Partial Correlation as a Tool for Mapping Functional-Structural Correspondence in Human Brain Connectivity

Francesca Santucci
NETWORKS Unit, IMT School for Advanced Studies, Lucca, Italy;
Computational Neuroimaging Lab, Biocruces-Bizkaia HRI, Barakaldo, Spain;
francesca.santucci@imtlucca.it

Antonio Jimenez-Marin
Computational Neuroimaging Lab, Biocruces-Bizkaia HRI, Barakaldo, Spain;

Andrea Gabrielli
Dipartimento di Ingegneria Civile, Informatica e delle Tecnologie Aeronautiche, Università degli Studi Roma Tre, Rome, Italy;
"Enrico Fermi" Research Center - CREF, Rome, Italy;
Istituto dei Sistemi Complessi - CNR, Rome, Italy;

Paolo Bonifazi
Computational Neuroimaging Lab, Biocruces-Bizkaia HRI, Barakaldo, Spain;
Ikerbasque: The Basque Foundation for Science, Bilbao, Spain;

Miguel Ibáñez-Berganza
NETWORKS Unit, IMT School for Advanced Studies, Lucca, Italy;

Tommaso Gili
NETWORKS Unit, IMT School for Advanced Studies, Lucca, Italy;

Jesus M. Cortes
Computational Neuroimaging Lab, Biocruces-Bizkaia HRI, Barakaldo, Spain;
Department of Cell Biology and Histology. University of the Basque Country, Spain;
Ikerbasque: The Basque Foundation for Science, Bilbao, Spain;

October 21, 2024

Abstract

Brain structure-function coupling has been studied in health and disease by many different researchers in recent years. Most of the studies have addressed functional connectivity matrices by estimating correlation coefficients between different brain areas, despite well-known disadvantages compared to partial correlation connectivity matrices. Indeed, partial correlation represents a more sensible model for structural connectivity since, under a Gaussian approximation, it accounts only for direct dependencies between brain areas. Motivated by this and following previous results by different authors, we investigate structure-function coupling using partial correlation matrices of functional magnetic resonance imaging (fMRI) brain activity time series under different regularization (a.k.a. noise-cleaning) algorithms. We find that, across different algorithms and conditions, partial correlation

provides a higher match with structural connectivity retrieved from Density Weighted Imaging data than standard correlation, and this occurs at both subject and population levels. Importantly, we also show that the precise regularization and thresholding strategy are crucial for this match to emerge. Finally, we assess neuro-genetic associations in relation to structure-function coupling, which presents promising opportunities to further advance research in the field of network neuroscience, particularly concerning brain disorders.

***K*eywords** FUNCTIONAL CONNECTIVITY, PARTIAL CORRELATION, PRECISION MATRIX, STRUCTURAL-FUNCTIONAL COUPLING, FMRI BOLD RESTING STATE, REGULARIZATION

1 Introduction

A fundamental problem in network neuroscience is understanding the relationship between Functional Connectivity (FC), accounting for similarity in the activation patterns between brain areas, and Structural Connectivity (SC), which accounts for the brain's anatomical connections, and that, for the particular case of magnetic resonance imaging-derived networks, reduces to white-matter streamlines connecting region pairs [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18]. Despite extensive research on this topic, the complex interplay between the numerous, rapidly evolving functional states and the relatively static underlying structural connectivity remains a major challenge. The intrinsic relationship suggests a many-to-one correspondence between FC and SC, whose comprehensive understanding remains elusive [7]. In other words, given SC, it is not possible to estimate FC or vice versa. Yet, it has been discovered that the underlying structure partially determines function and activity that, in turn, shapes the structure through the processes of neuromodulation and plasticity [7, 19, 20]. The investigation of the relationship between structure and function as a biomarker is generally referred to as structure-function coupling (SFC) [21]. This reciprocal interaction highlights the intricate and adaptive nature of neural networks, where structural elements provide the framework for functional dynamics, and FC continually refines and modulates SC.

The most common approach for studying SFC involves examining the correlation between the intensities (or *weights*) of the single connections (or *links*) of the structural and functional connectivity. Through this approach, several studies have uncovered that while structural and functional connections are positively correlated at rest [5], this correlation is not always consistent and exhibits variability across individuals [22, 23], as well as with age [19, 23], cognitive tasks [22], and different brain regions [24, 23]. Furthermore, it displays alterations in various brain-related disorders [25, 26, 27, 28, 29]. These findings highlight the complex and dynamic nature of SFC in the brain, offering promising avenues for further research into individual variability and clinical applications.

Nevertheless, as previously mentioned, the relationship between SC and FC does not follow a simple link-to-link mapping. For instance, two regions can be functionally connected without a direct structural connection, and SC evolves over much longer time scales than FC [20]. Therefore, another approach to studying SFC investigates the correspondence between structure and function at the module or aggregate level. In this multi-scale approach, brain network modules are compared at multiple resolutions, exploring the full set of nested partitions within a hierarchical tree. Applications of this procedure have revealed that FC and SC share a common structural modular architecture [30, 31, 9, 32, 33].

Another important aspect is that previous studies have estimated FC as the correlation matrix between pairs of functional time series of activity (usually blood oxygen level-dependent (BOLD) functional magnetic resonance imaging (fMRI) signals at rest) corresponding to different brain areas. An alternative method is estimating FC as the partial correlation (PC) of the time series. [34, 35, 36, 37, 38, 39, 40]. Under a Gaussian approximation, PC accounts for direct dependencies between brain areas only, implying that the absence of a link indicates their conditional independence. This makes the PC a more direct model for brain connectivity. [41, 42, 43, 36, 44, 45].

PC can be inferred as a *normalized* precision matrix, which is the inverse of the correlation matrix. Several advantages of PC over correlation have been shown. For instance, PC provides a higher link-wise match between SC and FC [45], exhibits reduced variance across subjects [46], and yields higher prediction scores for certain individual-level measures [39, 47]. However, the use of the PC is limited by the low accuracy of its statistical estimation in the small-sample limit, that is when the time series are short relative to the number of brain regions considered, which is often the case with BOLD fMRI time series [36, 45, 44, 48]. Indeed, the matrix inversion process amplifies the noise present in the correlation matrix [49]. To address this issue, various regularization, or noise-cleaning, methods have been proposed within network neuroscience for accurate inference of the correlation and precision matrices [35, 36, 46, 39]. Many of these methods introduce ℓ_1 or ℓ_2 penalty terms. The ℓ_1 penalty leads to sparse estimators, such as the Graphical Lasso (GLASSO) [50], while the penalty ℓ_2 results in Linear Shrinkage (LS) estimators, where matrices are not intended to be sparsified [51, 52, 44]. The GLASSO is, in general, a good approach for structure recovery [53], but it may not yield stable covariance coefficients since the ℓ_1 penalty used for sparsity often leads to unstable selection. Moreover, its costly

optimization is not suited for large-scale datasets. On the other hand, the Linear Shrinkage estimator is simple and fast to compute. It yields biased estimates that are more stable than the empirical covariance and is often recommended for functional connectivity [39].

Several studies have compared these methods based on criteria such as non-imaging features prediction power [39, 47, 40], stability across scans [54, 39, 40], retrieval of synthetic data generative models [55, 36, 40], and distribution-oriented measures such as test-set (or out-of-sample) likelihood [56, 48]. These studies suggest that the choice between sparse or shrinkage estimators may depend on the intended use and interpretation of FC.

In this context, some authors have investigated the dependence of the SFC on the FC inference method, including partial correlation [40, 45], but limited their analysis to the similarity between SC and FC at the level of the single links. Anyway, works such as [30, 9, 33, 31] have suggested that the SFC may be rooted in a common hierarchical modular organization rather than in a correspondence between the single connections.

Here, we go beyond both approaches by exploring the correspondence between SC and FC at the hierarchical aggregated level, inferring FC from both correlation and partial correlation matrices. Specifically, we show that: i) FC, when estimated from the regularized PC, exhibits greater similarity to SC in terms of modular structure, both at the subject and at the population level, and a smaller variance across subjects; ii) regularization is crucial to our results, which are robust with respect to the regularization methods most widely used in neuroscience; iii) thresholding, on top of the regularization, is fundamental for the emergence of the FC's hierarchical modular structure; iv) the SC-FC similarity reaches a maximum in correspondence with a specific partition of FC; v) the partition's modules are characterized by neurogenetic expression present in major diseases.

These findings advance our understanding of SFC in the brain, suggesting that the use of *regularized PC matrices* may provide a more accurate and stable representation of SC, enhancing the potential for clinical applications and insights into neurogenetic disease mechanisms.

2 Results

Connectivity matrices and cross-modularity for assessing the SC-FC correspondence SC and FC graphs (or connectomes) were obtained from data previously published in [31] and available at [57] for a population of $P = 136$ healthy participants (a methodological pipeline is given in Fig. S1). For each subject s, the data include diffusion-weighted imaging matrices and preprocessed voxel-level resting-state (rs) fMRI time series at various resolutions, among which we selected the one with $N = 183$ regions of interest (ROIs) and, for the time series, $T = 652$ time observations. In particular, we used diffusion-weighted matrices to extract a unique estimate of SC, while we used the rs-fMRI time series to derive the FC with two differentiated strategies: by estimating the correlation matrix $\mathbf{C}$ and the partial correlation matrix $\tilde{\mathbf{J}}$ of the time series, defined as $\tilde{J_{ij}} = \frac{J_{ij}}{\sqrt{J_{ii}J_{jj}}}$, which is the standardized version of the precision matrix $\mathbf{J} = \mathbf{C}^{-1}$. In the demeaned and standardized fMRI time series, contained in the $N \times T$ matrix $\mathbf{X}$ the ratio $q = N/T$ is not negligible, so inference methods that are better than the sample correlation $\mathbf{E} = \frac{\mathbf{X}\mathbf{X}^T}{T}$ are needed for accurately inferring $\mathbf{C}$ (thus, for its inverse), and here, we employed them by regularizing $\mathbf{C}$ with different algorithms. As we will show later, regularization is particularly relevant for the accurate inference of $\tilde{\mathbf{J}}$ [48].

Next, connectivity matrices $\mathbf{M}$ (where $\mathbf{M}$ contains the DWI data for SC and $M_{ij} = |\tilde{J}_{ij}|, |C_{ij}|$ for FC, with $|\cdot|$ indicating the absolute value), were cut at the so-called *percolation threshold* [58]: this procedure prescribes setting to zero all matrix elements smaller than the cutoff that would break the corresponding graph in two or more connected components (also see methods); interpreting these thresholded matrices as adjacency matrices, SC and FC graphs were generated (Fig.2). In the following, we will indicate the structural connectivity graph as $\mathcal{G}_S$ and the functional connectivity graph as $\mathcal{G}_F$, possibly distinguishing between the graphs inferred from $\mathbf{C}$ or $\tilde{\mathbf{J}}$ as $\mathcal{G}_F(\mathbf{C})$ and $\mathcal{G}_F(\tilde{\mathbf{J}})$, respectively. If no differently specified, we assume all graphs to be at the percolation threshold. Each pair of SC and FC graphs was compared at the module level for each subject in the dataset (Fig. 1) and at the population level (Fig. 3). To this aim, we applied a standard Hierarchical Agglomerative Clustering algorithm to obtain a sequence of nested partitions $P_{\mathcal{G}}(m)$ of graph $\mathcal{G}$ into $m = 2, \cdots, N$ modules, for both $\mathcal{G}_S$ and $\mathcal{G}_F$ separately. Subsequently, we computed the quality $Q_{\mathcal{G}}(m)$ of each partition in terms of Newman's modularity [59] and the agreement between each couple of partitions $P_{\mathcal{G}_S}(m)$ and $P_{\mathcal{G}_F}(m)$ in terms of adjusted Normalized Mutual Information (NMI) [60]. Notice that, here, we did not maximize Q but simply calculated it at all levels in the hierarchical tree. These three quantities allowed us to compute the **cross-modularity**

$$\chi_{\mathcal{G}_F,\mathcal{G}_S}(m) = \left(Q_{\mathcal{G}_F}(m) \cdot Q_{\mathcal{G}_S}(m) \cdot \nu_{\mathcal{G}_F,\mathcal{G}_S}(m)\right)^{1/3}, \tag{1}$$

which is a slightly different measure from the original one defined in [30]. The aim of cross-modularity is to quantify the reciprocal similarity of a couple of graphs in terms of their hierarchical modular structure while also taking into

account the quality of their individual partitions. This is done for all number of modules, so that a complete comparison of the whole hierarchy of nested partitions of the two graphs is provided. The range of $\chi(m)$ follows straightforwardly from those of Newman modularity and NMI: they both reach a maximum of 1 for a high-quality partition and a perfect match and are expected to vanish in case of random and unrelated partitions, respectively [61, 62, 60, 63, 64]. In this work, cross-modularity is always used to compare $\mathcal{G}_S$ with an estimate of $\mathcal{G}_F$; therefore, when referring, for the sake of shortness, to the "cross-modularity of $\mathcal{G}_F$," we will always mean the cross-modularity of $\mathcal{G}_F$ and the corresponding $\mathcal{G}_S$.

The cross-modularity $\chi(m)$ was then used for i) tuning the amount of regularization needed for the inference of $\mathcal{G}_F$ from time series data; ii) comparing the effects of inferring $\mathcal{G}_F$ as correlation or partial correlation matrices, as well as the impact of regularization. This evaluation was based on the similarity between the hierarchical nested structures of $\mathcal{G}_F$ and $\mathcal{G}_S$ at both the individual and population levels (Figs. 1a-c); iii) extracting the optimal partition from the best estimate of population-level FC graph (Fig. 1d). Finally, we made use of transcriptomic data from the Allen Human Brain Atlas [65]) and preprocessed in [31], in combination with information regarding the association between genes and diseases [66], to better motivate the biological interpretation of the different estimated partitions (Fig 4).

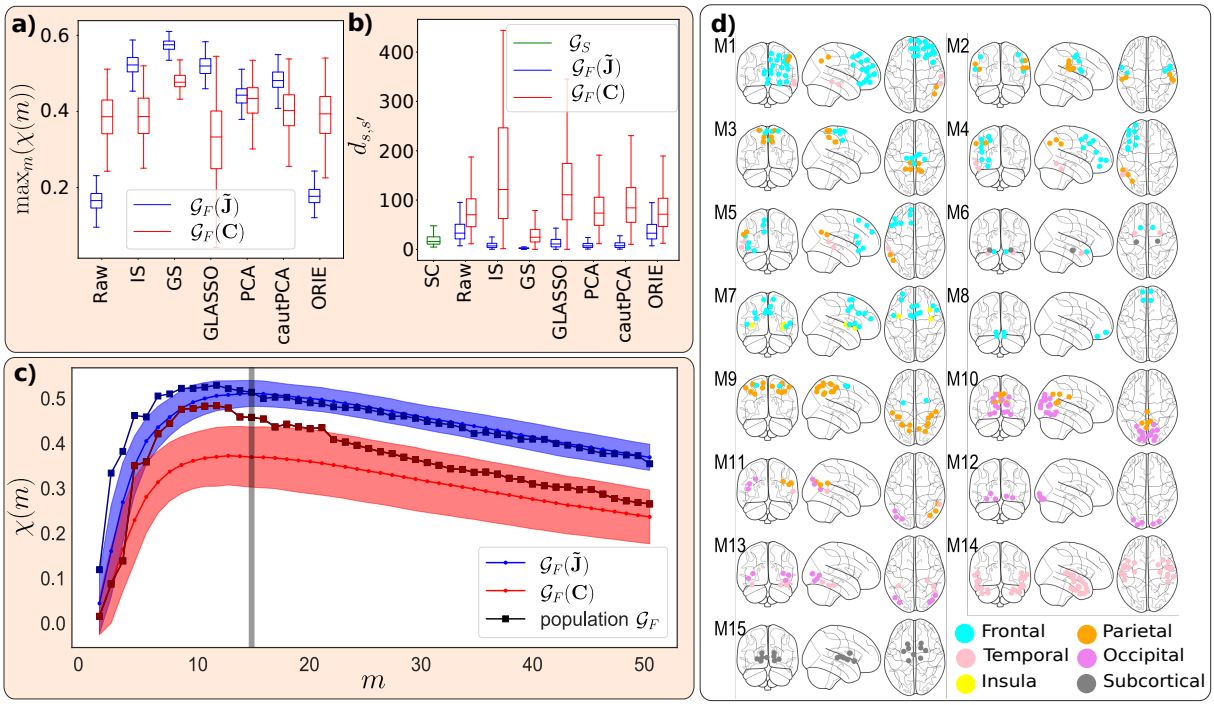

Figure 1: **Partial correlation enhances the FC-SC correspondence in the human brain**. Panels **a** and **b** show how two relevant measures such as (**a**) the cross-modularity and (**b**) the across-subject variance, measured in terms of spectral distance, change depending on whether $\mathcal{G}_F$ is estimated as $\mathbf{C}$ (red) or $\tilde{\mathbf{J}}$ (blue), and the relevance of using a regularization method ("raw" stands for no regularization). Most regularization strategies improve drastically the performance of $\mathcal{G}_F(\tilde{\mathbf{J}})$. More in detail, the boxplots represent the across-subject distribution of (**a**) $\max_m \chi(m)$ and (**b**) the spectral distances between each couple of subjects' $d_{s,s'}$; importantly, we have tuned the regularization in such a way to maximize the cross-modularity of the corresponding $\mathcal{G}_F$ (this is crucial for our specific dataset). Notice how the subjects' $\mathcal{G}_F(\tilde{\mathbf{J}})$ exhibit a systematic lower spectral distance, with respect to $\mathcal{G}_F(\mathbf{C})$, similarly to $\mathcal{G}_S$, that we have reported in green for comparison. The cross-modularity curves of $\mathcal{G}_F(\mathbf{C})$ and $\mathcal{G}_F(\tilde{\mathbf{J}})$, regularized with IS, are shown in subfigure **c**, both at the subject (the solid lines are the across-subject means and the shaded areas their standard-deviation) and at the population level (square-solid lines). The curves of $\mathcal{G}_F(\tilde{\mathbf{J}})$ are significantly higher than those of $\mathcal{G}_F(\mathbf{C})$ for almost all values of m and are characterized by a lower variance. Also, notice that the IS regularization strategy is such that the $\mathcal{G}_F(\mathbf{C})$ cross-modularity cannot be affected by it. The grey vertical line indicates the number of modules $m^* = 15$ that maximizes the cross-modularity of most subjects. We reported in subfigure **d** the partition of the population $\mathcal{G}_F(\tilde{\mathbf{J}})$ into m^* modules, drawing the ROIs belonging to each module as dots in the brain glass plots, colored according to the anatomical regions.

Partial correlation enhances the FC-SC correspondence in the human brain As represented in figure 1, $\mathcal{G}_F$ is more similar to $\mathcal{G}_S$ when inferred from regularized partial correlation than when inferred from regularized correlation in terms of hierarchical modularity, at the subject level, and this is robust across all regularization strategies that are most common in neuroscience. More in detail, with reference to figure 1a, we found that the cross-modularity $\chi(m)$ of $\mathcal{G}_F(\tilde{\mathbf{J}})$ was much lower than the one of $\mathcal{G}_F(\mathbf{C})$ when no regularization was applied. However, when regularization was applied, $\mathcal{G}_F(\tilde{\mathbf{J}})$ reached higher values of $\chi(m)$ for all regularization methods (except for the Optimally Rotationally Invariant Estimator (ORIE) that, despite providing the highest out-of-sample likelihood (see Fig. S3), did not produce any significant improvement to the cross-modularity and was consequently dropped in the rest of the analysis); on the other hand, the cross-modularity of $\mathcal{G}_F(\mathbf{C})$ did not exhibit significantly different variations after regularization (as expected). Moreover, $\mathcal{G}_F$ inferred from the regularized partial correlation produced a significantly higher $\chi(m)$ than the one inferred from the regularized correlation for methods such as Identity shrinkage (IS), Group Shrinkage (GS), Graphical Lasso (GLASSO), and Cautious PCA (cautPCA). In addition, FC inferred from partial correlation exhibited a smaller across-subject variance. A similar trend was found measuring the spectral distance (Fig. 1b) between each couple of subjects' matrices (see methods). We additionally observed that the cross modularity of $\mathcal{G}_F(\tilde{\mathbf{J}})$ is higher when regularizing with cautPCA, first introduced in [48], than when regularizing with the method known as *Eigenvalues Clipping* (indicated as PCA), of which cautPCA is a variant.

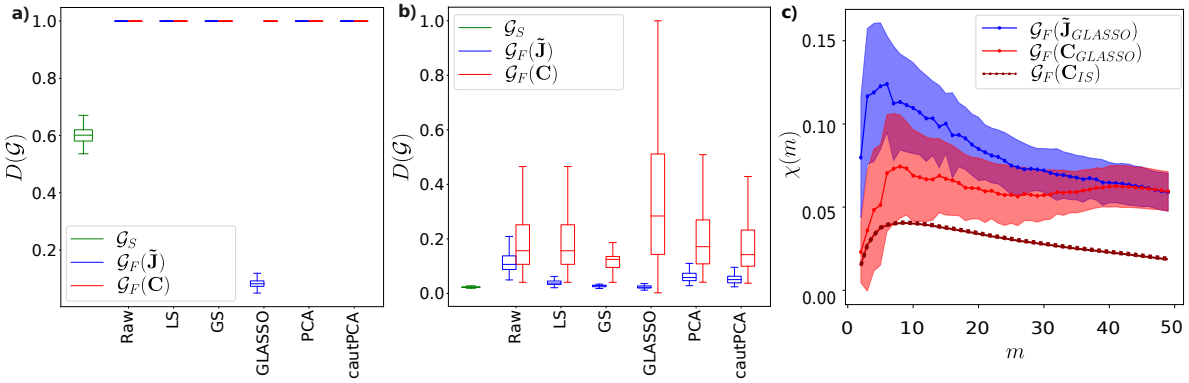

Figure 2: **Relevance of thresholding**. The boxplots **a,b** show the across-subject distribution of the single subjects' densities $D(\mathcal{G})$ of the SC (green) and FC graphs obtained from regularized $\mathbf{C}$ (red) and $\tilde{\mathbf{J}}$ (blue) matrices, without (**a**) and with (**b**) the pruning at the percolation threshold, for different regularization strategies (*"Raw"* stands for "no regularization"), and regularization parameter chosen to maximize the cross-modularity of the thresholded FC. We observe how only the partial correlation regularized with GLASSO provides an FC graph that is already sparse. We also reported in subfigure **c** the maximum cross-modularity achieved by G_F obtained from unthresholded matrices: they were inferred using GLASSO, tuned to maximize the cross-modularity of the corresponding unthresholded graphs. As a comparison, we added to the figure the average cross-modularity of $\mathcal{G}_F$ obtained from unthresholded correlation matrices regularized with IS (equal to the raw ones; we omit the cross-modularity of the corresponding $\mathcal{G}_F(\tilde{\mathbf{J}})$, which vanishes for most values of m). We observe that, while GLASSO provides a sparsification of the precision matrix, the cross-modularity is still much lower than the one that is obtained after further thresholding (compare with Fig. 1).

Thresholding is crucial for the partial correlation FC-SC correspondence We observed that the correspondence between $\mathcal{G}_S$ and $\mathcal{G}_F(\tilde{\mathbf{J}})$ only emerged provided that the latter was inferred from sparse partial correlation matrices, and otherwise resulted in a vanishing cross-modularity (Fig. 2). It is noteworthy that thresholding disrupts the positive definiteness of the matrices, rendering the thresholded connectivity matrices non-compliant with the mathematical definition of correlation matrices. Consequently, some researchers, including [44], advocate against the thresholding process. Instead, they recommend the sparsification induced by GLASSO to achieve sparsity, when needed. Nevertheless, we have confirmed that, although $\mathcal{G}_F(\tilde{\mathbf{J}})$, when regularized by GLASSO, is more similar to its non-thresholded counterpart compared to other regularization techniques due to its inherent sparsity, the percolation threshold still significantly enhances the cross-modularity of $\mathcal{G}_F(\tilde{\mathbf{J}})$ (Fig. 2).

Correspondence between the Shrinkage and LASSO approaches We emphasize that IS and GLASSO, the simplest representatives of the two broader regularization families most common in neuroscience (see methods), gave a similar cross-modularity of $\mathcal{G}_F(\tilde{\mathbf{J}})$. In other words, the FC graphs, extracted from partial correlation matrices regularized

with GLASSO and IS, are equally distant from $\mathcal{G}_S$ at the hierarchical modular level. Indeed, we checked that the two methods produce a similar partition of the corresponding population $\mathcal{G}_F(\tilde{\mathbf{J}})$ (Fig. S2).

Similarity of SC to the across-subject average FC An interesting observation is that GS provides a high cross-modularity and a low across-subject variance, explained by the fact that GS shrinks the estimate of $\mathbf{C}$ towards its across-subject average and the values of the shrinkage parameters (Fig. S3) that maximize $\chi(m)$ of $\mathcal{G}_F(\tilde{\mathbf{J}})$ (the similarity between $\mathcal{G}_F(\tilde{\mathbf{J}})$ and the same subject's $\mathcal{G}_S$) were such that the subjects' $\mathcal{G}_F(\mathbf{C})$ was very similar, and in some cases identical, to the average FC. This is related to the observation that the individual structural connectivities were very often more similar to the average FC than to the same subject's FC.

Cross-modularity curves identify the most representative partition The complete cross-modularity curves, both for individual subjects and at the population level, for $\mathcal{G}_F$ inferred from both correlation and partial correlation regularized with IS (we chose IS as reference regularization strategy, whenever the outcomes produced by the different regularization methods appeared qualitatively consistent), are shown in Fig. 1c (the analogous curves for GLASSO are also reported in the Supplementary Information). We found that in general (at both the individual subject and population levels, for $\mathcal{G}_F$ inferred from both regularized $\mathbf{C}$ and $\tilde{\mathbf{J}}$, and across all regularization methods), the cross-modularity $\chi(m)$ increases rapidly with the number of modules m when m is low. It exhibits a soft maximum when m is between 10 and 20 (though its exact position may slightly vary depending on the specific method used for inferring FC) and then decreases slowly as m increases. We stress that, as already mentioned, the cross-modularity is the product of two quantities that are expected to vanish in unstructured [62, 61] or unrelated graphs [64, 67]. Therefore, the fact that it presents a maximum is nontrivial, and the number of modules providing it can be used as a criterion for selecting the number of modules to partition the FC.

Accordingly, Fig. 1d presents a partition into 15 modules (the median across-subject argument of $\max_m \chi(m)$) of the population $\mathcal{G}_F(\tilde{\mathbf{J}})$, estimated with an IS regularization strategy. The corresponding partition produced by GLASSO, shown in Fig. S2, exhibits remarkable similarity. Figs. S4c-d display the overlap, computed as the Jaccard index, between each module and the Desikan-Killiany structural areas and the Yeo functional networks, respectively.

FC-SC correspondence at the population level Next, we showed the effect of the regularization on $\chi(m)$ at the population level. Again, we selected the simple IS as reference method. Figs. 3a,b reveal that the partial correlation provides a higher cross-modularity than the correlation, but the effect of regularization is less significant compared to the single-subject case. Nevertheless, we found that regularization improves the stability of the partition of the population $\mathcal{G}_F(\tilde{\mathbf{J}})$ across various sub-samplings of the dataset (see methods) and also across regularization parameters (Fig. S5). These stability effects, combined with cross-modularity data, guided us in selecting 0.8 as the optimal regularization parameter for the inference of the population $\mathcal{G}_F(\tilde{\mathbf{J}})$. Figs. 3c-e provide a representation of the population $\mathcal{G}_F(\mathbf{C})$, $\mathcal{G}_F(\tilde{\mathbf{J}})$ and $\mathcal{G}_S$. From a visual inspection, $\mathcal{G}_F(\tilde{\mathbf{J}})$ appears characterized by a higher modularity than $\mathcal{G}_F(\mathbf{C})$ and by a lower density, closer to that of $\mathcal{G}_S$.

Partition of the population FC and neurogenetic interpretation Finally, we characterized the modules of the population $\mathcal{G}_F(\tilde{\mathbf{J}})$ partition shown in Fig. 1 in terms of their participation in major brain-related disorders. In other words, we assessed whether some modules exhibit a significantly lower or higher expression of genes associated with a particular disease. To this end, we computed the across-modules Z-score of the median transcriptomic expression values for sets of genes associated with each brain-related disorder across the regions of interest in each module, following a procedure similar to that described in [31]. The results, shown in Fig. 4, reveal that module M15, corresponding to subcortical regions, is highly relevant for most disorders. It is positively associated with tumor-related and neurodegenerative diseases and negatively associated with psychiatric, substance abuse, and movement-related ones. Interestingly, module M6, which also contains two subcortical regions different from those present in M15, is linked to six diseases that are only weakly represented in M15. We additionally checked the stability of the population $\mathcal{G}_F(\tilde{\mathbf{J}})$ partition across different sub-samplings of the dataset. Indeed, we found an overall average NMI similarity score of 0.80 out of 1 between the main partition and those computed at each sampling (see methods). We also found that some modules are particularly persistent: notably, module M15 appears identical in all partitions (Fig. 4-b).

3 Discussion

In the study of brain structure-function coupling (SFC), functional connectivity (FC) inferred from partial correlation matrices shows a higher correspondence with structural connectivity (SC), even though most previous studies have only evaluated this correspondence at the single link level. In this work, we compared different methods to infer brain

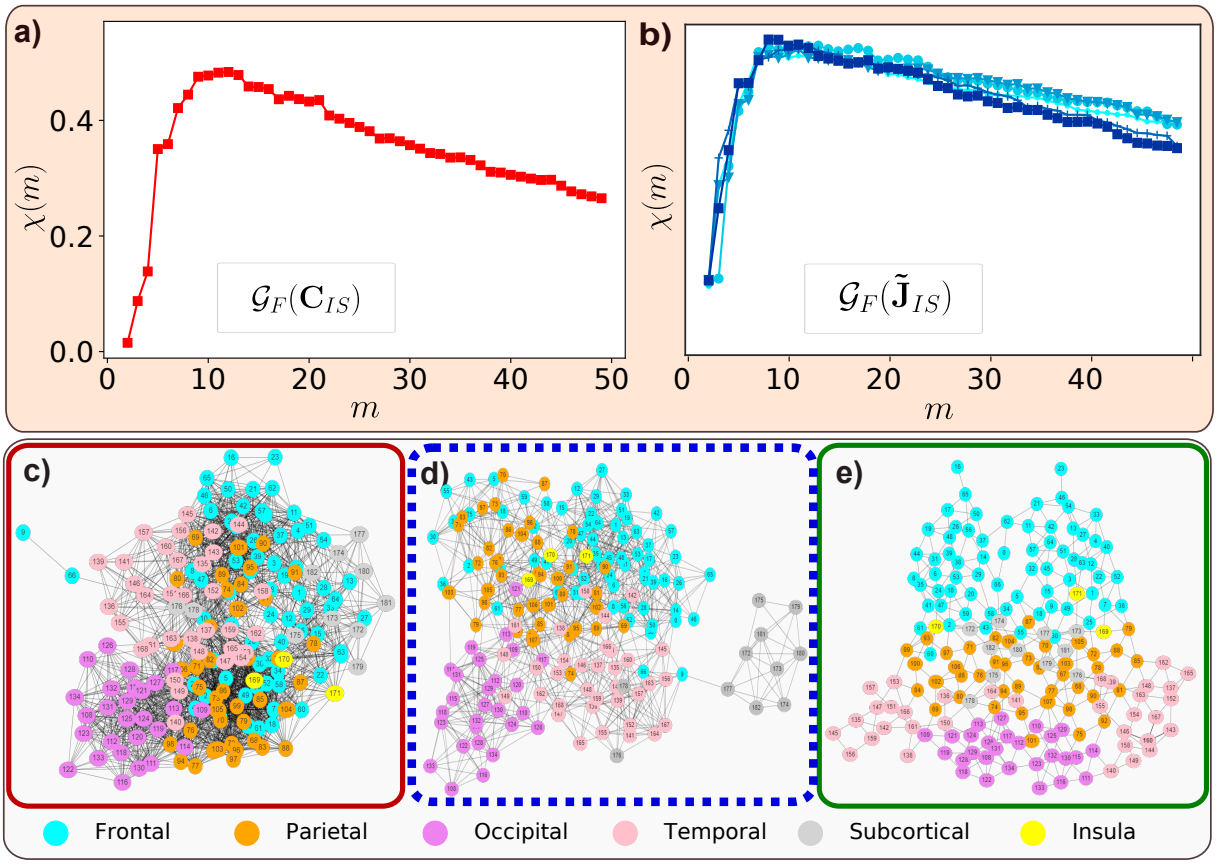

Figure 3: **Connectivity graphs, inferred from C and $\tilde{\mathbf{J}}$, at the population level**. The first row shows the population cross-modularity $\chi(m)$, for $\mathcal{G}_F(\mathbf{C})$ (**a**, red) and $\mathcal{G}_F(\tilde{\mathbf{J}})$ (**b**, blue), at the population level. The population matrices have been computed as the median of the single subject matrices, regularized with IS with different choices of the regularization parameters α. As single subject $\mathcal{G}_F(\mathbf{C})$ are not affected by the IS regularization, whose only effect is linearly shrinking the graphs' links weights, the population cross-modularity of $\mathcal{G}_F(\mathbf{C})$ does not depend on α. As for the population $\mathcal{G}_F(\tilde{\mathbf{J}})$, differently from the single subject case, the regularization does not produce a significant change in the cross-modularity curve. The second row reveals the aspect of the population brain graphs: $\mathcal{G}_F(\mathbf{C})$ (**c**, red), $\mathcal{G}_F(\tilde{\mathbf{J}})$ (**d**, blue) (regularized with IS and regularization parameter that slightly maximizes the respective $\max_m \chi(m)$), and $\mathcal{G}_S$ (**e**, green). A visual inspection appears in agreement with the rest of the results, suggesting that $\mathcal{G}_F(\tilde{\mathbf{J}})$ is closer to $\mathcal{G}_S$ than $\mathcal{G}_F(\mathbf{C})$

FC from resting-state BOLD time series. These methods include estimating FC from the correlation ($\mathbf{C}$) or the partial correlations ($\tilde{\mathbf{J}}$) of the time series, inferred through a regularization algorithm (Fig. 1), and further enhancing FC graph sparsity by setting the smallest matrix elements to zero, up to the so-called percolation threshold (Fig. 2).

Key for our study, while in statistics, regularization typically refers to removing noise that affects the sample correlation matrix when the number of observations is insufficient relative to the dimensionality of the time series, here, we use regularization for a conceptually different purpose: to obtain an estimate of the FC graph that closely matches, at an aggregate or hierarchical modularity level, the SC graph derived from diffusion-weighted imaging matrices. Indeed, the "best" estimate of the FC graph, in the sense of high aggregate similarity to SC does not, in principle, need to be obtained by thresholding the "best" estimate of the correlation matrix, in the sense of minimization of the out-of-sample error. For these reasons, we choose to tune the regularization methods so to maximize the cross-modularity of the corresponding $\mathcal{G}_F$.

Two families of regularization methods are most commonly used in neuroscience: Linear Shrinkage and Graphical Lasso. In this study, we also employed these methods, comparing them with two additional families: one derived from Principal Component Analysis and another based on Random Matrix Theory, known as the Optimal Rotationally

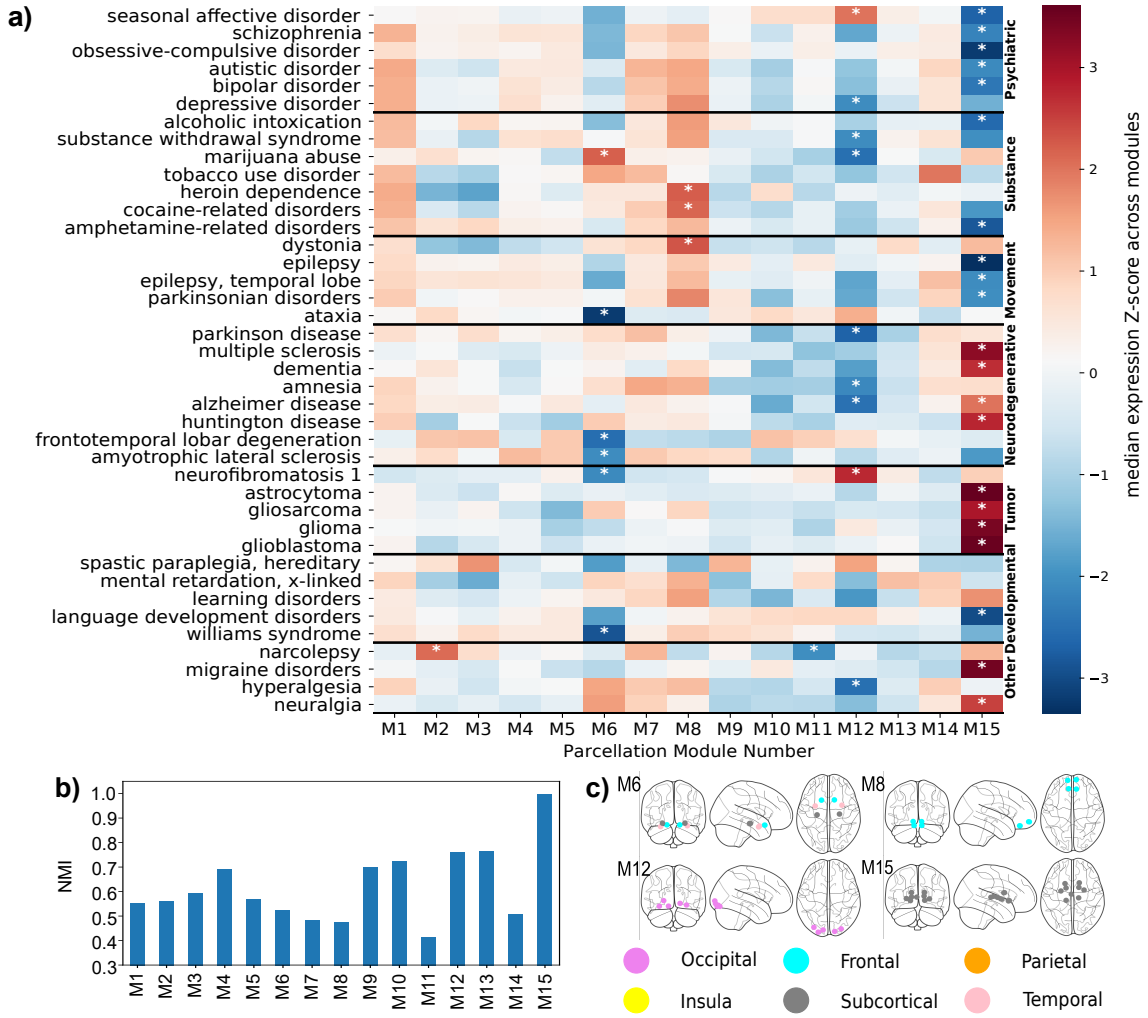

Figure 4: **Transcriptomic expression of major brain-related disorders across different modules**. For the modules of the same partition of the population $\mathcal{G}_F(\tilde{\mathbf{J}})$ illustrated in figure 1-d, we show **a** the across-modules Z-score of the median representation of the genes associated with each disorder. Modules with absolute values higher than 2 (meaning that the disease is significantly over or under-represented) are denoted by a white asterisk. Histogram **b** contains the persistence of the partition's modules across different samplings of the dataset's subjects; the result is expressed in terms of the average match (NMI) of each module with its most similar counterpart in each sampling's partition. Subfigure **c** shows the few modules significantly associated with three or more diseases. Notice the fundamental role of module M15, which is notably present across all samplings.

Invariant Estimator. For all cases, we evaluated cross-modularity, which measures the amount of SFC across different hierarchical module levels. The use of cross-modularity allowed us to i) fine-tune each regularization method to maximize the similarity between the corresponding FC and SC graphs architectures, ii) compare how different inference strategies affect the modular structure of the FC graph at both individual and population levels, and iii) determine an optimal partition of the population FC graph inferred from partial correlation matrices, which we further characterized in relation to brain-related diseases, using transcriptomic data. We have found that the partition maximizing cross-modularity of such FC graph differentiates well between subcortical and cortical structures. Additionally, the high association that occurs in subcortical modes compared to cortical modes highlights, in a way, the relevance of their involvement in most disorders, grouped here into seven categories of diseases (Fig. 4).

In general, and in line with findings from other authors [44], we found that inferring the FC graph from regularized partial correlation, rather than from correlation, enhances its similarity to the SC graph; the novelty is that we found that this occurred hierarchically across the modular structure both at the individual and population levels. In addition, we have found that, at the individual level, it reduces the variation across subjects (Fig. 1).

We also observed that applying a threshold to the corresponding matrix at the percolation cutoff is crucial for the FC graph inferred from partial correlation $\mathcal{G}_F(\tilde{\mathbf{J}})$ to exhibit a high degree of cross-modularity. However, some authors have argued that thresholding is not a principled approach, as it results in a non-positive definite matrix, which does not represent a valid covariance matrix and may not be invertible, thus preventing association with a Gaussian likelihood [44]. Consequently, they recommend using methods such as GLASSO to recover the adjacency matrix of $\mathcal{G}_F$ as a sparse partial correlation (a mathematically valid correlation matrix) whenever sparsity is required. Nevertheless, we confirmed that further thresholding a partial correlation matrix, even after applying GLASSO regularization, leads to $\mathcal{G}_F(\tilde{\mathbf{J}})$ with a significantly enhanced cross-modularity, comparable to that achieved with an Identity Shrinkage strategy (Fig 2).

Therefore, as a major result, we put emphasis on the relevance of thresholding at the percolation value, which helps to systematically eliminate weak connections, allowing the identification of the most robust structural-functional links in the brain. By applying this thresholding method, we aim to remove noise and focus on the core network architecture, which better represents the underlying connectivity patterns. As a matter of fact, we have observed that this approach enables more accurate mapping of the brain's modular organization, capturing both global and local connectivity features, which might be lost in traditional correlation-based analyses.

Additionally, we observed that Group Shrinkage provides the highest cross-modularity values but at the cost of losing subject-level information: the parameters yielding the highest cross-modularity values, often close to or equal to one, suggest that the structure at the individual subject level is closer to the population-level SC than to the same subject's structure. In any case, the noise-cleaning parameters providing the highest cross-modularity are high (towards the biased estimator) also for the methods IS, PCA, and cautPCA. Lastly, we discovered that cautPCA outperforms the Eigenvalue Clipping method, of which cautPCA was introduced in [48] as a variant, even though it is less effective than IS and GLASSO.

As previously noted, retrieving FC from time series data requires regularization; here, it was essential for accurately estimating $\mathcal{G}_F(\tilde{\mathbf{J}})$ for individual subjects, but less critical at the population level. Notice that the optimal amount of regularization crucially depends on the specific features of the time series, such as the ratio $q = \frac{N}{T}$ between the number of observations and the number of ROIs, their spatial correlation or their time correlation. For instance, with longer time series (as it typically occurs for M/EEG datasets), we would expect a reduced need for regularization. In our case, a high time correlation exists that affects the functional time series (Fig. S6), and, consequently, as demonstrated here, regularization for fMRI is highly recommendable.

4 Materials and Methods

4.1 Data

Subject connectivity data were downloaded from [57], previously processed in [31] from raw data from the multimodal dataset *Max Planck Institut Leipzig Mind-Brain-Body Dataset*, commonly referred to as LEMON [68]. These data include the preprocessed resting-state functional magnetic resonance imaging (rs-fMRI) hemodynamic blood oxygen level-dependent (BOLD) time series and the structural connectivity matrices of 136 healthy participants, 98 of them males, aged between 20 and 30 years old, already parcelled in 183 regions of interest (ROIs). The rs-time series include 652 time steps. Details about the preprocessing and the parcellation can be found in [31].

Transcriptomic expression data that contain the expression of genes within each ROI were also obtained from the dataset [31], where they were uploaded as a result of the preprocessing, using the *abagen* tool [69], of the open data from the Allen Human Brain Atlas (AHBA) [65]. In addition, we obtained the data relative to the association between genes and diseases from the archive associated with [66].

4.2 From data to brain Functional and Structural graphs

The SC and FC graphs were recovered from the respective connectivity matrices, interpreted as adjacency matrices: each matrix entrance can be read as the strength of the connection (*link weight*) between the corresponding couple of ROIs (*nodes*). For each subject s, we obtained the SC connectivity matrices from the Diffusion Weighted Imaging (DWI) data in the open dataset [31]. On the other hand, we computed multiple different estimates of the FC's connectivity matrix from the resting state BOLD time series: from their regularized correlation $\mathbf{C}^s_\mu(\alpha)$ and from its inverse $\mathbf{J}^s_\mu(\alpha) \equiv \mathbf{C}^s_\mu(\alpha)^{-1}$, whose elements were next normalized as $\tilde{J}^s_{\mu;ij} = \frac{J^s_{\mu;ij}}{\sqrt{J^s_{\mu;ii} J^s_{\mu;jj}}}$ (*partial correlation*), for each of the regularization methods μ that we took into account (see next section), with regularization parameter α; indeed, the comparison of the different inference procedures for the estimate of FC is a key topic of this work. The "raw," not

regularized, correlation matrix is the sample covariance of the fMRI time series; it corresponds to the sample correlation matrix when the time series are demeaned and standardized (null temporal averages and unit standard deviation) as, in the following, we will assume the data to be. Since there is no straightforward interpretation for negative link weights, we took both $\mathbf{C}$ and $\tilde{\mathbf{J}}$ matrices in absolute value.

Next, we cut both structural and functional matrices $\mathbf{M}$ (where $\mathbf{M}$ contains the DWI data for SC and $M_{ij} = |(\tilde{J}_\mu(\alpha))_{ij}|, |C_\mu(\alpha))_{ij}|$ for FC) at the *percolation threshold*, i.e., we set to zero all matrix entries smaller than the cutoff that would break the corresponding graph into two or more components (for more detail see the dedicated section). In the preliminary analysis, we verified that the way negative values are treated does not significantly affect the results since most of the negative elements are small in absolute value and are consequently removed in the thresholding step. At last, interpreting the thresholded matrices $\mathbf{M}$ as adjacency matrices, a SC graph $\mathcal{G}_S^s$ and different estimates of the FC graph $\mathcal{G}_F^s$ (i.e $\mathcal{G}_F(\mathbf{C}_\mu^s(\alpha))$ and $\mathcal{G}_F(\tilde{\mathbf{J}}_\mu^s(\alpha))$) were obtained for each subject s.

At the population level, we defined the population connectivity matrix as the across-subject median of the single-subject connectivity matrices.

4.3 Regularization methods

Given the single subject time series $\mathbf{X}$ (a $N \times T$ real matrix), demeaned and standardized, if T is finite or $q = N/T$ does not vanish, the sample correlation matrix $\mathbf{E} = \frac{\mathbf{X}\mathbf{X}^T}{T}$ and, even more, its inverse $\mathbf{J} = \mathbf{E}^{-1}$ are not good estimators of the true correlation and precision matrices [52, 70, 48]. Regularizing, or noise-cleaning, consists of proposing a matrix that is as similar as possible to the true correlation matrix (in particular, more than $\mathbf{E}$).

The regularization algorithms that we took into account in this work are far from being a complete list of all the methods present in literature [70]. Our small set of algorithms includes methods from the two most common approaches in neuroscience (i.e., the Linear Shrinkage approach, represented by the Identity and Group Shrinkage, and the sparse estimator approach, such as Graphical Lasso [44]), plus two algorithms that follow a Principal Component approach. All these algorithms were systematically evaluated in [48] on rs-fMRI BOLD signal and synthetic data, in terms of scores such as the element-wise distance $d(\mathbf{J}, \mathbf{J}^v)$ between the inferred precision $\mathbf{J}$ and the true precision $\mathbf{J}^v$, the test-set likelihood, pseudo-likelihood and completion-error (three quantities computed assuming a Normal Probability Distribution $\mathcal{N}(\mathbf{X}|\mathbf{C})$). Most of the regularization algorithms consist in cleaning the spectrum $\mathbf{\Lambda}$ of $\mathbf{E}$ while keeping the eigenvectors $\mathbf{U}$ unchanged (based on the assumption of no prior knowledge of the eigenvectors' structure) so that the correlation matrix cleaned with method μ has the general form $\mathbf{C}_\mu(\alpha) = \mathbf{U}\mathbf{\Lambda}(\alpha)\mathbf{U}^T$:
a,b) **Linear Shrinkage (LS)**, that introduces a ℓ_2 penalty term combining the sample estimator $\mathbf{E}$ with a target matrix $\mathbf{T}$ (independent from the data), such as the identity matrix $\mathbb{I}_N$ (in which case we dub the method **Identity Shrinkage (IS)**) or the across-subject average of the sample correlation matrix $\langle \mathbf{E} \rangle$ (we call this method **Group Shrinkage (GS)**). The cleaned matrix is $\mathbf{C}_{LS}(a) = (1-a)\mathbf{E} + a\mathbf{T}$ [49]. The Linear Shrinkage can have different interpretations. For example, it can be seen as a trade-off between bias and variance, as a shrinkage of the eigenvalues towards their grand mean [49], or also as an optimal Bayes estimator in the context of the Bayesian Random Matrix theory, choosing the minimum mean squared error as loss function and assuming a Gaussian distribution for the data and an Inverse-Wishart as prior distribution for the correlation matrix [70].
c) **Principal Component Analysis (PCA)**, (also known as *Eigenvalues Clipping* [70]) that consists in considering as significant only p eigenvalues of $\mathbf{E}$: $\hat{\lambda}_i = \lambda_i$ if $i \leq p$, else $\hat{\lambda}_i = \frac{\sum_{i<p} \lambda_i}{N-p}$.
d) **Cautious Principal Component Anlysis (cautPCA)**, a simple variant of PCA that was firstly introduced in [48], where it proved to slightly outperform PCA. In this case the spectrum is first changed as $\hat{\lambda}_i = \lambda_p \, \forall i \geq p$ and later rescaled such that $tr\,[\mathbf{C}] = N$.
e) **Graphical Lasso (GLASSO)**; in this case $\mathbf{C}$ and $\mathbf{J}$ are computed by maximization of the log-likelihood minus a ℓ_1 norm penalty term: $\mathbf{J}_{GLASSO}(a) = arg\max_J \{\ln N(\mathbf{X}|\mathbf{J}^{-1}) - a\sum_{i<j} |J_{ij}|\}$ [50]
f) **Optimally Rotationally Invariant Estimator (ORIE)**; first proposed in [71] and later extended in [72], the method consists of correcting the eigenvalues of the sample estimator $\mathbf{E}$ as $\lambda'_k = \frac{\lambda_k}{|1-q+qz_k s(z_k)|^2}$ where $s(z) = Tr(z\mathbb{I}_N - \mathbf{E})^{-1}$, and $z_k = \lambda_k - i\eta$ (i is the imaginary unit) [71, 72, 73, 70, 48]. This estimator minimizes, among all estimators that share the same eigenvectors of the sample correlation, the Euclidean distance from the true correlation matrix in the high-dimensional limit. The parameter η should be small, and such that $N\eta >> 1$ [73]. While a convenient choice is given by $\eta = N^{-1/2}$ [73, 72, 48], this may lead to values that are not small enough if N is not extremely large [73]. Therefore, following [48], we choose to cross-validate this parameter to maximize the validation-set likelihood.

Except for ORIE, the regularization methods that we took into account in this work depend on a tuning parameter that we generically call α, ranging from zero (corresponding to no regularization at all, in which case the correlation matrix estimate equates the sample correlation $\mathbf{E}$) to $\alpha = \alpha_{max}$ (maximum regularization): it coincides with the parameter a

of LS and GLASSO, and with $\frac{N-p}{N}$ for PCA and cautPCA. On the other hand, it is not related to the parameter η of ORIE that, for this reason, we treat differently.

The optimal value α^* is usually set by maximization of some score, for example, the log-likelihood $\ln(P(\mathbf{X}_v|\mathbf{C}_{tr}))$, evaluated, importantly, on a subset of data (the *validation-set* $\mathbf{X}_v$), different from the subset $\mathbf{X}_{tr}$ (*pure-training-set*) used to infer $\mathbf{C}_{tr}$ (so to avoid overfitting). A convenient method for doing this is by searching for the optimal value of α across a grid of values using the general procedure known as k-fold cross-validation (CV): the fraction of the data dedicated to learning a model (overall known as training-set, in contrast to the test-set used for the model's final evaluation) is randomly split into *k* smaller sets, or folds; at each of the *k* iterations, $k-1$ folds are taken as pure-training set, and the remaining one as validation-set; finally, each parameter's score is computed as the average of the iterations scores [48, 67]. Indeed, this is the procedure we followed to compute the value of α_L in figures S3.

Alternatively, some method-specific arguments for the choice of the parameter have been proposed. Ledoit and Wolf provided in [49] an analytical formula for computing the value of a that, under some conditions (such as a large number of observations and features), was proven to minimize the Frobenius distance between the corrected correlation matrix, in the Linear Shrinkage form, and the true one. In the context of PCA, the Minka criterion prescribed selecting the number of components that maximizes the training-set Bayesian evidence [48, 74].

Anyway, in this work, we selected the regularization parameters $\alpha_{\chi_{\mathcal{G}_F}}$ that provide the highest value of cross-modularity of the corresponding FC graph (notice that for each regularization method, there is only one parameter α_L that maximizes the validation-set likelihood of a correlation matrix, but two different parameters that maximize the cross-modularity of $\mathcal{G}_F(\mathbf{C})$ and $\mathcal{G}_F(\tilde{\mathbf{J}})$). This choice is linked to the values of the parameters that maximize the validation-set likelihood α_L, which here appeared anomalously low for most of the methods. This is particularly evident for the Linear Shrinkage methods, where we found α_L of the order of magnitude $10^{-5} - 10^{-4}$, while, on principle, we would expect values of the order of $\sim q \sim 10^{-1}$ (see Fig. S3). We interpreted this phenomenon as an artifact due to the high time correlation of our time series, causing an effective undersampling. Indeed, we checked on a second dataset [75], that the gap between the cross-modularity of $\mathcal{G}_F(\tilde{\mathbf{J}})$ inferred from α_L and $\alpha_{\chi_{\mathcal{G}_F(\tilde{\mathbf{J}})}}$ does not always appear (Fig. S6).

However, regularizing a correlation matrix in the statistical sense (i.e., maximizing the out-of-sample likelihood) does not, in principle, need to coincide with estimating a significant $\mathcal{G}_F(\tilde{\mathbf{J}})$ (inverting and standardizing the correlation and thresholding its absolute value). In fact, cross-modularity, a measure of the distance between $\mathcal{G}_F$ and $\mathcal{G}_S$, and validation-set likelihood, a measure of the prediction power of the correlation matrix from which $\mathcal{G}_F$ is extracted, are two completely different functions, so that optimizing α_L and $\alpha_{\chi_{\mathcal{G}_F}}$ may lead to different results. What are, if any, the conditions for them to coincide is a question that we propose to carefully investigate in a dedicated work. This is particularly evident in the case of the ORIE regularization method: the correlation matrices regularized with this method are among the top-performing algorithms in terms of distribution-based metrics such as test-set likelihood and completion error (Fig. S3), as previously noted in [48] at least for high enough T; yet, the corresponding $\mathcal{G}_F$ fails to accurately retrieve the $\mathcal{G}_S$ architecture.

Finally, the choice of cross-validating the parameter η of ORIE, in apparent contradiction with the approach that we adopted for tuning the other regularization methods, is motivated by the fact that η does not play the same role as the parameter α of the other estimators. In fact, we checked that the cross-modularity values do not vary significantly with η.

4.4 Percolation threshold

As mentioned, we obtained the connectivity graphs from the respective thresholded connectivity matrices. In particular, we thresholded the matrices at the percolation cutoff, meaning that we gradually removed the matrix elements in order of increasing weight, stopping just before the corresponding graph would break into more than one disconnected component. Among the many sparsification methods that have been proposed in neuroscience, this method has the advantage of providing a connected graph whose final density depends on the graph topology. Moreover, in [58], the application of the percolation threshold was shown to provide the optimal balance between the removal of noise and genuine information, maximizing the distance of FC from its randomized counterpart, and it was suggested that its application is critical for the extraction of the large-scale structure from the network.

4.5 Hierarchical Clustering

We compared the architectures of each pair of SC and FC graphs, both at the subject and population levels, based on their hierarchical clusterings.
Hierarchical, or agglomerative, clustering is an unsupervised method for finding communities in N-dimensional observation vectors. The algorithm produces a whole hierarchy of nested data partitions that can be represented as a

dendrogram. More in detail, we have made use of the hierarchical clustering algorithm implemented in the SciPy ([76]) module, with *cosine* metric and *weighted* method, taking as input the 1-D condensed distance matrix built from the connectivity graphs adjacency matrices, after having transformed the matrices elements $\mathbf{M}_{ij}$ as $\exp(-\mathbf{M}_{ij})$.
As a result of this step, we obtained a couple of sets of partitions $P_{\mathcal{G}_F}(m)$ and $P_{\mathcal{G}_S}(m)$ in m modules, $\forall m = 2, \cdots 50$, representing the hierarchical partitions of both Functional and Structural Connectivity graphs.

4.6 Cross-Modularity

To evaluate the similarity of the partitions of the Functional and Structural Connectivity, we introduced a metric that, given the partitions $P_{\mathcal{G}_F}(m)$ and $P_{\mathcal{G}_S}(m)$, simultaneously takes into account their individual quality and their reciprocal similarity, for each number of modules m. Therefore we define the cross-modularity $\chi(m)$ as

$$\chi_{\mathcal{G}_F,\mathcal{G}_S}(m) = \left(Q_{\mathcal{G}_F}(m) \cdot Q_{\mathcal{G}_S}(m) \cdot \nu_{\mathcal{G}_F,\mathcal{G}_S}(m)\right)^{1/3} \tag{2}$$

where $Q_{\mathcal{G}}(m)$ is the quality of partition $P_{\mathcal{G}}(m)$, measured in terms of the Newmann Modularity [59] (we measured it ignoring the links' weights), and $\nu_{\mathcal{G}_F,\mathcal{G}_S}(m)$ is the similarity of partitions $P_{\mathcal{G}_F}(m)$ and $P_{\mathcal{G}_S}(m)$, measured in terms of Adjusted Normalized Mutual Information (NMI) [60]. The cross-modularity is a variation of the homonym metric first introduced in [30]; there, a unique partition $P(m)$ was applied to both SC and FC graphs, and the agreement between the communities of the two graphs was then computed as the average agreement (measured as Sorensen index) between the couple of communities induced on SC and FC by each module of $P(m)$.

More in detail, Newman modularity is a measure of the quality of a particular partition of a graph into modules, defined in [59] as:

$$Q_{\mathcal{G}} = \frac{1}{2L}\sum_{i,j}^{N}(A_{ij} - \frac{d_i d_j}{2L})\delta(c_i, c_j) \tag{3}$$

where L is the number of edges, A the adjacency matrix, d_i the degree of i and $\delta(c_i, c_j)$ is 1 if i and j belong to the same community, else 0. This quantity is proportional to the number of intra-module edges minus its expected number in a network with equal degree sequence and random links. Consequently, modularity approaching 1 indicates a high-quality partition (high intra-cluster and low inter-cluster density), while it would vanish in the case of a random one [62].

Mutual Information (MI) is an information-theoretic tool that measures the amount of information shared by two partitions; its normalized and adjusted version (NMI), where the adjustment discards the matches due to chance, was proposed in [64]:

$$\text{AdjustedIndex} = \frac{\text{Index} - \text{ExpectedIndex}}{\text{MaxIndex} - \text{ExpectedIndex}} \tag{4}$$

so that this quantity vanishes in case of comparison of random partitions and reaches 1 for perfect coherence.

As a consequence, cross-modularity may achieve a maximum value of 1 in the case of high-quality partitions and a perfect match, while it vanishes in the case of unrelated partitions or if at least one of them is random.

In this work, we used the value of the peak of the cross-modularity $\max_m \chi(m)$ as a quality score for evaluating the different estimates of the FC, such as the regularized correlation and precision, for each regularization method, and, within the same method, for computing the best regularization parameter α^*. We have measured Newman modularity using the implementation in [61] and NMI using the implementation in [67].

4.7 Spectral distance

In addition to the single-subject cross-modularity, we also computed the spectral distance [77] of the connectivity graphs' adjacency matrices between each couple of subjects for both the SC graph and every estimate of the FC graph. In general, the spectral distance measures the difference between a couple of matrices in terms of the difference between their eigenvalues:

$$d(\mathbf{M}^{(1)}, \mathbf{M}^{(2)}) = \sum_{i}^{N} |\lambda_i^{(1)} - \lambda_i^{(2)}| \tag{5}$$

where $\lambda_i^{(j)}$, with $j = 1, 2$ and $i = 1, \cdots, N$ are the matrices $\mathbf{M}^{(1)}$ and $\mathbf{M}^{(2)}$ eigenvalues. To make fair the comparison between the spectral distance of the subjects' FC (which are bounded, in absolute value, between 0 and 1) and SC (which are expressed as the number of white matter streamlines) adjacency matrices, we normalized the values of the (thresholded) $\mathcal{G}_S$ adjacency matrix as $\frac{2\arctan(\cdot)}{\pi}$.

4.8 Characterization of the optimal FC brain partition in terms of brain-related disorders

Once obtained the optimal (in the sense of higher cross-modularity) partition of the optimal estimate of the population FC graph, we characterized its modules depending on their participation in 40 major brain-related disorders from 7 disease groups (Psychiatric disorders, Substance abuse, Movement disorders, Neurodegenerative diseases, Tumor conditions, Developmental disorders, and Others). Our method follows the procedure adopted in [31], work associated with the Zenodo dataset from which we extracted the transcriptomic data. For each disease and for each module, we computed the disease expression value as the median value across the genes associated with the disease and across the nodes contained in the module. The heatmap in figure 4 shows the across-modules Z-scores of the diseases expressions, with absolute values higher than 2 marked with an asterisk, and diseased grouped by the WHO categories (also obtained from archive [66]).

4.9 Robustness check of the optimal FC brain partition

We have checked the robustness of the optimal partition of the best estimate of the population FC graph by comparing it with different estimates inferred from many random sub-samplings of the dataset: 20 times, we randomly split the subject set into three groups, we computed from each group an estimate of the population $\mathcal{G}_F(\tilde{\mathbf{J}})$, with regularization method IS and regularization parameter $\alpha = 0.8$ and we partitioned it in 15 modules. Thus, we computed the match of each partition with the whole-dataset one (i.e., the partition shown in Fig 1) in terms of NMI, finding an overall score of 0.80. We additionally measured the robustness of each whole-dataset partition's module as its across-sampling average NMI similarity to the most similar module in the subsampling partition. The results, shown in figure 4-*b*, reveal that some modules, module M15 *in primis*, are particularly robust.

5 Technical Terms

Sample correlation matrix $\mathbf{E}$: the unbiased or "raw" estimator of the true correlation matrix. It is computed from the standardized time series contained in the $N \times T$ matrix $\mathbf{X}$ as $E = \frac{XX^T}{T}$.
Regularized Correlation Matrix $\mathbf{C}_\mu(\alpha)$: a de-noised estimate of the true correlation matrix of the standardized time series, using method μ and tuning parameter α.
Regularized Precision Matrix $\mathbf{J}_\mu(\alpha)$: the inverse of the regularized correlation matrix: $\mathbf{J}_\mu(\alpha) = \mathbf{C}_\mu^{-1}(\alpha)$.
Regularized Partial Correlation Matrix $\tilde{\mathbf{J}}_\mu(\alpha)$: the normalized version of the Precision matrix. In a Gaussian approximation, it accounts for direct dependencies only.
Structural Connectivity (SC) graph $\mathcal{G}_S$: the graph representing the anatomical connections between the brain areas. It is obtained from the thresholded Density Weighted Imaging data.
Functional Connectivity (FC) graph $\mathcal{G}_F$: the graph representing the activity patterns between brain areas. We use the notation $\mathcal{G}_F(\mathbf{C}_\mu(\alpha))$ or $\mathcal{G}_F(\tilde{\mathbf{J}}_\mu(\alpha))$ to specify whether the graph was inferred from the regularized correlation $\mathbf{C}_\mu(\alpha)$ or from the regularized partial correlation $\tilde{\mathbf{J}}_\mu(\alpha)$ matrix, taken in absolute value and thresholded.

References

[1] Karl J Friston. Functional and Effective Connectivity in Neuroimaging: A Synthesis. *Human Brain Mapping*, 2(1-2):56–78, 1994.

[2] Christopher J Honey, Rolf Kötter, Michael Breakspear, and Olaf Sporns. Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *Proceedings of the National Academy of Sciences*, 104(24):10240–10245, 2007.

[3] Martijn P. Van Den Heuvel, René C.W. Mandl, René S. Kahn, and Hilleke E. Hulshoff Pol. Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain. *Human Brain Mapping*, 30(10):3127–3141, 10 2009.

[4] Michael D. Greicius, Kaustubh Supekar, Vinod Menon, and Robert F. Dougherty. Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cerebral Cortex*, 19(1):72–78, 1 2009.

[5] C J Honey, O Sporns, L Cammoun, X Gigandet, J P Thiran, R Meuli, and P Hagmann. Predicting human resting-state functional connectivity from structural connectivity. *PNAS*, 2009.

[6] Christopher J. Honey, Jean Philippe Thivierge, and Olaf Sporns. Can structure predict function in the human brain?, 9 2010.

[7] Hae-Jeong Park and Karl Friston. Structural and Functional Brain Networks: From Connections to Cognition. *REVIEW SUMMARY*, 2013.

[8] Joaquin Goni, Martijn P. Van Den Heuvel, Andrea Avena-Koenigsberger, Nieves Velez De Mendizabal, Richard F. Betzel, Alessandra Griffa, Patric Hagmann, Bernat Corominas-Murtra, Jean Philippe Thiran, and Olaf Sporns. Resting-brain functional connectivity predicted by analytic measures of network communication. *Proceedings of the National Academy of Sciences of the United States of America*, 111(2):833–838, 2014.

[9] Olaf Sporns and Richard F. Betzel. Modular brain networks. *Annual Review of Psychology*, 67:613–640, 1 2016.

[10] Kanika Bansal, Johan Nakuci, and Sarah Feldt Muldoon. Personalized brain network models for assessing structure–function relationships, 10 2018.

[11] Bertha Vázquez-Rodríguez, Laura E. Suárez, Ross D. Markello, Golia Shafiei, Casey Paquola, Patric Hagmann, Martijn P. Van Den Heuvel, Boris C. Bernhardt, R. Nathan Spreng, and Bratislav Misic. Gradients of structure–function tethering across neocortex. *Proceedings of the National Academy of Sciences of the United States of America*, 116(42):21219–21227, 10 2019.

[12] Laura E. Suárez, Ross D. Markello, Richard F. Betzel, and Bratislav Misic. Linking Structure and Function in Macroscale Brain Networks, 4 2020.

[13] T. Sarwar, Y. Tian, B. T.T. Yeo, K. Ramamohanarao, and A. Zalesky. Structure-function coupling in the human connectome: A machine learning approach. *NeuroImage*, 226, 2 2021.

[14] Farnaz Zamani Esfahlani, Joshua Faskowitz, Jonah Slack, Bratislav Mišić, and Richard F. Betzel. Local structure-function relationships in human brain networks across the lifespan. *Nature Communications*, 13(1), 12 2022.

[15] Gustavo Deco, Viktor Jirsa, A R Mcintosh, Olaf Sporns, and Rolf Kö. Key role of coupling, delay, and noise in resting brain fluctuations. Technical report, 2024.

[16] Carmen Alonso-Montes, Ibai Diez, Lakhdar Remaki, Iñaki Escudero, Beatriz Mateos, Yves Rosseel, Daniele Marinazzo, Sebastiano Stramaglia, and Jesus M. Cortes. Lagged and instantaneous dynamical influences related to brain structural connectivity. *Frontiers in Psychology*, 6, 7 2015.

[17] Amor T. A, Russo R, Diez I, Bharath P, Zirovich M, Stramaglia S, Cortes J. M, de Arcangelis L, and Chialvo D. R. Extreme brain events: Higher-order statistics of brain resting activity and its relation with structural connectivity. *Europhysics Letters*, 111(6), 10 2015.

[18] S. Stramaglia, M. Pellicoro, L. Angelini, E. Amico, H. Aerts, J. M. Cortés, S. Laureys, and D. Marinazzo. Ising model with conserved magnetization on the human connectome: Implications on the relation structure-function in wakefulness and anesthesia. *Chaos*, 27(4), 4 2017.

[19] P. Hagmann, O. Sporns, N. Madan, L. Cammoun, R. Pienaar, V. J. Wedeen, R. Meuli, J. P. Thiran, and P. E. Grant. White matter maturation reshapes structural connectivity in the late developing human brain. *Proceedings of the National Academy of Sciences of the United States of America*, 107(44):19067–19072, 11 2010.

[20] Karla Batista-García-Ramó and Caridad Ivette Fernández-Verdecia. What we know about the brain structure-function relationship, 4 2018.

[21] Panagiotis Fotiadis, Linden Parkes, Kathryn A. Davis, Theodore D. Satterthwaite, Russell T. Shinohara, and Dani S. Bassett. Structure–function coupling in macroscale human brain networks, 10 2024.

[22] Alessandra Griffa, Enrico Amico, Raphaël Liégeois, Dimitri Van De Ville, and Maria Giulia Preti. Brain structure-function coupling provides signatures for task decoding and individual fingerprinting. *NeuroImage*, 250, 4 2022.

[23] Zijin Gu, Keith Wakefield Jamison, Mert Rory Sabuncu, and Amy Kuceyeski. Heritability and interindividual variability of regional structure-function coupling. *Nature Communications*, 12(1), 12 2021.

[24] Maria Giulia Preti and Dimitri Van De Ville. Decoupling of brain function from structure reveals regional behavioral specialization in humans. *Nature Communications*, 10(1), 12 2019.

[25] Rui Cao, Xin Wang, Yuan Gao, Ting Li, Hui Zhang, Waqar Hussain, Yunyan Xie, Jing Wang, Bin Wang, and Jie Xiang. Abnormal Anatomical Rich-Club Organization and Structural–Functional Coupling in Mild Cognitive Impairment and Alzheimer's Disease. *Frontiers in Neurology*, 11, 2 2020.

[26] Jonathan Tay, Marco Düring, Esther M.C. van Leijsen, Mayra I. Bergkamp, David G. Norris, Frank Erik de Leeuw, Hugh S. Markus, and Anil M. Tuladhar. Network structure-function coupling and neurocognition in cerebral small vessel disease. *NeuroImage: Clinical*, 38, 1 2023.

[27] Bo Hua, Xin Ding, Minghua Xiong, Fanyu Zhang, Yi Luo, Jurong Ding, and Zhongxiang Ding. Alterations of functional and structural connectivity in patients with brain metastases. *PLoS ONE*, 15(5), 5 2020.

[28] Angeliki Zarkali, Peter McColgan, Louise Ann Leyland, Andrew J. Lees, Geraint Rees, and Rimona S. Weil. Organisational and neuromodulatory underpinnings of structural-functional connectivity decoupling in patients with Parkinson's disease. *Communications Biology*, 4(1), 12 2021.

[29] Guangyao Liu, Weihao Zheng, Hong Liu, Man Guo, Laiyang Ma, Wanjun Hu, Ming Ke, Yu Sun, Jing Zhang, and Zhe Zhang. Aberrant dynamic structure–function relationship of rich-club organization in treatment-naïve newly diagnosed juvenile myoclonic epilepsy. *Human Brain Mapping*, 43(12):3633–3645, 8 2022.

[30] Ibai Diez, Paolo Bonifazi, Iñaki Escudero, Beatriz Mateos, Miguel A Muñoz, Sebastiano Stramaglia, and Jesus M Cortes. A novel brain partition highlights the modular skeleton shared by structure and function OPEN. *Nature Publishing Group*, 2015.

[31] Antonio Jimenez-Marin, Ibai Diez, Asier Erramuzpe, Sebastiano Stramaglia, Paolo Bonifazi, and Jesus M. Cortes. Open datasets and code for multi-scale relations on structure, function and neuro-genetics in the human brain. *Scientific Data*, 11(1), 12 2024.

[32] Richard F. Betzel, Maxwell A. Bertolero, Evan M. Gordon, Caterina Gratton, Nico U.F. Dosenbach, and Danielle S. Bassett. The community structure of functional brain networks exhibits scale-specific patterns of inter- and intra-subject variability. *NeuroImage*, 202, 11 2019.

[33] Maria Grazia Puxeddu, Joshua Faskowitz, Olaf Sporns, Laura Astolfi, and Richard F. Betzel. Multi-modal and multi-subject modular organization of human brain networks. *NeuroImage*, 264, 12 2022.

[34] Peter Fransson and Guillaume Marrelec. The precuneus/posterior cingulate cortex plays a pivotal role in the default mode network: Evidence from a partial correlation network analysis. *NeuroImage*, 42(3):1178–1184, 9 2008.

[35] Shuai Huang, Jing Li, Liang Sun, Jieping Ye, Adam Fleisher, Teresa Wu, Kewei Chen, and Eric Reiman. Learning brain connectivity of Alzheimer's disease by sparse inverse covariance estimation. *NeuroImage*, 50(3):935–949, 4 2010.

[36] Srikanth Ryali, Tianwen Chen, Kaustubh Supekar, and Vinod Menon. Estimation of functional connectivity in fMRI data using stability selection-based sparse partial correlation with elastic net penalty. *NeuroImage*, 59(4):3852–3861, 2 2012.

[37] Verónica Mäki-Marttunen, Ibai Diez, Jesus M. Cortes, Dante R. Chialvo, and Mirta Villarreal. Disruption of transfer entropy and inter-hemispheric brain functional connectivity in patients with disorder of consciousness. *Frontiers in Neuroinformatics*, 7, 11 2013.

[38] Daniele Poli, Vito Paolo Pastore, Sergio Martinoia, and Paolo Massobrio. From functional to structural connectivity using partial correlation in neuronal assemblies. *Journal of Neural Engineering*, 13(2):026023, feb 2016.

[39] Mehdi Rahim, Bertrand Thirion, and Gaël Varoquaux. Population shrinkage of covariance (PoSCE) for better individual brain functional-connectivity estimation. *Medical Image Analysis*, 54:138–148, 5 2019.

[40] Kirsten L Peterson, Ruben Sanchez-Romero, Ravi D Mill, and Michael W Cole. Regularized partial correlation provides reliable functional connectivity estimates while correcting for widespread confounding. *biorXiv*, 2023.

[41] Raymond Salvador, John Suckling, Christian Schwarzbauer, and Ed Bullmore. Undirected graphs of frequency-dependent functional connectivity in whole brain networks. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1457):937–946, 2005.

[42] Raymond Salvador, John Suckling, Martin R. Coleman, John D. Pickard, David Menon, and Ed Bullmore. Neurophysiological architecture of functional magnetic resonance images of human brain. *Cerebral Cortex*, 15(9):1332–2342, 9 2005.

[43] Guillaume Marrelec, Alexandre Krainik, Hugues Duffau, Mélanie Pélégrini-Issac, Stéphane Lehéricy, Julien Doyon, and Habib Benali. Partial correlation for functional brain interactivity investigation in functional MRI. *NeuroImage*, 32(1):228–237, 2006.

[44] G Varoquaux. Estimating brain functional connectivity and its variations from fMRI, 2019.

[45] Raphael Liégeois, Augusto Santos, Vincenzo Matta, Dimitri Van De Ville, and Ali H. Sayed. Revisiting correlation-based functional connectivity and its relationship with structural connectivity. *Network Neuroscience*, 4(4):1235–1251, 11 2020.

[46] Matthew R. Brier, Anish Mitra, John E. McCarthy, Beau M. Ances, and Abraham Z. Snyder. Partial covariance based functional connectivity computation using Ledoit-Wolf covariance regularization. *NeuroImage*, 121:29–38, 11 2015.

[47] Usama Pervaiz, Diego Vidaurre, Mark W. Woolrich, and Stephen M. Smith. Optimising network modelling methods for fMRI. *NeuroImage*, 211, 5 2020.

[48] Miguel Ibáñez-Berganza, Carlo Lucibello, Francesca Santucci, Tommaso Gili, and Andrea Gabrielli. Noise cleaning the precision matrix of short time series. *Physical Review E*, 108(2), 8 2023.

[49] Olivier Ledoit and Michael Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2):365–411, 2004.

[50] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 7 2008.

[51] Olivier Ledoit and Michael Wolf. Honey, I Shrunk the Sample Covariance Matrix. *The Journal of Portfolio Management*, 30(4):110 – 119, 2004.

[52] Gaël Varoquaux and R. Cameron Craddock. Learning and comparing functional connectomes across subjects. *NeuroImage*, 80:405–415, 10 2013.

[53] Shuai Huang, Jing Li, Liang Sun, Jun Liu, Teresa Wu, Kewei Chen, Adam Fleisher, Eric Reiman, and Jieping Ye. Learning Brain Connectivity of Alzheimer's Disease from Neuroimaging Data. Curran Associates, Inc., 2009.

[54] Amanda F. Mejia, Mary Beth Nebel, Anita D. Barber, Ann S. Choe, James J. Pekar, Brian S. Caffo, and Martin A. Lindquist. Improved estimation of subject-level functional connectivity using full and partial correlation with empirical Bayes shrinkage. *NeuroImage*, 172:478–491, 5 2018.

[55] Stephen M. Smith, Karla L. Miller, Gholamreza Salimi-Khorshidi, Matthew Webster, Christian F. Beckmann, Thomas E. Nichols, Joseph D. Ramsey, and Mark W. Woolrich. Network modelling methods for FMRI. *NeuroImage*, 54(2):875–891, 1 2011.

[56] Gaël Varoquaux, Alexandre Gramfort, Jean Baptiste Poline, and Bertrand Thirion. Brain covariance selection: better individual functional connectivity models using population prior. *Advances in neural information processing systems*, 23, 8 2010.

[57] Antonio Jimenez-Marin, Ibai Diez, Asier Erramuzpe, Sebastiano Stramaglia, Paolo Bonifazi, and Jesus Cortes, 2023.

[58] Carlo Nicolini, Giulia Forcellini, Ludovico Minati, and Angelo Bifone. Scale-resolved analysis of brain functional connectivity networks with spectral entropy. *NeuroImage*, 211, 5 2020.

[59] M E J Newman. Modularity and community structure in networks. Technical report, 2006.

[60] Nguyen Xuan Vinh, Julien Epps, and James Bailey. Information Theoretic Measures for Clusterings Comparison: Variants, Properties, Normalization and Correction for Chance. *Journal of Machine Learning Research*, 11:2837–2854, 2010.

[61] Aric A Hagberg, Daniel A Schult, and Pieter J Swart. Exploring Network Structure, Dynamics, and Function using NetworkX. In Ga\"el Varoquaux, Travis Vaught, and Jarrod Millman, editors, *Proceedings of the 7th Python in Science Conference*, pages 11–15, Pasadena, CA USA, 2008.

[62] P. Van Mieghem, X. Ge, P. Schumm, S. Trajanovski, and H. Wang. Spectral graph analysis of modularity and assortativity. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, 82(5), 11 2010.

[63] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[64] Simone Romano, Nguyen Xuan Vinh, James Bailey, and Karin Verspoor. Adjusting for Chance Clustering Comparison Measures. *Journal of Machine Learning Research*, 17:1–32, 2016.

[65] Michael J. Hawrylycz, Ed S. Lein, Angela L. Guillozet-Bongaarts, Elaine H. Shen, Lydia Ng, Jeremy A. Miller, Louie N. Van De Lagemaat, Kimberly A. Smith, Amanda Ebbert, Zackery L. Riley, Chris Abajian, Christian F. Beckmann, Amy Bernard, Darren Bertagnolli, Andrew F. Boe, Preston M. Cartagena, M. Mallar Chakravarty, Mike Chapin, Jimmy Chong, Rachel A. Dalley, Barry David Daly, Chinh Dang, Suvro Datta, Nick Dee, Tim A. Dolbeare, Vance Faber, David Feng, David R. Fowler, Jeff Goldy, Benjamin W. Gregor, Zeb Haradon, David R. Haynor, John G. Hohmann, Steve Horvath, Robert E. Howard, Andreas Jeromin, Jayson M. Jochim, Marty Kinnunen, Christopher Lau, Evan T. Lazarz, Changkyu Lee, Tracy A. Lemon, Ling Li, Yang Li, John A. Morris, Caroline C. Overly, Patrick D. Parker, Sheana E. Parry, Melissa Reding, Joshua J. Royall, Jay Schulkin, Pedro Adolfo Sequeira, Clifford R. Slaughterbeck, Simon C. Smith, Andy J. Sodt, Susan M. Sunkin, Beryl E. Swanson, Marquis P. Vawter, Derric Williams, Paul Wohnoutka, H. Ronald Zielke, Daniel H. Geschwind, Patrick R. Hof, Stephen M. Smith, Christof Koch, Seth G.N. Grant, and Allan R. Jones. An anatomically comprehensive atlas of the adult human brain transcriptome. *Nature*, 489(7416):391–399, 9 2012.

[66] Yashar Zeighami, Trygve E. Bakken, Thomas Nickl-Jockschat, Zeru Peterson, Anil G. Jegga, Jeremy A. Miller, Jay Schulkin, Alan C. Evans, Ed S. Lein, and Michael Hawrylycz. A comparison of anatomic and cellular transcriptome structures across 40 human brain diseases. *PLoS Biology*, 21(4), 4 2023.

[67] 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. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[68] Anahit Babayan, Miray Erbey, Deniz Kumral, Janis D. Reinelt, Andrea M.F. Reiter, Josefin Röbbig, H. Lina Schaare, Marie Uhlig, Alfred Anwander, Pierre Louis Bazin, Annette Horstmann, Leonie Lampe, Vadim V. Nikulin, Hadas Okon-Singer, Sven Preusser, André Pampel, Christiane S. Rohr, Julia Sacher, Angelika Thöne-Otto, Sabrina Trapp, Till Nierhaus, Denise Altmann, Katrin Arelin, Maria Blöchl, Edith Bongartz, Patric Breig, Elena Cesnaite, Sufang Chen, Roberto Cozatl, Saskia Czerwonatis, Gabriele Dambrauskaite, Maria Dreyer, Jessica Enders, Melina Engelhardt, Marie Michele Fischer, Norman Forschack, Johannes Golchert, Laura Golz, C. Alexandrina Guran, Susanna Hedrich, Nicole Hentschel, Daria I. Hoffmann, Julia M. Huntenburg, Rebecca Jost, Anna Kosatschek, Stella Kunzendorf, Hannah Lammers, Mark E. Lauckner, Keyvan Mahjoory, Ahmad S. Kanaan, Natacha Mendes, Ramona Menger, Enzo Morino, Karina Näthe, Jennifer Neubauer, Handan Noyan, Sabine Oligschläger, Patricia Panczyszyn-Trzewik, Dorothee Poehlchen, Nadine Putzke, Sabrina Roski, Marie Catherine Schaller, Anja Schieferbein, Benito Schlaak, Robert Schmidt, Krzysztof J. Gorgolewski, Hanna Maria Schmidt, Anne Schrimpf, Sylvia Stasch, Maria Voss, Annett Wiedemann, Daniel S. Margulies, Michael Gaebler, and Arno Villringer. Data descriptor: A mind-brain-body dataset of MRI, EEG, cognition, emotion, and peripheral physiology in young and old adults. *Scientific Data*, 6, 2019.

[69] Ross D Markello, Aurina Arnatkeviciute, Jean-Baptiste Poline, Ben D Fulcher, Alex Fornito, and Bratislav Misic. Standardizing workflows in imaging transcriptomics with the abagen toolbox. *eLIfe*, 2021.

[70] Joël Bun, Jean Philippe Bouchaud, and Marc Potters. Cleaning large correlation matrices: Tools from Random Matrix Theory, 1 2017.

[71] Olivier Ledoit and Sandrine Péché. Eigenvectors of some large sample covariance matrix ensembles. *Probability Theory and Related Fields*, 151(1):233–264, 10 2011.

[72] Joel Bun, Romain Allez, Jean Philippe Bouchaud, and Marc Potters. Rotational invariant estimator for general noisy matrices. *IEEE Transactions on Information Theory*, 62(12):7475–7490, 12 2016.

[73] Joel Bun, Jean-Philippe Bouchaud, and Marc Potters. My beautiful laundrette: Cleaning correlation matrices for portfolio optimization, 05 2016.

[74] Thomas P Minka. Automatic choice of dimensionality for PCA. MIT Press, 2000.

[75] Rossana Mastrandrea, Andrea Gabrielli, Fabrizio Piras, Gianfranco Spalletta, Guido Caldarelli, and Tommaso Gili. Organization and hierarchy of the human functional brain network lead to a chain-like core OPEN. 2017.

[76] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.

[77] Irena Jovanović and Zoran Stanić. Spectral distances of graphs. *Linear Algebra and Its Applications*, 436(5):1425–1435, 3 2012.

6 Supporting Information

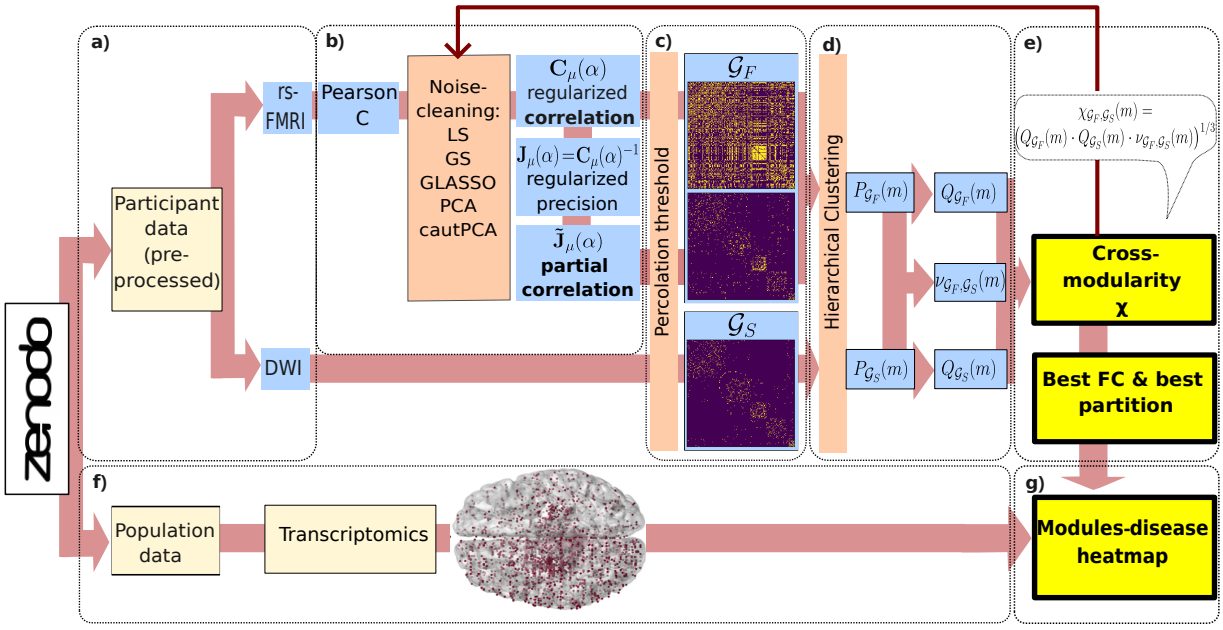

Figure S1: **Methodological sketch and pipeline**. **a)** Preprocessed data of 136 healthy subjects have been obtained from the open dataset [31], at [57]. Data include, for each subject, a Diffusion Weighted Imaging (DWI) matrix of $N = 183$ brain regions, and rs-fMRI BOLD time series of N brain regions and $T = 652$ time steps. **b)** Each subject FC may be estimated as the (regularized) Pearson correlation matrix $\mathbf{C}_\mu(\alpha)$ of the time series, where μ stands for the specific regularization method, and α for its tuning parameter, or as the partial correlation $\tilde{\mathbf{J}}$, the standardized version of the correlation inverse $\mathbf{J}_\mu(\alpha) = \mathbf{C}_\mu^{-1}(\alpha)$. **c)** We cut both functional and structural matrices, taken in absolute value, at the percolation threshold so that the corresponding $\mathcal{G}_F$ and $\mathcal{G}_S$ are sparse graphs. **d)** We computed the hierarchical clustering of both $\mathcal{G}_F$ and $\mathcal{G}_S$, obtaining a set of nested partitions $P_{\mathcal{G}_F}(m)$ and $P_{\mathcal{G}_S}(m)$, in m modules, for $m = 1, \cdots, N$, and, subsequently, their quality ($Q_{\mathcal{G}_F}(m)$ and $Q_{\mathcal{G}_S}(m)$) and their agreement $\nu_{\mathcal{G}_F,\mathcal{G}_S}(m)$. **e)** These three measures allow us to compute the cross-modularity $\chi_{\mathcal{G}_F,\mathcal{G}_S}(m) = (Q_{\mathcal{G}_F}(m) \cdot Q_{\mathcal{G}_S}(m) \cdot \nu_{\mathcal{G}_F,\mathcal{G}_S}(m))^{1/3}$, whose maximum can be used to tune the regularization methods and to identify a meaningful partition of the best estimate of the population FC. **f)** We made use of transcriptomic data extracted from [31] and data relative to the association between genes and a set of diseases extracted from [66] to evaluate the association between each of the modules of the optimal partition of the optimal estimation of FC and these diseases.

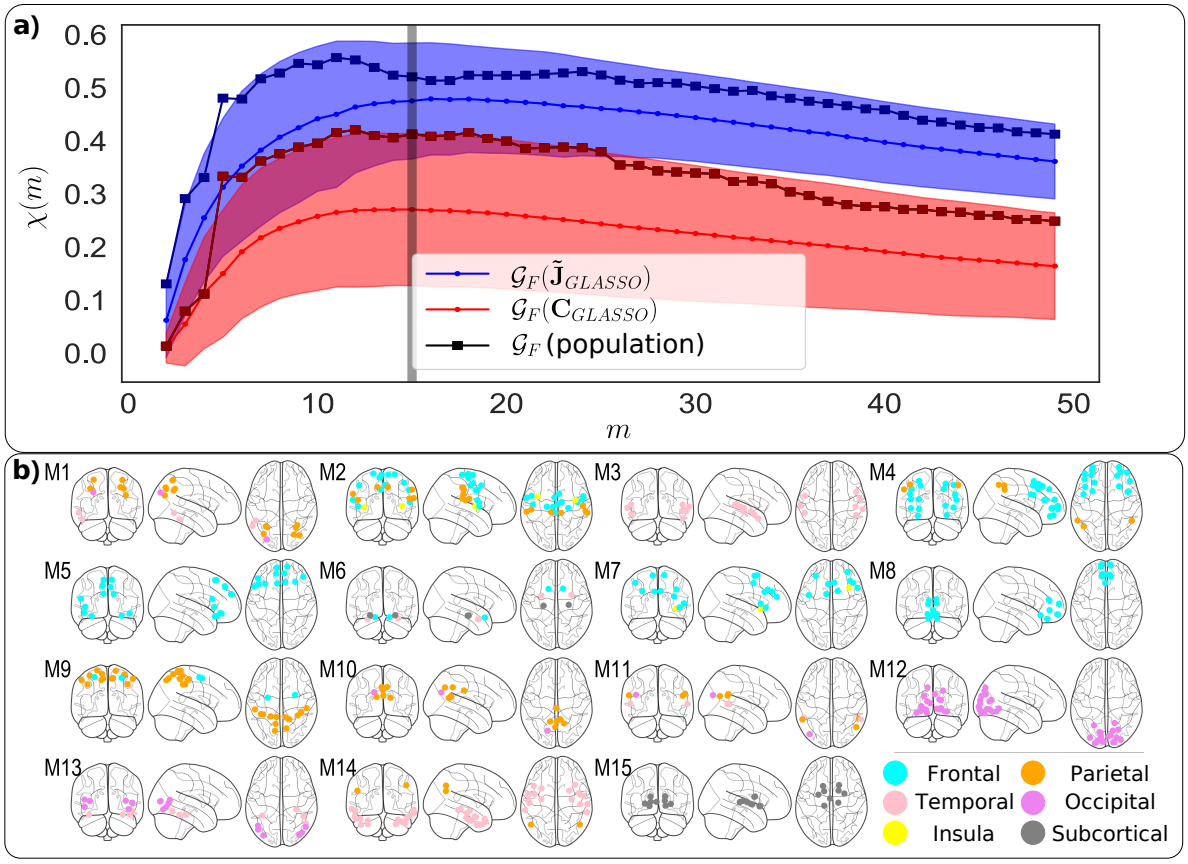

Figure S2: Analogous of figure 1, using **GLASSO as regularization strategy**. Most modules exhibit a remarkable similarity with the modules obtained regularizing with IS, with some of them (modules 6 and 15) being identical. Indeed, the two partitions have an NMI overlap score of 0.68 (on a scale where 1 stands for a perfect match and 0 is the expected score of two random partitions).

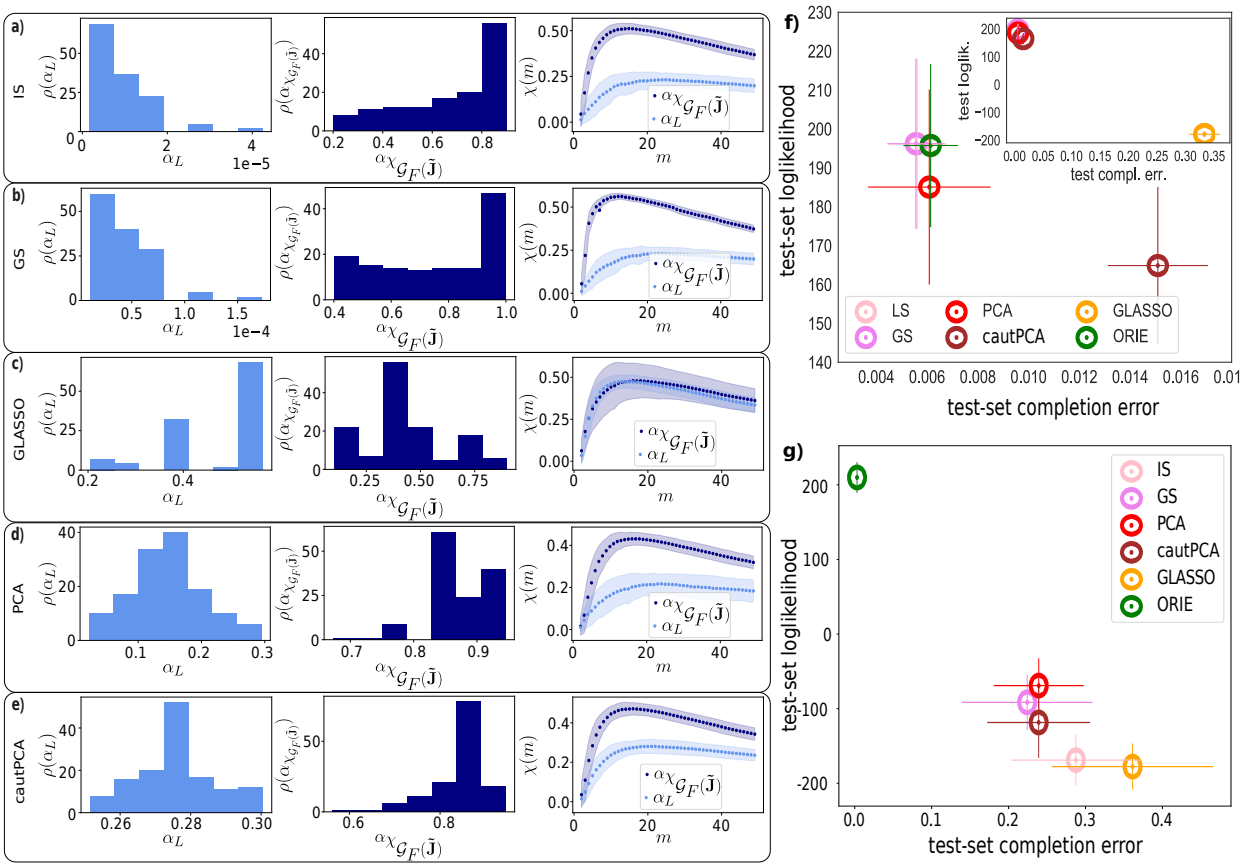

Figure S3: **Regularization tuning parameter**. This figure reveals a discrepancy between the regularization parameters α_L that maximize the validation-set likelihood and $\alpha_{\chi_{\mathcal{G}_F(\tilde{\mathbf{J}}_\mu)}}$ that maximize the cross-modularity of $\mathcal{G}_F(\tilde{\mathbf{J}}_\mu)$, at the single subject level, for all regularization methods μ (about the notation used in methods, the generic parameter α coincides with parameter a of the LS family and GLASSO, and with $\frac{N-p}{N}$ of the PCA family), except for GLASSO. We interpret the discrepancy as linked to anomalously small values of α_L due to the high time correlation of the time series, which induces effective undersampling. The first and second columns of rows **a-f** contain, respectively, the across-subject distributions of α_L and $\alpha_{\chi_{\mathcal{G}_F(\tilde{\mathbf{J}})}}$, while the third column contains the cross-modularity of the single subject $\mathcal{G}_F(\tilde{\mathbf{J}}_\mu(\alpha))$, obtained from the partial correlation matrix, in absolute value and thresholded at percolation, inferred choosing as regularization parameter α_L (light blue) and α_χ (dark blue). For the sake of completeness, we reported in the last row the scores of the regularized correlation (and hence precision) matrices in terms of test-set likelihood and completion error, with regularization parameter chosen as the one maximizing the validation-set likelihood (figure **f**) and the cross-modularity of $\mathcal{G}_F(\tilde{\mathbf{J}})$ (figure **g**).

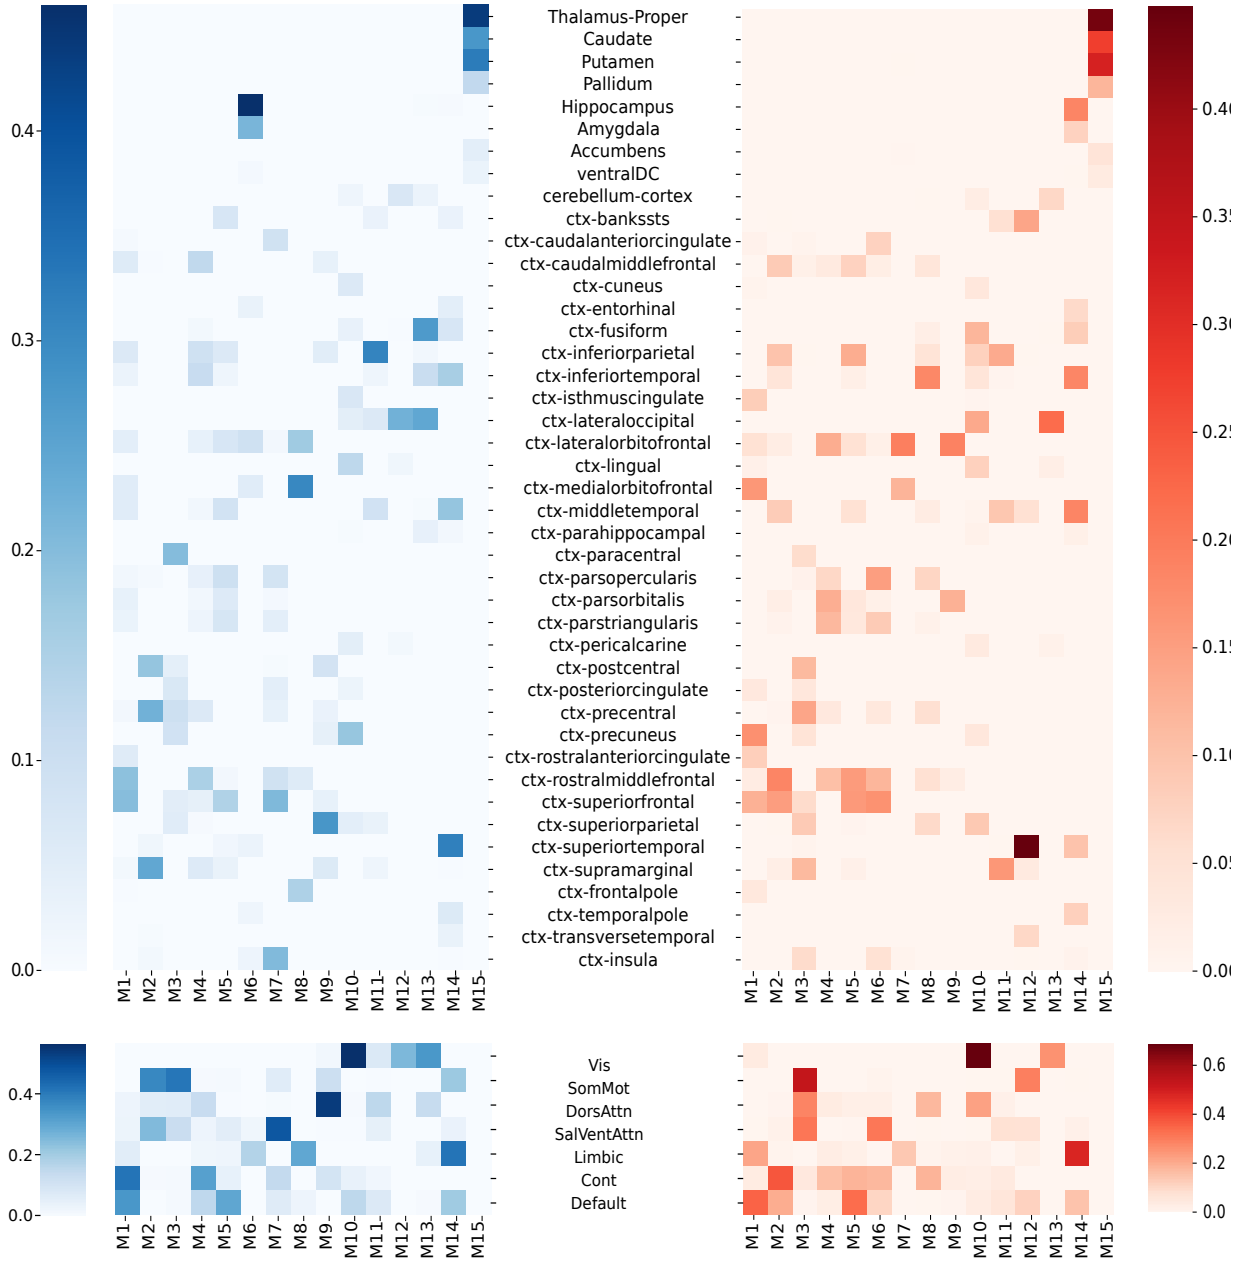

Figure S4: **Overlap between modules and brain areas**. This image shows the overlap between each module of both the population $\mathcal{G}_F(\mathbf{C})$ (red) and $\mathcal{G}_F(\tilde{\mathbf{J}})$ (blue), regularized with IS (in the case of $\mathcal{G}_F(\tilde{\mathbf{J}})$ the modules are the same as those shown in figure 1), the Desikan-Killany structural areas (subfigures *a*), and the Yeo functional areas (subfigures *b*). The overlap is measured in terms of the Jaccard index between each couple module-area's volumes. Notice the increased overlap of the $\mathcal{G}_F(\mathbf{C})$ with the Yeo functional areas.

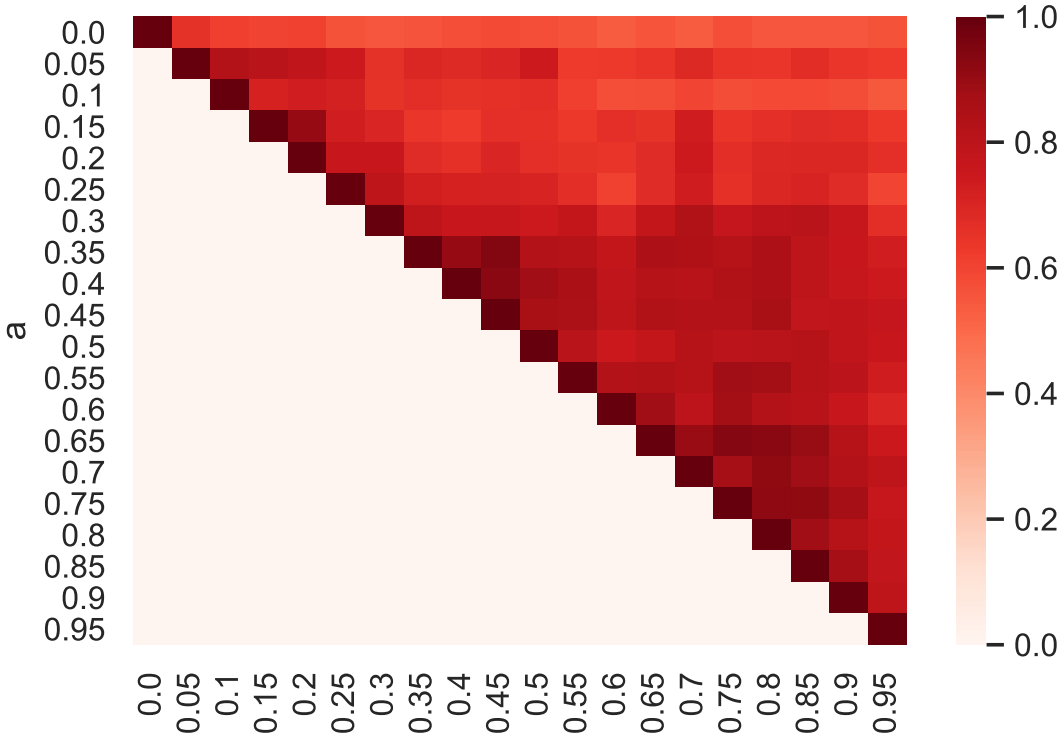

Figure S5: **Stability of the 15-modules-partition** of the population $\mathcal{G}_F(\tilde{\mathbf{J}}_{\mathbf{IS}}(\mathbf{a}))$, for different regularization parameters a. The heatmap contains the match (in terms of NMI) between the partitions for each couple of values of a.

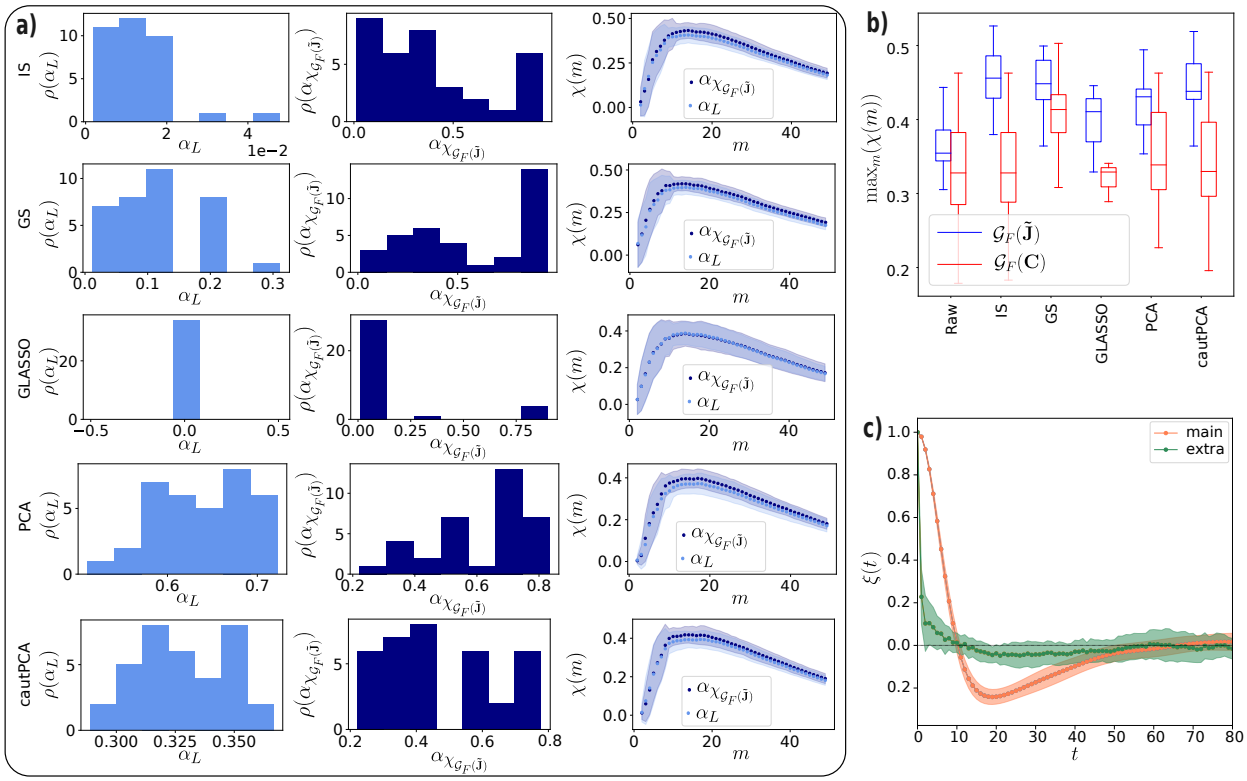

Figure S6: **Regularization tuning parameters** (subfigure **a**) and maximum cross-modularity (subfigure **b**) computed on a **supplementary extra dataset**, and time-correlations of the main and extra datasets (subfigure **c**). The extra dataset is composed of 35 subjects, with 116 ROIs and rs-BOLD time series of 180 time steps. Coherently with our main results, the cross-modularity of the FC graph retrieved from regularized partial correlation is significantly higher than one retrieved from correlation (subfig. **b**; also compare with Fig. 1) and the subjects' $\chi(m)$ curves exhibit a maximum at approximately the same value of m than the main dataset (subfig. **a**). On the other hand, in this case, we do not observe any significant gap between the cross-modularity of $\mathcal{G}_F(\tilde{\mathbf{J}}_\mu(\alpha))$ choosing the tuning parameter α_L that maximizes the validation-set likelihood and $\alpha_{\chi_{\mathcal{G}_F(\tilde{\mathbf{J}}_\mu)}}$ that maximizes the cross-modularity of $\mathcal{G}_F(\tilde{\mathbf{J}}_\mu)$ (subfig. **a**; the only bar in the histogram relative to GLASSO's α_L corresponds to an optimal value of 0.01). We additionally observe that the ratio between the number of observations and the number of brain areas of the extra dataset is over twice ($q = 0,64$) the one of the main dataset ($q = 0,28$), so that, in principle, we would expect larger regularization parameters. Indeed, this is the case for the parameters α_L, while the parameters $\alpha_{\chi_{\mathcal{G}_F(\tilde{\mathbf{J}})}}$ appear in the same order of magnitude as those of the main dataset (compare with Fig S3). This may be due to a higher time correlation affecting the main dataset (subfig. *c*; for each subject, we have computed the time correlation as $\xi(t) = \frac{1}{T-t}\sum_{t'}^{T-t} \vec{x}(t') \cdot \vec{x}(t'+t)$, where $\vec{x}(t)$ is time t time series observation. The solid lines represent the across-subject average, and the shaded areas represent the across-subject standard deviation.)

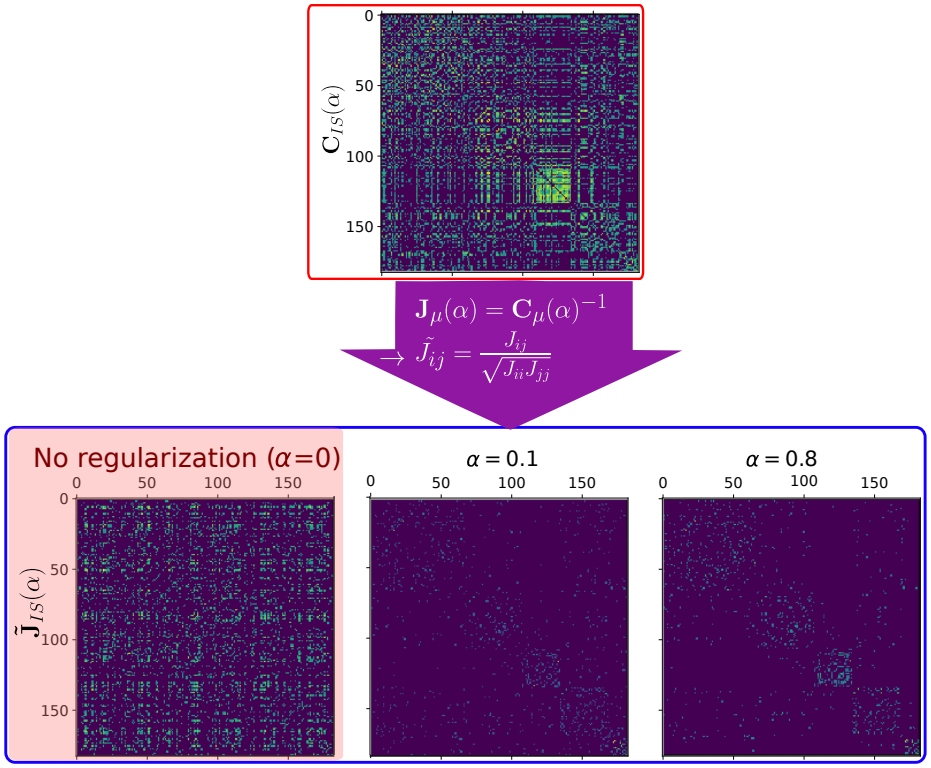

Figure S7: Example of the **effect of IS regularization on $\mathcal{G}_F(\mathbf{C})$ and $\mathcal{G}_F(\tilde{\mathbf{J}})$**, for a single representative subject. The goal of the image is to illustrate how regularization can produce little or no change on $\mathcal{G}_F(\mathbf{C}_\mu(\alpha))$, and yet deeply affect the corresponding $\mathcal{G}_F(\tilde{\mathbf{J}}_\mu(\alpha))$. In particular, IS only acts on the ratio of the diagonal and out-of-diagonal correlation matrix elements so that the corresponding FC graph, which discards the self-loops, only changes in links' weights by a scale factor. On the other hand, the diagonal elements of a correlation matrix affect all elements of the corresponding partial correlation matrix so that the final $\mathcal{G}_F(\tilde{\mathbf{J}}_{\mathbf{IS}}(\alpha))$ is remarkably modified by the IS regularization. First row: subject's 0 correlation matrix's out-of-diagonal elements. The same matrix corresponds to all choices of $\alpha < 1$, apart from a scale factor, adopting an IS regularization strategy. For three values of α, the corresponding partial correlation matrices' out-of-diagonal elements are shown in the second row.