\Lambda = (\lambda_{ij})_{i,j}$, with $0 \leq \lambda_{ij} \leq 1$.

of the hyperparameter Λ , (ii) estimation of the edge set E , and (iii) estimation of the edge weights or signed connectivity strengths Ω .

Step (i) will be discussed in the next section. Assuming for the moment that Λ is known, the last two steps are simply:

(ii) The edges e_{ij} in \hat{E} correspond to those entries of $\hat{\Theta}_E$ for which $|\hat{\theta}_{ij}| > 0$.

(iii) Given the estimated edge set, we estimate edge strengths $\Omega = (\omega_{ij})_{i,j}$ as the maximum likelihood estimates restricted to the edge set \hat{E} of the graph $\hat{\mathcal{G}}$, based on the sample correlation matrix \mathbf{S} . Our estimator reads

$$\hat{\Theta}_\Omega = \arg \min_{\Theta_{|\hat{\mathcal{G}}|} > 0} \varphi_\Omega(\Theta) \stackrel{\text{def.}}{=} -\log \det(\Theta) + \text{trace}(\Theta \mathbf{S}) \quad (2)$$

which is over the space of positive-definite, symmetric matrices with a given support set \hat{E} . The restriction to the graph prohibits the use of straightforward matrix inversion, and instead we use iterative proportional scaling (IPS), iteratively solving plain maximum likelihood models alternated over maximal cliques of the graph $\hat{\mathcal{G}}$ [20]. The obtained solution is an unbiased maximum likelihood estimator conditional on the support set of the graph, hence the only bias in our final model is the pruning of the edges. Given new data (a subject, session, or condition), we can infer the corresponding functional connectome.

3.3 Choosing the Prior Atlas

The choice of the regularizer Λ is essential to optimize the bias/variance compromise. A first approach consists in choosing $\Lambda = \lambda(\mathbf{1}\mathbf{1}^\top - \text{Id}_p)$. As in [21], this makes the problem low-dimensional. Hence, we use a repeated random jackknifing cross-validation scheme. As a score function, we opt for the negative log-likelihood, which is our main objective function. However, the extrapolation of that single parameter λ is an issue, since a given λ value corresponds to different cardinality of \hat{E} for different datasets (even if ratios such as λ/λ_{\max} are considered, where λ_{\max} is the critical regularizer that yields a trivial solution). Instead we propose to accumulate the support estimates across subjects and sessions. The intuition behind this approach is that noise in the data can cause instabilities in the graph lasso estimate. However, these instabilities are not reproducible across subjects, and the population distribution of the presence of edges is reliable. This approach is related to the “stability selection” strategy, also developed for sparse Gaussian graphical models [22]. Consider q subjects, each subject having 2 separate scan sessions ².

Estimation of $\lambda^{(i)}$ s. For each subject i , we use a grid search to estimate $\lambda^{(i)}$ as that λ in $\Lambda = \lambda(\mathbf{1}\mathbf{1}^\top - \text{Id}_p)$ for which $\hat{\Theta}_\Omega^{(i)}$ maximizes the likelihood generalization. The generalization capability of the likelihood can be approximated

² We consider the back-to-back recorded left-right and right-left phase encoding as a single session. To marginalize the effect of the phase encoding, we always sample both encodings evenly. Thus, if we state that n samples are jackknifed from a session, we actually mean that $n/2$ samples are jackknifed from each phase encoding associated with that session.

by using n samples from one session as a train session to estimate $\Theta_{E^{(i)}}^{(i)}$; $E^{(i)}$ is then taken as the support in Eq. (2) to estimate $\Theta_{\Omega^{(i)}}^{(i)}$ from n samples. We then use the n' complementary samples in the latter session to compute the likelihood of $\widehat{\Theta}_{\Omega^{(i)}}^{(i)}$. This is repeated k times, and we report $\lambda^{(i)}$, as that λ that maximizes the average likelihood. In this study, we have used $n = 200$, $n' = 600$, $k = 5$, and for each repetition the session used for training has been chosen randomly.

Estimation of Λ . Using the optimal $\lambda^{(i)}$ for subject i , we re-estimate $E^{(i)}$ from $n'' = 400$ samples, using Eq. (1) for each session, yielding $\widehat{E}^{(i,1)}$ and $\widehat{E}^{(i,2)}$. Suppose that all subjects (and sessions) are drawn from a single population with P_{ij} the marginal probability of observing edge e_{ij} . We simply estimate these probabilities by using maximum likelihood under marginal Bernoulli distributions on the edges, based on the $\{E^{(i,1)}, E^{(i,2)}\}$ estimates. We then define the penalization parameter Λ as $\mathbb{1}\mathbb{1}^T - \mathbf{P}$, where $\mathbf{P} \stackrel{\text{def.}}{=} (P_{ij})_{i,j}$. Note that, by construction, \mathbf{P} is a symmetric matrix with ones on its diagonal.

Λ can then be used on a new subject as a penalty in the penalized maximum likelihood estimation Eq. (1). Importantly, it appears as a penalty, not a constraint, and it is thus a “soft” prior. In other words, it does not force the existence or absence of a connections; it simply promotes these choices.

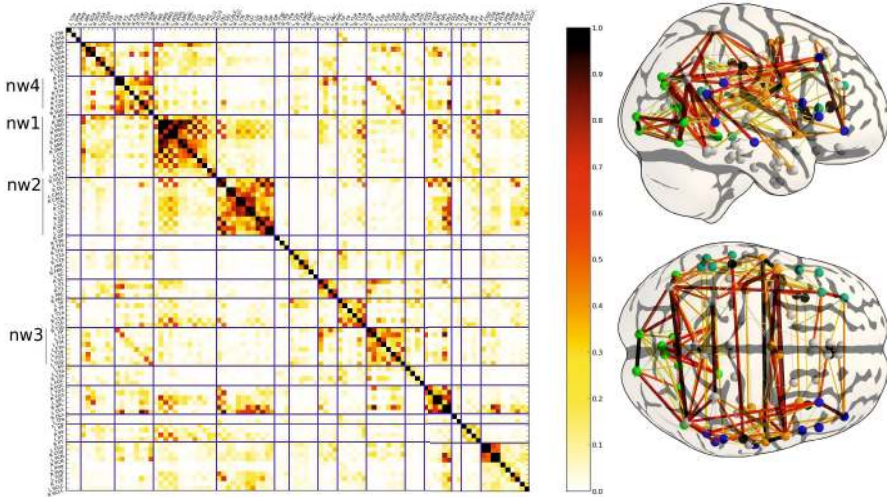


Fig. 1. Prior on the functional connectivity, noted \mathbf{P} in the text: the coefficients in the matrix represent the frequencies of edge detections. ROI labels and their abbreviations, as well as a visualization of all networks can be found in an addendum to a preprint of this manuscript at <http://hal.inria.fr/hal-00991124>.

4 Resulting Atlas and Discussion

The probability estimates in \mathbf{P} give a global view of the connectome in our training set of 46 subjects. It is important to note that the structure of \mathbf{P} reflects the sparsity assumptions, which means that some connections have indeed never been ‘observed’ over the entire set of subjects, see Fig. 1. This connectivity prior matrix is very structured; it displays a checkerboard pattern: regions are most likely to be connected to a region in the same hemisphere than the corresponding one in the opposite hemisphere. In addition, visualized as a brain graph, it features well-known characteristics of brain connectivity, such as strong homologous inter-hemispheric or fronto-parietal connections.

To visualize the latent network structure, we re-organized the matrix \mathbf{P} using hierarchical clustering with complete linkage, choosing the number of clusters that minimizes the silhouette criterion [23]. Clustering on probabilities of edges in the support of a graph gives FC networks (see Fig. 2) similar to those found in group analyses based on independent component maps, see, e.g., [5]. This suggests that the FC atlas has indeed captured valuable information from our reference population.

Other than having an interpretable graph representation, it should be emphasized that our approach provides a readily-usable subject-specific estimator, alleviating the need of a search over the parameter space (λ) for a new subject. As any medical-imaging atlases, our model incorporates accumulated knowledge. In contrast to group analyses that need a complete re-computation of the graph or the components upon incorporation of a new subject [6,24], we only need to update a single matrix, namely the estimator \mathbf{P} . Moreover, based on the observations of [10], we conjecture that the incorporation of the structured atlas-based prior will yield more accurate estimation of functional connectivity than traditional graph lasso estimators that enforce a uniform ℓ_1 shrinkage on all connections.




Fig. 2. Lateral and top views of the four networks containing the most nodes. From left to right we distinguish a motor network, a visual network, as well as two lateralised TPJ-to-IFG networks (TPJ: temporo-parietal junction – IFG: Inferior Frontal Gyrus).

5 Conclusions

Sparsifying penalties are the most promising tools to learning functional connectomes from noisy and scarce fMRI data. Injecting structure in this prior, favoring connections that are known to be likely, can greatly improve the ability of the estimation procedures to recover real connections. We have introduced a procedure to learn an atlas of brain functional connections by accumulating knowledge across subjects, associated with a subject-specific estimator of a sparse Gaussian graphical model that quantifies functional connectivity. Future work includes the validation of the atlas-based strategy for individual connectivity estimation and subjects connectivity comparison.

Acknowledgements. Data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. We acknowledge funding from the NiConnect project grant.

References

1. Friston, K.J., Frith, C.D., Liddle, P.F., Frackowiak, R.S.J.: Functional connectivity: the principal-component analysis of large (PET) data sets. *Journal of Cerebral Blood Flow and Metabolism* 13, 5–14 (1993)
2. Menon, V.: Developmental pathways to functional brain networks: emerging principles. *Trends in Cognitive Sciences* (2013)
3. Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., Beckmann, C.F., Nichols, T.E., Ramsey, J.D., Woolrich, M.W.: Network modelling methods for fMRI. *NeuroImage* 54(2), 875–891 (2011)
4. Hayasaka, S., Laurienti, P.J.: Comparison of characteristics between region-and voxel-based network analyses in resting-state fMRI data. *NeuroImage* 50(2), 499–508 (2010)
5. Damoiseaux, J.S., Rombouts, S.A.R.B., Barkhof, F., Scheltens, P., Stam, C.J., Smith, S.M., Beckmann, C.F.: Consistent resting-state networks across healthy subjects. *Proceedings of the National Academy of Sciences* 103(37), 13848–13853 (2006)
6. Smith, S.M., Miller, K.L., Moeller, S., Xu, J., Auerbach, E.J., Woolrich, M.W., Beckmann, C.F., Jenkinson, M., Andersson, J., Glasser, M.F., Van Essen, D.C., Feinberg, D.A., Yacoub, E.S., Ugurbil, K.: Temporally-independent functional modes of spontaneous brain activity. *Proceedings of the National Academy of Sciences* 109(8), 3131–3136 (2012)
7. Marrelec, G., Krainik, A., Duffau, H., Pélégriani-Issac, M., Lehericy, S., Doyon, J., Benali, H.: Partial correlation for functional brain interactivity investigation in functional MRI. *NeuroImage* 32(1), 228–237 (2006)
8. Ma, S., Calhoun, V.D., Phlypo, R., Adali, T.: Dynamic changes of spatial functional network connectivity in healthy individuals and schizophrenia patients using independent vector analysis. *NeuroImage* 90, 196–206 (2014)

9. Rubinov, M., Sporns, O.: Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage* 52(3), 1059–1069 (2010); *Computational Models of the Brain*
10. 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)
11. Ng, B., Varoquaux, G., Poline, J.B., Thirion, B.: Implications of inconsistencies between fMRI and dMRI on multimodal connectivity estimation. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) *MICCAI 2013, Part III. LNCS*, vol. 8151, pp. 652–659. Springer, Heidelberg (2013)
12. Wang, J.H., Zuo, X.N., Gohel, S., Milham, M.P., Biswal, B.B., He, Y.: Graph theoretical analysis of functional brain networks: Test-retest evaluation on short- and long-term resting-state functional MRI data. *PLoS One* 6(7), e21976 (2011)
13. Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M.: The minimal preprocessing pipelines for the human connectome project. *NeuroImage* 80, 105–124 (2013); *Mapping the Connectome*
14. Smith, S.M., Beckmann, C.F., Andersson, J., Auerbach, E.J., Bijsterbosch, J., Douaud, G., Duff, E., Feinberg, D.A., Griffanti, L., Harms, M.P., Kelly, M., Laumann, T., Miller, K.L., Moeller, S., Petersen, S., Power, J., Salimi-Khorshidi, G., Snyder, A.Z., Vu, A.T., Woolrich, M.W., Xu, J., Yacoub, E., Ugurbil, K., Van Essen, D.C., Glasser, M.F.: Resting-state fMRI in the human connectome project. *NeuroImage* 80, 144–168 (2013); *Mapping the Connectome*
15. Behzadi, Y., Restom, K., Liau, J., Liu, T.T.: A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *NeuroImage* 37(1), 90–101 (2007)
16. Desikan, R.S., Ségonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, B.T., Albert, M.S., Killiany, R.J.: An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage* 31(3), 968–980 (2006)
17. Deligianni, F., Varoquaux, G., Thirion, B., Sharp, D.J., Ledig, C., Leech, R., Rueckert, D.: A framework for inter-subject prediction of functional connectivity from structural networks. *IEEE Transactions on Medical Imaging* (August 2013)
18. Friedman, J., Hastie, T., Tibshirani, R.: Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* 3, 432–441 (2008)
19. Banerjee, O., El Ghaoui, L., d’Aspremont, A.: Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *Journal of Machine Learning Research* 9, 485–516 (2008)
20. Lauritzen, S.L.: *Graphical Models*. Oxford Statistical Science, vol. 17. Clarendon Press (1996)
21. Mazumder, R., Hastie, T.: The graphical lasso: new insights and alternatives. *Electronic Journal of Statistics* 6, 2125–2149 (2012)
22. Meinshausen, N., Bühlmann, P.: Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72(4), 417–473 (2010)
23. Rousseeuw, P.J.: Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics* 20, 53–65 (1987)
24. Ng, B., Varoquaux, G., Poline, J.B., Thirion, B.: A novel sparse group Gaussian graphical model for functional connectivity estimation. In: *Information Processing in Medical Imaging, Asilomar, États-Unis* (June 2013)