Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

Dynamics of functional network organization through graph mixture learning

Ilaria Ricchi^{a,b,c,1,*}, Anjali Tarun^{a,b,1}, Hermina Petric Maretic^c, Pascal Frossard^c, Dimitri Van De Ville^{a,b}

^a Institute of Bioengineering, École Polytechnique Fédérale de Lausanne (EPFL), Geneva, 1202, Switzerland

^b Department of Radiology and Medical Informatics, University of Geneva, Geneva, 1202, Switzerland

^c School of Engineering, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, 1015, Switzerland

ARTICLE INFO

Keywords: Dynamic functional connectivity Structure and function Task fMRI Resting-state (Meta)states

ABSTRACT

Understanding the organizational principles of human brain activity at the systems level remains a major challenge in network neuroscience. Here, we introduce a fully data-driven approach based on graph learning to extract meaningful repeating network patterns from regionally-averaged timecourses. We use the Graph Laplacian Mixture Model (GLMM), a generative model that treats functional data as a collection of signals expressed on multiple underlying graphs. By exploiting covariance between activity of brain regions, these graphs can be learned without resorting to structural information. To validate the proposed technique, we first apply it to task fMRI with a known experimental paradigm. The probability of each graph to occur at each time-point is found to be consistent with the task timing, while the spatial patterns associated to each epoch of the task are in line with previously established activation patterns using classical regression analysis. We further on apply the technique to resting state data, which leads to extracted graphs that correspond to well-known brain functional activation patterns. The GLMM allows to learn graphs entirely from the functional activity that, in practice, turn out to reveal high degrees of similarity to the structural connectome. The Default Mode Network (DMN) is always captured by the algorithm in the different tasks and resting state data. Therefore, we compare the states corresponding to this network within themselves and with structure.

Overall, this method allows us to infer relevant functional brain networks without the need of structural connectome information. Moreover, we overcome the limitations of windowing the time sequences by feeding the GLMM with the whole functional signal and neglecting the focus on sub-portions of the signals.

1. Introduction

Functional magnetic resonance imaging (fMRI) is a unique tool to probe the functional architecture of the human brain. Specifically, spontaneous fluctuations of blood-oxygenation-level dependent (BOLD) signals have shown to be synchronised between brain regions during resting state (RS) despite the absence of task or external stimuli (Biswal et al., 1995; Smith et al., 2009; 2011). A repertoire of functional networks have been identified in healthy (Beckmann et al., 2005; Shirer et al., 2012; Thomas Yeo et al., 2011) and clinical populations (Liao et al., 2014; Tagliazucchi et al., 2010; Zöller et al., 2019). These functional networks include both sensory regions and higher-level cognitive ones, such as the default-mode network (DMN), which generally shows reduced activity during an externally-oriented task (Greicius et al., 2003), and becomes more engaged during internal mentation (Andrews-Hanna, 2012; McCormick and Telzer, 2018).

The interest in understanding the intrinsic functional organization of the human brain has motivated many new methods to deal with RS. Conventional functional connectivity (FC) constitutes a single whole-brain functional connectome by measuring Pearson correlation between pairs of regionally-averaged time courses (Bullmore and Sporns, 2009). FC is thus computed using the entire RS scan, intrinsically assuming a "stationary" relationship, but several "dynamic" extensions that acknowledge temporal fluctuations of FC have been proposed (Bolton et al., 2020; Chang and Glover, 2010; Hutchison et al., 2013; Preti et al., 2017). One popular approach is the sliding-window technique (Elton and Gao, 2015b; Kucyi and Davis, 2014; Madhyastha and Grabowski, 2014), where time courses are segmented into temporal windows so that a time-dependent FC matrix can be obtained. Then, further analysis by graph metrics (Betzel et al., 2016; Sizemore and Bassett, 2018), or di-

* Corresponding author.

¹ These authors contributed equally.

https://doi.org/10.1016/j.neuroimage.2022.119037.

Received 25 May 2021; Received in revised form 29 December 2021; Accepted 23 February 2022 Available online 25 February 2022. 1053-8119/© 2022 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/)





E-mail address: ilaria.ricchi@epfl.ch (I. Ricchi).

mensionality reduction methods such as singular value decomposition (SVD) (Leonardi et al., 2013), k-means clustering (Allen et al., 2014), or hierarchical clustering (Yang et al., 2014), can be applied to extract the most relevant brain patterns. The timescale of these dynamically-occurring FC patterns is limited by the temporal window length, and their spatial specificity by the nature of FC that characterizes interacting activity between brain regions. For this reason, Bayesian and probabilistic methods represent good examples of extracting repeating activity patterns from whole-brain data (Smith et al., 2011; Varoquaux et al., 2010; Vidaurre et al., 2017). In neuroimaging, these approaches have been explored with the aim of applying unsupervised learning to estimate different hidden temporal states; e.g., Hidden Markov Models (HMM) (Eavani et al., 2013) to perform an activation network analysis (Taghia et al., 2018; Vidaurre et al., 2018; 2018; 2019; 2016).

While dynamic FC patterns are generated by looking at fluctuations in second-order statistics (i.e., correlation), alternatively, instantaneous activity-based brain states can access shorter timescales and directly explain the empirical BOLD time courses.

In this work, we propose a new framework to estimate multiple functional states, using a recently introduced Bayesian-based generative model that is the Graph Laplacian Mixture Model (GLMM), suitable for mixed signals expressed on different graphs (Maretic et al., 2018; Petric Maretic and Frossard, 2020). Several graph learning techniques (Dong et al., 2019; Mateos et al., 2019) have been recently proposed in order to infer a meaningful graph structure from data. Due to the instantaneous nature of these methods, the choice of crucial parameters such as window lengths (Leonardi and Van De Ville, 2015) is eliminated.

General approaches for learning graphs to represent data structure started with Dong et al. (2016) who focuses on graph Laplacian matrix inference, which enforces data smoothness on the inferred graph (Kalofolias, 2016). Using a more efficient solution, dictionary-based methods assume that signals can be modeled as a sparse combination of localised graph dictionary atoms (Petric Maretic et al., 2017; Thanou et al., 2017). Multiple graph inference methods include works on time-varying graphs (Kalofolias et al., 2017; Yamada et al., 2019), where temporal signals reflect the structure of several graphs, each of which is active in a pre-defined time period. The recent work of Gan et al. (2019) imposes that the sparsity pattern on all inferred graphs should be similar through a Bayesian prior. All of these works focus either on inferring only one graph or already have a predefined set of signals for each graph that needs to be inferred. Such a strong assumption is not realistic in RS fMRI data, where a wide array of dynamic functional networks are known to occur. By appealing to GLMM (Petric Maretic and Frossard, 2020), we simultaneously learn multiple graph structures (represented by graph Laplacian matrices) and, separately, nodal activation patterns. Additionally, along with these two elements, brain signals are decomposed into clusters with a certain probability. These three elements form together what we call here state.

Interestingly, when we validate this approach with task fMRI data, we observe the probabilities reflecting the dynamic of epochs of the experimental paradigms, even though these are completely unknown to the method – which is anyway the case for RS fMRI data. Different from other HMM frameworks, we have neither a Markovian assumption nor any constraint on the temporal axis. By being more flexible, we still are able to capture a meaningful dynamic.

We also show that the extracted states consist of brain areas that are consistent with previously observed regions implicated in the corresponding task. We then apply the GLMM to RS data and obtain the most prevalent states governing spontaneously interacting brain areas.

Overall, this approach allows us to revisit the relationship between brain function and underlying structure, which is one of the fundamental questions in neuroscience (Atasoy et al., 2016; Betzel et al., 2016; Goni et al., 2014; Gu et al., 2015; Hutchison et al., 2013; Preti and Van De Ville, 2019; Tarun et al., 2020). The traditional approaches to estimate structural connectivity from FC, do not account for fluctuations of brain activity and thus how the structure-function link is exploited differently over time. Therefore, here we show that the estimated graph Laplacian matrices reveal indicative similarities with SC. Moreover, we observe that the most notable brain pattern that consistently is part of states in all tasks, as well as in RS, is the default mode network (DMN). Different studies have tried to capture how the DMN differs in different neurological conditions (Nair et al., 2020; Ren et al., 2020; Shine et al., 2015; Starck et al., 2013). That is why we finally focus our attention to the similarities and differences of the DMN patterns estimated during the rest epochs across all task paradigms and during resting state. The question consequently arises as to how the differences in connectivity structure that give rise to various DMN graphs are related to the brain's underlying anatomical structure.

Finally, we leverage on a dissimilarity score to index the level of cognitive association: the more dissimilar the score is, the more it indicates that that network belongs to a higher cognitive level task, suggesting an appropriate method of comparison with respect to classical Pearson correlations.

2. Materials and methods

2.1. Data and preprocessing

We use MRI data from the Human Connectome Project (HCP). The MRI acquisition protocols have been extensively described and presented elsewhere (Glasser et al., 2013). In particular, we used 50 subjects (see Supplementary Appendix A.6 for IDs) consisting of 4 sessions of RS scans (1200 volumes each, a total of 4800 frames), and 2 sessions each of task fMRI data (i.e., working memory, relational memory, social, language, emotion, and motor tasks). Functional volumes underwent the standard pre-processing steps (Van Dijk et al., 2010). All functional images were first realigned to the mean functional volume for each participant. The realigned volumes were registered to the structural T1 data using rigid-body registration (SPM12, https://www.fil.ion.ucl.ac.uk), and were detrended (i.e., constant, linear, quadratic) to remove signal drifts. Then, the images were smoothed using a Gaussian kernel with FWHM equal to 6mm. Finally, we used the Automated Anatomical Labeling (AAL, 90 regions) atlas that was re-sliced to native fMRI space to parcellate fMRI volumes and compute regionally averaged fMRI BOLD signals. The structural and diffusion-weighted MRI data of each subject were downloaded from the HCP and were processed using MRtrix3 (http://www.mrtrix.org/). We used single shell (b=3000) multi-tissue to estimate the response function, while fiber orientation distribution functions were computed using constrained spherical deconvolution of order 8. Tractogram generation was performed using deterministic tractography with about 2×10^7 output streamlines and was seeded from the white matter. Fiber density was used as metric to define the individual SCs and were computed by dividing the total number of fibers that connect each of the pairwise regions of the AAL atlas with both the mean fiber length and the average of the sizes of the two regions considered. The normalization with respect to the region sizes is done to ensure that the strength of the connection is not biased towards the size of the ROIs. The final SC matrix was obtained by averaging all SC matrices of all subjects.

2.2. Graph Laplacian mixture model

The GLMM is a generative model assuming that the observations belong to different types of signals with different underlying graph structures (Petric Maretic and Frossard, 2020). The graphs are unknown and are modeled by their graph Laplacian matrix; i.e., \boldmath $L = \boldmath D - \boldmath A$, where \boldmath D is a diagonal matrix of node degrees, and \boldmath A is the weighted adjacency matrix. \unboldmath math

The model estimation problem wants to fit the observed data to recover signal clusters, as well as the associated activation patterns and graph structures. Specifically, observed fMRI signals are grouped into clusters defined by different unknown brain activation patterns. The model recovers these clusters, characterized by graph Laplacians together with the probability of each cluster to occur. These graph Laplacians bring information on the brain networks activation patterns (estimating means), functional brain connectivity structure (Laplacians), and dynamics (probability of occurrence as a function of time).

Formally, each of the *M* observed signals $x_m \in \mathbb{R}^N$ belongs to exactly one cluster *k* represented by the graph Laplacian $L_k \in \mathbb{R}^{N \times N}$ and mean $\mu_k \in \mathbb{R}^N$. A binary latent variable $z_m \in \mathbb{R}^K$ has exactly one non-zero value, which denotes the cluster *k* that x_m belongs to. A probability α_k defines a prior probability distribution of x_m belonging to cluster *k*, namely $p(z_{m,k} = 1) = \alpha_k, \forall m$.

Finally, the graph Laplacian $L_k \in \mathbb{R}^{N \times N}$ models smooth changes in signals on the corresponding graph. Large edge weight values in L_k thus capture pairs of vertices that change their values in similar ways. These connections can be seen as partial correlations between two vertices in a certain state (Dempster, 1972).

Under these assumptions, signals in each cluster *k* follow a Gaussian distribution determined through the graph Laplacian¹ $x \sim \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{L}_k^{-1})$ (Dong et al., 2016): with $L_k = D_k - W_k$.

$$p(\boldsymbol{x}_m | \boldsymbol{z}_{m,k} = 1) = p(\boldsymbol{x}_m | \boldsymbol{\mu}_k, \boldsymbol{L}_k) = \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{L}_k^{-1})$$
(1)

Recall that the cluster of each signal is a priori unknown. Marginalising over latent variables z denoting which cluster the signals belongs to, we have:

$$p(\mathbf{x}_m) = \sum_{z_m} p(z_m) p(\mathbf{x}_m | z_m)$$
(2)

$$=\sum_{k=1}^{K} p(z_{m,k}=1)p(\mathbf{x}_m|z_{m,k}=1)$$
(3)

$$=\sum_{k=1}^{K} \alpha_k \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{L}_k^{-1}), \tag{4}$$

s.t.
$$L_k \in \mathcal{L}, \forall k$$
 (5)

$$\sum_{k=1}^{K} \alpha_k = 1,\tag{6}$$

$$\alpha_k > 0, \forall k \tag{7}$$

Here (5) ensures that all L_k 's are valid Laplacian matrices, $\mathcal{L} = \{L | L_{i,j} = L_{j,i} \leq 0, \forall i \neq j \& \sum_{j=1}^{N} L_{i,j} = 0, \forall i\}$. Eqs. (6) and (7) ensure that α defines a valid probability measure.

2.2.1. GLMM algorithm

Given *M* observed *N*-dimensional signals in the data matrix $X \in \mathbb{R}^{N \times M}$, we want to recover the parameters of our generative model (4). To do so, we will look at the maximum a posteriori (MAP) estimate for our parameters: probabilities $\alpha = \alpha_1, ..., \alpha_K$, means $\mu = \mu_1, ..., \mu_K$ and graph Laplacians $L = L_1, ..., L_K$. Namely, we assume the data has been sampled independently from the distribution ((4)) defined through the graph Laplacians. In addition, we take into account the constraints on the graph structure given in (5), as well as possible prior information on the graphs (such as sparsity), and we maximise over the a-posteriori distribution of our model:

$$\arg \max_{\alpha,\mu,L} \ln p(\alpha,\mu,L|X)$$
(8)

$$\propto \arg \max_{\alpha,\mu,L} \ln p(X|\alpha,\mu,L)p(L)$$
(9)

$$= \arg \max_{\boldsymbol{\alpha}, \boldsymbol{\mu}, \boldsymbol{L}} \ln \prod_{m=1}^{M} p(\boldsymbol{x}_{m} | \boldsymbol{\alpha}, \boldsymbol{\mu}, \boldsymbol{L}) p(\boldsymbol{L})$$
(10)

$$= \arg \max_{\boldsymbol{\alpha},\boldsymbol{\mu},\boldsymbol{L}} \ln \prod_{m=1}^{M} \sum_{k=1}^{K} \alpha_k \mathcal{N}(\boldsymbol{x}_m | \boldsymbol{\mu}_k, \boldsymbol{L}_k^{-1}) p(\boldsymbol{L}_k)$$
(11)

$$= \arg \max_{\boldsymbol{\alpha}, \boldsymbol{\mu}, \boldsymbol{L}} \sum_{m=1}^{M} \ln \sum_{k=1}^{K} \alpha_k \mathcal{N}(\boldsymbol{x}_m | \boldsymbol{\mu}_k, \boldsymbol{L}_k^{-1}) p(\boldsymbol{L}_k).$$
(12)

This problem does not have a closed form solution. It could be simplified through posterior probabilities $\gamma \in \mathbb{R}^{M \times K}$, with $\gamma_{m,k}$ modeling the probability that the signal x_m belongs to cluster k:

$$\gamma_{m,k} = p(z_{m,k} = 1 | \boldsymbol{x}_m, \boldsymbol{\mu}_k, \boldsymbol{L}_k)$$
(13)

$$=\frac{p(z_{m,k}=1)p(\mathbf{x}_m|z_{m,k}=1,\boldsymbol{\mu}_k,\boldsymbol{L}_k)}{\sum_{l=1}^{K}p(z_{m,l}=1)p(\mathbf{x}_m|z_{m,l}=1,\boldsymbol{\mu}_l,\boldsymbol{L}_l)}$$
(14)

$$=\frac{\alpha_k \mathcal{N}(\mathbf{x}_m | \boldsymbol{\mu}_k, \boldsymbol{L}_k^{-1})}{\sum_{l=1}^K \alpha_l \mathcal{N}(\mathbf{x}_m | \boldsymbol{\mu}_l, \boldsymbol{L}_l^{-1})}$$
(15)

The parameters α , μ and *L* can now be estimated iteratively using an expectation maximisation (EM) algorithm, in which the graph Laplacians *L* are estimated with a graph learning scheme.

It is worth noting that the graph learning step includes the graph priors imposed through $p(L_k)$, including the sparsity prior which helps control the sparsity/density of the solution and is directly enforced through the following regularisation term:

$$-\beta_{1,k} tr(\mathbf{1}^T log(\operatorname{diag}(\boldsymbol{L}_k))) + \beta_{2,k} \|\boldsymbol{L}_k\|_{F.off}^2.$$
(16)

The term diag(L_k) is a vector with the diagonal values (node degrees) from L_k , and $||L_k||_{F,off}^2$ is the Frobenius norm of the off-diagonal values in L_k . The graph learning problem naturally provides sparse solutions, while increasing $\beta_{1,k}$ strengthens graph connectivity, and increasing $\beta_{2,k}$ promotes density. For a detailed derivation and analysis of this part of the algorithm, we refer to the appendix.

2.3. Application to fMRI data

The general workflow of this work is summarized in Fig. 1. We build a data matrix X that contains the timecourses of all AAL 90 regions. The timecourses of all sessions and all subjects are then concatenated together to form a huge data matrix Y. This is then fed to the GLMM algorithm, which learns the several graph Laplacians (L_k), means (μ_k) for each graph and signal clustering probabilities (γ). The means and graph Laplacians will be representing the whole population, while the probabilities could be interpreted subject- specific because a probability is assigned to each time point. We then validate the performance of the proposed framework using task fMRI as reference ground truth, by demonstrating that the timing of the task paradigms are captured by the averaged γ -values across all subjects. In order to have a timescale comparable to the experimental timing, we average the probabilities γ) across subjects.

We can visualize the graph representation of each state by associating to each node a color that reflects its mean activation (encoded in the estimated μ) and each brain region is connected with a weighted adjacency matrix computed from the inferred Laplacians. The activation patterns are also associated to a probability of occurrence thanks to their clustering probabilities (γ). This allows us to associate each state to a task experimental phase.

The hyper-parameters of the model are optimized with a grid-search by splitting the data in half and running the GLMM in parallel in the two groups and matching exclusively the activation patterns (μ_k). More

 $^{^{1}}$ Note that L_{k}^{-1} here denotes a pseudo-inverse of the graph Laplacian L_{k}



Fig. 1. General workflow: from fMRI signals to the extraction of states. (A) Mean BOLD signals from four sessions of fMRI recordings are computed within each region of the AAL90 atlas and are concatenated together to form the subject data matrix X, whose dimensionality is NxM where M represents the timecourse and N corresponds to the number of regions. The final data (Y) is obtained by concatenating 50 HCP subjects. (B) Plate notation of the generative model for extracting functional states and their corresponding estimated graphs represented through the Laplacians. (C) The output of the algorithm are K-number of states, their corresponding graph Laplacians, and the probability that each state would occur at a particular time-point. In the above example, we show K = 3.

details are provided in Supplementary Appendix A.2. The optimization of the number of clusters (K) is also guided by the consistency of the obtained clustering probabilities with the task paradigm. We cannot expect K to correspond to the number of tasks, since we cannot discount the possibility for a state to correspond to more than one task epoch. For that reason, we evaluate the results by iteratively changing K according to the concordance of the estimated γ with respect to the experimental task paradigm. In doing so, we observe that the set of meaningful network patterns is the same regardless of the proposed number of clusters K. Furthermore, varying K in such a way results in better overall accuracy, suggesting that the number of optimal clusters might be different from the number estimated through more traditional methods. We denote this method as K- γ *itero-homogeneity*. See Supplementary Material for further information (Appendix A.2.1, Fig. A2).

On the other hand, RS data lacks a ground truth, so the number of clusters K has been chosen based on the optimized silhouette measure and the consensus clustering procedure Monti et al. (2003), which is a resampling-based method for optimal class discovery. K has eventually been varied according to the procedure mentioned above, in order to capture multiple networks. In practice, changing the optimal number of clusters does not seem to affect the final estimation, instead, it opens the possibility of inferring more or fewer networks.

2.4. Comparison of learned functional graphs to brain structure

The GLMM not only solves a clustering problem but additionally estimates a direct correlation matrix represented by the Laplacian and, differently from a Gaussian Mixture Model, the GLMM performs an implicit dimensionality reduction while leading to more interpretable results. An important benefit of the algorithm compared to other clustering methods is indeed the estimation of the *graph Laplacian*, directly estimating its inverse L_k^{-1} . Not only does this help in obtaining more meaningful brain patterns, but it also conveys much more information and details about our clusters, allowing a direct comparison of the functional connectivity matrix with the structural connectome.

Since the DMN is strongly emerging in the state of all datasets being analyzed, they have been compared to each other with canonical metrics (Pearson correlation) focusing mainly on the activation patterns (μ). The DMNs used for comparison are the ones extracted from both task and RS data. While in RS, we rely simply on the visualization of the activation patterns of the brain areas involved, in the task data we can verify the occurrence of the corresponding state dynamic looking at the γ probabilities and matching the corresponding rest epoch of the experiment (that are basically the rest periods between task blocks). This gives us an additional validation of the assessed brain network.

Additionally, we also compare connectivity matrices of each DMNstate with SC derived from DW-MRI by using a spectral euclidean difference. The comparison is done by first decomposing the weighted adjacency matrices (A) of each graph derived from the Laplacians (A = D - L) into their constituent eigenvalues and eigenvectors. After a normalization step of both the functional matrix A and the structural one, the eigenspectrum of A is transformed using the Procrustes algorithm (Goodall, 1991; Kendall, 1989) to match the ordering of the eigenvectors of SC. The output of the Procrustes transform is the rotation matrix that is used to retrieve the transformed (functional) eigenvalues. In particular, given the rotation matrix Z generated by the transformation, the rotated eigenvalues are computed as follows:

$$\hat{\Lambda} = ZEW, \tag{17}$$

where W is the weighted adjacency matrix of the functional network taken into consideration and E is the vector of its original eigenvalues.

The spectral Euclidean difference reflects the degree of similarity between the learned graphs in each task and the anatomy defined by the SC. Thus, each network will have a certain score of distance that indicates how that function couples with structure: the smaller the metric, the closer the estimated graph to the structure. Subsequently, the scores have been normalized and sorted to see how the multiple DMNs estimated from the different tasks and RS differ from the structure.

3. Results

3.1. Estimated network timecourses are consistent with the timing of task paradigms

The GLMM algorithm enables the recovery of three important elements for the whole pull of subjects: (1) means or centers of activity; (2) connectivity matrices in the form of graph Laplacians; (3) temporal activity profile of each network in terms of their likelihood to occur at each time point. The proposed framework extracts consistent patterns of brain activity for each of the considered tasks, along with a temporal profile that is strikingly in agreement with the experimental paradigm, even though the GLMM is not given any information on the timings. Fig. 2 displays the estimated timecourses for six selected tasks overlaid with task conditions, which are distinguished by the background colors. Each task has rest epochs in between the task blocks and these are consistently associated with corresponding Default Mode Networks A4. The choice of the number of states visualized is driven by the K- γ iterohomogeneity, as explained in the section above 2.3.

States 2 and 3 of the Language task capture moments when subjects undergo the Story and Math epochs, respectively, while State 1 corresponds to the resting epoch of the task. For the Social task, State 3 consistently captures the RS epoch. However, the method does not distinguish the conditions *Mental* and *Random* since the plotted likelihood (γ) of State 3 corresponds to both task epochs. On the other hand, State 1 captures the transition between the rest and task epochs. The same observation can be made in the other tasks, in particular, the motor, relational, and working memory. In each of these tasks, some states capture transitions between epochs and uniquely identify whether it is a transition between rest and task condition or a transition between two different task conditions; i.e., State 2 in Relational Memory captures transitions between the conditions Relation and Match, while State 4 captures transitions between rest and any of the two tasks. Another consistent state that is always present regardless the number of states estimated is the State 1 of the Motor task that seems to correspond to the tongue movement in the experimental paradigm.

3.2. GLMM captures activation patterns corresponding to each task and consistent meta-analytic characteristics

In order to better interpret the states and maximize the concordance of the estimated γ with respect to the experimental task paradigm, the optimal has been set to the number of task conditions. Figs. 3(A) and (B) display a representative set of GLMM results for the Language and Emotion tasks, respectively. Along with the γ probabilities, already shown before, GLMM also recovers a brain graph represented with nodes whose color corresponds to the means (μ) estimated. Active nodes (brain areas) represented in red correspond to more positive μ values. These values represent the mean activation of each node in terms of signals. The size of the nodes is proportional to their degree and the edges are the weighted edges estimated from the Laplacians inferred. While the means give information on which node is active the most, the graph structure on top of this gives the interactions between those activations, namely how connected two different activations are.

The states for RS epochs of both tasks reveal a spatial pattern that corresponds to the DMN network. Even though they have the same coactivation patterns, they differ in the way they are connected. The regions implicated in State 2 (encoding the condition *Math*) include the parietal areas (superior and inferior) and the frontal region (e.g., middle frontal gyrus, opercular part of the inferior frontal gyrus). These regions are well in-line with the established associated areas corresponding to numbers and calculations (Arsalidou et al., 2018). On the other hand, active areas in State 3 (condition *Story*) are the hippocampus, frontal, and the bilateral superior and anterior temporal cortex, consistent with previously observed regions implicated with story processing tasks (Barch et al., 2013).

Meanwhile, states 2 and 3 of the *Emotion* task correspond to the *Fear* and *Neutral* conditions, respectively, with strong differences in terms of the regions activated. In particular, *Fear* triggers activations in the bilateral central gyrus, superior occipital gyrus, and the parietal cortices. These regions cover the somatosensory cortex, which is responsible not only for processing sensory information from various parts of the body, but also for emotional processing, including generation of emotional states and emotion regulation (Bufalari et al., 2007; Dolan, 2002; Kropf et al., 2019). Meanwhile, the *Neutral* state shows activation of visual regions that are strongly related to shape (Anzai et al., 2007; Hegdé and Van Essen, 2000). Unexpectedly, we have found relatively low means in the amygdala, a well-known region that is typically activated in emotional matching paradigms (Hariri et al., 2000; Preckel et al., 2019).

Another relevant brain pattern is the one associated to the *Motor* task and in particular to the tongue movements. The first State is not only capturing perfectly the dynamic but also highlighting activations of the operculum, insula, thalamus, medial and posterior cingulate in line with literature (Corfield et al., 1999; Xiao et al., 2017).

3.3. Extracted states during resting state

After validating the outcome of the proposed framework when applied to task fMRI using the known experimental task paradigm as the ground truth, we apply the method to RS fMRI. The RS brain networks that GLMM is capable of capturing are consistent with literature (van den Heuvel and Hulshoff Pol, 2010). Fig. 4 (A) displays the estimated brain networks together with the average likelihood to occur across subjects with its standard deviation.

Fig. 4 (B) shows the probability (γ) of each state to occur at each time-point for one representative subject, giving a proxy of the state's activity profiles in the same order presented in (A). As expected, we observe the DMN to be highly occurring. We have found a state that clearly covers areas of the visual cortex, another state that shows regions corresponding to the auditory network and some portions of the frontoparietal cortex, and another one that contains the bilateral temporal cortices and the insula, which is analogous to the salience network. Additionally, we also have found activations in the left and right hemispheres to be separately clustered.

Hence, GLMM uses an existing pool of empirical fMRI data to approximate graphs that, interestingly, arise to be meaningful brain networks consistent with the literature (van den Heuvel and Hulshoff Pol, 2010).

3.4. Comparison of learned graphs reminiscent of the DMN

We have found that states corresponding to the rest epochs of the different tasks always bear striking similarity to the DMN, as expected. To understand the nature of the recovered DMN-related states, we compute Pearson correlation between the coactivation patterns between these DMN-related means (μ) for each task as well as RS, which have an important topological meaning, since each node represent a brain area.



Fig. 2. Estimated timecourses with respect to selected task paradigms. The γ values are plotted over the experimental paradigm of Language (A), Emotion (B), Social (C) and Relational Memory (D) tasks. The black signal corresponds to the probability of belonging to a specific *state*.

Fig. 4(C) displays the correlation between the extracted DMN-related states from all tasks and RS fMRI. The axes are sorted from lowest to highest total correlation per row. A separation between low-level sensory-motor (*e.g.*, language, motor, rest) *versus* high-level cognition (*e.g.* working memory, relational memory, social and emotion).

These co-activation patterns are all highly correlated but *Language*, *Relational Memory and Rest* DMNs present some negative correlations with the other tasks DMNs. The reader is pointed to figure A4 in Supplementary Material to visually assess these differences.

3.5. Comparison of learned functional DMN graphs with structure

Here we want to compare the DMN-related states to structure by simple visualization and applying some similarity or dissimilarity metrics. As mentioned above, these states are related to the rest epoch of the

different tasks datasets and the DMN-state extracted from RS. Fig. 5 (A) shows the group-averaged SC across all subjects considered in this work, and Fig. 5 (B) displays the weighted adjacency matrix computed as A = D - L from the estimated graph Laplacian matrix of the DMN network of the Resting State. The adjacency matrix is visually more sparse than the SC. Moreover, a direct comparison (i.e., Pearson correlation across connections) between the state matrices and SC reveals similarity values within the range of r = 0.48 - 0.63, as shown in Fig. 5(C). It is also noteworthy that, compared to the conventional FC-SC relationship, where FC is obtained by correlating inter-regional BOLD timecourses using either Pearson correlation or partial correlation, the correlation with SC is much higher for GLMM-based state matrices. It has to be noticed that they are not statistically different within methods, namely the correlation of each DMN with SC is similar in all the methods.

By computing the dissimilarity spectral scores for each DMN and sorting them in ascending order, it emerges that there is a trend of "externally directed" or less introspective tasks. This score is normalized to enhance the visualization and it is only representative of the gradient, not of the significant difference.

Low cognitive level tasks, such as Motor and Resting State and Language, are expected to have low scores of dissimilarity with structure, hence are more similar to structure. Conversely, more introspective and "internally directed" tasks, such as Relational Memory, Social and Emotion, differ more from structure. Fig. 5(D) wants to highlight the gradient and order of the DMN under consideration to assess a potential decoupling index with the aim of creating a framework for potential application in meta-analyses (Preti and Van De Ville, 2019) (see Supplementary Materials Appendix A.5).

Fig. 5 can be complemented with the visualization of the DMNs in Appendix A.4 to advance reasonable conclusions. In fact, even though the identified brain pattern in the rest epochs across all tasks contains the key regions of the DMN, the corresponding connectivity structure of the state matrix varies, suggesting task-specific mechanisms in the way brain areas connect.



Fig. 3. States corresponding to each epoch of the (A) Language task and (B) Emotional task paradigms. Each node corresponds to each brain region in the AAL atlas. The colors denote the values of the extracted clusters (means μ), the edges denote the connectivity derived from the estimated graph Laplacians, and the size the nodes indicate the degree of each node's connections. The rest epochs of the two tasks show regions of the default mode network.

3.6. Extracted states during resting state

After validating the outcome of the proposed framework when applied to task fMRI using the known experimental task paradigm as the ground truth, we apply the method to RS fMRI. The RS brain networks that GLMM is capable of capturing are consistent with literature (van den Heuvel and Hulshoff Pol, 2010). Fig. 4 (A) displays the estimated brain networks together with the average likelihood to occur across subjects with its standard deviation.

Fig. 4 (B) shows the probability (γ) of each state to occur at each time-point for one representative subject, giving a proxy of the state's activity profiles in the same order presented in (A). As expected, we observe the DMN to be highly occurring. We have found a state that clearly covers areas of the visual cortex, another state that shows regions corresponding to the auditory network and some portions of the frontoparietal cortex, and another one that contains the bilateral temporal cortices and the insula, which is analogous to the salience network. Additionally, we also have found activations in the left and right hemispheres to be separately clustered.

Hence, GLMM uses an existing pool of empirical fMRI data to approximate graphs that, interestingly, arise to be meaningful brain networks consistent with the literature (van den Heuvel and Hulshoff Pol, 2010).

3.7. Comparison of learned graphs reminiscent of the DMN

We have found that states corresponding to the rest epochs of the different tasks always bear striking similarity to the DMN, as expected. To understand the nature of the recovered DMN-related states, we compute Pearson correlation between the coactivation patterns between these DMN-related means (μ) for each task as well as RS, which have an important topological meaning, since each node represent a brain area. Fig. 4(C) displays the correlation between the extracted DMN-related states from all tasks and RS fMRI. The axes are sorted from lowest to highest total correlation per row. A separation between low-level sensory-motor (*e.g.*, language, motor, rest) *versus* high-level cognition (*e.g.* working memory, relational memory, social and emotion).

These co-activation patterns are all highly correlated but Language, Relational Memory and Rest DMNs present some negative correlations



Fig. 4. Extracted states from resting-state (A) Six states corresponding to the RS data. Values in percentage show the mean occurrences of each state with its standard deviations across all 50 subjects considered. (B) Example activity profile of each state for an example subject in the same order presented in (A). (C) Spatial correlation of estimated means corresponding to DMN extracted from resting state and all DMNs from task fMRI data.

with the other tasks DMNs. The reader is pointed to figure A4 in Supplementary Material to visually assess these differences.

3.8. Comparison of learned functional DMN graphs with structure

Here we want to compare the DMN-related states to structure by simple visualization and applying some similarity or dissimilarity metrics.

As mentioned above, these states are related to the rest epoch of the different tasks datasets and the DMN-state extracted from RS.

Fig. 5 (A) shows the group-averaged SC across all subjects considered in this work, and Fig. 5 (B) displays the weighted adjacency matrix computed as A = D - L from the estimated graph Laplacian matrix of the DMN network of the Resting State. The adjacency matrix is visually more sparse than the SC. Moreover, a direct comparison (i.e., Pearson correlation across connections) between the state matrices and SC reveals similarity values within the range of r = 0.48 - 0.63, as shown in Fig. 5(C). It is also noteworthy that, compared to the conventional FC-SC relationship, where FC is obtained by correlating inter-regional BOLD timecourses using either Pearson correlation or partial correlation, the correlation with SC is much higher for GLMM-based state matrices. It has to be noticed that they are not statistically different within methods, namely the correlation of each DMN with SC is similar in all the methods.

By computing the dissimilarity spectral scores for each DMN and sorting them in ascending order, it emerges that there is a trend of "externally directed" or less introspective tasks. This score is normalized to enhance the visualization and it is only representative of the gradient, not of the significant difference.

Low cognitive level tasks, such as Motor and Resting State and Language, are expected to have low scores of dissimilarity with structure, hence are more similar to structure. Conversely, more introspective and "internally directed" tasks, such as Relational Memory, Social and Emotion, differ more from structure. Fig. 5(D) wants to highlight the gradient and order of the DMN under consideration to assess a potential decoupling index with the aim of creating a framework for potential application in meta-analyses Preti and Van De Ville (2019) (see Supplementary Materials Appendix A.5.

Fig. 5 can be complemented with the visualization of the DMNs in A4 to advance reasonable conclusions. In fact, even though the identified brain pattern in the rest epochs across all tasks contains the key regions of the DMN, the corresponding connectivity structure of the state matrix varies, suggesting task-specific mechanisms in the way brain areas connect.

4. Discussion

4.1. General findings

We have proposed the GLMM framework for explaining brain activity based on a generative model. Specifically, we were able to obtain A. Structural Connectome (SC)

C. Correlation between FC measures and SC



Fig. 5. Weighted adjacency matrix computed from the learned Laplacian and its relation to the structural connectome. (A) Normalized group-averaged SC matrix corresponding to all subjects considered in this work. (B) Normalized adjacency matrix extracted from the Laplacian matrix of the DMN extracted from Resting State. (C) Similarity between SC and different FC measures: using (1) GLMM-based adjacency extracted from learned graphs and classical measures of FC using (2) Pearson Correlation and (3) Partial Correlation. (D) Percentage dissimilarity scores for networks of a specific task, relative to the DMN score of that same task.

states that are characterized by: (1) spatial pattern of regional activation levels; (2) underlying graph that describes the interactions of the regions; (3) likelihood over time.

First, we validated the approach by demonstrating that the extracted patterns are consistent with established neurophysiological descriptions corresponding to the tasks. We also showed that the probability of these states to occur at each time-point is able to capture the timing of the task paradigms, even though no temporal information is given to the GLMM. Surprisingly, also the hyperparameters are optimized by focusing solely on the activation patterns. Therefore the algorithm has no clue on the timings.

Setting the number of states to a number that is comparable with the task epochs allows us to have simpler visualization and clearer interpretation of each state with respect to the task paradigm, keeping in mind that the one state might correspond to a combination of task epochs and *vice versa*. Interestingly, when the states are significantly more than the number of task epochs, the most prominent activation patterns were the ones characterized by the activation of medial prefrontal cortex, posterior cingulate cortex and angular gyrus - that are commonly associated to the Default Mode Network (Buckner and DiNicola, 2019). Moreover, their corresponding temporal dynamic was always characterized by a high likelihood (γ) during the rest epochs of task fMRI data.

Secondly, we applied the approach to RS data, which revealed contributions of some of the well-known RS networks, such as the DMN, visual, auditory/attention, and salience networks. Finally, we took advantage of the estimated graph Laplacian matrices to understand the interactions of the regions, and how the strength of these interactions relates to the underlying brain structure obtained from DW-MRI. We showed that the adjacency matrices computed from the graph Laplacians bear closer similarity to SC than conventionally defined FC matrices obtained by Pearson or partial correlations of time-courses. We conclude that the graphs estimated by the GLMM are generally more correlated to structure even if the direct input is simply the empirical fMRI data, without any additional structural information.

Since a DMN-related state was found in all datasets, we have compared the spatial patterns to assess how they correlate to each other. Sorting the correlation values, a distinction between low- and high-level cognitive tasks was revealed. It is interesting to notice that the correlations can be explained by visually looking at which brain area has positive or negative activations (μ) (Fig. A4). In particular, the DMN estimated from Language seems to have low correlations with the DMNs estimated from the other tasks because the nodes corresponding to inferior frontal and temporal gyrus, hippocampus, amygdala and calcarine fissure have opposite sign values.

Comparison of the matrices using the spectral approach also showed a similar trend of low-level and high-level cognitive *tasks*. Comparing the DMN-related graphs to the SC is beneficial to evaluate how the co-activation patterns are supported in terms of connectivity, considering the estimated Laplacians that is one of the main benefits of using GLMM.

One limitation of our approach might be the concatenation in time of all subjects' data, therefore we assume that parameters are constant in the whole population, which is a clear over- simplification knowing that individual fingerprinting can be reliably achieved on the basis of RS fMRI recordings Finn et al. (2015). Nonetheless, we are able to construct a set of maps that represent the entire population and this could be helpful as a normative model of healthy individuals to assess patients in a clinical setting, knowing that there are functional differences in the brain network related to autism Nair et al. (2020), Alzheimer's Disease (Grieder et al., 2018) and other neurological conditions (Liao et al., 2014; Starck et al., 2013; Tagliazucchi et al., 2010; Zöller et al., 2019). Moreover, the subject specificity can be taken into account when looking at the time dynamic of the γ probabilities, since temporality plays also a key role in the human brain fingerprinting (Van De Ville et al., 2021).

4.2. Estimated Graph Laplacians conform with structural connectivity from DW-MRI

Pearson correlation has been widely deployed to measure FC between different brain regions (Chang and Glover, 2010; Friston et al., 1994; Honey et al., 2009; Horn et al., 2014; Hutchison et al., 2013; Preti et al., 2017; Supekar et al., 2010). The simplicity of the approach, however, comes with a number of well-known limitations: it only reflects an averaged measure of the association between regions at all timepoints, leaving out the dynamic nature of brain function (Preti et al., 2017). Moreover, it does not provide evidence of a direct relationship between pairwise regions without possible confounding effects that is introduced by a relationship to a third region due to the inherent transitive nature of the metric. To address this issue, partial correlation has been shown to improve the results by regressing out the effects of other variables (Deligianni et al., 2013; Liégeois et al., 2020; Smith et al., 2011; Wang et al., 2016). The GLMM approach directly estimates a Laplacian matrix to characterize the connectivity profile of a state; e.g., the GLMM adjacency matrix in Fig. 5(B) is much sparser than the FC matrix, even sparser than the SC, because of the regularization which controls sparsity. We also show in Fig. 5(C) that the GLMM adjacency matrices bear higher similarity to SC than FC variants. Unlike classical FC, GLMM adjacency matrices include direct interactions between regions within a network and do not include indirect correlations which are dependant on other regions of interest. This phenomena occurs because the GLMM, built on Gaussian Graphical models, provably have this property (Uhler, 2017). Moreover, the GLMM method is a dynamic approach, whereby each GLMM matrix corresponds to particular timepoints (i.e., each task epoch turns out to correspond to a specific state, and consequently to a specific matrix). It has been shown in previous work (Liégeois et al., 2016) that there are fluctuations in similarity between FC and SC when sliding-window analysis is applied instead of static FC.

As shown in Fig. 5 (C), the GLMM state matrices have closer similarity to the SC. This could be explained by the fact that these are state-specific graphs whose edges reflect pairwise regions that show synchronous changes. In this case, since we are specifically comparing the graphs corresponding to the DMN in rest, and in all tasks, the edges of the learned graphs reveal how signals defined on the regions comprising the DMN simultaneously fluctuates with all others. Therefore, these graphs recovered by the GLMM offer a more accurate description of functional connectivity networks. We surmise that their close similarity with SC support previous findings that brain activity is shaped by the anatomical backbone on which it manifests Cabral et al. (2011); Honey et al. (2009); Liégeois et al. (2016).

From the standard deviation of how the different DMNs are correlated to the SC, it can be noticed that there is no statistical difference within the same FC method chosen, but a similar trend is consistently found applying the spectral dissimilarity and running the GLMM with different conditions (hyper-parameters).

The GLMM provides a good tool to describe functional connectivity of a whole-population with meaningful comparisons to structure. Therefore, in future works this approach could be applied to a pull of patients to better assess these differences.

4.3. Distance between functional states and structural connectome reflects a gradient along tasks

The dissimilarity score based on spectral distance can be interpreted as a function-structure decoupling index of each state. In particular, each DMN-related state is associated with a value that indicates how *far away* it is from structure. This metric of function-structure *distance* was capable of separating low- from high-level tasks. As shown in previous works (McCormick and Telzer, 2018), task engagement plays a role in the regulation of DMN in order to perform goal-directed behaviors. Even though subjects are resting, they are involved with different levels of engagement.

Generally, there has always been an interest in disentangling the role of the DMN in different tasks (Elton and Gao, 2015a). Here, we focus on how these DMN-related states connect functionally, and their distance with anatomy as reflected by the SC. It is interesting to notice that structure-function decoupling allows distinguishing the different tasks. Thus, as shown in Elton and Gao (2015a), the DMN may play a great role in internal and external tasks through a flexible coupling with task-relevant brain areas.

However, it is important to highlight that this metric is introduced as a simple tool of sorting the (dis)similarity of functional networks with respect to structure. We acknowledge that this metric is extremely dependent on the choice of SC and FC generation and no statistical difference is claimed: the scores are normalized for a better visualization. Nonetheless, the trend highlighted is also found using traditional Pearson correlation, as shown in Fig. 5(C). Moreover, this gradient can also be explained visually, looking at the density of the connections.

4.4. Methodological perspectives

Dynamic analyses of functional imaging data during rest and tasks have been going on for a decade (Chang and Glover, 2010). Since their inception, several methodological advances have been introduced to probe the functional organization of the brain from a dynamical point of view. A thorough review of the existing methodological tools (Hutchison et al., 2013; Preti et al., 2017) has classified existing approaches into four distinct groups of methods: (i) sliding-window correlations, (ii) frame-wise analyses, (iii) state modeling, and (iv) temporal modeling. We consider our proposed technique to be an integration of state and temporal modeling, whereby the GLMM approach provides states characterized in terms of activity and connectivity, as well as the state time course, akin to how Hidden Markov Models (HMM) (Stevner et al., 2019; Vidaurre et al., 2017) are able to capture various brain states and their likelihood to occur at each time-point. Unlike HMMs, however, GLMM does not assume any Markovian process and, differently from a Gaussian Mixture Model (GMM), it extracts the states directly from the averaged BOLD data within parcellated brain regions without a dimensionality reduction step such as PCA, that are often needed to apply GMMs. The reason is the implicit dimensionality reduction that occurs in GLMM by imposing a Laplacian structure on the inferred graphs (Petric Maretic and Frossard, 2020). This model has been shown to outperform standard clustering methods on high-dimensional tasks, even when intuitively there is a priori no inherent graph structure in the data, achieving better clustering accuracy. The reader is directed to Petric Maretic and Frossard (2020) for more details and proofs. Additionally, the inferred Laplacians add a strong layer of interpretability to our findings and offer multiple possibilities for further analysis and

understanding of brain networks. The extracted individual graph structures describe the interactions of regions comprising the states. Apart from the enhanced visual understanding of interactions between the brain regions, these graphs enable further comparison of different brain networks, as shown in Figs. 4 and 5.

Meanwhile, the extraction of the activity time course is a relevant aspect that this method highlights: in conventional methods, a functional graph is a representation of the functional connectome, which reflects the statistical inter-dependency between brain signals across all brain regions (nodes). The main issue with this approach is that one needs to define time windows and the size of these windows may bias the analysis (Hindriks et al., 2016; Leonardi et al., 2013). Furthermore, the input to the GLMM algorithm is the concatenated raw fMRI signals, and thus the experimental paradigm is completely unknown to the algorithm. Nevertheless, the γ values (probabilities of belonging to a cluster) impressively manage not only to capture the experimental conditions of the tasks, but also to extract network activity patterns that are consistent with previously established knowledge using classical regression analyses.

A particularity of the model is the imposition of a special structure on the Gaussian, bringing several benefits. Namely, as shown in Petric Maretic and Frossard (2020), the model is very robust to a highdimensional setting, when compared to standard mixtures of Gaussians. Furthermore, the graph Laplacian matrices obtained with this method offer a high level of interpretability. Therefore, the method can be applied directly to atlas-based time courses, and does not need to resolve to any prior dimensionality reduction, providing results which can be directly connected to the atlas under consideration, making the results more explicable.

4.5. Technical limitations

While the proposed method permits to capture more information about estimated brain networks, it still suffers from some limitations, mostly due to the oversimplification of a very complex problem. Firstly, the method makes the assumption that each fMRI signal corresponds to exactly one graph, i.e., one brain network. When fitting the model, it does return probabilities of the signal belonging to each of the states, but that is largely different from assuming that one signal actually originated as a combination of several neurological networks, a phenomena that could very well be present in practice due to temporal overlap at the hemodynamic timescale. Furthermore, the method does not specify any time constraints. Even though the experimental paradigm is completely unknown to the algorithm, the γ values (probabilities of belonging to a state) manage to capture the timing of the experimental conditions. This observation helped us to validate the method, showing that even with no temporal information, meaningful states are obtained.

Finally, additional information from the literature could be incorporated into the graph inference problem, so that more accurate and reliable findings can be obtained. These limitations can be addressed in future work.

5. Conclusion

This study presents a new framework for uncovering dynamic representations of brain activity. The extracted state graphs are sparse and mainly capture connections supported by the underlying SC.

The degree of dissimilarity between DMN-related states and SC may allow to advance a distinction between tasks, with a trend that could be replicated by running the algorithm multiple times in different conditions. Even though the majority of the DMN functional connectivity matrices are not statistically different, the spectral distance can be a good approximation of the matricial information that the FC carries and can be an interesting decoupling metric to investigate behavioral specialization through meta-analyses in future works. Overall, our findings validate the potential of the proposed technique in providing a meaningful representation of brain activity. Future works could focus on application of our method in clinical settings for population studies, or focusing the attention on a more subject-specific analysis.

Correspondence

Correspondence and requests for materials should be addressed to Ilaria Ricchi (email: ilaria.ricchi@epfl.ch).

Declaration of Competing Interest

The authors declare that they have no competing financial interests.

Credit authorship contribution statement

Ilaria Ricchi: Software, Validation, Formal analysis, Writing – original draft, Writing – review & editing. **Anjali Tarun:** Conceptualization, Formal analysis, Data curation, Writing – original draft, Writing – review & editing, Supervision. **Hermina Petric Maretic:** Methodology, Conceptualization, Writing – original draft, Writing – review & editing, Supervision. **Pascal Frossard:** Methodology, Writing – review & editing, Supervision. **Dimitri Van De Ville:** Conceptualization, Writing – review & editing, Supervision.

Acknowledgements

This work was supported by the Swiss National Science Foundation under the Project Grant 205321-163376. We thank Dr. Maria Giulia Preti for providing guidelines on the meta-analysis with Neurosynth.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.neuroimage.2022.119037

References

- Allen, E.A., Damaraju, E., Plis, S.M., Erhardt, E.B., Eichele, T., Calhoun, V.D., 2014. Tracking whole-brain connectivity dynamics in the resting state. Cerebral Cortex 24 (3), 663–676. doi:10.1093/cercor/bhs352.
- Andrews-Hanna, J.R., 2012. The Brain's default network and its adaptive role in internal mentation. The Neuroscientist 18 (3), 251–270. doi:10.1177/1073858411403316.
- Anzai, A., Peng, X., Van Essen, D.C., 2007. Neurons in monkey visual area V2 encode combinations of orientations. Nat. Neurosci. 10 (10), 1313–1321. doi:10.1038/nn1975.
- Arsalidou, M., Pawliw-Levac, M., Sadeghi, M., Pascual-Leone, J., 2018. Brain areas associated with numbers and calculations in children: meta-analyses of fmri studies. Dev Cogn Neurosci 30 (January 2017), 239–250. doi:10.1016/j.dcn.2017.08.002.
- Atasoy, S., Donnelly, I., Pearson, J., 2016. Human brain networks function in connectomespecific harmonic waves. Nat Commun 7 (1), 10340. doi:10.1038/ncomms10340.
- Barch, D.M., Burgess, G.C., Harms, M.P., Petersen, S.E., Schlaggar, B.L., Corbetta, M., Glasser, M.F., Curtiss, S., Dixit, S., Feldt, C., Nolan, D., Bryant, E., Hartley, T., Footer, O., Bjork, J.M., Poldrack, R., Smith, S., Johansen-Berg, H., Snyder, A.Z., Van Essen, D.C., 2013. Function in the human connectome: task-fMRI and individual differences in behavior. Neuroimage 80, 169–189. doi:10.1016/j.neuroimage.2013.05.033.
- Beckmann, C.F., DeLuca, M., Devlin, J.T., Smith, S.M., 2005. Investigations into resting-state connectivity using independent component analysis. Philosophical Transactions of the Royal Society B: Biological Sciences 360 (1457), 1001–1013. doi:10.1098/rstb.2005.1634.
- Betzel, R.F., Fukushima, M., He, Y., Zuo, X.-N., Sporns, O., 2016. Dynamic fluctuations coincide with periods of high and low modularity in resting-state functional brain networks. Neuroimage 127, 287–297. doi:10.1016/j.neuroimage.2015.12.001.
- Biswal, B.F., Yetkin, Z., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magn Reson Med 34, 537–541. doi:10.1002/mrm.1910340409.
- Bolton, T.A., Morgenroth, E., Preti, M.G., Van De Ville, D., 2020. Tapping into multifaceted human behavior and psychopathology using fmri brain dynamics. Trends Neurosci. xx (xx), 1–14. doi:10.1016/j.tins.2020.06.005.
- Buckner, R., DiNicola, L., 2019. The brain's default network: updated anatomy, physiology and evolving insights. Nat Rev Neurosci 20, 432–441. doi:10.1038/s41583-019-0212-7.

- Bufalari, I., Aprile, T., Avenanti, A., Di Russo, F., Aglioti, S.M., 2007. Empathy for pain and touch in the human somatosensory cortex. Cerebral Cortex 17 (11), 2553–2561. doi:10.1093/cercor/bhl161.
- Bullmore, E., Sporns, O., 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. Nat. Rev. Neurosci. 10 (3), 186–198. doi:10.1038/nrn2575.
- Cabral, J., Hugues, E., Sporns, O., Deco, G., 2011. Role of local network oscillations in resting-state functional connectivity. Neuroimage 57 (1), 130–139. doi:10.1016/j.neuroimage.2011.04.010.
- Chang, C., Glover, G.H., 2010. Time-frequency dynamics of resting-state brain connectivity measured with fmri. Neuroimage 50 (1), 81–98. doi:10.1016/j.neuroimage.2009.12.011.
- Corfield, D.R., Murphy, K., Josephs, O., Fink, G.R., Frackowiak, R.S.J., Guz, A., Adams, L., Turner, R., 1999. Cortical and subcortical control of tongue movement in humans: a functional neuroimaging study using fmri. J Appl Physiol 86 (5), 1468–1477. doi:10.1152/jappl.1999.86.5.1468.
- Deligianni, F., Varoquaux, G., Thirion, B., Sharp, D.J., Ledig, C., Leech, R., Rueckert, D., 2013. A framework for inter-subject prediction of functional connectivity from structural networks. IEEE Trans Med Imaging 32 (12), 2200–2214. doi:10.1109/TMI.2013.2276916.
- Dempster, A.P., 1972. Covariance selection. Biometrics 157-175.
- Dolan, R. J., 2002. Neuroscience and psychology: Emotion, cognition, and behavior. 10.1126/science.1076358

Dong, X., Thanou, D., Frossard, P., Vandergheynst, P., 2016. Learning Laplacian matrix in smooth graph signal representations. IEEE Trans. Signal Process. 64 (23), 6160–6173. Dong, X., Thanou, D., Rabbat, M., Frossard, P., 2019. Learning graphs from data: a signal

representation perspective. IEEE Signal Process Mag 36 (3), 44–63.

- Eavani, H., Satterthwaite, T.D., Gur, R.E., Gur, R.C., Davatzikos, C., 2013. Unsupervised learning of functional network dynamics in resting state fMRI. Inf Process Med Imaging (23) 426–437. doi:10.1007/978-3-642-38868-2_36.
- Elton, A., Gao, W., 2015. Task-positive functional connectivity of the default mode network transcends task domain. J Cogn Neurosci. (27(12):2369-81) doi:10.1162/jocn_a_00859.
- Elton, A., Gao, W., 2015. Task-related modulation of functional connectivity variability and its behavioral correlations. Hum Brain Mapp 36 (8), 3260–3272. doi:10.1002/hbm.22847.
- Finn, E.S., Shen, X., Scheinost, D., Rosenberg, M.D., Huang, J., Chun, M.M., Papademetris, X., Constable, R.T., 2015. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. Nat. Neurosci. 18 (11), 1664– 1671. doi:10.1038/nn.4135.
- Friston, K.J., Jezzard, P., Turner, R., 1994. Analysis of functional mri time-series. Hum Brain Mapp 1 (2), 153–171. doi:10.1002/hbm.460010207.
- Gan, L., Yang, X., Narisetty, N., Liang, F., 2019. Bayesian joint estimation of multiple graphical models. In: Advances in Neural Information Processing Systems, pp. 9799–9809.
- 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., 2013. The minimal preprocessing pipelines for the human connectome project. Neuroimage 80, 105–124. doi:10.1016/j.neuroimage.2013.04.127.
- Goni, J., van den Heuvel, M.P., Avena-Koenigsberger, A., Velez de Mendizabal, N., Betzel, R.F., Griffa, A., Hagmann, P., Corominas-Murtra, B., Thiran, J.-P., Sporns, O., 2014. Resting-brain functional connectivity predicted by analytic measures of network communication. Proceedings of the National Academy of Sciences 111 (2), 833– 838. doi:10.1073/pnas.1315529111.

Goodall, C., 1991. Procrustes methods in the statistical analysis of shape. Journal of the Royal Statistical Society 53 (2), 285–339.

- Greicius, M.D., Krasnow, B., Reiss, A.L., Menon, V., 2003. Functional connectivity in the resting brain: anetwork analysis of the default mode hypothesis. Proceedings of the National Academy of Sciences 100 (1), 253–258. doi:10.1073/pnas.0135058100.
- Grieder, M., Wang, D.J.J., Dierks, T., Wahlund, L.-O., Jann, K., 2018. Default mode network complexity and cognitive decline in mild alzheimer's disease. Front Neurosci 12, 770. doi:10.3389/fnins.2018.00770.
- Gu, S., Pasqualetti, F., Cieslak, M., Telesford, Q.K., Yu, A.B., Kahn, A.E., Medaglia, J.D., Vettel, J.M., Miller, M.B., Grafton, S.T., Bassett, D.S., 2015. Controllability of structural brain networks. Nat Commun 6 (1), 8414. doi:10.1038/ncomms9414.
- Hariri, A., Bookheimer, S., Mazziotta, J., 2000. Modulating emotional responses: effects of a neocortical network on the limbic system. Neuroreport 11 (1), 43–48. doi:10.1038/s41598-018-37163-9.
- Hegdé, J., Van Essen, D.C., 2000. Selectivity for complex shapes in primate visual area V2. J. Neurosci. 20 (5), 1–6. doi:10.1523/jneurosci.20-05-j0001.2000.
- Hindriks, R., Adhikari, M., Murayama, Y., Ganzetti, M., Mantini, D., Logothetis, N., Deco, G., 2016. Can sliding-window correlations reveal dynamic functional connectivity in resting-state fmri? Neuroimage 127, 242–256. doi:10.1016/j.neuroimage.2015.11.055.
- Honey, C.J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.P., Meuli, R., Hagmann, P., 2009. Predicting human resting-state functional connectivity from structural connectivity. Proceedings of the National Academy of Sciences 106 (6), 2035–2040. doi:10.1073/pnas.0811168106.
- Horn, A., Ostwald, D., Reisert, M., Blankenburg, F., 2014. The structural-functional connectome and the default mode network of the human brain. Neuroimage 102 (P1), 142–151. doi:10.1016/j.neuroimage.2013.09.069.
- Hutchison, R., Womelsdorf, T., Allen, E., Bandettini, P., Calhoun, V., Corbetta, M., Della Penna, S., Duyn, J., Glover, G., Gonzalez-Castillo, J., Handwerker, D., Keilholz, S., Kiviniemi, V., Leopold, D., de Pasquale, F., Sporns, O., Walter, M.M.C, 2013. Dynamic functional connectivity: promise, issues, and interpretations. Neuroimage doi:10.1016/j.neuroimage.2013.05.079.

- Kalofolias, V., 2016. How to learn a graph from smooth signals. In: Artificial Intelligence and Statistics, pp. 920–929.
- Kalofolias, V., Loukas, A., Thanou, D., Frossard, P., 2017. Learning time varying graphs. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, pp. 2826–2830.
- Kendall, D.G., 1989. A survey of the statistical theory of shape. Statistical Science 4 (2), 87–120. doi:10.1214/ss/1177012582.
- Kropf, E., Syan, S.K., Minuzzi, L., Frey, B.N., 2019. From anatomy to function: the role of the somatosensory cortex in emotional regulation. Rev Bras Psiquiatr 41 (3), 261–269. doi:10.1590/1516-4446-2018-0183.
- Kucyi, A., Davis, K.D., 2014. Dynamic functional connectivity of the default mode network tracks daydreaming. Neuroimage 100, 471–480. doi:10.1016/j.neuroimage.2014.06.044.
- Leonardi, N., Richiardi, J., Gschwind, M., Simioni, S., Annoni, J.-m., Schluep, M., Vuilleumier, P., Van De Ville, D., 2013. Principal components of functional connectivity: anew approach to study dynamic brain connectivity during rest. Neuroimage 83, 937– 950. doi:10.1016/j.neuroimage.2013.07.019.
- Leonardi, N., Van De Ville, D., 2015. On spurious and real fluctuations of dynamic functional connectivity during rest. Neuroimage 104, 430–436. doi:10.1016/j.neuroimage.2014.09.007.
- Liao, W., Zhang, Z., Mantini, D., Xu, Q., Ji, G.J., Zhang, H., Wang, J., Wang, Z., Chen, G., Tian, L., Jiao, Q., Zang, Y.F., Lu, G., 2014. Dynamical intrinsic functional architecture of the brain during absence seizures. Brain Structure and Function 219 (6), 2001– 2015. doi:10.1007/s00429-013-0619-2.
- Liégeois, R., Ziegler, E., Phillips, C., Geurts, P., Gómez, F., Bahri, M.A., Yeo, B.T.T., Soddu, A., Vanhaudenhuyse, A., Laureys, S., Sepulchre, R., 2016. Cerebral functional connectivity periodically (de)synchronizes with anatomical constraints. Brain Structure and Function 221 (6), 2985–2997. doi:10.1007/s00429-015-1083-y.
- Liégeois, R., Santos, A., Matta, V., Van De Ville, D., Sayed, A.H., 2020. Revisiting correlation-based functional connectivity and its relationship with structural connectivity. Network Neurosci. 0 (0), 1–17. doi:10.1162/netn_a_00166.
- Madhyastha, T.M., Grabowski, T.J., 2014. Age-related differences in the dynamic architecture of intrinsic networks. Brain Connect 4 (4), 231–241. doi:10.1089/brain.2013.0205.
- Maretic, H.P., El Gheche, M., Frossard, P., 2018. Graph heat mixture model learning. In: 2018 52nd Asilomar Conference on Signals, Systems, and Computers. IEEE, pp. 1003–1007.
- Mateos, G., Segarra, S., Marques, A.G., Ribeiro, A., 2019. Connecting the dots: identifying network structure via graph signal processing. IEEE Signal Process Mag 36 (3), 16–43.
- McCormick, E., Telzer, E., 2018. Contributions of default mode network stability and deactivation to adolescent task engagement. Sci Rep 8 (18049). doi:10.1038/s41598-018-36269-4.
- Monti, S., Tamayo, P., Mesirov, J., Golub, T., 2003. Consensus clustering: a resamplingbased method for class discovery and visualization of gene expression microarray data. Mach Learn 52 (1–2), 91–118. doi:10.1023/A:1023949509487.
- Nair, A., Jolliffe, M., Lograsso, Y., Bearden, C.E., 2020. A review of default mode network connectivity and its association with social cognition in adolescents with autism spectrum disorder and early-onset psychosis. Front Psychiatry 11 (614). doi:10.3389/fpsyt.2020.00614.
- Petric Maretic, H., Frossard, P., 2020. Graph Laplacian mixture model. IEEE Trans. Signal Inf. Process. Networks (c) doi:10.1109/TSIPN.2020.2983139.
- Petric Maretic, H., Thanou, D., Frossard, P., 2017. Graph learning under sparsity priors. In: ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings, pp. 6523–6527. doi:10.1109/ICASSP.2017.7953413.
- Preckel, K., Trautwein, F.M., Paulus, F.M., Kirsch, P., Krach, S., Singer, T., Kanske, P., 2019. Neural mechanisms of affective matching across faces and scenes. Sci Rep 9 (1), 1–10. doi:10.1038/s41598-018-37163-9.
- Preti, M.G., Bolton, T.A., Van De Ville, D., 2017. The dynamic functional connectome: state-of-the-art and perspectives. Neuroimage 160 (December), 41–54. doi:10.1016/j.neuroimage.2016.12.061.
- Preti, M.G., Van De Ville, D., 2019. Decoupling of brain function from structure reveals regional behavioral specialization in humans. Nat Commun 10 (1), 4747. doi:10.1038/s41467-019-12765-7.
- Ren, H., Zhu, J., Su, X., Chen, S., Zeng, S., Lan, X., Zou, L.-Y., Sughrue, M.E., Guo, Y., year=2020. Application of structural and functional connectome mismatch for classification and individualized therapy in alzheimer disease. Front Public Health 8, 720. doi:10.3389/fpubh.2020.584430.
- Shine, J., Muller, A., O'Callaghan, C., Hornberger, G., Halliday, G., Lewis, S.J., 2015. Abnormal connectivity between the default mode and the visual system underlies the manifestation of visual hallucinations in parkinson's disease: a task-based fmri study. npj Parkinson's Dis. 1 (15003). doi:10.1038/npjparkd.2015.3.
- Shirer, W.R., Ryali, S., Rykhlevskaia, E., Menon, V., Greicius, M.D., 2012. Decoding subject-Driven cognitive states with whole-Brain connectivity patterns. Cerebral Cortex 22 (1), 158–165. doi:10.1093/cercor/bhr099.
- Sizemore, A.E., Bassett, D.S., 2018. Dynamic graph metrics: tutorial, toolbox, and tale. Neuroimage 180, 417–427. doi:10.1016/j.neuroimage.2017.06.081.
- Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R., Beckmann, F.C., 2009. Correspondence of the brain's functional architecture during activation and rest. PNAS 106 (31), 13040– 13045. doi:10.1073/pnas.0905267106.
- Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., Beckmann, C.F., Nichols, T.E., Ramsey, J.D., Woolrich, M.W., 2011. Network modelling methods for FMRI. Neuroimage 54 (2), 875–891. doi:10.1016/j.neuroimage.2010.08.063.
- Starck, T., Nikkinen, J., Rahko, J., Remes, J., Hurtig, T., Haapsamo, H., Jussila, K., Kuusikko-Gauffin, S., Mattila, M.-L., Jansson-Verkasalo, E., Pauls, D., Ebeling, H., Moilanen, I., Tervonen, O., Kiviniemi, V., 2013. Resting state fmri reveals a default

mode dissociation between retrosplenial and medial prefrontal subnetworks in asd despite motion scrubbing, Front Hum Neurosci 7, 802. doi:10.3389/fnhum.2013.00802.

- Stevner, A.B., Vidaurre, D., Cabral, J., Rapuano, K., Nielsen, S.F., Tagliazucchi, E., Laufs, H., Vuust, P., Deco, G., Woolrich, M.W., Van Someren, E., Kringelbach, M.L., 2019. Discovery of key whole-brain transitions and dynamics during human wakefulness and non-REM sleep. Nat Commun 10 (1), 1035. doi:10.1038/s41467-019-08934-3.
- Supekar, K., Uddin, L.Q., Prater, K., Amin, H., Greicius, M.D., Menon, V., 2010. Development of functional and structural connectivity within the default mode network in young children. Neuroimage 52 (1), 290–301. doi:10.1016/j.neuroimage.2010.04.009.
- Taghia, J., Cai, W., Ryali, S., Kochalka, J., Nicholas, J., Chen, T., Menon, V., 2018. Uncovering hidden brain state dynamics that regulate performance and decision-making during cognition. Nat Commun 9 (2505). doi:10.1038/s41467-018-04723-6.
- Tagliazucchi, E., Balenzuela, P., Fraiman, D., Chialvo, D.R., 2010. Brain resting state is disrupted in chronic back pain patients. Neurosci. Lett. 485 (1), 26–31. doi:10.1016/j.neulet.2010.08.053.
- Tarun, A., Behjat, H., Bolton, T., Abramian, D., Van De Ville, D., 2020. Structural mediation of human brain activity revealed by white-matter interpolation of fmri. Neuroimage 213, 116718. doi:10.1016/j.neuroimage.2020.116718.
- Thanou, D., Dong, X., Kressner, D., Frossard, P., 2017. Learning heat diffusion graphs. IEEE Trans. Signal Inf. Process. Networks 3 (3), 484–499.
- Thomas Yeo, B.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., Fisch, B., Liu, H., Buckner, R.L., 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J. Neurophysiol. 106 (3), 1125–1165. doi:10.1152/jn.00338.2011.
- Uhler, C., 2017. Gaussian graphical models: An algebraic and geometric perspective. 1707.04345.
- Van De Ville, D., Farouj, Y., Preti, M.G., Liégeois, R., Amico, E., 2021. When makes you unique: temporality of the human brain fingerprint. Sci Adv 7 (42), eabj0751. doi:10.1126/sciadv.abj0751.
- van den Heuvel, M.P., Hulshoff Pol, H.E., 2010. Exploring the brain network: a review on resting-state fmri functional connectivity. Eur. Neuropsychopharmacol. 20 (8), 519– 534. doi:10.1016/j.euroneuro.2010.03.008.
- Van Dijk, K.R.A., Hedden, T., Venkataraman, A., Evans, K.C., Lazar, S.W., Buckner, R.L., 2010. Intrinsic functional connectivity as a tool for human connectomics: theory, properties, and optimization. J. Neurophysiol. 103 (1), 297–321. doi:10.1152/jn.00783.2009.

- Varoquaux, G., Gramfort, A., Poline, J.B., Thirion, B., 2010. Brain covariance selection: better individual functional connectivity models using population prior. In: Zemel, R., Shawe-Taylor, J. (Eds.), Advances in Neural Information Processing Systems. John Lafferty, Vancouver, Canada. https://hal.inria.fr/inria-00512451
- Vidaurre, D., Abeysuriya, R., Becker, R., Quinn, A.J., Alfaro-Almagro, F., Smith, S.M., Woolrich, M.W., 2018. Discovering dynamic brain networks from big data in rest and task. Neuroimage 180, 646–656. doi:10.1016/j.neuroimage.2017.06.077.
- Vidaurre, D., Hunt, L.T., Quinn, A.J., Hunt, B.A.E., Brookes, M.J., Nobre, A.C., Woolrich, M.W., 2018. Spontaneous cortical activity transiently organises into frequency specific phase-coupling networks. Nat Commun 9 (2987). doi:10.1038/s41467-018-05316-z.
- Vidaurre, D., Hunt, L.T., Quinn, A.J., Hunt, B.A.E., Brookes, M.J., Nobre, A.C., Woolrich, M.W., 2019. Discovery of key whole-brain transitions and dynamics during human wakefulness and non-rem sleep. Nat Commun 10 (1035). doi:10.1038/s41467-019-08934-3.
- Vidaurre, D., Quinn, A.J., Baker, A.P., Dupret, D., Tejero-Cantero, A., Woolrich, M.W., 2016. Spectrally resolved fast transient brain states in electrophysiological data. Neuroimage 126, 81–95. doi:10.1016/j.neuroimage.2015.11.047.
- Vidaurre, D., Smith, S.M., Woolrich, M.W., 2017. Brain network dynamics are hierarchically organized in time. Proc. Natl. Acad. Sci. U.S.A. 114 (48), 12827–12832. doi:10.1073/pnas.1705120114.
- Wang, Y., Kang, J., Kemmer, P.B., Guo, Y., 2016. An efficient and reliable statistical method for estimating functional connectivity in large scale brain networks using partial correlation. Front Neurosci 10 (MAR). doi:10.3389/fnins.2016.00123.
- Xiao, F.-L., Gao, P.-y., Qian, T.-y., Sui, B.-b., Xue, J., Zhou, J., Lin, Y., 2017. Cortical representation of facial and tongue movements: a task functional magnetic resonance imaging study. Clin Physiol Funct Imaging 37 (3), 341–345. doi:10.1111/cpf.12304.
- Yamada, K., Tanaka, Y., Ortega, A., 2019. Time-varying graph learning based on sparseness of temporal variation. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, pp. 5411–5415.
- Yang, Z., Craddock, R.C., Margulies, D.S., Yan, C.-G., Milham, M.P., 2014. Common intrinsic connectivity states among posteromedial cortex subdivisions: insights from analysis of temporal dynamics. Neuroimage 93 (P1), 124–137. doi:10.1016/j.neuroimage.2014.02.014.
- Zöller, D., Sandini, C., Karahanoğlu, F.I., Padula, M.C., Schaer, M., Eliez, S., Van De Ville, D., 2019. Large-scale brain network dynamics provide a measure of psychosis and anxiety in 22q11.2 deletion syndrome. Biological Psychiatry: Cognitive Neuroscience and Neuroimaging 4 (10), 881-892. doi:10.1016/j.bpsc.2019.04.004.