# A framework for inter-subject prediction of functional connectivity from structural networks

Fani Deligianni, Gaël Varoquaux, Bertrand Thirion, David J. Sharp, Christian Ledig, Robert Leech, Daniel Rueckert

# A framework for inter-subject prediction of functional connectivity from structural networks

Fani Deligianni*† , Gaël Varoquaux‡ , Bertrand Thirion‡, David J. Sharp§, Christian Ledig†, Robert Leech§, and Daniel Rueckert†

* Imaging and Biophysics Unit, Institute of Child Health, University College London, UK

‡INRIA, Saclay-Île-de-France, Parietal team, France - CEA/ DSV/ I2BM/ Neurospin/ LNAO

§ C3NL, The Centre for Neuroscience, The Division of Experimental Medicine, Imperial College London, UK

† BioMedia Group, Department of Computing, Imperial College London, UK

*Corresponding author - f.deligianni@ucl.ac.uk

*Abstract*—**Functional connections between brain regions are supported by structural connectivity. Both functional and structural connectivity are estimated from in-vivo MRI and offer complementary information on brain organisation and function. However, imaging only provides noisy measures, and we lack a good neuroscientific understanding of the links between structure and function. Therefore, inter-subject joint modeling of structural and functional connectivity, the key to multimodal biomarkers, is an open challenge. We present a probabilistic framework to learn across subjects a mapping from structural to functional brain connectivity. Expanding on our previous work [1], our approach is based on a predictive framework with multiple sparse linear regression. We rely on the randomized LASSO to identify relevant anatomo-functional links with some confidence interval. In addition, we describe resting-state (rs)-fMRI in the setting of Gaussian graphical models, on the one hand imposing conditional independences from structural connectivity and on the other hand parameterizing the problem in terms of multivariate autoregressive models. We introduce an intrinsic measure of prediction error for functional connectivity that is independent of the parameterization chosen and provides the means for robust model selection. We demonstrate our methodology with regions within the default mode and the salience network as well as, atlas-based cortical parcellation.**

*Index Terms*—**structural brain connectivity, functional brain connectivity, statistical associations, predictive modeling**

## I. Introduction

Descriptions of brain function rely on complementary principles of functional specialization and functional integration, also referred as connectivity. Functional specialization refers to the localization of particular aspects of cognitive processes to brain regions and is studied via the well-established framework of brain mapping. However, it must be complemented by the study of interactions between brain regions [2]. These interactions are mediated by physical connections –brain fibers– and appear in the functional signal via so-called *functional connectivity*. The large-scale organizations reflected by both types of connectivity are the focus of intense research [3], [4], [5], [6]. Indeed, connectivity and network integration appear as markers of brain states such as consciousness [7], [8] or high-level cognition [9], as well as brain pathologies, including schizophrenia [10], [11], attention deficit hyperactivity disorder (ADHD) [12], autism [13], Alzheimer disease (AD) [14] and traumatic brain injury (TBI) [15], [16].

Magnetic Resonance Imaging (MRI) is central to the study of brain networks because it allows non-invasive, in-vivo measures of both functional and structural connections. Diffusion Weighted MRI (DWMRI) captures local tissue properties that can be used to estimate neuronal fibers via tractography techniques that output structural networks [17], [18], [19]. On the other hand, functional MRI (fMRI) measures hemodynamic, metabolic activity. Beyond its typical use for mapping task-related regions, it reveals intrinsic functional connectivity via spontaneous fluctuations of brain activity [20], in *resting-state* fMRI (rs-fMRI) experiments. This observed functional connectivity stems from the correlated patterns of activity between interacting regions [2], [21]. Combining DWMRI and rs-fMRI can provide a whole-brain connectivity description that reflects structure and function [22], [23], [24], [25].

Integration of structural and functional brain connectivity from multimodal imaging holds the promise of improved markers of brain function and malfunction [26]. The first step towards this objective is to build reliable anatomo-functional models that extract meaningful relationships in healthy subjects. This would allow the development of a sound methodology to draw conclusions on the statistical significance of the results and the corresponding confidence intervals. The key challenges to face are that: *i)* both rs-fMRI and DWMRI are noisy, indirect measures of function and structure, respectively. Their relationship is underpinned by the fact that structural connectivity is the carrier of the signal which is transmitted from one region to another, *ii)* linking functional connectivity to anatomical connectivity is a highly multivariate statistical problem due to the number of possible neuronal pathways. Most previous approaches have circumvented this challenge by neglecting indirect effects, based on the simplifying assumptions that an observed functional connection is underpinned solely by the structural connection between the two corresponding brain regions [27], [28]. On the contrary, empirical studies [29], [27], [30] have reported that there is significant functional connectivity mediated from one region to another via indirect pathways.

The coupling between function and structure has been also studied with neurocomputational models [3]. These build from the hypothesis that dynamical resting states emerge from the interplay between large-scale brain structure and oscillatory neural dynamics at the local level. Connectivity in these networks is enforced based on long-range structural connections and conduction delays proportional to their length. This simulation work sheds light on models of intrinsic functional activity directly linked to the underlying neurobiology. However, this approach is not well suited to estimating subject specificities from imaging data, and the authors do not quantify precisely the agreement between empirical and simulated functional connectivity.

To our knowledge the first probabilistic framework that jointly models functional and structural connectivity interactions has been introduced recently in Venkataraman et al. [31]. They model the observed correlation between two areas as a random variable with parameters dependent on both the latent functional and structural connectivity. The limits of this model are that it treats each connection independently and it ignores the variance of connectivity strength across connections.

Functional connectivity, inter-subject studies are usually based on coefficient-level comparison of correlation matrices. The underlying variability model is often expressed as a univariate additive linear model on the covariance matrix, which encodes both the group variability as well as the subject specific contributions. However, a major limitation of this approach is that it models separately the variation of the different matrix coefficients. Correlation matrices have a certain structure: they are by construction positive definite. Constructing a matrix with coefficients chosen independently leads, with high probability, to a non-positive definite matrix, hence, it does not yield any consistent signal model. Therefore, it is not possible to generalize the learned model to new subjects [32].

In this paper, we introduce a joint predictive model of brain structural and functional connectivity. In the absence of any specific knowledge, we represent the link between structure and function with multiple sparse linear regressions. An important specificity of our approach is to account for indirect connections, both by predicting partial correlations rather than marginal correlations in the functional signal, and by learning a many-to-one mapping from structural to functional connections. Unlike previous studies [31], we use prediction error, *i.e.* out-of-sample modeling error, as an Occam razor to select the best performing model while limiting model complexity and corresponding overfit. We quantify modeling error independently of the choice of parameterization of functional connectivity, *i.e.* correlation or partial correlation. Functional connectivity is considered as a multivariate statistic, rather than a collection of independent correlations and hence the predicted functional-connectivity matrices are constrained to the set of observable correlation matrices, *i.e.* symmetric positive definite (SPD) matrices.

This paper extends our previous work [1] in several ways. We introduce a model selection framework based on nested cross-validation, which relates prediction error and the sparsity of the Gaussian graphical model. Furthermore, we discuss

model identification, that is how to estimate with some confidence the relevant anatomo-functional links. Finally, we present a detailed evaluation of our model with both atlas-based ROIs and regions derived based on known functional networks such as the default mode network (DMN) and the salience network (SN). These parcellations provide intuitive brain network descriptions and allow to relate our results with existing literature [16].

The paper is organized as follows. In section II, we introduce our predictive modeling framework. In section III, we present a joint model linking the two modalities and we discuss model selection and model identification. In section IV, we discuss the overall validation design that includes both the validation of the model and the qualitative representation of anatomo-functional connections. In section V we give empirical results that show how our model performs in terms of predictive accuracy and we relate qualitatively our identification results with current literature. Finally, in sections VI and VII we discuss the strengths and weaknesses of our approach and the evidence supporting it.

*Notations:* We write vectors with bold letters, $\boldsymbol{a} \in \mathbb{R}^n$, matrices with capital bold letters, $\boldsymbol{A} \in \mathbb{R}^{n \times n}$, and we denote $\|\boldsymbol{A}\|$ the operator norm, $\|\boldsymbol{A}\|^2 = \sum_{i,j=1}^{n} \boldsymbol{A}_{ij}^2$. $\boldsymbol{I}_n$ is the identity matrix of $\mathbb{R}^{n \times n}$ and $\boldsymbol{O}_n$ the zero matrix. When considering a variable with values across a group of subjects, we denote the corresponding subject-specific variable with an exponent: $\boldsymbol{A}^s$. Quantities estimated from the data at hand are written $\widehat{\boldsymbol{A}}$. $\boldsymbol{A}^{-1}$ is the matrix inverse of $\boldsymbol{A}$, $\boldsymbol{A}^\mathsf{T}$ is the transposed matrix, and $\boldsymbol{A}^{-\mathsf{T}}$ is the inverse transposed. Finally, we denote $\mathcal{Sym}_n$ the set of $n \times n$ symmetric matrices and $\mathcal{Sym}_n^+$ the set of symmetric definite positive (SPD) matrices.

## II. PREDICTIVE MODELING OF FUNCTIONAL CONNECTIVITY

In this section we introduce our predictive modeling settings: using statistical learning to infer a link between anatomical and functional connectivity. We start by reviewing the statistical learning framework we rely on and the corresponding modeling choices it entails. Then we expose the specificities of our learning problem. Briefly, modeling challenges stem from *i)* the fact that both connectivity modalities need to be described as multivariate objects, which gives rise to structured output prediction problems *ii)* the small sample size compared to the complexity of the objects, which leads to seeking simpler models via a sparsity description of the problem and learning procedures that can include penalization.

### A. Casting the problem as a statistical learning task

*Problem setting:* In a population of $S$ subjects, we aim to infer the link between the connections between brain regions estimated by tractography of DWMRI and the observed synchronization in the brain activity observed via rs-fMRI. Given a set of $n$ ROIs, we represent the set of anatomical connections for a subject $s$ as a connectivity matrix $\boldsymbol{A}^s \in \mathcal{Sym}_n$. The corresponding brain activity in the ROIs is summarized by $n$ time series of length $t$, $\boldsymbol{F}^s \in \mathbb{R}^{n \times t}$. We are interested in explaining the correlation structure of the observed rs-fMRI

time-series $\boldsymbol{F}^s$ from the subject's anatomical connectivity matrix $\boldsymbol{A}^s$. For this purpose, we propose to learn the link between $\boldsymbol{A}^s$ and the covariance matrix of the fMRI data $\boldsymbol{\Sigma}^s \in \mathbb{R}^{n \times n}$.

The typical population size used in clinical studies is small compared to the number of possible connections that grows as $\frac{1}{2}n^2$. As a result, the statistical estimation of a link across subjects between anatomical and functional connection is ill-posed, and it will need to resort to regularization, *i.e.* injecting sparsity constrains. We adapt the empirical risk minimization framework [33], [34], in which sparsity, or other regularization strategies, can easily be added via penalization terms. This framework gives theoretical bounds on the performance of learning algorithms and it can be used to compare different models as well as to select the best model parameters.

*The statistical learning framework:* Briefly, the goal is to learn a mapping $f$ from a space $\mathcal{A}$ describing anatomical connectivity matrices to a second space $\mathcal{F}$ describing functional connectivity, such that for a new subject, given the anatomical connectivity $\boldsymbol{A}^{\text{test}}$, the function $f$ yields a prediction for the functional connectivity $\boldsymbol{\Sigma}^{\text{pred}} = f(\boldsymbol{A}^{\text{test}})$ close to the subject's observed functional connectivity $\boldsymbol{\Sigma}^{\text{test}}$. In our case, the prediction function $f$ expresses a multi-modal model of brain connectivity, parameterized by a vector of model parameters $\boldsymbol{\theta}$. Learning the model from the data is performed by adjusting these model parameters to minimize the sum of a *loss*, *i.e.* a term quantifying the goodness of fit of the model to the observed data, and a regularization term that controls for overfitting by penalizing more complex models and can account for prior information on the parameters. For instance, the celebrated LASSO estimator [35] relies on these principles: the square loss quantifies the error in regression settings, and the $\ell_1$ penalty brings sparsity to the solution.

In general, the model is learned from the data by solving the following optimization problem:

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} \sum_s l\big(\boldsymbol{y}^s, f_{\boldsymbol{\theta}}(\boldsymbol{x}^s)\big) + p(\boldsymbol{\theta}), \qquad (1)$$

where $\boldsymbol{x}^s$ are features extracted from the input space, here structural connectivity, and $\boldsymbol{y}^s$ are the corresponding observed features of the output space, here functional connectivity.

In the empirical risk minimization framework, specifying our learning problem entails the choice of a predictive model with the associated function $f_{\boldsymbol{\theta}}$, a loss $l$, and a penalization $p$. Several considerations come into play to guide these choices. First, the prediction function $f_{\boldsymbol{\theta}}$ constitutes the forward model and thus reflects our understanding of the mechanisms linking $\boldsymbol{x}$ and $\boldsymbol{y}$. The simplest and most parsimonious assumption here is the linear model. With regards to this inter-modality mapping, prior work has shown that the strength of correlations observed in the resting-state signal between a pair of regions correlates with a measure of their anatomical connectivity [30]. It is thus reasonable to stipulate a linear relation between measures of functional and anatomical connectivity. We describe the anatomo-functional mapping in details in section III-A.

Second, the loss should be well-suited to quantify errors for the output space. In particular, it must be consistent with its distribution. For instance, as we will see, the square loss –

Euclidean distance– is ill-suited to estimate the error between predicted functional connectivity and estimated functional connectivity. However, it leads to well-known estimators and efficient optimization algorithms. Therefore, we will formulate the learning problem on a parameterization well suited for the square loss. Third, the penalization is used to encode prior knowledge on the problem, *i.e.* specify which model to prefer out of a family of models fitting the data equally well. In the proposed framework this prior knowledge reflects sparsity.

Since, we have chosen to formalize this learning problem as a structured-output multivariate regression, suitable for predicting multiple interdependent variables: rather than predicting each functional connection separately, we seek to predict a complete matrix with a structure compatible with our output space. We show how to achieve this in section III-B. Statistical consistency is an important property for our learning framework to be useful as an inference tool. A learning strategy is consistent if its prediction error converges to the expected error in the large sample limit. Empirical risk minimizers achieve such consistency under weak assumptions, such as convexity of the loss and stability of the estimates under resampling [36], [37]. We demonstrate evidence that these properties are fulfilled in the results section.

*Parameterization of Structural and Functional Connectivity :* A statistical framework linking functional and anatomical connectivity across subjects is underpinned by variability models for anatomical and functional connectivity. A common practice to investigate variations of anatomical connectivity across subjects is to use linear models on a tract-based measure of fractional anisotropy [38]. While the relationship between local changes in fractional anisotropy and a measure of connectivity between regions is subject to debate, modeling inter-individual differences in anatomical connectivity using independent tract-level statistics is a well-accepted paradigm [39]. For our purposes, such an approach implies that the off-diagonal coefficients of our anatomical connectivity matrix $\boldsymbol{A}^s$ are good candidates for the input variable to our learning problem. Given that linear relations are the most parsimonious hypothesis from a statistical standpoint and have been widely used to relate a measure of structural connectivity to observed functional connectivity [38], [39], [30], we will base our approach on a linear model from structural connectivity to functional connectivity.

Another challenge of predicting functional brain connectivity is that it can be described by many different parameters. Most often, the functional connectivity between a set of regions is described by the covariance matrix, $\boldsymbol{\Sigma} \in \mathbb{R}^{n \times n}$, between the time series of the mean activation in the different regions of interest. However, this characterisation reflects signal transmission from one region to another via indirect links. Partial correlation reflects the level of interaction between two regions after removal of the common influences from all the other regions. Therefore, it is a closer description of direct signal transmission from one area to another. In fact, Fransson and Marrelec [40] reported that partial correlations identified between-regions describe interaction better than correlations based on rs-fMRI. Several studies [41], [32], [1] have shown that the precision matrix, the inverse of the covariance,

$K = \Sigma^{-1}$, has adequate properties for the estimation of brain functional networks. Note that the partial correlation matrix $P$ is estimated from the precision matrix $K$ with normalization to have unit diagonal and reversing the sign of the off-diagonal elements. We describe how to obtain a suitable precision matrix in section III-A.

## III. JOINT MODELING OF ANATOMIC AND FUNCTIONAL CONNECTIVITY

In this section, we introduce a joint generative model of fMRI time series and anatomical connections. This model is used to reformulate the learning problem into an easier problem with output variables in a vector space. We also discuss the model selection: comparing different models based on predictive power. Finally, we use state-of-the-art sparse recovery techniques to identify the most robust associations between function and structure and quantify confidence for these measurements.

### A. Generative model

Sparsity is introduced in two levels. Firstly, the structural support removes noisy and irrelevant connections in a unified manner across structure and function. This constrains the conditional dependencies between the time-series, within each subject, and improves the estimation of partial correlation. Secondly, sparse linear regression, across subjects, links each functional connection with a subset of structural connections and it allows data-driven extraction of multi-stage pathways.

We use the square loss only to estimate the fit of each functional connection independently, whereas for the whole functional connectivity matrices we discuss losses suitable for SPD matrices, later in this section.

*Single-subject generative model for fMRI:* In this paragraph, we model the observed fMRI data for a single subject. We temporarily drop the subject exponent on the matrices.

The rs-fMRI data are described by the correlation matrices of the activity in the different ROIs. The simplest probabilistic model with a given order-two moment is the multivariate Gaussian model. A graphical model representation of such distribution gives a picture that can be linked to brain connectivity. Following [41], we assume a multivariate autoregressive model (MAR) to describe the generative process of fMRI time series. If $\boldsymbol{f}(\tau) \in \mathbb{R}^n$ is the multivariate vector giving the fMRI signal in the $n$ ROIs at time $\tau$, $\boldsymbol{f}(\tau + 1)$ is given as:

$$\boldsymbol{f}(\tau + 1) = \boldsymbol{T}\,\boldsymbol{f}(\tau) + \boldsymbol{e}(\tau + 1), \qquad (2)$$

with $\boldsymbol{e}$ additive Gaussian noise between variables with zero mean and identity covariance, and $\boldsymbol{T} \in \mathbb{R}^{n \times n}$ a matrix specifying the connections between variables that can be understood as transition probabilities.

We consider the resting-state brain activity as a stationary process. Eq. (2) is then written as a zero-lag process:

$$\boldsymbol{F} = \boldsymbol{T}\,\boldsymbol{F} + \boldsymbol{E} \quad \Leftrightarrow \quad \boldsymbol{F} = (\boldsymbol{I}_n - \boldsymbol{T})^{-1}\boldsymbol{E}, \qquad (3)$$

where $\boldsymbol{F} \in \mathbb{R}^{n \times t}$ is the observed signal matrix for $t$ time points, and $\boldsymbol{E} \in \mathbb{R}^{n \times t}$ the corresponding noise matrix. Thus,

the covariance of the observed time series, $\boldsymbol{\Sigma}$, is given by:

$$\begin{aligned} \boldsymbol{\Sigma} \;&=\; \frac{1}{t}\boldsymbol{F}\,\boldsymbol{F}^{\mathsf{T}} \;=\; (\boldsymbol{I}_n - \boldsymbol{T})^{-1}\mathrm{cov}\boldsymbol{E}\,(\boldsymbol{I}_n - \boldsymbol{T})^{-\mathsf{T}} \quad (4) \\ &=\; \left((\boldsymbol{I}_n - \boldsymbol{T})^{\mathsf{T}}(\boldsymbol{I}_n - \boldsymbol{T})\right)^{-1}, \qquad\qquad (5) \end{aligned}$$

as $\mathrm{cov}\boldsymbol{E} = \boldsymbol{I}_n$. In the context of multivariate Gaussian, $\boldsymbol{\Sigma}^{-1}$, the inverse covariance is called the precision matrix $\boldsymbol{K}$. We call $\boldsymbol{B} = \boldsymbol{I}_n - \boldsymbol{T}$ the *interaction matrix*. $\boldsymbol{B}$ can be easily linked to region-to-region connectivity via $\boldsymbol{T}$ and gives a full parameterization of our model as a directed Gaussian graphical model. It follows from Eq. (5) that $\boldsymbol{B}$ is a matrix square root of the precision matrix of the observed time series:

$$\boldsymbol{K} = \boldsymbol{\Sigma}^{-1} = \boldsymbol{B}^{\mathsf{T}}\boldsymbol{B}. \qquad (6)$$

*Anatomical to functional model:* We use a multivariate linear model between all the anatomical connectivity values and the coefficients of the matrix $\boldsymbol{B}$. Formally, we introduce vector representations of the matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ simply by taking all their coefficients above the diagonal:

$$\begin{aligned} \boldsymbol{x}^s &= \mathrm{Vec}(\boldsymbol{A}^s) = \{\boldsymbol{A}_{i,j}^s, i > j\}, \qquad \boldsymbol{x}^s \in \mathbb{R}^{1/2\,n\,(n-1)}, \\ \boldsymbol{y}^s &= \mathrm{Vec}(\boldsymbol{B}^s) = \{\boldsymbol{B}_{i,j}^s, i > j\}, \qquad \boldsymbol{y}^s \in \mathbb{R}^{1/2\,n\,(n-1)}. \end{aligned}$$

The model is then written

$$\forall k \in \{1 \dots \tfrac{1}{2}n\,(n-1)\} \qquad \boldsymbol{y}_k^s = \boldsymbol{\beta}_k\,\boldsymbol{x}^s, \qquad (7)$$

where $\boldsymbol{\beta}_k$ are the coefficients relating the functional connection $\boldsymbol{y}_k$ to the whole-brain anatomical connectivity[1], $\boldsymbol{x}$. The interpretation of this model is that each coefficient of the matrix $\boldsymbol{B}$ is the transition probability between regions, which is a linear function of the various anatomical connections.

Note that, unlike $\boldsymbol{K}$, $\boldsymbol{B}$ is not constrained to be positive-definite. In addition to the simple interpretation of the coefficients of $\boldsymbol{B}$, this property is a reason for choosing $\boldsymbol{B}$ as a parameterization of functional connectivity for our predictive model. Indeed, we can predict all coefficients of $\boldsymbol{y}$ separately without any constraints.

### B. Estimation of the joint model

Here we give the learning strategy to estimate the model described above. The different steps of the full estimation procedure are illustrated in Fig. 1.

*Covariance selection for the fMRI precision matrix:* In Gaussian models, conditional independence between variables, also known as the Markov structure of the model, is given by the zeros in the precision matrix. Imposing a conditional independence structure, *i.e* estimating from the fMRI data a sparse precision matrix, reduces the small-sample estimation error present in the empirical covariance matrix [32], [42]. This is known in statistics as *covariance selection* [43]. Indeed, for a fully-connected graphical model, the number of functional connectivity parameters can be much greater than the number of samples: $t < \frac{1}{2}n\,(n-1)$. This results in a large estimation error on the sample covariance matrix [44], [32].

---

[1] $\boldsymbol{\beta}_k$ also contain a coefficient representing an offset, or *intercept*, that we omit from the equation, as can be conflated into the vector of coefficients $\boldsymbol{\beta}_k$ by adding a constant term to the design matrix $\{\boldsymbol{x}^s\}$.

Fig. 1. Processing pipeline for predicting functional connectivity from structural brain connectivity:

**1.** Tracts are transformed into a connectivity measure between ROIs

**2.** On each subject, the fMRI signal is extracted for each ROI

**3.** A T-test on anatomical connectivity matrices across subjects gives the graph of connected regions common to the group

**4.** Using this graph, a sparse inverse covariance is estimated from the fMRI time series with the IPS algorithm

**5.** An approximated minimum degree ordering is found for graph $G$

**6.** Using this ordering, the interaction matrix for each subject is computed from the Cholesky decomposition of the subject's fMRI precision matrix

**7.** The interaction matrix is rescaled according to Eq. 8

**8.** A multivariate linear predictor is learned from the subjects' anatomical matrices to the interaction matrices
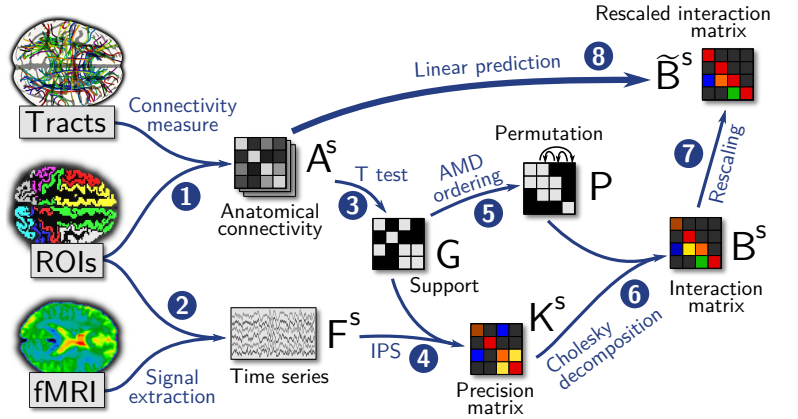
To impose a conditional independence structure of the fMRI model, we use anatomical connections that are not significantly positive across subjects [1]: given the anatomical connectivity matrices across the various subjects, $\{\boldsymbol{A}^s, s = 1 \ldots S\}$, for each pair of regions $(i, j)$, we perform a T-test to test $\boldsymbol{A}^s_{i,j} \neq 0$ in our population of subjects. The output of this procedure is a set of candidate connections, that can be represented as a non-valued graph $\mathcal{G}$, or equivalently a symmetric matrix $\boldsymbol{G} \in \mathcal{S}ym_n$ of zeros and ones. We call this matrix the *support* of our functional and anatomical connectivity graphs.

The maximal likelihood estimate of $\boldsymbol{K}^s$ for each subject can then be computed using the iterative proportional scaling algorithm (IPS) [44], taking as an input this support $\boldsymbol{G}$, and the observed time series for the subject $\boldsymbol{F}^s$.

*Estimation of the interaction matrix $\boldsymbol{B}^s$:* We use the Cholesky decomposition of the precision matrix $\boldsymbol{K}^s$ to estimate $\boldsymbol{B}^s$. The Cholesky decomposition is not invariant by a permutation of the rows and columns of the input matrix $\boldsymbol{K}$. We use the approximate minimum degree (AMD) ordering [45], which provides a permutation of $\boldsymbol{K}^s$ favoring a sparser Cholesky decomposition. This permutation depends only on the support of the matrix $\boldsymbol{K}^s$, $\boldsymbol{G}$, and thus it is identical across subjects. Note that the scaling of the fMRI activity recording in the different ROIs is generally not deemed related to structural information. For this reason, most investigators study correlation, and not covariance, which corresponds to imposing a certain scaling to the columns of $\boldsymbol{B}$. For this purpose, we rescale the diagonal of $\boldsymbol{B}$ to ones with the following rule:

$$\tilde{\boldsymbol{B}}^s = \boldsymbol{B}^s \, \text{diag}(\boldsymbol{B}^s)^{-1}, \qquad (8)$$

where $\text{diag}(\boldsymbol{B}^s)$ is the diagonal matrix made of only the coefficients on the diagonal of $\boldsymbol{B}^{s\,2}$. In simpler terms, we consider that the partial-correlations, and not the elements of the inverse covariance, are the quantities that can be mapped to structural connectivity across subjects.

*Anatomical to functional regression:* The binary matrix $\boldsymbol{G}$ imposes a common support across subjects to anatomical

---

<sup>2</sup>Note that another possible rescaling strategy would correspond to multiplying $\boldsymbol{B}^s$ on the left, rather than on the right, in Eq. (8). As can be seen from Eqs. (3) and (6), this amounts to setting the covariance of the noise term $\boldsymbol{E}$ in the MAR and it imposes that our matrix $\tilde{\boldsymbol{B}}^s$ reflects only the structure and not the innovation terms of the MAR. We do not discuss this option further, as it gives poorer prediction in the experiments below.

and functional connectivity. However, inside this support, these two quantities are modeled with multivariate linear model, specified by Eq. 7. Without any prior information, this linear model is ill-posed as there are many more candidate anatomical connections than subjects: we are in high-dimensional settings. Therefore we use the LASSO [35], which is a sparse $\ell_1$-penalized regression:

$$\hat{\boldsymbol{\beta}}_k = \underset{\boldsymbol{\beta}}{\arg\min} \left( \frac{1}{2} \|\boldsymbol{Y}_k - \boldsymbol{\beta}\boldsymbol{X}\|^2 + \lambda \|\boldsymbol{\beta}\|_1 \right), \qquad (9)$$

where $\boldsymbol{X}$ and $\boldsymbol{Y}$ are the multi-subject concatenation of the input and output vectors: $\boldsymbol{X} = \{\boldsymbol{x}^s, s = 1 \ldots S\}$ and $\boldsymbol{Y} = \{\boldsymbol{y}^s, s = 1 \ldots S\}$. LASSO offers two major advantages over ordinary least square regression that are very useful in modeling brain connectivity. Firstly, it improves prediction by setting some coefficients to zero. This results in removing noisy and irrelevant variables and thus reducing the total variance. Secondly, it allows the selection of the most relevant variables and thus it links each functional connection with a subset of structural connection in a data driven way. Under certain assumptions, if the true model is sparse, the correct predictors can be identified even when the number of variables is higher than the number of observation [46].

### C. Model selection and identification

Here we discuss how to select the model parameters and the anatomo-functional links that we can reliably consider as non zero.

*Choice of model parameters:* The model and associated estimation procedure presented above takes two parameters: the significance threshold $\alpha$ on the inter-subject test of anatomical connectivity between pairs of ROIs, and the amount of penalization $\lambda$ on the LASSO regression from anatomical to functional connectivity. The parameters set the sparsity of the model at two different levels: $\alpha$ set the joint sparsity of the anatomical and functional graphs, while $\lambda$ sets the sparsity of the relation from anatomy to function.

We set a different value of $\lambda$ for each connection to be predicted. For this, given $\boldsymbol{X}$ and $\boldsymbol{Y}_k$ in as in Eq. 9, we apply 5-fold cross-validation, choosing $\lambda$ to maximize the explained variable of unseen samples of $\boldsymbol{Y}_k$. We denote this procedure of selecting the $\lambda$-value by cross-validation fitting a LASSO on the full data as *LassoLarsCV*. The *LassoLarsCV* is

a parameter-free algorithm. Note that this cross-validation can be considered as nested, since it is performed in an internal loop of the general estimation algorithm. We set $\alpha$ to maximize predictive power of our model on unseen data using random hold-out cross validation with 10% held out.

*Losses of functional connectivity:* The Euclidean distance is not suitable to quantify differences between functional graphs. $\mathcal{S}ym_p^+$ can be parameterized as a Riemannian manifold using an intrinsic metric [47], which is suitable to build a full statistical framework on $\mathcal{S}ym_p^+$ [48], [49]:

$$d_{AI}(\boldsymbol{C}, \boldsymbol{D})^2 = \text{tr}\big(\log \boldsymbol{C}^{-\frac{1}{2}} \boldsymbol{D} \boldsymbol{C}^{-\frac{1}{2}}\big)^2 \qquad (10)$$

This metric is invariant under affine scaling and inversion of the matrices, and it is thus equivalent for a wide range of parameterization for functional connectivity. In particular, if $\boldsymbol{D}$ is the predicted matrix, and $\boldsymbol{C}$ the ground truth, it gives the same prediction error on the correlation matrices and on the precision matrices. The metric $d_{AI}$ cannot be used with non-structured output predictors, such as the independent prediction of each coefficient of the functional connectivity matrix, because they cannot guarantee that the prediction will be positive definite. To compare all models, whether they give SPD predictions or not, with a prediction error independent of the parametrization of functional connectivity, we develop the metric given in Eq. (11) around the target matrix, $\boldsymbol{\Xi}^{\text{target}}$, that can be a correlation, precision, partial correlation matrix and it is by construction positive definite:

$$d(\boldsymbol{\Xi}^{\text{pred}}, \boldsymbol{\Xi}^{\text{target}}) = \|(\boldsymbol{\Xi}^{\text{target}})^{-1}(\boldsymbol{\Xi}^{\text{target}} - \boldsymbol{\Xi}^{\text{pred}})\| \qquad (11)$$

For our generative model of the fMRI signal, a natural measure of distance between the predicted fMRI covariance and the observed one is the $KL$ divergence of these models. Under the assumption, that fMRI time series are well-described by a multivariate normal model parameterized by a covariance matrix $\boldsymbol{\Sigma}$, the relevant measure of divergence between two models, a predicted and a target model, is the *Kullback-Leibler divergence*[3]:

$$\mathcal{KL}(\boldsymbol{\Sigma}^{\text{pred}}, \boldsymbol{\Sigma}^{\text{target}}) = \frac{1}{2}\bigg(\log\det(\boldsymbol{\Sigma}^{\text{pred}}) - \log\det(\boldsymbol{\Sigma}^{\text{target}})$$
$$+ \text{tr}\big((\boldsymbol{\Sigma}^{\text{pred}})^{-1}\boldsymbol{\Sigma}^{\text{target}}\big)\bigg) \qquad (12)$$

Unlike $d_{AI}$ or $d$, it applies only to covariance matrices, and is not invariant, or close to, by scaling or re-parameterization.

*Model identification:* Once we have selected the optimal model in terms of predictive performance, we seek to identify the important anatomo-functional non-zero coefficients $\beta_k$. For this purpose, we use the randomized LASSO, a generalization of LASSO with better sparse recovery properties [50]. The randomized LASSO estimate is computed by solving the LASSO problem with random weights $W_j$ within specified

---
[3]Note that neither $d$ nor the KL divergence are distances, as they are not symmetric. They however are suitable to define a loss.

---

**Algorithm 1:** Anatomo-functional model estimation

**Input:** $\{\boldsymbol{F}^s \in \mathbb{R}^{n \times t}, s = 1, \dots, S\}$, the time series for each subject; $\{\boldsymbol{A}^s \in \mathbb{R}^{n \times n}, s = 1, \dots, S\}$, the anatomical connectivity matrices for each subject; $\alpha$, the threshold on the p-value for group-level anatomical connectivity.

**Output:** $\boldsymbol{G} \in \{0,1\}^{n \times n}$ the support of functional connectivity graphs; $\boldsymbol{P} \in \{0,1\}^{n \times n}$, permutation of the ROIs' ordering; $\{\beta_k \in \mathbb{R}^{n \times n}\}$ the anatomo-functional link coefficients[b].

1: $\boldsymbol{G} \leftarrow \boldsymbol{O}_n$ (initialize empty support)
2: **for** i=1 to n, j=1 to n, where T-test($\{\boldsymbol{A}_{i,j}^s\} \neq 0) > \alpha$ **do**
3:     $\boldsymbol{G}_{i,j} \leftarrow 1$
4: **end for**
5: $\boldsymbol{P} \leftarrow$ AMD($\boldsymbol{G}$) (Approximate Minimum Degree ordering)
6: **for** s=1 to S **do**
7:     $\boldsymbol{K}^s \leftarrow$ IPS($\boldsymbol{F}^s, \boldsymbol{G}$) (Iterative Proportional Fitting)
8:     $\boldsymbol{B}^s \leftarrow$ Cholesky($\boldsymbol{P}\,\boldsymbol{K}^s$)
9:     $\tilde{\boldsymbol{B}}^s \leftarrow \boldsymbol{B}^s \text{diag}(\boldsymbol{B}^s)^{-1}$ (scale $\boldsymbol{B}^s$ according to Eq. 8)
10: **end for**
11: $\boldsymbol{X} = \{\text{Vec}(\boldsymbol{A}^s)\}; \quad \boldsymbol{Y} = \{\text{Vec}(\tilde{\boldsymbol{B}}^s)\}$
12: **for** i=1 to n, j=1 to n, such that $\boldsymbol{G}_{i,j} \neq 0$ **do**
13:     $\beta_k \leftarrow$ LassoLarsCV($\boldsymbol{X}, \boldsymbol{Y}_k$); $\quad k = i + nj$
14: **end for**

[b]Note that as the matrices are symmetric, we only need to compute at most $^1/_2\,n\,(n-1)$ coefficients. We write the full computations to simplify the exposition of the algorithm.

---

bounds[4]:

$$\hat{\beta}_k = \underset{\beta}{\text{argmin}}\bigg(\|\boldsymbol{Y}_k - \beta\,\boldsymbol{X}\|^2 + \lambda \sum_{j=1}^{N} \frac{|\beta_j|}{W_j}\bigg) \qquad (13)$$

This randomized penalization regression is solved many times. The probability that a functional connection is related to a anatomical connection is then given by the fraction of times the coefficient is selected during the repetitions. The benefit is that the identification of relevant anatomo-functional links is not sensitive to the parameters of the regression equations. Note that we have one parameter for each functional connection, $y_k$. For the randomized LASSO, we use the penalty parameter $\lambda_k$ selected previously by cross-validation with the *LassoLarsCV*.

## IV. EXPERIMENTS

### A. Multi-subject and multi-modal brain imaging dataset

Brain connectivity analysis was performed in 41 normal adults (24 males, age $29.07 \pm 9.54$ years). The data were acquired at a 3T Philips scanner of the Computational, Cognitive and Clinical Neuroimaging Laboratory (C3NL), Imperial College London, UK. Resting-state fMRI data acquisition was performed with eyes closed. Subjects were instructed neither to fall asleep, nor to think of something in particular. rs-fMRI: T2*-weighted gradient EPI sequence, TR/TE=2000/30, 31 ascending slices with thickness $3.25mm$, gap $0.75mm$, voxel size $2.5 \times 2.5 \times 4mm$, flip angle 90º , FOV $280 \times 220 \times 123mm$, matrix $112 \times 87$. DWI: 64 non-collinear directions, in 72 slices, slice thickness $2mm$, FOV $224mm$, matrix 128x128, voxel

---
[4]$W_j$ chosen randomly equal to .5 or 1, as recommended by [50]. In practice, the results are not very sensitive to the randomization strategy used to draw $W_j$.

size $1.75 \times 1.75 \times 2 mm^3$, b-value 1000 $s/mm^2$. High resolution T1-weighted whole-brain structural images were also obtained in all subjects.

FSL was used for image pre-processing of both DWI and rs-fMRI images [51]. This involved eddy current correction of DWI and motion correction as well as spatial smoothing and whitening of rs-fMRI images. Brain extraction was performed originally with FSL. Multi-resolution bias correction with N4ITK [52] was applied to T1s, B0s and the middle fMRI volumes to improve the robustness of the non-rigid registration tools [53].

BOLD fluctuations are profound in gray matter, while tractography is more reliable in delineating white matter fibers. Hence, we are interested in defining cortical ROIs that are located in gray matter. There are usually two competing approaches in defining brain ROIs: *i)* atlas-based ROIs and *ii)* connectivity-based ROIs. We tested our methodology with ROIs from both approaches.

**Atlas-based ROIs:** Atlas-based ROIs are accurately defined in T1 space using a robust label propagation technique [54]. Note that labels of the original atlases were subdivided into their white matter (WM), gray matter (GM) and cerebrospinal fluid (CSF) portions using FSL FAST tissue classification [55]. The individual atlases were propagated using transformations calculated with a non-rigid registration procedure based on free-form deformations (FFD) [53]. The propagated label sets are then merged into a consensus probabilistic segmentation using a locally weighted fusion strategy [56]. These probabilistic labels were then used as spatial priors in a subsequent refinement step, where a probabilistic intensity model is solved using the Expectation-Maximization (EM) algorithm. For this procedure we used our previously published method as described in [57]. These automatically computed labels are then transferred from T1 space to both fMRI and B0 space using intra-subject, non-rigid registration [53]. In total we used 54 cortical regions.

**Connectivity-based ROIs:** Connectivity-based ROIs are predefined in MNI space and they are derived from a multi-subject independent component analysis (ICA) [58]. Subsequently, specific regions were selected by an expert review of the resting-state literature. An effort was made to split large regions into smaller and the boundaries were manually refined so there is no overlap. We identify ROIs that belong to the default mode network (DMN): *i)* posterior cingulate cortex (PCC), *ii)* medial prefrontal cortex (mPFC), *iii)* superior parietal regions (SupPar), *iv)* inferior parietal regions (InferPar), *v)* occipital regions (Occ) and *vi)* thalamus (Thal). We also identify ROIs that belong to the salience network (SN): *i)* anterior insula (AI), *ii)* presupplementary motor area (preSMA) and *iii)* middle frontal regions (MidFront). These regions are propagated to native fMRI and B0 space with non-rigid registration. Left and right hemispheric parts are identified based on each subject atlas-based segmentation. We also use the subject GM mask to keep only voxels within GM. We used a total of 18 regions.

To construct corresponding functional networks the fMRI signal was averaged across voxels within each area. Confound parameters were regressed out with linear regression: the mean CSF and white-matter signal, as well as the six motion correction parameters estimated with FEAT, FSL [51].

To build the anatomical connectivity matrices $A^s$, the tracks between regions are identified based on the standard ball and stick model available in FSL [51], [59]. However, measurements of connection probability are difficult to interpret as the probability measure is very sensitive to noise in the data, as well as the size and separation of the ROIs [23], [24]. In our work, the strength along these tracks is determined as the integral of the local diffusion anisotropy. This is estimated based on the diffusive transfer between voxels using the orientation distribution function (ODF). The ODFs are calculated directly from all samples on the volume fraction of diffusion, allowing direct correspondence with the underlying fiber model [23], [24].

### B. Model validation by prediction performance

Here, we use cross-validation to compare our model to other approaches based on prediction performance across subjects. Specifically, we revisit the different choices relevant for the construction of the model: *i)* structured prediction using a Cholesky decomposition, *ii)* imposing a common graph support in both modalities *iii)* ordering the MAR with the AMD ordering, and *iv)* scaling the innovation terms to 1 –Eq. (8). The various alternatives tested are:

**MAR AMD:** the model presented in section III estimated with Algorithm 1.

**MAR Random-order:** the same model and estimation strategy, but choosing random orders rather than an approximate minimum degree ordering before performing the Cholesky, and averaging the prediction for many of these random choices.

**Dense:** the same model and estimation strategy, but rather than using the IPS to estimate a sparse inverse covariance, we estimate a dense precision matrix using the Ledoit-Wolf covariance shrinkage [60]. This approach guaranties a more accurate and well-conditioned covariance estimate under little assumptions and with no parameters to tune.

**Unstructured:** independent prediction of each entry of the precision matrix. The precision matrix has been estimated based on the Ledoit-Wolf covariance shrinkage.

Finally, to set the filling factor of anatomical and functional connectivity graphs that gives the optimum prediction performance, we vary the threshold $\alpha$ on the T-test so that the percentage of selected connections ranges from 10% to 90%.

In Fig. 2 we show a qualitative example of how a predicted interaction matrix ($B^S$), Fig. 2a, is transformed to a precision matrix, Fig. 2b, by multiplying with its transpose and subsequently it is inverted to derive a correlation matrix, Fig. 2c, in the AMD model. The derived correlation matrix, Fig. 2c, is subsequently compared to the 'ground truth' correlation, Fig. 2f, with each of the measures described above: *i)* $d_{AI}$, *ii)* $d$ and *iii)* $KL$. The 'ground truth' correlation is estimated as the correlation coefficient of each pair of average time-series

(a) Predicted interaction matrix      (b) Corresponding predicted precision matrix      (c) Corresponding predicted correlation matrix

(d) Training interaction matrix      (e) Structural Connectivity      (f) Target correlation matrix
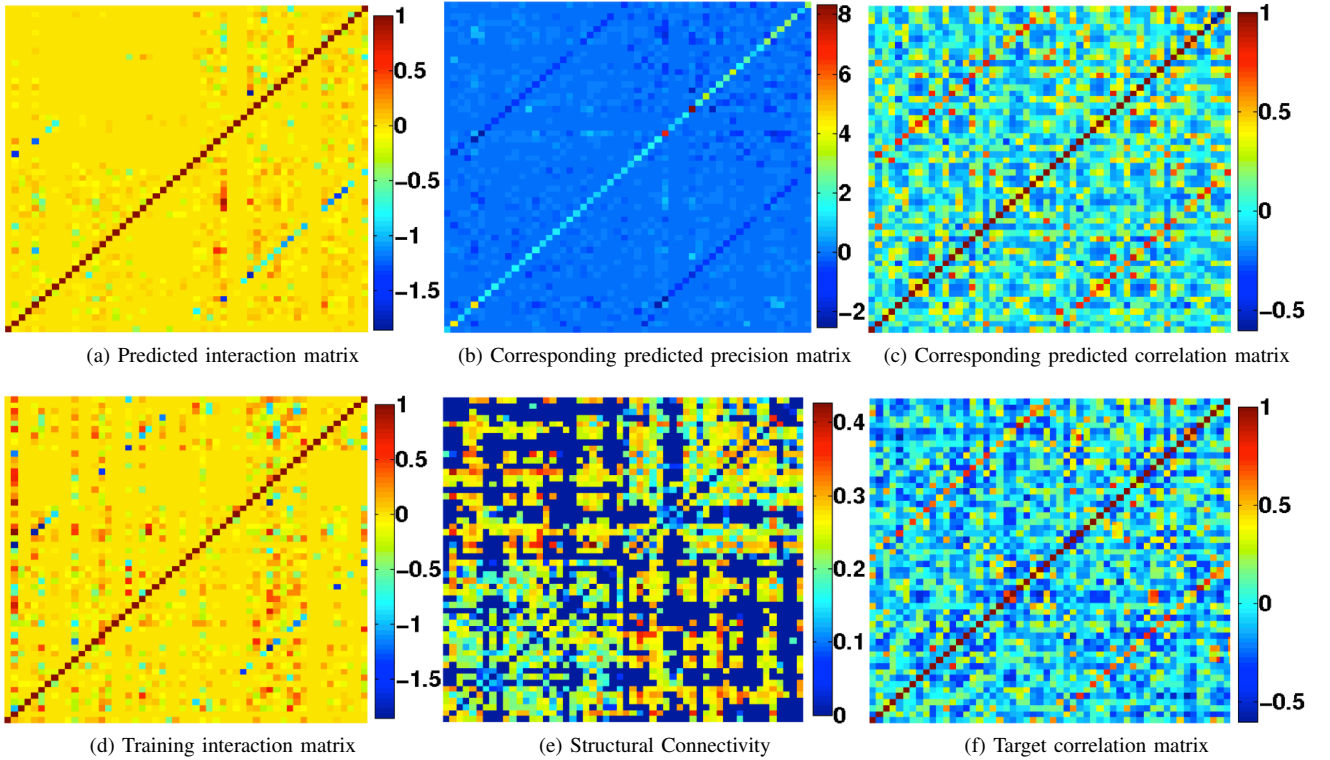
Fig. 2. We show a qualitative example of how in the AMD model a) a predicted interaction matrix is transformed to b) a precision matrix by multiplying with its transpose and subsequently it is inverted to derive c) a correlation matrix. The derived correlation matrix Fig. 2c is subsequently compared to the 'ground truth' correlation Fig. 2a with each of the measures described above: *i)* $d_{AI}$, *ii)* $d$ and *iii)* $KL$. We also show an example of d) an interaction matrix used for training the model as well as e) the structural connectivity matrix used for prediction of the interaction matrix a). The structural connectivity is derived with 70% support. In f) we show the 'ground truth' correlation, which is estimated as the correlation coefficient of each pair of averaged time-series within each ROI once the WM, CSF and motion parameters have been regressed out. Note that the functional-connectivity matrices display diagonal features in the middle of the upper and lower part of the matrices. These are due to homologous inter-hemispheric connections; in other words, each region tends to be highly correlated with the corresponding region in the opposite hemisphere.

within each ROI once the WM, CSF and motion parameters have been regressed out. We also show an example of an interaction matrix used for training the model, Fig. 2d, as well as the structural connectivity matrix used in this example for prediction, Fig. 2e. We follow the same procedure for all the models that give SPD predictions. For the unstructured model we predict the precision matrix directly and subsequently it is inverted to derive the corresponding correlation matrix. Note that precision matrices inherently reflect partial correlation between ROIs and are more suited for the association of structure with function.

*C. Identification of relevant anatomo-function connections*

Once the best performing model has been selected using cross-validated predictive power, we are interested in identifying which structural connections are determined given a functional connection. This is important to interpret the results, and possibly draw neuroscientific or medical conclusions in an applicative context. Thus, for the model estimation we use randomized LASSO across all 41 subjects. This results in a sparse array with rows that correspond to functional connections and columns that correspond to the structural connections and the probability to have been selected. Although, there is a profound lack of ground truth data, there are numerous functional studies that draw conclusions about functional

and structural connectivity among regions implicated in well known networks such as the DMN and the SN. The SN is thought to regulate dynamically other networks, such as the DMN and it is implicated in impaired inhibition and slow information processing speed.

## V. RESULTS

*A. Model selection*

Fig. 3 summarises the results of model selection with both atlas-based ROIs, Fig. 3a-3c and connectivity-based ROIs, Fig. 3d-3f. Each column demonstrates the prediction performance as it is measured with the intrinsic metrics $d_{AI}$ introduced in Eq. (10), $d$ as in Eq. (11) and the Kullback-Leibler divergence, $\mathcal{KL}$, Eq. (12), respectively. Mean values and standard deviation is estimated for each scenario with cross-validation, which consists of 100 randomly selected triplets (left-out subjects) out of the 41 subjects. Each plot shows the performance of various models explored. The performance of the unstructured model is only estimated with the $d$ metric, since the prediction is non-SPD.

The prediction performance of the proposed model, (MAR AMD) in section III outperforms all the other approaches for all the different support levels. The results are robust across different types of parcellations and different performance measures. In the MAR Random Order 100 random permutations
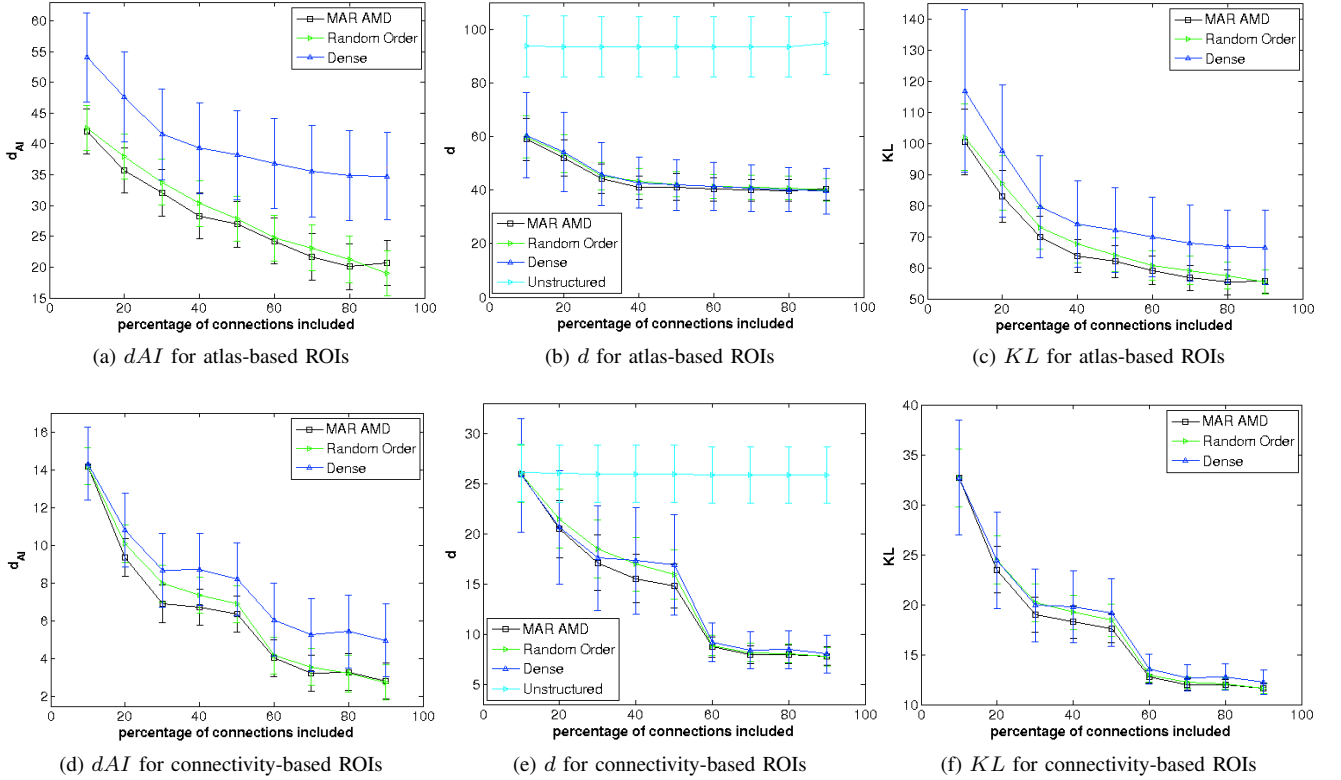
Fig. 3. Model selection and prediction performance. a-c) summarise the results of model selection and prediction performance with atlas-based ROIs. d-f) summarise the results of model selection and prediction performance with connectivity-based ROIs. Each column demonstrates the prediction performance as it measured with the intrinsic metrics $d_{AI}$ introduced in Eq. (10), $d$ introduced in Eq. (11) and the Kullback-Leibler divergence, $\mathcal{KL}$, Eq. (12), respectively. Each plot shows the performance of the following models: *i)* MAR AMD, *ii)* MAR Random Order, *iii)* Dense AMD. The performance of the unstructured model is only estimated with the $d$ metric, since the prediction is non-SPD.

| MAR AMD - Random Order | | 10% | 20% | 30% | 40% | 50% | 60% | 70% | 80% | 90% |
|---|---|---|---|---|---|---|---|---|---|---|
| | 54 | 3.3e-18 | 1.2e-49 | 1.9e-42 | 1.1e-47 | 2.7e-16 | 2.2e-07 | 2.5e-28 | 3.4e-30 | 2.0e-04 |
| | 18 | 4.0e-05 | 2.6e-48 | 1.5e-49 | 1.4e-44 | 9.0e-43 | 4.2e-05 | 6.4e-29 | 2.4e-03 | 8.9e-04 |

TABLE I

WILCOXON SIGNED-RANK TEST (PAIRED, TWO-SIDED TEST) ON THE DIFFERENCE BETWEEN THE MATCHED SAMPLES IN THE MAR AMD AND RANDOM ORDER APPROACHES, AS THEY MEASURED WITH $d_{AI}$. THE TABLE VALUES INDICATE THE PROBABILITY OF OBSERVING THE GIVEN RESULT UNDER THE NULL HYPOTHESIS.

for the interaction matrix $K$ are estimated and applied to each of the 100 randomly selected triplets, respectively. Here, we test whether random orders of the interaction matrix would significantly affect the performance of the model. Inter-subject variability accounts for most of the magnitude of the error bars. Therefore, to demonstrate that there is substantial difference between the MAR-AMD model and the Random Order, we estimate the Wilcoxon signed rank test as shown in Table I. The results reveal statistically significant differences between the two approaches for all support levels.

The unstructured model performs significantly worse than the others. In figures 3b, 3e, the loss appears to be almost constant due to the scaling. For both atlas-based ROIs and connectivity-based ROIs, it follows a similar pattern as the rest of the models: For example, with atlas-based ROIs, the difference between the predicted functional connectivity and the 'ground truth' is 93.65 for 10% support, it gradually falls to 93.4 for 80% support and it raises again to 94.7 for 90% support.

For the multiple-penalized regression we use a LASSO implementation in scikit-learn toolbox based on the Lars algorithm [61], which is computationally very efficient. This involves an intrinsic cross-validation to estimate the optimum $\lambda$. To examine the distribution of $\lambda$ across all connections and cross-validations, we plot the histogram of $\lambda$ values for different support levels at Fig. 4. We note that the optimal $\lambda$ value is consistently close to $0.015$ for most of the connections. This holds for every support level, though here we present only the histogramms for 50%, 60%, 70% and 80% connections included Fig. 4.

With a Markov structure of 90% the structural connectivity matrices are extremely ill-conditioned with several entries close to zero. In this case, the estimation of the prediction fails on average 10 out of in total 100 cross-validation trials. The results presented here are by averaging the successful cross-validation trials. For the remaining support levels there is no drop-out. The results suggest that for both ROIs defined either based on atlas-based segmentation or drawn from known
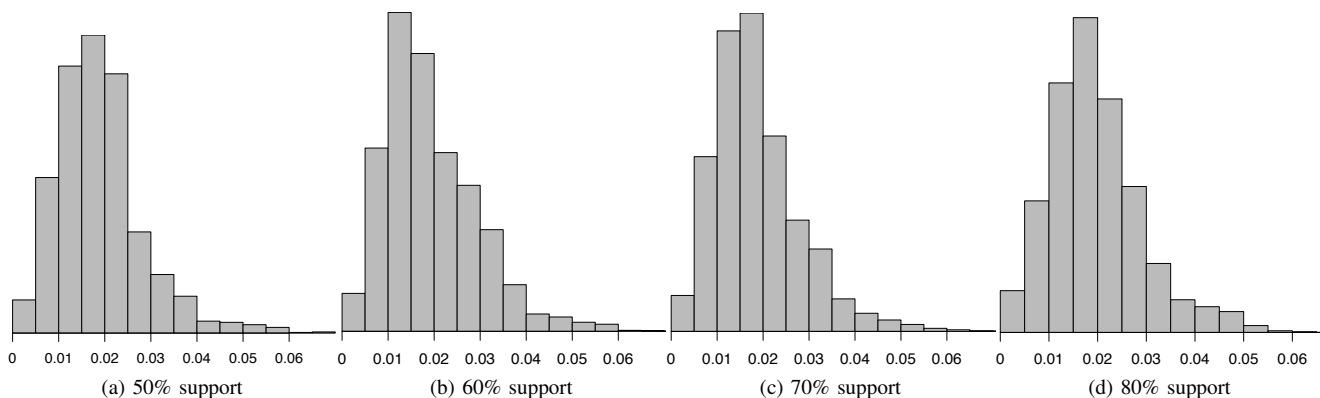
Fig. 4. Histograms of the $\lambda$ values across all cross-validations and connections for a) 50% support, b) 60% support, c) 70% support and d) 80% support. We note that the optimal $\lambda$ value is consistently close to 0.015 for most of the connections.
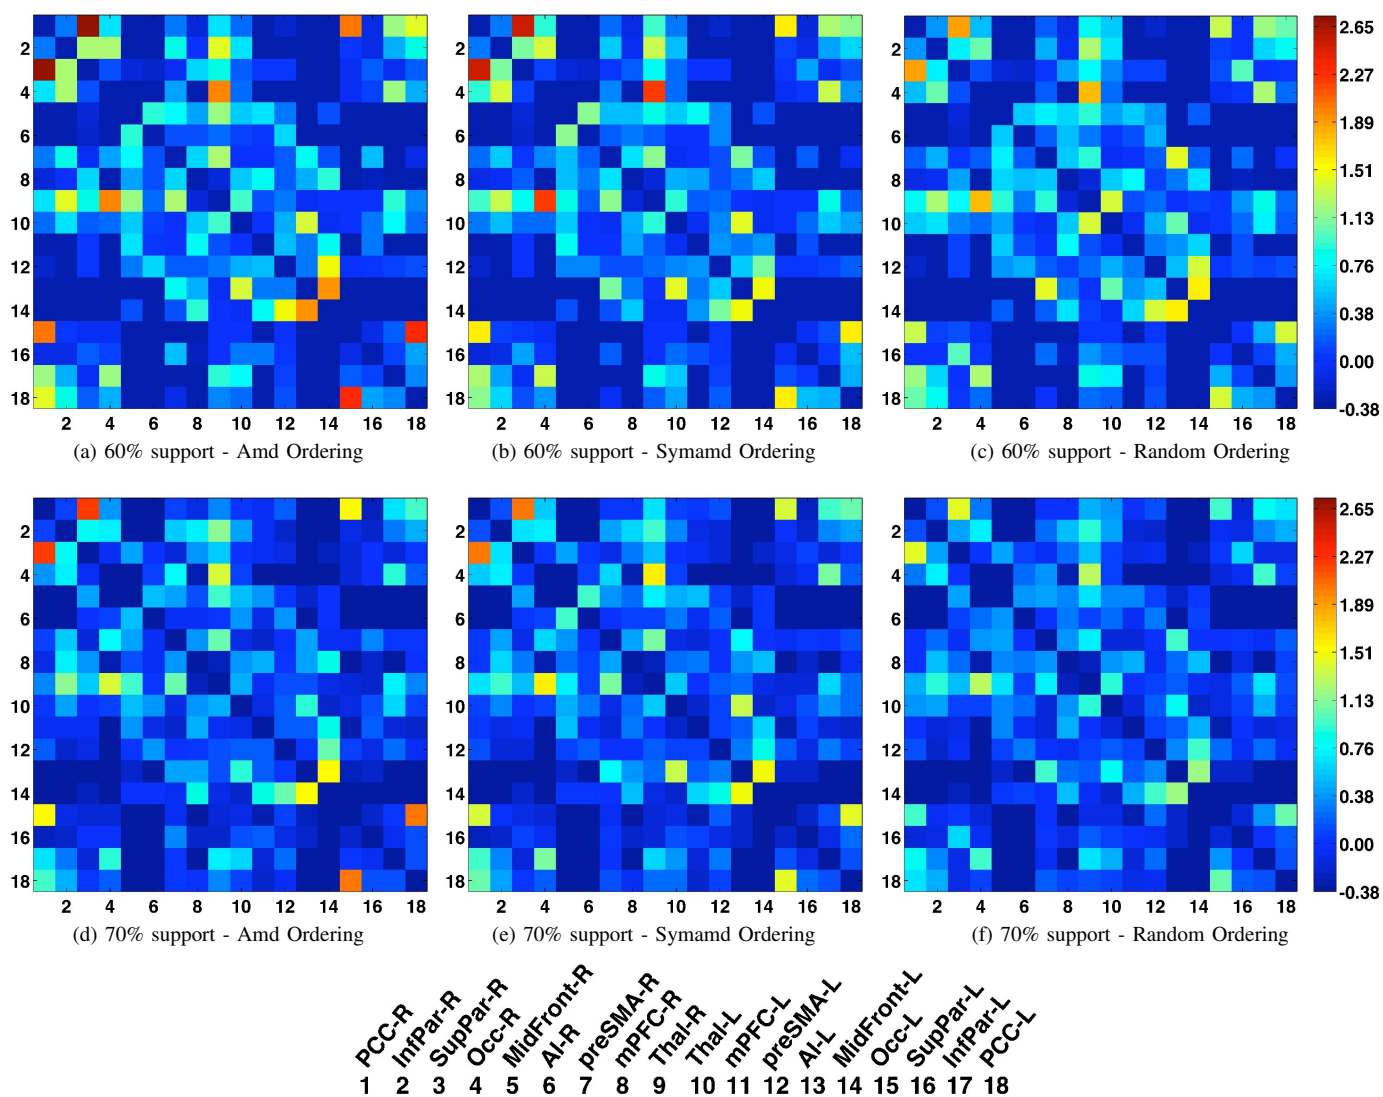


Fig. 5. Z-scores of the structural connections related to the functional connection between the left occipital lobe and the right PCC. Results are based on model identification and bootstrap by drawing the subjects with replacement over 5000 iterations. The z-scores are shown for 60% and 70% support at the top and bottom row, respectively. Three different strategies have been adapted for the ordering of the precision matrix: a,d) AMD, b,e) SYMAMD and c,f) Random Ordering.

networks a Markov structure of around $60-70\%$ achieves the higher performance across all models and metrics.

## B. Connections identification

We run model identification across all subjects for 60% and 70% of the connections included, to examine how the identified connections based on the proposed data-driven approach are related to current evidence in neuroimaging and neuroscience literature. Note that these support levels minimize the prediction error in Fig. 3. Also, the model selection demonstrated that the MAR AMD model outperforms the others. We now utilize it to identify the structural connections closely related to each functional connection. Furthermore, we use bootstrap by drawing the subjects with replacement over 5000 iterations to obtain the corresponding z-scores, which results in a matrix of $153 \times 18 \times 18$ for the connectivity-based ROIs and a matrix of $1431 \times 54 \times 54$ for the atlas-based ROIs. Fig. 5 shows one slice of the z-scores for the connectivity-based ROIs that depicts the structural connections related to the functional connection between the left occipital lobe and the right PCC.

To assess how sensitive the model identification process is to the ordering of the precision matrix, we show results from three different strategies: AMD Fig. 5a, 5d, SYMAMD Fig. 5b, 5e and Random Ordering Fig. 5c, 5f. Both AMD and SYMAMD estimate an ordering that aims to sparsify the cholesky decomposition. They provide similar but slightly perturbed ordering sequences given the same support. On the other hand, for the Random Ordering, we estimate a random ordering for each bootstrap iteration. The results demonstrate that similar key connections are identified for all the models. However, the sparse models result in significantly higher confidence intervals given the same sample size.

Subsequently, we plot the brain connections with the higher z-scores for 60% and 70% support, Fig. 6. Fig. 6 is divided in a left part (left two columns) and a right part (right two columns). The left part shows the overall structural connections linked to the function within the DMN and the SN. The right part of the figure is focused on structural connections linked to the functional connectivity between the left part of occipital lobe and the right PCC. Each column presents identification results with the same support. Each row shows the 3D brain model from top, right and left view, respectively. Each connection is represented as a 3D tube between a pair of regions with diameter proportional to the average probability over all bootstrap iterations. Connections are thresholded with only 6% with the highest z-scores shown here, which corresponds to $1.25$ and $0.95$ for 60% and 70% support levels, respectively.

The thalamus, preSMA and superior parietal regions appear as hubs of the underlying networks when all connections are taken into consideration. When we examined the functional connection between right PCC and left occipital lobe direct structural connections emerge as well as other pathways via homologous regions (left PCC) and the thalamus. Structural connections between the anterior insula and middle frontal gyrus as well as the anterior insula and the thalamus are also highlighted. Note that SN tracts to the right anterior insula have been found to predict default mode network function after traumatic brain injury in a task-related paradigm [16].

Connection identification results for the atlas-based ROIs are shown in Fig. 7. The top row shows the structural connections with the highest 2% z-score values that predict the functional connection between the left occipital lobe and the right posterior part of the cingulate gyrus (PCC). At the bottom row, we show the structural connections with the 2% highest z-scores related to the functional connection between the left and right frontal lobes. The value of z-score threshold is around $1.6$ for both examples. We note that for connectivity-based ROIs pathways emerge that involve subcortical regions, whereas for the cortical atlas-based ROIs pathways emerge that involve inter-hemispheric connections between homologous regions.

## VI. DISCUSSION

Both DWI and rs-fMRI are indirect and noisy measures of structure and function, respectively. Therefore, it is not clear how inter-subject variability in functional connectivity is related to structural connectivity and to what degree it may result from disease progression or imply it. Furthermore, there are inconsistencies between independent structural and functional connectivity studies, which may result from differences in patient groups, imaging and analysis methods [4]. Multi-modal integrative approaches of connectivity are potentially useful tools in understanding brain function and for investigating the impact of pathology on brain networks [62], [63].

## A. Blending brain modeling with high-dimensional statistics

Integrating functional and structural connectivity under a predictive framework is a very challenging statistical task since, the dimensionality of both the input and the output space increase as $n^2$, where $n$ is the number of ROIs. With 54 ROIs, there are up to 1431 coefficients to learn, for each of the 1431 functional connections, which is to be put in perspective with the typical number of subjects in a study, seldom more than a hundred. Previous anatomo-functional studies have relied on simplification of the problem, either by predicting distributed variations [64], or by only studying known univariate associations between a given functional connection and the corresponding direct structural connection [30]. Here, we learn indirect effects using a many-to-one mapping from structural to functional connections. In addition, to achieve a simple description of the mapping, we rely on node-level linear models. For this purpose, we propose a generative model that is well-suited to the high-dimensionality of the data: *i)* in the context of Gaussian graphical models we impose a Markov independence structure on functional connectivity using structural connectivity and *ii)* we employ the randomized LASSO to select the structural connections predicting a given functional connection. In addition, we assess the goodness of the fit of our model and set its different parameters using prediction performance on unseen data, the standard strategy for model selection in the context of supervised learning. In this context, overfitting is controlled by the external cross-validation loop. A bias is introduced, but in high-dimensional

(a) 60% support     (b) 70% support     (c) 60% support     (d) 70% support

(e) 60% support     (f) 70% support     (g) 60% support     (h) 70% support

(i) 60% support     (j) 70% support     (k) 60% support     (l) 70% support

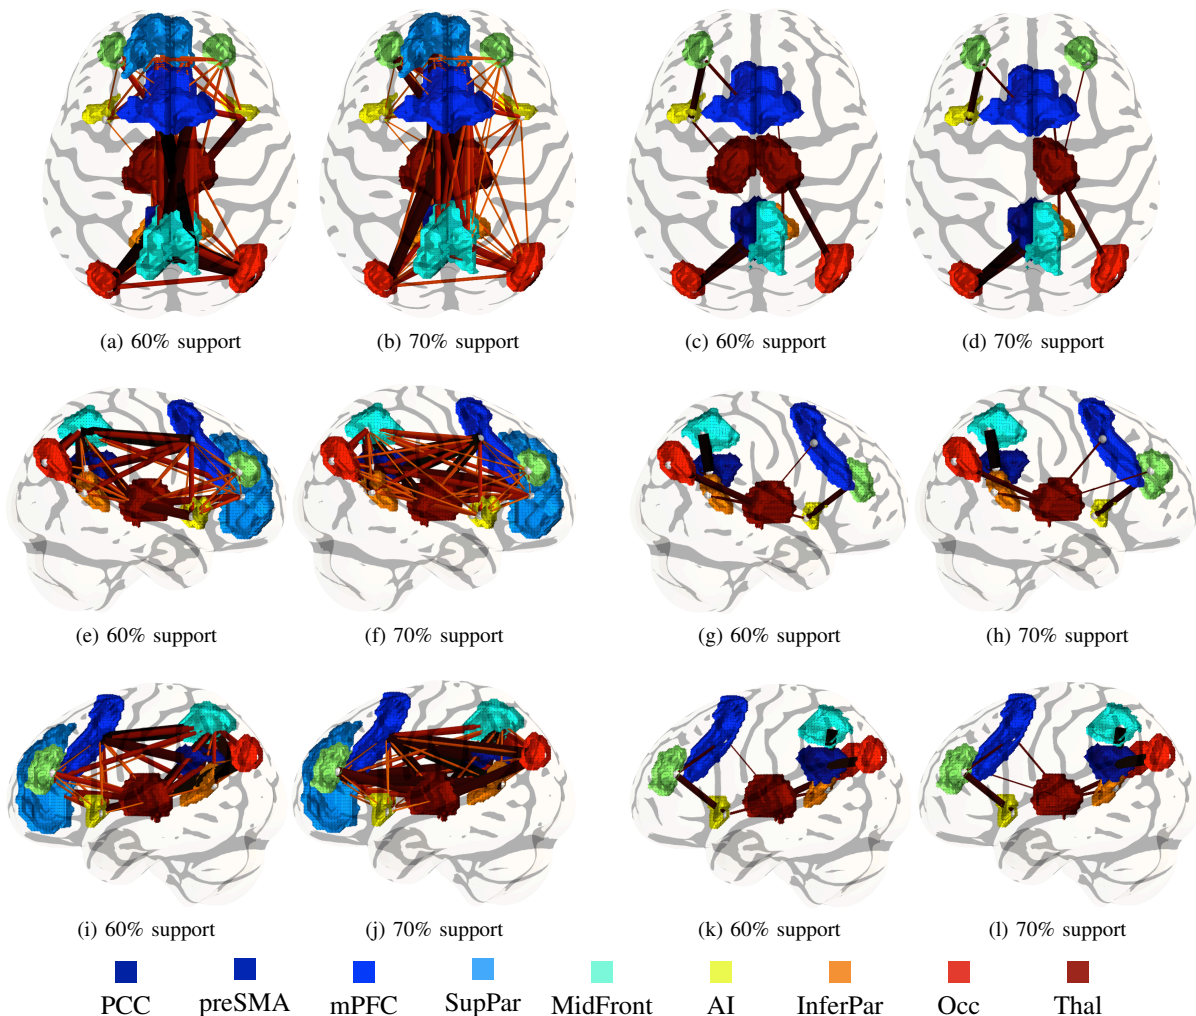PCC    preSMA    mPFC    SupPar    MidFront    AI    InferPar    Occ    Thal

Fig. 6. Connectivity associations between structure and function in the underlying network for 60% and 70% of connections included. The figure is divided in left and right part. The left part shows the overall structures linked to the function within the DMN and the SL. The thalamus, preSMA and superior parietal regions appear as hubs of the underlying networks. The right part of the figure is focused on structural connections linked to the functional connectivity between the left occipital lobe and the right pre-cuneus. The results highlight direct connections between pre-cuneus and occipital lobe as well as other pathways. Each column presents identification results with the same support. Each row shows the 3D brain model from top, right and left view, respectively. Each region is associated with a unique color. Each connection is represented as a 3D tube between a pair of regions with diameter proportional to the probability of the connection. Connections are thresholded with only 20% with the highest probability shown here.

statistics, bias-variance tradeoffs are necessary. Importantly, our model evaluation is based on an intrinsic metric suitable to quantify distance between symmetric positive definite matrices $\mathcal{S}ym_n^+$. This metric is invariant under affine transformations and inversion, and enables us to compare different parameterizations of functional connectivity, such as correlation, or partial correlations.

Describing resting-state fMRI with a Gaussian graphical model provides an underlying generative process [32], [42], [41], which can assist in gaining a deeper understanding of the mechanisms of functional connectivity, compared to purely descriptive models as [30]. We found empirically that enforcing the independence structure based on structural connectivity [1] improves the prediction of the model. From the high-dimensional statistical point of view, this procedure gives a better conditioned covariance matrix; in addition it is consistent with the simple picture that absence of fiber between regions should correspond to absence of functional connections.

We assume a multivariate linear relationship between anatomical and functional connectivity. In fact, the interaction matrix reflects signal transition probabilities between regions, which are considered as a linear function of the related anatomical connections. Neuroscientific findings that support this come from neuronal models that associate axonal diameter, myelination and fiber integrity to conduction delays and wiring cost. For example, metabolic cost is often assumed proportional to the total volume of fibers [65], indicating a multiple linear relation between function and structure.

To find associations between structure and function we use sparse models. We rely on the randomized LASSO rather than the LASSO, as it improves upon the LASSO in a number of ways. Firstly, the randomized LASSO gives a reliable detection of important features with less stringent conditions than the LASSO [50]. Secondly, it assigns a selection score for each structural connection that reflects the probability

| (a) OccL-PCCR | (b) OccL-PCCR | (c) OccL-PCCR |

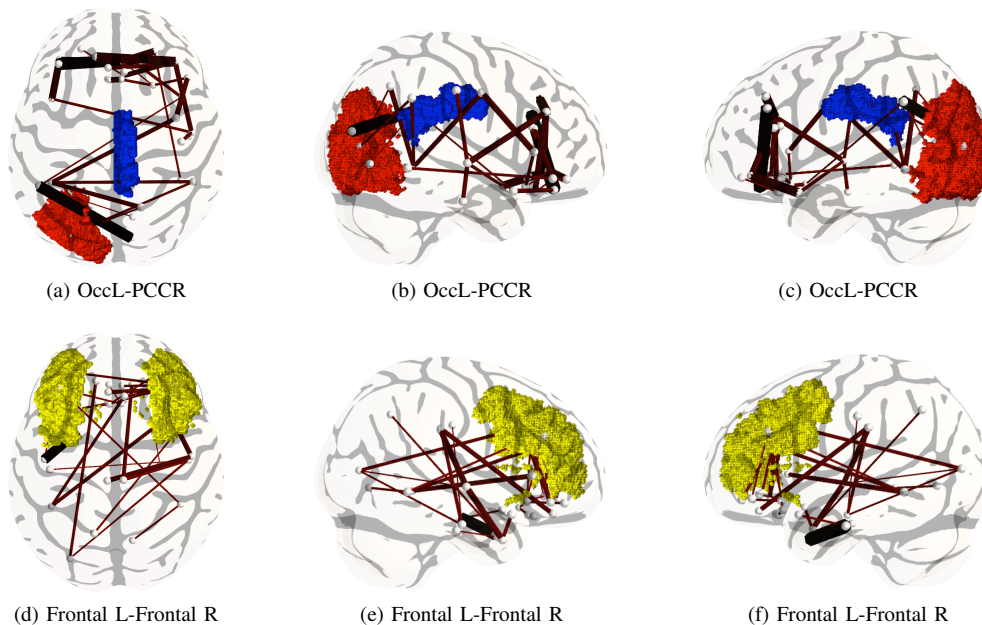| (d) Frontal L-Frontal R | (e) Frontal L-Frontal R | (f) Frontal L-Frontal R |

Fig. 7. Connection identification results for the atlas-based ROIs. Top row shows the structural connections with the highest 2% z-score values that predict the functional connectin between the left Occipital lobe and the right posterior part of the cingulate gyrus (PCC). At the bottom row, we show the structural connections with the 2% highest z-scores related to the functional connection between the left and right frontal lobes. The value of z-score threshold is around 1.6 for both examples. Both examples are estimated with support 60%.

of it to be selected and can give a good control on false detections. This is a significant step towards interpreting the results compared to the standard LASSO estimate that gives no indication on the statistical significance of the selected features. Another benefit of the randomized LASSO is that it decreases the dependence of the selected coefficients on the initial choice of regularization parameter, $\lambda$. Note that in our approach $\lambda$ is selected independently for each functional connection by cross-validation.

### B. Remaining challenges to neuroscientific interpretation

*a) Node ordering:* Predicting a brain-wide functional connectivity structure is a challenging problem as it boils down to predicting second-order statistics of the signal, such as correlations or partial-correlations, which cannot be well modeled via a linear model. For this reason, we use a parameterization based on multivariate autoregressive models, that can be linearly related to the observed functional signal, as in equations (2) and (3). Such a first-order parameterization requires to choose an ordering in the variables, *i.e.* the ROIs. Our functional-connectivity model shares similarities with structural equation models used for brain connectivity in fMRI [66]. In these models, the structure is stipulated and, as cycles pose challenges to the estimation, directed acyclic graphs are used, which amounts to choosing a node ordering. Choosing an ordering is not satisfactory, as the brain is best described by top-down and bottom-up connections forming feedback loops. However, we believe that it is a necessary implication of using node-centric linear models. These offer the strong benefit of simple interpretation. In practice, we use the approximate minimum degree ordering, that leads to sparser autoregressive models. Empirically, this choice yields a better prediction performance than model averaging on random orderings. Our results indicate that random ordering eventually converges to similar results obtained from AMD ordering but it requires a larger sample size.

*b) Non-stationarity in rs-fMRI:* In this work, the precision matrix, inverse of covariance is used to describe the statistical characteristics of the fMRI signal under the assumption of stationarity. In other words, this captures the steady-state contribution to the connectivity. There is evidence that in addition to this contribution, neural networks fluctuate from one brain state to another. Hence, future work should study non-stationarity in more detail. However, anatomical connectivity reflects a wiring, static over the duration of an experiment, and is thus most likely linked to steady-state functional connectivity. A promising future direction is the combined acquisition of rs-fMRI and electrophysiological data EEG that have temporal resolution of milliseconds and can reliably detect profound changes in brain states, such as sleep, task related and so on. Nevertheless, our proposed method highlights relationships between structure and function that could potentially change due to brain plasticity and development. In this perspective of long-term changes in brain networks that affect both structure and function significantly, modeling of short-term non-stationarity becomes less critical. Future studies should investigate changes in this anatomo-functional relationship in disease and relate them to existing knowledge and neuroscientific evidence.

*c) Confounding spatial variability:* Inter-subject differences in connectivity, structural or functional, may also arise from differences in the spatial location and the extent of the corresponding brain regions [67], [41]. In this paper, we investigated either subject-specific ROIs defined on purely anatomical criteria, or functional ROIs from the literature. Further work to improve the neuroscientific interpretability of

our method calls for better definitions of ROIs, to reject this confounding source of variability. For instance, our approach could benefit from connectivity-based segmentation, to define the ROIs exploiting patterns of cortico-cortical connections of cortical areas [68], [69]. While most studies so far assume that parcellation based solely on anatomical information reflects functional specialization [22], [30], [1], [64], [70], functional areas extracted from fMRI are not completely defined by folding patterns or cytoarchitectural boundaries [71]. In addition, they appear in different locations and varying size across subjects [72].

*d) From anatomy to function:* Using the anatomical constraints enables a maximum-likelihood estimate of functional connectivity. Estimation in a Bayesian framework with continuous anatomical support would yield a maximum-a-posteriori estimate where the functional connectivity would be biased by the anatomical connectivity. This bias would not be transferable to the new data, and thus would hinder generalization of the anatomo-functional model. Furthermore, joint estimation by iteratively refining both the anatomical connectivity matrix and the functional connectivity matrix would be interesting, drawing from the linear link that we model between a parameterization of functional connectivity and anatomical connectivity. However, the focus of this paper is inter-subject predictive modeling, for which it is not clear that joint estimation would be beneficial. It would introduce a coupling between the two modalities due to the estimation procedure that would add up to the natural coupling that we are trying to measure.

## VII. CONCLUSION

We have presented a probabilistic framework that explains inter-subject variability of functional connectivity based on structural connectivity, associating each functional connection with structural connections. This association is data-driven and can reveal interactions between function and structure that remain unknown. Our proposed approach outperforms a number of possible variants in terms of prediction error on two different choices of ROIs. Although, several prior studies measured localized differences between structure and function, they are based on prior hypothesis and in isolation of the rest of the brain [29], [15]. To our knowledge, this is the first attempt to develop a full-brain predictive model.

Our contribution provides a systematic method to highlight networks and provide neuroscientists and clinicians with an exploratory tool with interpretable and reliable parameter estimates. It is a step forward in learning multimodal biomarkers: extracting underlying anatomo-functional networks that support the differences between diseased and healthy populations. Predictive models open the door to better data-driven tracking of brain diseases progression. Indeed, there is a plethora of possible associations and using domain knowledge is challenging as the imaging modalities provide only very indirect measures of connectivity. In addition, we lack ground truth neuroscientific knowledge. In these settings, data-driven approaches based on statistical learning can fill the gaps.

## REFERENCES

[1] F. Deligianni, G. Varoquaux, B. Thirion, E. Robinson, D. J. Sharp, A. D. Edwards, and D. Rueckert, "A probabilistic framework to infer brain functional connectivity from anatomical connections," *IPMI*, p. 296, 2011.

[2] K. Friston, "Functional and effective connectivity in neuroimaging: a synthesis," *Hum Brain Mapp*, vol. 2, p. 56, 1994.

[3] J. Cabral, E. Hugues, O. Sporns, and G. Deco, "Role of local network oscillations in resting-state functional connectivity," *NeuroImage*, vol. 57, p. 130, 2011.

[4] E. Bullmore and O. Sporns, "Complex brain networks: graph theoretical analysis of structural and functional systems," *Nat Rev Neurosci*, vol. 10, p. 186, 2009.

[5] O. Sporns, G. Tononi, and R. Kotter, "The human connectome: a structural description of the human brain," *PLoS Comput Biol*, vol. 1, p. e42, 2005.

[6] O. Sporns, G. Tononi, and G. Edelman, "Theoretical neuroanatomy: relating anatomical and functional connectivity in graphs and cortical connection matrices," *Cereb Cortex*, vol. 10, p. 127, 2000.

[7] F. Ferrarelli, M. Massimini, S. Sarasso, A. Casali, B. A. Riedner, G. Angelini, G. Tononi, and R. A. Pearce, "Breakdown in cortical effective connectivity during midazolam-induced loss of consciousness," *P Natl Acad Sci Usa*, vol. 107, p. 2681, 2010.

[8] G. Tononi, "An information integration theory of consciousness," *BMC neuroscience*, vol. 5, 2004.

[9] J. Richiardi, H. Eryilmaz, S. Schwartz, P. Vuilleumier, and D. Van De Ville, "Decoding brain states from fMRI connectivity graphs," *NeuroImage*, vol. 56, p. 616.

[10] J. Burns, "An evolutionary theory of schizophrenia: Cortical connectivity, metarepresentation, and the social brain," *Behavioral and Brain Sciences*, vol. 27, p. 831, 2004.

[11] A. Garrity, G. Pearlson, K. McKiernan, D. Lloyd, K. Kiehl, and V. Calhoun, "Aberrant" default mode" functional connectivity in schizophrenia," *Am J Psychia*, vol. 164, p. 450, 2007.

[12] K. Konrad and S. Eickhoff, "Is the ADHD brain wired differently? A review on structural and functional connectivity in attention deficit hyperactivity disorder," *Hum Brain Mapp*, vol. 31, p. 904, 2010.

[13] L. Uddin and V. Menon, "The anterior insula in autism: Under-connected and under-examined," *Neuroscience and Biobehavioral Reviews*, vol. 33, p. 1198, 2009.

[14] K. Wang, T. Jiang, M. Liang, L. Wang, L. Tian, X. Zhang, K. Li, and Z. Liu, "Discriminative analysis of early alzheimer's disease based on two intrinsically anti-correlated networks with resting-state fMRI," *MICCAI*, vol. 4191, pp. 340–347, 2006.

[15] D. Sharp, C. Beckmann, R. Greenwood, K. Kinnunen, V. Bonnelle, X. D. Boissezon, J. Powell, S. Counsell, M. Patel, and R. Leech, "Default mode network functional and structural connectivity after traumatic brain injury," *Brain*, vol. 134, p. 2233, 2011.

[16] V. Bonnelle, T. E. Ham, R. Leech, K. M. Kinnunen, M. A. Mehta, R. J. Greenwood, and D. J. Sharp, "Salience network integrity predicts default mode network function after traumatic brain injury," *P Natl Acad Sci Usa*, vol. 109, p. 4690, 2012.

[17] D. K. Jones, "Diffusion MRI: Theory, methods, and applications," p. 624, 2010.

[18] H. Johansen-Berg, T. Behrens, M. Robson, I. Drobnjak, M. Rushworth, J. Brady, S. Smith, D. Higham, and P. Matthews, "Changes in connectivity profiles define functionally distinct regions in human medial frontal cortex," *P Natl Acad Sci Usa*, vol. 101, p. 13335, 2004.

[19] T. Behrens, M. Woolrich, M. Jenkinson, H. Johansen-Berg, R. Nunes, S. Clare, P. Matthews, J. Brady, and S. Smith, "Characterization and propagation of uncertainty in diffusion-weighted MR imaging," *Magnet Reson Med*, vol. 50, p. 1077, 2003.

[20] S. M. Smith, K. L. Miller, S. Moeller, J. Xu, E. J. Auerbach, M. W. Woolrich, C. F. Beckmann, M. Jenkinson, J. Andersson, M. F. Glasser, D. C. V. Essen, D. A. Feinberg, E. S. Yacoub, and K. Ugurbil, "Temporally-independent functional modes of spontaneous brain activity," *P Natl Acad Sci Usa*, vol. 109, p. 3131, 2012.

[21] C. F. Beckmann, M. DeLuca, J. T. Devlin, and S. M. Smith, "Investigations into resting-state connectivity using independent component analysis." *Philos Trans R Soc Lond B Biol Sci*, vol. 360, p. 1001, 2005.

[22] P. Hagmann, L. Cammoun, X. Gigandet, R. Meuli, C. J. Honey, V. J. Wedeen, and O. Sporns, "Mapping the structural core of human cerebral cortex," *PLoS Biol*, vol. 6, p. e159, 2008.

[23] Y. Iturria-Medina and E. Canales-Rodriguez, "Characterizing brain anatomical connections using diffusion weighted MRI and graph theory," *NeuroImage*, 2007.

[24] E. Robinson, A. Hammers, A. Ericsson, A. Edwards, and D. Rueckert, "Identifying population differences in whole-brain structural networks: a machine learning approach," *NeuroImage*, vol. 50, p. 910, 2010.

[25] E. Robinson, M. Valstar, and A. Hammers, "Multivariate statistical analysis of whole brain structural networks obtained using probabilistic tractography," *MICCAI*, vol. 5241, p. 486, 2008.

[26] D. Zhang and M. E. Raichle, "Disease and the brain's dark energy," *Nat Rev Neurol*, vol. 6, p. 15, 2010.

[27] M. D. Greicius, K. Supekar, V. Menon, and R. F. Dougherty, "Resting-state functional connectivity reflects structural connectivity in the default mode network," *Cerebral Cortex*, vol. 19, p. 72, 2009.

[28] M. van den Heuvel, R. Mandl, J. Luijges, and H. H. Pol, "Microstructural organization of the cingulum tract and the level of default mode functional connectivity," *J Neurosci*, vol. 28, p. 10844, 2008.

[29] J. Damoiseaux and M. Greicius, "Greater than the sum of its parts: a review of studies combining structural connectivity and resting-state functional connectivity," *Brain Struct Funct*, vol. 213, p. 525, 2009.

[30] C. Honey, O. Sporns, L. Cammoun, X. Gigandet, J. Thiran, R. Meuli, and P. Hagmann, "Predicting human resting-state functional connectivity from structural connectivity," *P Natl Acad Sci Usa*, vol. 106, p. 2035, 2009.

[31] A. Venkataraman, Y. Rathi, M. Kubicki, C. Westin, and P. Golland, "Joint modeling of anatomical and functional connectivity for population studies," *IEEE Trans. Med. Imag.*, vol. 31, p. 164, 2012.

[32] G. Varoquaux, A. Gramfort, J. B. Poline, and B. Thirion, "Brain covariance selection: better individual functional connectivity models using population prior," *NIPS*, 2010.

[33] V. Vapnik, "An overview of statistical learning theory," *IEEE Trans Neural Networks*, vol. 10, p. 988, 1999.

[34] A. Barron, L. Birgé, and P. Massart, "Risk bounds for model selection via penalization," *Probab. Theory Related Fields*, vol. 113, p. 301, 1999.

[35] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J Royal Statist Soc B*, vol. 58, p. 267, 1996.

[36] S. Mukherjee, P. Niyogi, T. Poggio, and R. Rifkin, "Learning theory: stability is sufficient for generalization and necessary and sufficient for consistency of empirical risk minimization," *Adv Comput Math*, vol. 25, p. 161, 2006.

[37] F. Bach and E. Moulines, "Non-asymptotic analysis of stochastic approximation algorithms for machine learning," in *NIPS*, 2011, p. 451.

[38] S. Smith, M. Jenkinson, H. Johansen-Berg, D. Rueckert, T. Nichols, C. Mackay, K. Watkins, O. Ciccarelli, M. Cader, P. Matthews *et al.*, "Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data," *Neuroimage*, vol. 31, p. 1487, 2006.

[39] D. Jones, "Challenges and limitations of quantifying brain connectivity in vivo with diffusion MRI," *Imaging*, vol. 2, p. 341, 2010.

[40] P. Fransson and G. Marrelec, "The precuneus/posterior cingulate cortex plays a pivotal role in the default mode network: Evidence from a partial correlation network analysis," *Neuroimage*, vol. 42, p. 1178, 2008.

[41] S. Smith, K. Miller, G. Salimi-Khorshidi, M. Webster, C. Beckmann, T. Nichols, J. Ramsey, and M. Woolrich, "Network modelling methods for fMRI," *Neuroimage*, vol. 54, p. 875, 2011.

[42] G. Varoquaux, A. Gramfort, J. B. Poline, and B. Thirion, "Markov models for fMRI correlation structure: is brain functional connectivity small world, or decomposable into networks?" *Journal of Physiology - Paris*, vol. 106, pp. 212–221, 2012.

[43] A. Dempster, "Covariance selection," *Biometrics*, vol. 28, p. 157, 1972.

[44] S. Lauritzen, *Graphical models*. Oxford University Press, USA, 1996.

[45] P. Amestoy, T. Davis, I. Duff *et al.*, "An approximate minimum degree ordering algorithm," *SIAM J Matrix Anal A*, vol. 17, p. 886, 1996.

[46] D. Donoho, "For most large underdetermined systems of linear equations the minimal $\ell_1$-norm solution is also the sparsest solution," *Comm. Pure Appl. Math.*, vol. 59, p. 797, 2006.

[47] W. Förstner and B. Moonen, "A metric for covariance matrices," *Qua vadis geodesia*, p. 113, 1999.

[48] X. Pennec, P. Fillard, and N. Ayache, "A Riemannian framework for tensor computing," *Int J Comp Vision*, vol. 66, p. 41, 2006.

[49] C. Lenglet, M. Rousson, R. Deriche, and O. Faugeras, "Statistics on the manifold of multivariate normal distributions: Theory and application to diffusion tensor MRI processing," *J Math Imag Vis*, vol. 25, p. 423, 2006.

[50] N. Meinshausen and P. Buhlmann, "Stability selection," *J Roy Statist Soc B*, vol. 27, p. 417, 2010.

[51] S. Smith, M. Jenkinson, M. Woolrich, C. Beckmann, T. Behrens, H. Johansen-Berg, P. Bannister, M. D. Luca, I. Drobnjak, D. Flitney, R. Niazy, J. Saunders, J. Vickers, Y. Zhang, N. D. Stefano, J. Brady, and P. Matthews, "Advances in functional and structural MR image analysis and implementation as FSL," *NeuroImage*, vol. 23, p. 208, 2004.

[52] N. J. Tustison, B. B. Avants, P. A. Cook, Y. Zheng, A. Egan, P. A. Yushkevich, and J. C. Gee, "N4itk: Improved n3 bias correction," *IEEE Trans. Med. Imag.*, vol. 29, p. 1310, 2010.

[53] D. Rueckert, L. Sonoda, C. Hayes, and D. Hill, "Non-rigid registration using free-form deformations: application to breast MR images," *IEEE Trans. Med. Imag.*, vol. 18, p. 712, 1999.

[54] R. Heckemann, J. Hajnal, P. Aljabar, D. Rueckert, and A. Hammers, "Automatic anatomical brain MRI segmentation combining label propagation and decision fusion," *NeuroImage*, vol. 33, p. 115, 2006.

[55] Y. Zhang, M. Brady, and S. Smith, "Segmentation of brain MR images through a hidden markov random field model and the expectation-maximization algorithm," *IEEE Trans. Med. Imag.*, vol. 20, p. 45, 2001.

[56] X. Artaechevarria, A. Barrutia, and C. Ortiz-de Solorzano, "Combination strategies in multi-atlas image segmentation: Application to brain MR data," *Adv NIPS*, vol. 28, pp. 1266–1277, 2009.

[57] C. Ledig, R. Wolz, P. Aljabar, J. Lötjönen, R. A. Heckemann, A. Hammers, and D. Rueckert, "Multi-class brain segmentation using atlas propagation and EM-based refinement," in *ISBI*, 2012, p. 896.

[58] G. Varoquaux, S. Sadaghiani, P. Pinel, A. Kleinschmidt, J. B. Poline, and B. Thirion, "A group model for stable multi-subject ICA on fMRI datasets," *NeuroImage*, vol. 51, p. 288, 2010.

[59] T. Behrens, M. Woolrich, M. Jenkinson, H. Johansen-Berg, R. Nunes, S. Clare, P. Matthews, J. Brady, and S. Smith, "Characterization and propagation of uncertainty in diffusion-weighted MR imaging," *Magnet Reson Med*, vol. 50, p. 1077, 2003.

[60] O. Ledoit and M. Wolf, "A well-conditioned estimator for large-dimensional covariance matrices," *J Multivar Anal*, vol. 88, p. 365, 2004.

[61] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine Learning in Python ," *J Mach Learn Res*, vol. 12, p. 2825, 2011.

[62] T. Jiang, Y. Liu, F. Shi, N. Shu, B. Liu, J. Jiang, and Y. Zhou, "Multimodal magnetic resonance imaging for brain disorders: Advances and perspectives," *Brain Imaging and Behavior*, vol. 2, p. 249, 2008.

[63] M. Ewers, R. A. Sperling, W. E. Klunk, M. W. Weiner, and H. Hampel, "Neuroimaging markers for the prediction and early diagnosis of alzheimer's disease dementia," *Trends in Neurosciences*, vol. 34, p. 430, 2011.

[64] F. Deligianni, E. Robinson, C. Beckmann, D. Sharp, A. Edwards, and D. Rueckert, "Inference of functional connectivity from structural brain connectivity," *ISBI*, p. 1113, 2010.

[65] D. Chklovskii, T. Schikorski, and C. Stevens, "Wiring optimization in cortical circuits," *Neuron*, vol. 34, pp. 341–7, Jan 2002.

[66] A. McIntosh and F. Gonzalez-Lima, "Structural equation modeling and its application to network analysis in functional brain imaging," *Hum Brain Map*, p. 2, 1994.

[67] A. Zalesky, A. Fornito, and E. Bullmore, "Network-based statistic: Identifying differences in brain networks," *NeuroImage*, vol. 53, p. 1197, 2010.

[68] R. Passingham, K. Stephan, and R. Kotter, "The anatomical basis of functional localization in the cortex," *Nat Rev Neurosci*, vol. 3, p. 606, 2002.

[69] S. Jbabdi, M. W. Woolrich, and T. E. J. Behrens, "Multiple-subjects connectivity-based parcellation using hierarchical Dirichlet process mixture models," *NeuroImage*, vol. 44, p. 373, 2009.

[70] F. Deligianni, E. Robinson, C. Beckmann, D. Sharp, D. Edwards, and D. Rueckert, "Inference of functional connectivity from direct and indirect structural brain connections," *ISBI*, 2011.

[71] B. Yeo, F. Krienen, J. Sepulcre, M. Sabuncu *et al.*, "The organization of the human cerebral cortex estimated by intrinsic functional connectivity," *J Neurophysio*, vol. 106, p. 1125, 2011.

[72] G. Varoquaux, A. Gramfort, F. Pedregosa, V. Michel, and B. Thirion, "Multi-subject dictionary learning to segment an atlas of brain spontaneous activity," in *IPMI*, 2011, p. 562.

[73] P. Ramachandran and G. Varoquaux, "Mayavi: 3d visualization of scientific data," *Computing in Science Engineering*, vol. 13, p. 40, 2011.