Voxel-Wise Brain Graphs from Diffusion MRI: Intrinsic Eigenspace Dimensionality and Application to Functional MRI

Hamid Behjat*[1,2,✉], Anjali Tarun*[1], David Abramian[3], Martin Larsson[4], and Dimitri Van De Ville[1,5]

[1]Neuro-X Institute, Ecole Polytechnique Federale de Lausanne, Lausanne, Switzerland
[2]Department of Biomedical Engineering, Lund University, Lund, Sweden
[3]Department of Biomedical Engineering, Linköping University, Linköping, Sweden
[4]Centre for Mathematical Sciences, Lund University, Lund, Sweden
[5]Department of Radiology and Medical Informatics, University of Geneva, Geneva, Switzerland
*H. Behjat and A. Tarun provided equal contribution. ✉ Corresponding author; e-mail: hamid.behjat@epfl.ch

Structural brain graphs are conventionally limited to defining nodes as gray matter regions from an atlas, with edges reflecting the density of axonal projections between pairs of nodes. Here we explicitly model the entire set of voxels within a brain mask as nodes of high-resolution, subject-specific graphs. We define the strength of local voxel-to-voxel connections using diffusion tensors and orientation distribution functions derived from diffusion MRI data. We study the graphs' Laplacian spectral properties on data from the Human Connectome Project. We then assess the extent of inter-subject variability of the Laplacian eigenmodes via a procrustes validation scheme. Finally, we demonstrate the extent to which functional MRI data are shaped by the underlying anatomical structure via graph signal processing. The graph Laplacian eigenmodes manifest highly resolved spatial profiles, reflecting distributed patterns that correspond to major white matter pathways. We show that the intrinsic dimensionality of the eigenspace of such high-resolution graphs is only a mere fraction of the graph dimensions. By projecting task and resting-state data on low-frequency graph Laplacian eigenmodes, we show that brain activity can be well approximated by a small subset of low-frequency components. The proposed graphs open new avenues in studying the brain, be it, by exploring their organisational properties via graph or spectral graph theory, or by treating them as the scaffold on which brain function is observed at the individual level.

Brain graph | diffusion MRI | functional MRI | graph signal processing | spectral graph theory

1. Introduction

Magnetic resonance imaging (MRI) has provided an effective means to map the brain's anatomical scaffold, using diffusion MRI, and, in parallel, to track brain neural activity using functional MRI (fMRI). Extensive datasets that include diffusion and functional data on the same set of subjects, such as the Human Connectome Project (HCP) (Essen et al., 2013), have been made freely available, and with such accessibility, various methodological developments that aim to integrate the two modalities have emerged.

Computational neuroimaging has successfully adopted graph theory, creating a new field of interdisciplinary research called *network neuroscience* (Bassett and Sporns, 2017). Structural and functional *connectomes* are independently defined, and consequently analyzed using graph theory measures to provide a better understanding of their organizing network principles (Bullmore and Sporns, 2009). There is also an increasing interest in deciphering how underlying brain anatomy supports the emergence of spatially and temporally varying distributed patterns of functional activity (Mišić et al., 2015; Iraji et al., 2019; Suárez et al., 2020). In this perspective, it is fitting to consider advancing network neuroscience to a more unifying analysis approach that accounts for the interplay between brain structure and function. The study of signal propagation on structural connectomes (Vézquez-Rodríguez et al., 2020; Weninger et al., 2022) is an example avenue of research that is gaining momentum. Leveraging principles from the recently emerged field of graph signal processing (GSP) (Shuman et al., 2013; Ortega et al., 2018), an alternative framework is taking form, in which functional data are interpreted as functions defined atop of a graph that describes the morphological or wiring structure of the brain, and, in turn, processed using spectral methods that are informed by the underlying brain structure.

GSP generalizes principles from classical discrete signal processing for time series to data defined on irregular domains. GSP has found numerous applications across multiple domains—see e.g. (Ortega et al., 2018) for a recent review, and in particular within neuroimaging, examples include: brain state decoding (Petrantonakis and Kompatsiaris, 2018; Ghoroghchian et al., 2020; Cattai et al., 2021), brain signal denoising (Einizade and Sardouie, 2022), brain activation mapping (Behjat et al., 2015; Abramian et al., 2021; Behjat et al., 2021), source localization (Hyde et al., 2019), diagnosing neuropathology (Itani and Thanou, 2021), tracking fast spatiotemporal cortical dynamics (Glomb et al., 2020; Ruéqueralt et al., 2021), brain fingerprinting and task decoding (Griffa et al., 2022) via quantifying the degree of coupling between brain function and structure (Preti and Van De Ville, 2019), identifying dynamically evolving populations of neurons (Charles et al., 2022), deciphering signatures of attention switching (Huang et al., 2018), manifesting white matter pathways that mediate cortical activity (Tarun et al., 2020), and elucidating perturbations of consciousness induced by brain injury or drugs (Atasoy

et al., 2017; Luppi et al., 2022).

The majority of existing applications of GSP in functional brain imaging are limited to region-wise analyses, where regions are defined using a priori brain atlases having 100 to 1000 regions. Motivated by the promising results from existing region-wise GSP studies in linking brain structure and function, and novel recent methods for high-resolution connectomics (Mansour-L et al., 2021) and activation mapping (Zhao et al., 2022), it is fitting to further explore the benefits of large-scale brain graphs at the resolution of *voxels*. Here we present models to derive the strength of connection between adjacent voxels using diffusion MRI data, one based on tensors estimated from diffusion tensor imaging (DTI) (Basser et al., 1994), and the other based on diffusion orientation distribution functions (ODF) estimated from high angular resolution diffusion imaging (HARDI) data (Tuch, 2004); we consider the two signal representation settings to make the methodology applicable to a larger set of available diffusion MRI data. Using data of 100 subjects from the HCP (Essen et al., 2013), we validate the graphs via studying their nodal and spectral measures. We then probe the intrinsic dimensionality of their eigenspace using a Procrustes validation scheme that characterizes inter-subject variability. Finally, we demonstrate the relevance of such high-spatial-resolution voxel-wise graphs within a GSP setting, particularly through studying the energy spectral density of resting-state and task fMRI data on these graphs. We conclude the paper by discussing potential future research avenues in using voxel-wise graphs, in particular, in studying the interaction between brain structure and function.

2. Materials and Methods

2.A. Datasets

We used MRI data from the publicly available HCP dataset—the 100 unrelated subjects, WU-Minn Consortium (Essen et al., 2013). MRI acquisition protocols of the dataset and preprocessing guidelines for diffusion MRI are extensively described elsewhere (Glasser et al., 2013). We used the minimally preprocessed diffusion and anatomical data. The resting-state and task fMRI scans of each subject were realigned to their mean images, and were registered and resampled onto the diffusion data through rigid-body registration using SPM[1]. Two signal reconstruction methods were applied to the diffusion data: (1) DTI tensor fitting using FSL[2], and (2) ODF estimations using DSI Studio[3].

2.B. Graphs and their spectra

Let $\mathcal{G} := (\mathcal{N}, \mathbf{A})$ denote an undirected, single-connected, weighted graph, consisting of a node set $\mathcal{N}$, where $|\mathcal{N}| = N$, and a symmetric $N \times N$ weighted adjacency matrix $\mathbf{A}$, wherein any of its nonzero elements a_{ij} represent the weight of an edge (i, j) in the graph. The normalized graph Laplacian (Chung, 1997) is defined as

$$\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}, \quad (1)$$

where $\mathbf{I}$ denotes the identity matrix, and $\mathbf{D}$ denotes the graph degree matrix, which is diagonal with elements $d_{i,i} = \sum_{j=1}^{N} a_{ij}$. $\mathbf{L}$ can be diagonalized as $\mathbf{L} = \mathbf{U}\Lambda\mathbf{U}^\top$, where $\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, ..., \mathbf{u}_N]$ is an orthonormal matrix stacking the eigenvectors $\mathbf{u}_i$, $i = 1, \ldots, N$, and Λ is a diagonal matrix stacking the corresponding eigenvalues $\Lambda_{i,i} = \lambda_i$, which are real and non-negative due to symmetry and positive semi-definiteness of $\mathbf{L}$. Without loss of generality, we assume that the diagonal elements in Λ, and the corresponding columns in $\mathbf{U}$, are sorted based on the magnitude of the eigenvalues, *i.e.,* $\forall i, j$ if $i < j$ then $\Lambda_{i,i} \leq \Lambda_{j,j}$. As such, the graph Laplacian eigenvalue set satisfies $\{0 = \lambda_1 \leq \lambda_2 \cdots \leq \lambda_N := \lambda_{\text{max}} \leq 2\}$, where the upper bound is guaranteed due to the use of the normalized Laplacian matrix (Chung, 1997). This set defines the Laplacian *spectrum* of the graph, and the eigenvector set $\{\mathbf{u}_i\}_{i=1,\ldots,N}$ defines an orthonormal basis that spans the $\mathbb{R}^N$ space of vectors defined on the nodes of the graph; in the following, we occasionally refer to the Laplacian eigenvectors also as *eigenmodes*, a nomenclature commonly used in the neuroimaging community.

The eigenvalues of a graph Laplacian carry a notion of frequency, which is directly linked to the extent of spatial saliency manifested by their corresponding eigenvectors. To understand this link, a metric known as total variation (TV) (Sandryhaila and Moura, 2014) can be computed for each eigenvector, or more precisely, for any given graph signal. A graph signal defined on the nodes of a graph can be represented as a vector $\mathbf{x} \in \mathbb{R}^N$, where the i-th element, $\mathbf{x}[i]$, is the signal value at the i-th node of the graph. For a given graph signal $\mathbf{x}$, the TV of $\mathbf{x}$ is defined as $TV(\mathbf{x}) = \mathbf{x}^\top \mathbf{L}\mathbf{x}$, a measure that quantifies the extent of variation observed in $\mathbf{x}$ relative to the underlying graph structure. Given that the eigenvectors are orthonormal, *i.e.,* $\mathbf{u}_i^\top \mathbf{u}_i = 1$, and that $\mathbf{L}\mathbf{u}_i = \lambda_i \mathbf{u}_i$, it follows that the TV of Laplacian eigenvectors reduces to $TV(\mathbf{u}_i) = \mathbf{u}_i^\top \mathbf{L}\mathbf{u}_i = \lambda_i$, showing that the variability of each Laplacian eigenvector is reflected by the associated eigenvalue, or in other words, that Laplacian eigenvectors associated to larger eigenvalues reflect a greater extent of spatial variability. Alternatively, spatial variability of graph signals/eigenmodes can be quantified via a measure of zero-crossings(Petrantonakis, 2021; Petrantonakis and Kompatsiaris, 2022). In particular, we define a weighted zero-crossing measure as (Abramian et al., 2021)

$$ZC(\mathbf{u}_k) = \frac{1}{2}\sum_{i \neq j} a_{ij} H(-\mathbf{u}_k[i]\mathbf{u}_k[j]), \quad (2)$$

where $H(\cdot)$ is the Heaviside step function and a_{ij} is the edge weight that connects voxels v_i and v_j. The higher

[1] https://www.fil.ion.ucl.ac.uk/spm/software/spm12
[2] https://fsl.fmrib.ox.ac.uk
[3] https://dsi-studio.labsolver.org

the zero-crossing metric, the greater the associated variability in the eigenvectors' spatial patterns.

2.C. Spectral decomposition of graph signals

For a given graph signal **x**, its spectral representation, commonly referred to as the graph Fourier transform (GFT) of **x**, is given as

$$\widetilde{\mathbf{x}} = \mathbf{U}^{\top}\mathbf{x}. \quad (3)$$

The signal can be perfectly recovered through the inverse GFT operation as, $\mathbf{x} = \mathbf{U}\widetilde{\mathbf{x}} = \sum_{i=1}^{N} \widetilde{\mathbf{x}}[i]\mathbf{u}_i$. As such, any given graph signal can be seen as a linear combination of the orthonormal set of Laplacian eigenvectors. In particular, $|\widetilde{\mathbf{x}}[i]|^2$ gives the energy spectral density of the signal associated to the i-th eigenmode. Given a set of graph signals $\mathcal{X} = \{\mathbf{x}_k\}_{k=1,\ldots,S}$ (e.g. graph signals derived from individual time frames of a given fMRI session) we compute the ensemble energy spectral density (EESD) of the lower end of the spectrum of $\mathcal{X}$ as

$$EESD(i) = \frac{1}{S}\sum_{k=1}^{S} \underline{\widetilde{\mathbf{x}}}_k[i], \quad i = 1, \ldots, C, \quad (4)$$

where C denotes a desired cutoff index specifying the number of lower end spectral indices to be studied ($C = 1000$ in this study), and $\underline{\mathbf{x}}_k$ denotes the demeaned and normalized version of $\mathbf{x}_k$ obatined as

$$\underline{\mathbf{x}}_k = (\mathbf{x}_k - \mathbf{u}_1^{\top}\mathbf{x}_k\mathbf{u}_1)/||\mathbf{x}_k - \mathbf{u}_1^{\top}\mathbf{x}_k\mathbf{u}_1||_2, \quad (5)$$

which ensures $|\underline{\widetilde{\mathbf{x}}}_k[1]|^2 = 0$ and $\sum_{i=1}^{N} |\underline{\widetilde{\mathbf{x}}}_k[i]|^2 = 1$.

2.D. Brain graph design

For each subject, we define a weighted brain graph characterized by a node set $\mathcal{N} = \{1, \ldots, N\}$ defined based on the set of voxels that fall within the subject's brain mask, covering gray matter (GM), white matter (WM) and cerebrospinal fluid (CSF), representing a 3D mesh arrangement. In particular, each node i is associated to a voxel, denoted v_i, with coordinates (x_i, y_i, z_i). The graph edges are defined based on the adjacency of voxels within the Moore neighborhood cubic lattice of size $3 \times 3 \times 3$ and $5 \times 5 \times 5$, where the latter size is only used in the ODF-based design. For the $5 \times 5 \times 5$ design, voxels in the outer layer that fall in parallel to the voxels within the inner layer were excluded, enabling encoding of connections to a maximum of 98 voxels/directions in the neighborhood of each focal voxel, whereas the $3 \times 3 \times 3$ design enables encoding connections to a maximum of 26 different directions. As such, the 5-connectivity design trades localization for better angular resolutions. With this definition of edges, each node i entails a neighborhood set of the indices of nodes in $\mathcal{N}$ that are adjacent to it, denoted $\mathcal{N}_i$, where $|\mathcal{N}_i| \leq 26$ and ≤ 98 for the 3- and 5-connectivity designs, respectively.

We define the edge weights based on a measure of *inter-voxel fiber coherence* across all pairs of adjacent voxels. In particular, to make the presented method applicable to a wider range of available diffusion MRI data, we present two edge weighting schemes using two signal representation models, one using diffusion tensors and one using diffusion ODFs; we denote the resulting edge weighting schemes as the DTI-based and ODF-based methods, respectively. The tensor and ODF models both aim to represent structural information on the intra-voxel axon fiber arrangement. The diffusion tensor model can be seen as a multivariate Gaussian describing the distribution of fiber bundle alignment, whereas the diffusion ODF model defines the radial projection of the diffusion function, providing an estimate of the empirical distribution of water diffusion.

Let $\mathbf{r}_{ij}$ denote the vector pointing from the center of the voxel v_i to the center of the voxel v_j. Let $\{p(i, \mathbf{r}_{ij})\}_{j \in \mathcal{N}_i}$ denote estimates of the extent of diffusion at voxel v_i in directions $\{\mathbf{r}_{ij}\}_{j \in \mathcal{N}_i}$. In the following, we first present a DTI-based and an ODF-based approach to estimate $\{\{p(i, \mathbf{r}_{ij})\}_{j \in \mathcal{N}_i}\}_{i=1}^{N}$. We then use these estimates to define the graph edge weights.

2.D.1. DTI-based quantification of diffusion orientation. DTI is a model-based method for reconstructing the diffusion signal from diffusion MRI. The assumption in DTI is that the diffusion pattern follows the shape of a 3D ellipsoid. The molecular displacement of water at voxel i in the direction $\mathbf{r}_{ij}$ can be approximated by a 3D Gaussian distribution with real symmetric diffusion tensor $\mathbf{T}$ as the covariance matrix:

$$p(i, \mathbf{r}_{ij}) = \frac{1}{\sqrt{(2\pi)^3|\mathbf{T}|}} \exp(-\frac{1}{2}\mathbf{r}_{ij}^{\top}\mathbf{T}^{-1}\mathbf{r}_{ij}), \quad (6)$$

where $|\mathbf{T}|$ is the determinant of the diffusion tensor. The calculation of $p(i, \mathbf{r}_{ij})$ requires a discretization step that guarantees a one-to-one mapping between the (continuous) multivariate Gaussian model and the (discrete) weighting of nodes in the brain graph. For the $3 \times 3 \times 3$ neighborhood encoding scheme, for a given voxel v_i, the set of values $\{p(i, \mathbf{r}_{ij})\}_{j \in \mathcal{N}_i}$ can be arranged into $3 \times 3 \times 3$ discrete representation, which mimics the structure of a 3D finite impulse response (FIR) filter. As such, the problem of obtaining $\{p(i, \mathbf{r}_{ij})\}_{j \in \mathcal{N}_i}$ can be alternatively seen as that of obtaining the coefficients of an FIR filter. We provide the details of this procedure in the Supplementary Material.

2.D.2. ODF-based quantification of diffusion orientation. Unlike diffusion tensors, ODFs do not follow a specific model and shape. Thus, a one-to-one discretization of a continuous function similar to that presented for the DTI model cannot be applied. Here we build on the construction previously presented by Iturria-Medina (Iturria-Medina et al., 2007). Within standard spherical coordinates, parametrized by $(r, \theta, \phi,)$, let $O_i(\hat{u})$ denote the ODF associated to voxel v_i with its center

of coordinate being the voxel's center, with $\hat{u}(\theta,\phi) = (sin(\theta)cos(\phi), sin(\theta)sin(\phi), cos(\theta))^\top$ denoting the unit direction vector. Given $O_i(\hat{u})$, a measure of the extent of diffusion at voxel v_i along direction $\mathbf{r}_{ij}$ can be obtained as

$$p(i,\mathbf{r}_{ij}) = \frac{1}{\Omega_{ij}} \int_{\Omega_{ij}} O_i^n(\hat{u}) d\Omega, \quad \textbf{(7)}$$

where Ω_{ij} denotes a given solid angle around $\mathbf{r}_{ij}$, $d\Omega = sin(\theta)d\theta d\phi$ denotes the infinitesimal solid angle element; in particular, a solid angle of $4\pi/26$ and $4\pi/98$ is used for the 3- and 5-connectivity schemes, respectively. The exponent $n > 0$ is a desired power factor that is used to sharpen the ODF given the limited degree to which diffusion ODFs can differentiate fiber orientations (Jones et al., 2013). As such, $p(i,\mathbf{r}_{ij})$ gives an average measure of the surface area of the ODF within a spherical cap defined by the solid angle. Given a discrete representation of $O_i(\hat{u})$ in form of N_o samples, denoted $\{O_{i,k}\}_{k=1}^{N_o}$, along N_o spherical directions from the center of voxel v_i, (7) can be approximated as

$$p(i,\mathbf{r}_{ij}) \approx \frac{1}{|\mathcal{D}_{ij}|} \sum_{k \in \mathcal{D}_{ij}} O_{i,k}^n, \quad \textbf{(8)}$$

where $\mathcal{D}_{ij}$ denotes a subset of direction indices $\{1,\ldots,N_o\}$ whose associated set of directions fall within Ω_{ij}. The normalization by the cardinality of $\mathcal{D}_{ij}$ is due to the difference in the number of ODF samples that fall within the solid angle subtended along the different neighborhood directions.

2.D.3. Brain graph edge weights. The graph edge weights a_{ij} are defined by using the estimates of diffusion orientation at the associated voxels v_i and v_j, i.e., $p(i,\mathbf{r}_{ij})$ and $p(j,\mathbf{r}_{ji})$, as well as the strength of anisotropy at those voxels. In particular, for a given voxel v_i, let $F(v_i)$ and $Q(v_i)$ denote the voxel's fractional anisotropy (FA) and quantitative anisotropy (QA), respectively. FA can be calculated directly from the eigenvalues of the diffusion tensor (Alexander et al., 2007) and QA pertains to the amount of diffusion anisotropy along the fiber orientation as originally defined by Yeh et. al. (Yeh et al., 2010). Using these measures, we define the graph edge weights a_{ij} as

$$a_{ij} = \overbrace{\frac{P_{\text{mag}}(v_i)P_{\text{mag}}(v_j)}{\alpha^2}}^{\text{magnitude}} \overbrace{\left(\frac{p(i,\mathbf{r}_{ij})}{\beta_i} + \frac{p(j,\mathbf{r}_{ji})}{\beta_j}\right)}^{\text{orientation}}, \quad \textbf{(9)}$$

where $\alpha = \max_k\{P_{\text{mag}}(v_k)\}$, $\beta_k = 2\max_{l \in \mathcal{N}_k}\{p(k,\mathbf{r}_{kl})\}$ and $P_{\text{mag}}(v_i)$ denotes the magnitude of anisotropy at voxel i, defined as

$$P_{\text{mag}}(v_i) = \begin{cases} F(v_i), & \text{DTI-based design} \\ Q(v_i), & \text{ODF-based design.} \end{cases} \quad \textbf{(10)}$$

The connectivity structure of the graph is then characterized in $\mathbf{A}$, such that $a_{ij} > 0$, given as in (9), if nodes i and j are connected through an edge, and $a_{ij} = 0$ if otherwise.

The *orientation* term in (9) gives a measure of diffusion orientation coherence between the two connected voxels, whereas the anisotropies give a *magnitude* measure that is useful in delineating tissues; the use of anisotropies is in contrast to using probabilistic tissue maps as used in (Iturria-Medina et al., 2007), enabling the design of the graphs using only the diffusion data. Given two adjacent voxels that exhibit highly coherent diffusion orientations, a large weight is associated to the connection only if the two voxels also exhibit notably large anisotropies. This interplay between the orientation term and magnitude term enables, for example, to prevent associating large weight to an edge between a WM voxel and a CSF voxel. Furthermore, the normalizations incorporated in the definition, i.e., the α and β_k terms, ensure having an unbiased definition of weights relative to the structure of the diffusion tensors/ODFs across the brain, and, mathematically, they impose bounds on the orientation and magnitude terms—both terms bounded to $[0,1]$, which in turn results in a_{ij} also being bounded to $[0,1]$.

2.E. Group-level eigenmodes

The voxel-wise nature of the studied graphs renders their size excessively large, with $7.8 \pm 0.7 \times 10^5$ nodes across the 100 subjects considered. The sheer size of the voxel-wise graphs impedes computing the full eigendecomposition of the graph Laplacian, and, therefore, we compute and study the first leading 1000 eigenvectors corresponding to the lowest spectral frequencies. To preserve subtle subject-specific spatial details, we construct all graphs in the native space of each subject's diffusion data. To enable inter-subject comparison of eigenmodes, the DARTEL normalization algorithm (Ashburner, 2007) implemented in SPM12 was used to define a group-level template coordinate space, based on the group's T1-weighted MRI data. This results in a structural T1-weighted template as well as a set of transformation maps per subject. Each subject's eigenmodes are then transformed into the template space using the subject-specific transformations, resulting in inter-subject spatially aligned eigenmodes.

2.F. Consistent inter-subject ordering of eigenmodes

The ordering of DARTEL-normalized eigenmodes is not necessarily consistent across subjects, including sign ambiguity and linear combinations between modes with close eigenvalues. To obtain a consistent ordering of the eigenmodes, and enable inter-subject comparison of individual eigenmodes, we used the Procrustes transform (Goodall, 1991), which finds the optimal rotation, translation, and/or reflection between two linear subspaces. We implemented a scheme of Procrustes transformation, where the subspace is defined by the first K DARTEL-normalized eigenmodes of any M subset of subjects. Let $\mathbf{u}_{i,m}$ denote the i-th eigenmode of subject

m, and let $X_{K,m} = [\mathbf{u}_{1,m}\ \mathbf{u}_{2,m}\ \cdots\ \mathbf{u}_{K,m}]$.

Algorithm 1: group-level matching of eigenmodes
$X_{avg} \leftarrow X_{K,1}$
for $j = 1$ **to** J **do**
  for $m = 1$ **to** M **do**
    $X_{K,m} \leftarrow$ Reorder columns in $X_{K,m}$ such that the resulting permuted matrix best matches X_{avg}
  end for
  $X_{avg} \leftarrow$ average $\{X_{K,1}, \ldots, X_{K,M}\}$
end for

The *reordering* is done based on Procrustes transformation estimates, and J denotes the optimal number of iterations to ensure that X_{avg} does not remain biased towards its initial value, i.e., $X_{K,1}$. The columns of the resulting $\{X_{K,m}\}_{m=1}^{M}$ are reordered in such a way that they optimally match each other and can thus be compared across the subjects. Although the Procrustes transformation removes much of the variance that is common between subjects, it cannot discount for subject-specific details; in the following section we define a metric that quantifies the extent of remaining inter-subject structural variability.

2.G. Quantification of the extent of inter-subject structural variability as encoded in brain graphs

The precision of the Procrustes transformation can be evaluated by quantifying the cosine similarity between corresponding eigenmodes of different subjects, after DARTEL normalization and Procrustes transformation. For any pair of subjects, a symmetric cosine similarity matrix is obtained, where the deviation of the off-diagonal elements of the matrix from zero quantifies inter-subject structural variability. If the set of eigenmodes of two subjects has been ideally matched, the cosine similarity matrix should be the identity matrix. Therefore, to quantify the mismatch between a pair of subjects m and n based on their first K eigenmodes, we define a measure of the extent of inter-subject structural variability, termed *Procrustes error*, via computing the cosine similarity between pairs of eigenmodes as

$$E_{m,n}(K) = \frac{1}{2}\sqrt{\sum_{i=1}^{K}\sum_{\substack{j=1\\ j\neq i}}^{K}\left(\frac{\mathbf{u}_{i,m}^{\top}\mathbf{u}_{j,n}}{\|\mathbf{u}_{i,m}\|\,\|\mathbf{u}_{j,n}\|}\right)^{2}}. \quad \textbf{(11)}$$

We used this error term to determine J in Algorithm 1 and also to compare the different graph designs.
To validate the performance of the Procrustes transformation, we implemented a bootstrap scheme, which successively applies the transformation on the first K eigenmodes of two randomly chosen subjects; the implementation is summarized in the following algorithm:

Algorithm 2: Procrustes validation
for $K = 100$ **to** 1000 **step** 100 **do**
  for $j = 1$ **to** 1000 **do**
    $\{m, n\} \leftarrow$ randomly select 2 values $\in \{1, \ldots, 100\}$

    $\{X_{K,m}, X_{K,n}\} \leftarrow$ **Algorithm 1** on $\{X_{K,m}, X_{K,n}\}$

    $\epsilon_j \leftarrow E_{m,n}(K)$
  end for
  $\mu_K \leftarrow$ mean of $\{\epsilon_j\}_{j=1,\ldots,1000}$
  $\sigma_K \leftarrow$ standard deviation of $\{\epsilon_j\}_{j=1,\ldots,1000}$
end for

When comparing two brain graph designs, the design that results in a higher μ_K, with reasonably small σ_K, is interpreted as capturing more subject-specific structural features.

2.H. Spectral decomposition of fMRI data

As proof-of-concept of the applicability of the proposed whole-brain, voxel-wise brain graphs, we evaluated the extent to which brain fMRI data are spatially shaped by the underlying brain structure as encoded by the graphs. In particular, we constructed fMRI graph signals from six functional tasks as well a resting-state acquisition, across 100 subjects. Each fMRI time frame was transformed in to a single graph signal. This was done by extracting the fMRI voxels associated to the graph vertices, i.e., voxels that fall within each subject's brain mask, arranging them as a vector, ordered based on the order of vertices as reflected in each subject's graph adjacency matrix. As such, for each subject, a time-evolving series of graph signals were obtained from each of the subjects' task or resting-state 4D fMRI volumes. We studied the energy spectral density of the extracted graph signals associated to the first 1000 spectral indices of the ODF-3 graph. To serve as a null, we also evaluated the energy spectral density of synthesized shuffled fMRI signals—obtained through random permutation of voxel indices of each fMRI volume to destroy spatial order, as well as white Gaussian noise signals.

3. Results

Figure 1 shows the degree distributions of the graphs across the different tissue types. The connectivity strength is highest in WM nodes compared to GM and CSF nodes. The degree distribution of ODF-5 graphs is a shifted version of that of the ODF-3 graphs towards a higher degree, reflecting the larger number of connections possible with the ODF-5 design. It can be observed that the nodal degrees within gray matter and CSF almost coincide for the DTI design whereas this is much less pronounced in the ODF designs (compare Figure 1(A) with Figs. 1(B) and (C)); this is more pronounced for the ODF-3 design, which reflects that the ODF-3 design bears larger differences across tissue types compared to the DTI-3 design. Figure 1(D) shows the distribution of nodes

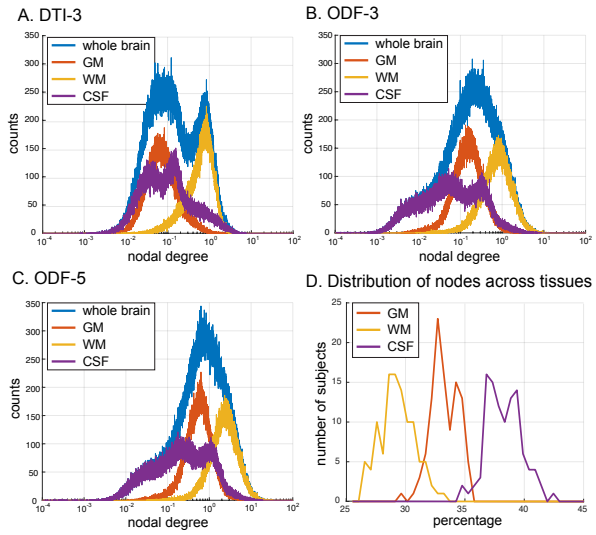

Figure 1. Degree distribution of the (A) DTI-3, (B) ODF-3, and (C) ODF-5 voxel-wise brain graphs for a representative subject. (D) Distribution of the number of nodes in each tissue type across all subjects.

for the different tissue types across all subjects; the median number of nodes for GM, WM, and CSF were 254 299, 221 964, and 294 028, respectively, with standard deviations 25 322, 28 579, and 25 742, respectively.

3.A. Spectral Comparison

Figure 2(A) shows the first Laplacian eigenmodes obtained from the three brain graphs of a representative subject. Noting that the first eigenmode of $\mathbf{L}$ is a function of the graph nodal degrees—$\mathbf{u}_1 = \mathbf{D}^{1/2}\mathbf{1}$ where $\mathbf{1}$ denotes the constant function that assumes the value of 1 on each node, the spatial pattern manifested by the first eigenmodes is a corroboration of the results shown in Figure 1(A)-(C), demonstrating that the distinction between tissue types naturally arises from the assignment of the connectivity weights in the brain graph, in particular through the assignment of the magnitude term in (9). Moreover, the first eigenmode manifests a specific profile of local tissue structure, in which higher values reflect voxels/regions that are more strongly connected to their surrounding neighbourhood, particularly observed at regions of less ambiguous fiber structure, e.g. within the corpus callosum The second and third eigenmodes shown in Figure 2(B) manifest global morphological organization of the brain, contrasting the posterior and anterior, and the left and right brain regions, respectively. The next eigenmodes exhibit a greater extent of spatial variability and localized information.

Figure 3(A) shows the lower-end graph spectra in which the rate of increase in the first 1000 eigenvalues are shown. The ODF-5 graph has relatively larger eigenvalues than the ODF-3 and DTI-3 graphs. Given that the eigenvalues entail a notion of spatial saliency, the increase in local connectivity in the ODF-5 implies that the associated eigenmodes have greater degrees of freedom, and as such, spatial saliency can become higher.

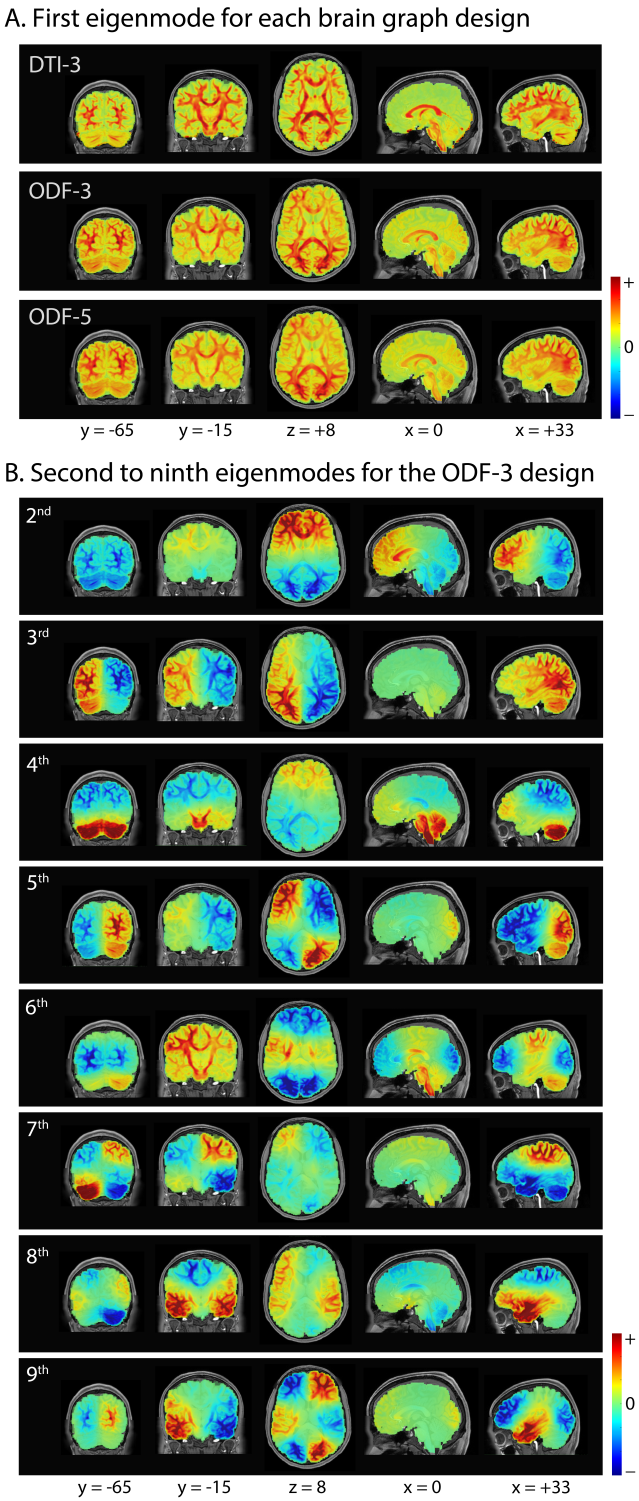

Figure 2. (A) First Laplacian eigenmode of 3-connectivity DTI, 3-connectivity ODF and 5-connectivity ODF brain graphs of a representative subject. (B) The next lowest frequency eigenmodes corresponding to the 3-connectivity ODF brain graph.

The spatial saliency of the eigenmodes' can be quantified by computing their weighted zero-crossing, cf. (2). Figure 3(B) shows a trend that is consistent with that of the eigenvalues, thereby confirming the general notion that higher indexed eigenmodes encompass a larger extent of spatial saliency.

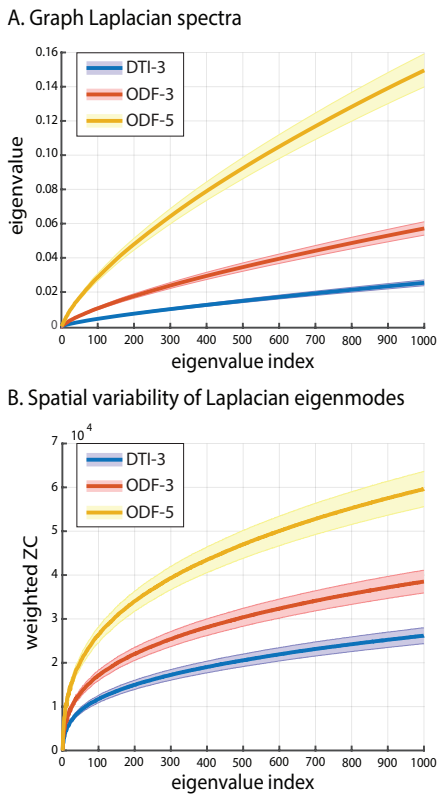

Figure 3. (A) Lower-end eigenvalues of DTI-3, ODF-3 and ODF-5 graphs, consisting of their first 1000 eigenvalues. (B) weighted zero-crossing the corresponding eigenmodes; cf. (2); solid lines show the mean and shades show the standard deviation across the 100 subjects.

3.B. Procrustes Validation

We evaluated the inherent inter-subject variability encoded in two voxel-wise brain graph designs via the Procrustes transform, which finds the optimal configuration that matches single subject eigenmodes to an averaged set. This step, however, cannot account for subject-specific fiber pathways encoded in the eigenmodes, as manifested by the cosine similarity analysis of pairs of subjects. The cosine similarity matrices of two sets of eigenmodes before and after applying Procrustes transformation are shown in Figure 4. The diagonal structure of the cosine similarity matrix associated to the transformed eigenmodes shows the effectiveness of the transformation in aligning eigenmodes of the same neuroanatomical spatial nature. The insets show traces of flipped signs and unordered eigenmodes in the original vectors, which were corrected after Procrustes transformation. The precision of the Procrustes analysis improves after several rounds of the transformation; see Figure S1 in Supplementary Materials.

Figure 5 shows the Procrustes error when different subset of eigenmodes are used, for each of the three graph designs, reflecting the extent of inter-subject structural variability captured by the eigenmodes. Furthermore, to serve as a null for comparing and validating the brain graphs, we synthesized 100 random orthonormal vectors of the same dimension as each of the subject-specific

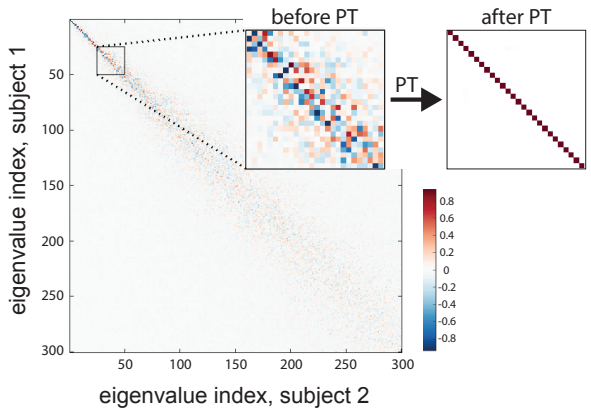

Figure 4. Cosine similarity between the first 300 eigenmodes before and after Procrustes transformation (PT) for two representative subjects.

eigenmodes, applied DARTEL normalization, and then subjected the resulting vectors to Procrustes validation. All three brain graphs show a decreasing trend, and they get closer to that of the null upon reaching higher K-values. Despite the differences in the Procrustes errors across the three brain graphs for low values of K, all errors eventually converge to a single point for high K. In addition, the ODF-based designs show higher Procrustes errors than the DTI-based design, across K. The ODF-5 design slightly outperformed the ODF-3 design, suggesting the benefit of using the larger neighborhood in encoding fiber orientations with better angular resolution.

3.C. Spectral decomposition of fMRI

Figure 6(A) shows the ensemble energy spectral densities of the fMRI. The energy spectral density of the fMRI data is characterized by a power-law behavior. The lowest frequency component eigenmodes capture the majority of the energy content (spatial variability), which is about two orders of magnitude greater than what is captured by the 1000th eigenmode. Moreover, a notable dispersion is observed in the energy profiles of approximately the first 100 spectral indices across the different

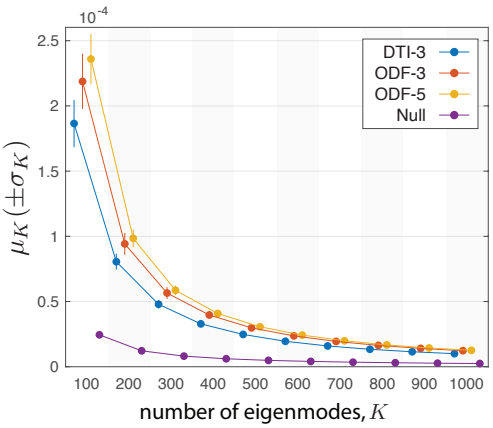

Figure 5. Quantification of the extent of inter-subject structural variability captured by the eigenmodes, cf. Algorithm 2. The four markers (DTI-3, ODF-3, ODF-5, and Null) within each vertical shade are associated to the same K.

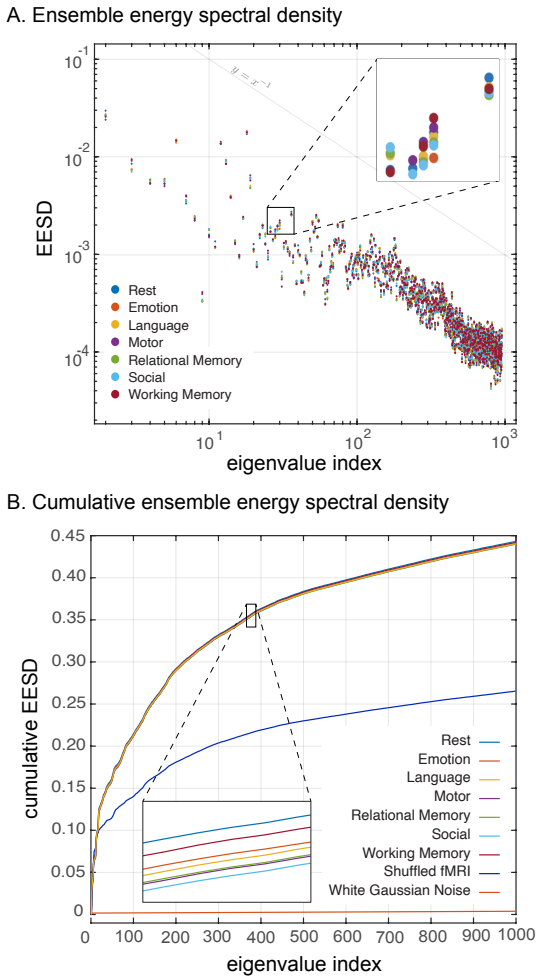

Figure 6. (A) Ensemble average energy spectral density of fMRI data of 100 subjects for the first 1000 eigenmodes of ODF-3 brain graph. (B) Same as in (A) but showing the cumulative values as well as a comparison to shuffled fMRI data and white Gaussian noise.

tasks, whereas for the higher spectral indices, the profiles are more closely packed, following a steady power-law drop. This observation is more apparent by inspecting the cumulative ensemble energy (CEE) profiles, see Figure 6(B); results are shown also for synthesized shuffled fMRI signals (cf. Section H) and Gaussian noise signals. In particular, the CEE profiles of task and rest fMRI sharply differ from that of Gaussian noise as well as shuffled fMRI data. For Gaussian noise, each eigenmode captures a fraction of the total energy equivalent to approximately $1/N$, where N is the number of the graph nodes, whereas shuffled fMRI still entails the distribution of values as in the real fMRI data, but lacks the exquisite spatial dependencies manifested in fMRI data that is linked to the underlying structure. The CEE profiles of fMRI data show that functional brain activity is expressed preferentially by lower-frequency components; approximately 85% of the total signal energy content[4] captured by the first 1000 eigenvectors, wherein the contributions from the first eigenmode are equal to zero due to that the data were demeaned, cf. Section **??**.

[4]For each graph signal, the first eigenmode captured approximately 70% of the total signal energy. In the demeaning step, cf. Section **??**, the contribution from the first eigenmode was regressed out. As such, baed on Figure 6(B), the total amount of energy captured by the first 1000 eigenmodes amounts to approximately 85% of total signal energies of the original signals, i.e., $70 + 0.45 \times 30$.

4. Discussion

Voxel-wise brain graphs enable overcoming several limitations associated to existing region-wise brain graphs. For example, they obviate the need for cortical parcellation, and as such, downstream analysis will not be affected by the choice of parcellation scheme (de Reus and Van den Heuvel, 2013). Moreover, analyses on local-encoding voxel-wise graphs prevents variations in results as a function of the algorithm that is employed to approximate the number of WM tracts for region-wise graphs (Maier-Hein et al.). Lastly, given that the proposed graphs are constructed at the native diffusion space and encompass the whole brain, tissue segmentation and subsequent transformation to a template space is not needed, which can be challenging especially in populations that exhibit a complex mixture of brain structural deficits.

Results on fMRI show that despite the high dimensionality of voxel-wise graphs, brain functional activity can be well approximated by merely a small subset of their low-frequency Laplacian harmonics, whereas in contrast for region-wise brain graphs a larger subset of their total number of Laplacian harmonics is required (Atasoy et al., 2017; Preti and Van De Ville, 2019; Glomb et al., 2020). This observation shows that functional brain maps are smooth relative to the underlying fiber architecture and tissue profile morphologies, which, on the one hand, is consistent with energy profile of fMRI graph signals on tissue-specific, voxel-wise graphs (Behjat et al., 2015, 2021), and on the other hand, can be linked to the decreasing trend observed in the Procrustes validation errors (see Figure 5), where increasing K-values reduce the error close to that of a randomly generated graph. This energy pattern is reminiscent of a power-law behaviour, which is interesting in light of the evidence of scale-free behaviour in human brain activity (Van De Ville et al., 2010; Ciuciu et al., 2012) observed in both temporal and spatial scales. Moreover, a practical implication of such energy profiles is that the lower-spectral-end energy content of fMRI data on whole-brain voxel-wise graphs has the potential to provide signatures of mental activity, similar to that observed for cerebral cortex gray matter graphs (Behjat and Larsson, 2020), which substantially reduces the computational burden associated to diagonalization of the graph Laplacian. Alternatively, to study the energy profile across the spectrum, a filter design scheme that adapts to the ensemble signal content can be used to efficiently partition the spectrum (Behjat et al., 2016), which can be implemented in a computationally efficient manner (Behjat and Van De Ville, 2019), obviating the need to even compute individual eigenvectors.

Voxel-wise brain graphs hold the potential to open new research avenues to study the brain. One avenue is to study the graphs from a pure structural perspective, using spectral graph theoretical measures that have been used to e.g. discriminate auditory gyri subtypes (Maghsadhagh et al., 2021), or to perform subject identification and characterization of hemispheric asymmetries (Wachinger et al., 2016). A second, more interesting, avenue of research is to employ voxel-wise brain graphs within the context of relating brain structure to function. The proposed voxel-wise graphs can be leveraged to perform whole-brain anatomically-informed spatial filtering and interpolation of fMRI data, operations that are inherent within numerous fMRI processing pipelines; e.g. spatial smoothing to enhance whole-brain fMRI activation mapping, as done using tissue-specific designs in gray matter (Behjat et al., 2015, 2021) and white matter (Abramian et al., 2021; Zhao et al., 2022). Moreover, it yet remains to be studied how functional connectivity (FC) and their associated measures can be extended to accurately integrate structural information. FC has often been associated to Euclidean distance (Alexander-Bloch et al., 2013), whereas by using the Laplacian eigenmodes of the proposed graph, functional distance can be better interpreted in relation to the underlying brain structure. That is, functional variations that are captured using low-frequency eigenmodes are smooth with respect to long-distance white matter bundles, whereas localized and short-distance functional associations are expected to be dominated by higher frequency components. Lastly, the exquisite voxel-wise scale of the proposed graphs can enable assessing the extent to which brain structural-functional relations hold at spatially finer mesoscales (Mansour-L et al., 2021); e.g. by using graph Slepians (Van De Ville et al., 2017; Petrovic et al., 2019; Georgiadis et al., 2021) or variants of localized graph filter banks (Shuman, 2020; Isufi et al., 2022) and spectral transforms (Ghandehari et al., 2021; de Loynes et al., 2022; Tay, 2022), focus can be placed on a particular subset of nodes, thus, providing a finer level of analytical resolution than that provided by conventional region-wise graphs.

5. Conclusion

Two methods for constructing voxel-wise brain graphs from diffusion MRI data were studied. Through a Procrustes validation scheme that reflects inter-subject structural differences, it was shown that low-frequency eigenmodes of such high spatial resolution graphs reflect the highest amount of structural information from diffusion MRI. This finding was corroborated by the manifested energy spectral density of functional signals showing the preferential expression of human brain activity onto lower frequency components. Overall, the presented results signify the capability of voxel-wise brain graphs' eigenmodes in capturing anatomically-constrained functional variations that are specific to different cognitive tasks. By treating voxel-wise brain graphs as the scaffold on which brain function is observed, they hold the potential to open new research avenues to study the brain, in particular, enabling the development of novel GSP methods to study the interplay between brain structure and function, in health and disease.

Acknowledgements

This work was supported in part by the Swiss National Science Foundation under Grant 205321-163376 and in part by the Swedish Research Council under Grant 2018-06689. A preliminary version of this work was presented in (Tarun et al., 2019).

References

Abramian, D., Larsson, M., Eklund, A., Aganj, I., Westin, C.-F., and Behjat, H. Diffusion-informed spatial smoothing of fMRI data in white matter using spectral graph filters. *Neuroimage*, 237:118095, 2021.

Alexander-Bloch, A. F., Vértes, P. E., Stidd, R., Lalonde, F., Clasen, L., Rapoport, J., Giedd, J., Bullmore, E. T., and Gogtay, N. The Anatomical Distance of Functional Connections Predicts Brain Network Topology in Health and Schizophrenia. *Cereb. Cortex*, 23(1): 127–138, 2013.

Alexander, A. L., Lee, J. E., Lazar, M., and Field, A. S. Diffusion tensor imaging of the brain. *Neurotherapeutics*, 4(3):316–329, 2007.

Ashburner, J. A fast diffeomorphic image registration algorithm. *Neuroimage*, 38(1):95 – 113, 2007.

Atasoy, S., Roseman, L., Kaelen, M., Kringelbach, M. L., Deco, G., and Carhart-Harris, R. L. Connectome-harmonic decomposition of human brain activity reveals dynamical repertoire re-organization under LSD. *Scientific Reports*, (1):17661, 2017.

Basser, P., Mattiello, J., and Lebihan, D. Estimation of the Effective Self-Diffusion Tensor from the NMR Spin Echo. *Journal of Magnetic Resonance, Series B*, (3):247–254, mar 1994.

Bassett, D. S. and Sporns, O. Network neuroscience. *Nature Neuroscience*, (3):353–364, 2017.

Behjat, H. and Larsson, M. Spectral characterization of functional MRI data on voxel-resolution cortical graphs. In *Proc. IEEE Int. Symp. Biomed. Imaging*, pages 558–562. IEEE, 2020.

Behjat, H. and Van De Ville, D. Spectral design of signal-adapted tight frames on graphs. In *Vertex-Frequency Analysis of Graph Signals*, pages 177–206. Springer, 2019.

Behjat, H., Leonardi, N., Sörnmo, L., and Van De Ville, D. Anatomically-adapted graph wavelets for improved group-level fMRI activation mapping. *Neuroimage*, 123:185–199, 2015.

Behjat, H., Richter, U., Van De Ville, D., and Sörnmo, L. Signal-adapted tight frames on graphs. *IEEE Trans. Signal Process.*, 64(22):6017–6029, 2016.

Behjat, H., Westin, C.-F., and Aganj, I. Cortical surface-informed volumetric spatial smoothing of fMRI data via graph signal processing. In *Proc. IEEE Int. Conf. Eng. Med. Biol. Soc.*, pages 3804–3808. IEEE, 2021.

Behjat, H., Aganj, I., Abramian, D., Eklund, A., and Westin, C.-F. Characterization of spatial dynamics of fMRI data in white matter using diffusion-informed white matter harmonics. In *Proc. IEEE Int. Symp. Biomed. Imaging*, pages 1586–1590. IEEE, 2021.

Bullmore, E. and Sporns, O. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat. Rev. Neurosci.*, (3):186–198, 2009.

Cattai, T., Scarano, G., Corsi, M.-C., Bassett, D. S., Fallani, F. D. V., and Colonnese, S. Improving J-divergence of brain connectivity states by graph Laplacian denoising. *IEEE Trans. Signal Info. Process. Netw.*, 7:493–508, 2021.

Charles, A. S., Cermak, N., Affan, R. O., Scott, B. B., Schiller, J., and Mishne, G. GraFT: graph filtered temporal dictionary learning for functional neural imaging. *IEEE Trans. Image Process.*, 31:3509–3524, 2022.

Chung, F. *Spectral graph theory*. AMS, Providence, RI., 1997.

Ciuciu, P., Varoquaux, G., Abry, P., Sadaghiani, S., and Kleinschmidt, A. Scale-free and multifractal time dynamics of fMRI signals during rest and task. *Front. Neurosci.*, 3 JUN(June):1–18, 2012.

de Loynes, B., Navarro, F., and Olivier, B. Localized Fourier analysis for graph signal processing. *Applied and Computational Harmonic Analysis*, 57:1–26, 2022.

de Reus, M. A. and Van den Heuvel, M. P. The parcellation-based connectome: limitations and extensions. *Neuroimage*, 80:397–404, 2013.

Einizade, A. and Sardouie, S. H. A unified approach for simultaneous graph learning and blind separation of graph signal sources. *IEEE Trans. Signal Info. Process. Netw.*, 8:543–555, 2022.

Essen, D. V., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., and Ugurbil, K. The WU-Minn human connectome project: An overview. *Neuroimage*, pages 62–79, 2013.

Georgiadis, K., Adamos, D. A., Nikolopoulos, S., Laskaris, N., and Kompatsiaris, I. Covariation informed graph Slepians for motor imagery decoding. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 29:340–349, 2021.

Ghandehari, M., Guillot, D., and Hollingsworth, K. Gabor-type frames for signal processing on graphs. *Journal of Fourier Analysis and Applications*, 27(2):1–23, 2021.

Ghoroghchian, N., Groppe, D. M., Genov, R., Valiante, T. A., and Draper, S. C. Node-centric graph learning from data for brain state identification. *IEEE Trans. Signal Info. Process. Netw.*, 6:120–132, 2020.

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., and Jenkinson, M. The minimal preprocessing pipelines for the Human Connectome Project. *Neuroimage*, pages 105–124, 2013.

Glomb, K., Queralt, J. R., Pascucci, D., Defferrard, M., Tourbier, S., Carboni, M., Rubega, M., Vulliemoz, S., Plomp, G., and Hagmann, P. Connectome spectral analysis to track EEG task dynamics on a subsecond scale. *Neuroimage*, 221:117137, 2020.

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

Griffa, A., Amico, E., Liégeois, R., Van De Ville, D., and Preti, M. G. Brain structure-function coupling provides signatures for task decoding and individual fingerprinting. *Neuroimage*, 250:118970, 2022.

Huang, W., Bolton, T. A. W., Medaglia, J. D., Bassett, D. S., Ribeiro, A., and Van De Ville, D. A Graph Signal Processing Perspective on Functional Brain Imaging. *Proc. IEEE*, 106(5):868–885, 2018.

Hyde, D. E., Peters, J., and Warfield, S. K. Multi-resolution graph based volumetric cortical basis functions from local anatomic features. *IEEE Trans. Biomed. Eng.*, 66(12):3381–3392, 2019.

Iraji, A., Deramus, T. P., Lewis, N., Yaesoubi, M., Stephen, J. M., Erhardt, E., Belger, A., Ford, J. M., McEwen, S., Mathalon, D. H., et al. The spatial chronnectome reveals a dynamic interplay between functional segregation and integration. *Hum. Brain Mapp.*, 40(10):3058–3077, 2019.

Isufi, E., Gama, F., Shuman, D. I., and Segarra, S. Graph filters for signal processing and machine learning on graphs. *arXiv preprint arXiv:2211.08854*, 2022.

Itani, S. and Thanou, D. Combining anatomical and functional networks for neuropathology identification: A case study on autism spectrum disorder. *Medical Image Analysis*, 69:101986, 2021.

Iturria-Medina, Y., Canales-Rodríguez, E., Melie-García, L., Valdés-Hernández, P., Martínez-Montes, E., Alemán-Gómez, Y., and Sánchez-Bornot, J. Characterizing brain anatomical connections using diffusion weighted MRI and graph theory. *Neuroimage*, (3): 645–660, 2007.

Jones, D. K., Knösche, T. R., and Turner, R. White matter integrity, fiber count, and other fallacies: the do's and don'ts of diffusion mri. *Neuroimage*, 73:239–254, 2013.

Luppi, A. I., Vohryzek, J., Kringelbach, M. L., Mediano, P. A., Craig, M. M., Adapa, R., Carhart-Harris, R. L., Roseman, L., Pappas, I., Peattie, A. R., et al. Distributed harmonic patterns of structure-function dependence orchestrate human consciousness. *BioRxiv*, pages 2020–08, 2022.

Maghsadhagh, S., da Rocha, J. L. D., Benner, J., Schneider, P., Golestani, N., and Behjat, H. A discriminative characterization of Heschl's gyrus morphology using spectral graph features. In *Proc. IEEE Int. Conf. Eng. Med. Biol. Soc.*, pages 3577–3581. IEEE, 2021.

Maier-Hein, K. H., Neher, P. F., Houde, J.-C., Côté, M.-A., Garyfallidis, E., and et al. The challenge of mapping the human connectome based on diffusion tractography. *Nat. Commun.*, (1):1349.

Mansour-L, S., Tian, Y., Yeo, B. T. T., Cropley, V., and Zalesky, A. High-resolution connectomic fingerprints: Mapping neural identity and behavior. *Neuroimage*, 229:117695, 2021.

Mišić, B., Betzel, R. F., Nematzadeh, A., Goni, J., Griffa, A., Hagmann, P., Flammini, A., Ahn, Y.-Y., and Sporns, O. Cooperative and competitive spreading dynamics on the human connectome. *Neuron*, 86 (6):1518–1529, 2015.

Ortega, A., Frossard, P., Kovačević, J., Moura, J. M. F., and Vandergheynst, P. Graph signal processing: Overview, challenges, and applications. *Proc. IEEE*, 106(5):808–828, 2018.

Petrantonakis, P. C. and Kompatsiaris, I. Single-trial NIRS data classification for brain–computer interfaces using graph signal processing. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 26(9):1700–1709, 2018.

Petrantonakis, P. C. and Kompatsiaris, I. Fast feature extraction from large scale connectome data sets using zero crossing counts over graphs. In *Proc. Eur. Signal Process. Conf.*, pages 927–931. IEEE, 2022.

Petrantonakis, P. C. Higher order crossings analysis of signals over graphs. *IEEE Signal Process. Lett.*, 28:837–841, 2021.

Petrovic, M., Bolton, T. A. W., Preti, M. G., Liégeois, R., and Van De Ville, D. Guided graph spectral embedding: Application to the C. elegans connectome. *Network Neuroscience*, 3(3):807–826, 2019.

Preti, M. G. and Van De Ville, D. Decoupling of brain function from structure reveals regional behavioral specialization in humans. *Nat. Commun.*, (1):4747, 2019.

Rué-Queralt, J., Glomb, K., Pascucci, D., Tourbier, S., Carboni, M., Vulliémoz, S., Plomp, G., and Hagmann, P. The connectome spectrum as a canonical basis for a sparse representation of fast brain activity. *NeuroImage*, 244:118611, 2021.

Sandryhaila, A. and Moura, J. M. F. Discrete Signal Processing on Graphs: Frequency Analysis. *IEEE Trans. Signal Process.*, (12): 3042–3054, 2014.

Shuman, D. I., Narang, S. K., Frossard, P., Ortega, A., and Vandergheynst, P. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Process. Magazine*, (3):83–98, 2013.

Shuman, D. I. Localized spectral graph filter frames: A unifying framework, survey of design considerations, and numerical comparison. *IEEE Signal Process. Mag.*, 37(6):43–63, 2020.

Suárez, L. E., Markello, R. D., Betzel, R. F., and Misic, B. Linking structure and function in macroscale brain networks. *Trends Cogn. Sci.*, 24(4):302–315, 2020.

Tarun, A., Abramian, D., Behjat, H., and Van De Ville, D. Graph spectral analysis of voxel-wise brain graphs from diffusion-weighted MRI. In *Proc. IEEE Int. Symp. Biomed. Imaging*, pages 159–163, 2019.

Tarun, A., Behjat, H., Bolton, T., Abramian, D., and Van De Ville, D. Structural mediation of human brain activity revealed by white-matter interpolation of fMRI. *Neuroimage*, 213:116718, 2020.

Tay, D. B. Spectral mappings for graph wavelets. *IEEE Trans. Signal Process.*, 70:3107–3122, 2022.

Tuch, D. S. Q-ball imaging. *Mag. Reson. Med.*, 52(6):1358–1372, 2004.

Van De Ville, D., Britz, J., and Michel, C. M. EEG microstate sequences in healthy humans at rest reveal scale-free dynamics. *Proc. Natl Acad. Sci.*, 107(42):18179–18184, 2010.

Van De Ville, D., Demesmaeker, R., and Preti, M. G. When Slepian meets Fiedler: Putting a focus on the graph spectrum. *IEEE Signal Process. Lett.*, 24(7):1001–1004, 2017.

Vézquez-Rodríguez, B., Liu, Z.-Q., Hagmann, P., and Misic, B. Signal propagation via cortical hierarchies. *Network Neuroscience*, 4(4): 1072–1090, 2020.

Wachinger, C., Salat, D., Weiner, M., Reuter, M., and Initiative, A. D. N. Whole-brain analysis reveals increased neuroanatomical asymmetries in dementia for hippocampus and amygdala. *Brain*, 139(12): 3253–3266, 2016.

Weninger, L., Srivastava, P., Zhou, D., Kim, J. Z., Cornblath, E. J., Bertolero, M. A., Habel, U., Merhof, D., and Bassett, D. S. Information content of brain states is explained by structural constraints on state energetics. *Physical Review E*, 106(1):014401, 2022.

Yeh, F. C., Wedeen, V. J., and Tseng, W. Y. I. Generalized q-sampling imaging. *IEEE Trans. Med. Imag.*, 29(9):1626–1635, 2010.

Zhao, Y., Gao, Y., Zu, Z., Li, M., Schilling, K. G., Anderson, A. W., Ding, Z., and Gore, J. C. Detection of functional activity in brain white matter using fiber architecture informed synchrony mapping. *bioRxiv*, 2022.

Supplementary Materials

Iterative Procrustes analysis

The precision of the Procrustes analysis improves after several rounds of the transformation. The evolution of the Procrustes error for each successive iterations of Procrustes transformation, evaluated through the term μ_K introduced in Algorithm 2, is displayed in Fig. S1(A). After the first iteration of Procrustes transformation, a large is observed in the error (about 60% change); after three iterations (for $J = 3$, 94% drop of error), we obtained the final set of group-averaged eigenmodes, the first few of which are displayed in Fig. S1(B). Major WM features that are relatively smoother than subject-specific bundles are preserved, whereas subject-specific local structures are lost.

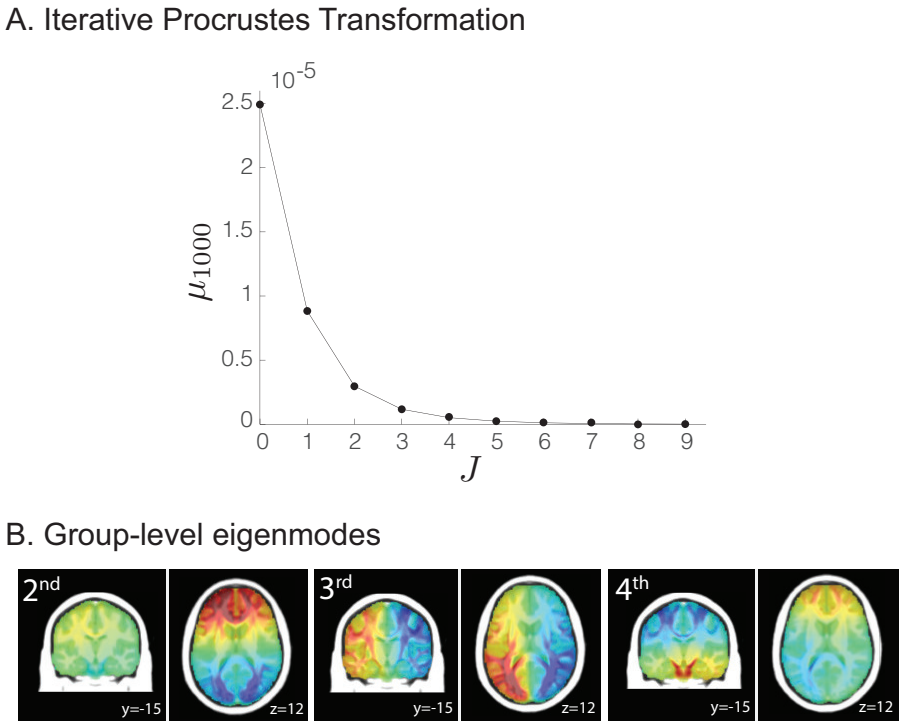

Figure S1. (A) Residual error after successive PTs, cf. Algorithm 1. (B) The second, third, and fourth group-averaged eigenmode.

Cartesian discretisation of a Gaussian via FIR filter design

The molecular displacement of water at voxel i in the direction $\mathbf{r}_{ij}$ can be approximated by a 3D Gaussian distribution with real symmetric diffusion tensor $\mathbf{T}$ as the covariance matrix:

$$p(i, \mathbf{r}_{ij}) = \frac{1}{\sqrt{(2\pi)^3 |\mathbf{T}|}} \exp(-\frac{1}{2} \mathbf{r}_{ij}^\top \mathbf{T}^{-1} \mathbf{r}_{ij}), \tag{S1}$$

where $|\mathbf{T}|$ is the determinant of the diffusion tensor. The calculation of $p(i, \mathbf{r}_{ij})$ requires a discretization step that guarantees a one-to-one mapping between the (continuous) multivariate Gaussian model and the (discrete) weighting of nodes in the brain graph. Note that the Fourier transform of the Gaussian function in (S1) is another Gaussian function. For the $3 \times 3 \times 3$ neighborhood encoding scheme, for a given voxel v_i, the set of values $\{p(i, \mathbf{r}_{ij})\}_{j \in \mathcal{N}_i}$ can be arranged into $3 \times 3 \times 3$ discrete representation, which mimics the structure of a 3D finite impulse response (FIR) filter. Let $\mathbf{h}[\cdot, \cdot, \cdot]$ denote an FIR filter, characterized by it's frequency response $H(\omega)$ as

$$H(\omega) = \sum_{v=-1}^{1} \sum_{l=-1}^{1} \sum_{k=-1}^{1} \mathbf{h}[v, l, k] e^{-j\omega v} e^{-j\omega l} e^{-j\omega k}. \tag{S2}$$

The coefficients $\mathbf{h}$ can then be obtained by a one-to-one mapping from the continuous domain to the discrete domain through the matching of corresponding frequencies. The frequency matching is done by considering the cartesian discretization of a Gaussian using its Fourier domain representation. For simplicity of notation, we present the discretization process for a 1D Gaussian, which can be readily extended to higher dimensions. The 1D Gaussian is given as

$$G(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2} x^2}, \tag{S3}$$

where σ^2 is a constant which describes the variance of the curve and $\frac{1}{\sqrt{2\pi\sigma^2}}$ is a normalization factor. Its Fourier transform is given by

$$G(\omega) = e^{-2\pi^2\sigma^2\omega^2}. \tag{S4}$$

Obtaining a 3-sample discretization of the Gaussian using it's Fourier domain representation is a problem that can be equivalently seen as finding the coefficients of a 3-tap FIR filter, $\mathbf{h}$, matching the Frequency response given in $G(\omega)$ with that of the filter $\mathcal{H}(\omega)$ given as

$$H(\omega) = \sum_{n=-1}^{n=1} \mathbf{h}[n]e^{-j\omega n} = \mathbf{h}[-1]e^{j\omega} + \mathbf{h}[0] + \mathbf{h}[1]e^{-j\omega}. \tag{S5}$$

Noting that $e^x = \sum_{k=0}^{\infty} x^k/k!$, Taylor series expansion of $G(\omega)$ and $H(\omega)$ up to the second order gives

$$G(\omega) \approx 1 - 2\pi^2\sigma^2\omega^2 + \frac{1}{2}(-2\pi^2\sigma^2\omega^2)^2 \tag{S6}$$

$$\begin{aligned} H(\omega) \approx\ & \mathbf{h}[-1] + \mathbf{h}[0] + \mathbf{h}[1] + (\mathbf{h}[-1] - \mathbf{h}[1])j\omega \\ & - \frac{1}{2}(\mathbf{h}[-1] + \mathbf{h}[1])\omega^2. \end{aligned} \tag{S7}$$

Matching the coefficients of the polynomials in (S6) and (S7), a system of linear equations is obtained as

$$\begin{aligned} 1 &= \mathbf{h}[-1] + \mathbf{h}[0] + \mathbf{h}[1] \\ 0 &= \mathbf{h}[-1] - \mathbf{h}[1] \\ 2\pi^2\sigma^2 &= \tfrac{1}{2}(\mathbf{h}[-1] + \mathbf{h}[1]), \end{aligned}$$

where the solution gives the desired FIR impulse response as

$$\mathbf{h} = [2\pi^2\sigma^2, 1 - 4\pi^2\sigma^2, 2\pi^2\sigma^2]. \tag{S8}$$

In the case of the 3D multivariate Gaussian density, we have $G(\mathbf{x}) = (\sqrt{(2\pi)^3|\Sigma|})^{-1}e^{-\frac{1}{2}\mathbf{x}^\top\Sigma^{-1}\mathbf{x}}$, wherein the exponent is a quadratic form in the 3D vector variable $\mathbf{x}$ covering $x-$, $y-$, and $z-$directions, and σ^2 is replaced by a 3×3 invertible covariance matrix Σ that is equivalent to the tensor $\mathbf{T}$ in (S1). The goal is then to find the coefficients of a $3\times3\times3$ FIR filter by matching its frequency response $\mathcal{H}(\omega)$, as given in (S2), to the Fourier transform of the 3D Gaussian, which can be derived in a similar way as done for the 1D case above.