Network alignment and similarity reveal atlas-based topological differences in structural connectomes

Matteo Frigo, Emilio Cruciani, David Coudert, Rachid Deriche, Emanuele Natale, Samuel Deslauriers-Gauthier

Network alignment and similarity reveal atlas-based topological differences in structural connectomes

Matteo Frigo[1,3], Emilio Cruciani[2,3], David Coudert[2], Rachid Deriche[1], Emanuele Natale[2], and Samuel Deslauriers-Gauthier[1]

[1]Université Côte d'Azur, Inria, France
{matteo.frigo,rachid.deriche,samuel.deslauriers-gauthier}@inria.fr
[2]Université Côte d'Azur, Inria, CNRS, I3S, France
{emilio.cruciani,david.coudert,emanuele.natale}@inria.fr
[3]These authors contributed equally to this work.

Abstract

Brain atlases are central objects in network neuroscience, where the interactions between different brain regions are modeled as a graph called connectome. In structural connectomes, nodes are parcels from a predefined cortical atlas and edges encode the strength of the axonal connectivity between regions measured via diffusion Magnetic Resonance Imaging (MRI) tractography. Herein, we aim at providing a novel perspective on the evaluation of brain atlases by modeling it as a network alignment problem, with the goal of tackling the following question: given an atlas, how robustly does it capture the network topology across different subjects? To answer such a question, we introduce two novel concepts arising as natural generalizations of previous ones. First, the graph Jaccard index (GJI), a graph similarity measure based on the well-established Jaccard index between sets; the GJI exhibits natural mathematical properties that are not satisfied by previous approaches. Second, we devise WL-align, a new technique for aligning connectomes obtained by adapting the Weisfeiler-Lehman (WL) graph-isomorphism test. We validated the GJI and WL-align on data from the Human Connectome Project database, inferring a strategy for choosing a suitable parcellation for structural connectivity studies. Code and data are publicly available.

1 Introduction

Due to the immense complexity of the brain, it is impossible to gain any insight into its global operation without simplifying assumptions. One such assumption, which has been widely used by neuroscientists, is that the brain, and in particular the cortical surface, can be divided into distinct and homogeneous areas. Of course the definition of homogeneous areas greatly depends on one's point of view, which has led to a plethora of brain parcellations. For example, the cortical surface has been subdivided based on its cytoarchitecture (Brodmann, 1909), gyri (Desikan et al., 2006), functional organization (Schaefer et al., 2017), axonal connectivity (Gallardo et al., 2018b), and combinations of these and other features (Glasser et al., 2016). There is also significant evidence that cortical regions vary in shape, size, number, and location across subjects and even across individual tasks, making the existence of a single canonical atlas unlikely. In addition to studying the characteristics of specific brain regions defined by a parcellation, there has been a growing interest in their relationship and interactions, an emerging field known as connectomics. In this context, the focus is shifted from understanding how information is segregated in the brain to how it is integrated. For example, through diffusion Magnetic Resonance Imaging (MRI) tractography, structural connections between brain areas can be recovered. The result is a network whose nodes correspond to cortical regions and whose edge weights represent the strength of the structural connectivity between pairs of regions. A similar network can also be built from resting state functional MRI yielding a functional, rather than structural, network. These brain networks, which encode the structural and functional connections of the brain, are referred to as connectomes (Sporns et al., 2005; Hagmann, 2005). Given functional or structural connectomes, their features can be compared across subjects and populations to link network changes to pathology or to further increase our understanding of its organization. An underlying assumption is that a correspondence exists between nodes of the network across subjects, a condition which is usually satisfied by using a group parcellation (Parisot et al., 2015; Gallardo et al., 2018b). The drawback of this strategy is that it ignores any subject specific changes in cortical organization and reduces the specificity of the results. An alternative approach is to construct a mapping between the nodes of the network prior to the comparison, therefore allowing the use of subject specific atlases. To our knowledge, this approach has never been investigated in the field of network neuroscience.

The construction of a mapping between network nodes corresponds to what is known in various fields as *network alignment* or *graph matching* (Barak et al., 2019) and has been applied in fields other than neuroscience (Singh et al., 2008; Conte et al., 2004; Ayache and Faverjon, 1987; Korula and Lattanzi, 2014). Graph alignment solutions (called *alignments*) correspond to a permutation of the labels of the nodes of a graph which maximizes its similarity to a second graph. There is no standard way to measure

the quality of its solutions (Bayati et al., 2013). This is also reflected in the neuroimaging literature, where various measures of similarity between brain networks are used (Chung et al., 2017; Becker et al., 2018; Osmanlıoğlu et al., 2019; Deslauriers-Gauthier et al., 2020; Villareal-Haro et al., 2020). In the context of connectomics, a graph alignment is a reordering of the labels of the nodes of a brain network that maximizes its similarity with a second one while preserving the topology. Describing the brain network through its connectivity (a.k.a. adjacency) matrix, permutations of the node labels correspond to identical permutations of the rows and columns of the connectivity matrix. This problem is distinct from the brain atlas correspondence and parcel matching problems (Mars et al., 2016; Gallardo et al., 2018a). The main difference is that there the permutation acts only on the rows of the connectivity matrix because they find correspondences between *connectivity fingerprints* that rely on external features. Conversely, graph alignment does not rely on any external information and uses only information contained in the topology of the graphs.

The complexity of finding the optimal alignment between two graphs using brute force is exponential in the number of nodes. It is therefore intractable even for the smallest of brain networks, which typically have 50 cortical regions. Spectral methods are a popular approach to the alignment problem (Nassar et al., 2018; Feizi et al., 2019; Hayhoe et al., 2019), despite being subject to limitations (Wilson and Zhu, 2008). Modern machine learning paradigms exploit deep learning techniques fo finding an alignment (Liu et al., 2016; Li et al., 2018; Heimann et al., 2018), however they make use of partially available information about the alignment (Liu et al., 2016), or lack explainability and interpretability.

We first introduce, in Section 2, the *graph Jaccard index* (GJI), a natural objective function for the network alignment problem. For a given alignment, the GJI rewards correct matches while simultaneously penalizing mismatches, overcoming shortcomings of previous approaches (Feizi et al., 2019).

We then propose, in Section 2.3, a new graph alignment heuristic, the *Weisfeiler-Lehman Alignment* (WL-align), based on a weighted variant of the Weisfeiler-Lehman algorithm for graph isomorphism (Weisfeiler and Leman, 1968). WL-align is amenable to concrete interpretability in terms of local network structure around each node (Figure 3) and can be integrated with other heuristics.

After reviewing alternative approaches, we compare WL-align against FAQ (Vogelstein et al., 2015), another efficient brain-alignment heuristic which is solely based on network structure.

2 Theory

A brain network is characterized as an edge-weighted graph $G = (V, E)$, where each of the n nodes represents a brain region and each weight w_{ij} encodes the strength of the connection between regions i and j. The graph G can always be considered as complete, given that if an edge (i, j) is not in G it

can be associated to a null weight $w_{ij} = 0$. The matrix that encodes in position (i, j) the weight of the edge w_{ij} between nodes i and j is called adjacency matrix of G and is denoted as $\mathrm{Adj}(G)$. In the context of connectomics (Sporns et al., 2005; Hagmann, 2005), the adjacency matrix is also known as *connectivity* matrix. In this work we consider only networks with non-negative edge weights. For structural connectomes this does not impose any special preprocessing, since they are usually constructed using streamline count, length, or weights which are already non-negative. However, functional connectomes can contain negative entries because they are typically based on the correlation of resting state functional MRI signals. A practical solution, already used in other studies (Deslauriers-Gauthier et al., 2020), is to threshold the connectomes, therefore replacing negative entries by zeros.

2.1 Brain alignment

In order to compare two networks it is of fundamental importance to establish a correspondence between the nodes of the two graphs. Given two networks $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ respectively of n_1 and n_2 nodes, it is possible to define an injective map[1] $m : V_1 \to V_2$ that is called *graph matching* or *network alignment*. An edge $(u, v) \in E_1$ is correctly *matched* by m if $(m(u), m(v)) \in E_2$ and both edges have the same weight. Notice that a graph matching that matches all edges corresponds to an injective graph homomorphism. In the context of connectomics we will refer to m also as a *brain alignment*. A simple representation of this function is that of a matching matrix P_m of dimension $n_2 \times n_1$ (with $n_2 \geq n_1$) defined as

$$(P_m)_{ij} = \begin{cases} 1 & \text{if } m(j) = i, \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

Notice that if $n_1 = n_2$ then P_m is a permutation matrix. If m is an isomorphism between G_1 and G_2, then the transformation between the adjacency matrices of the two graphs is fully characterized by the matching matrix and is given by

$$\mathrm{Adj}(G_1) = P_m^{\intercal}\, \mathrm{Adj}(G_2)\, P_m. \tag{2}$$

2.2 Quality of brain alignments

Once a brain alignment is identified, its quality can be assessed by evaluating the (dis)similarity of the resulting two networks. On a lexical note, we remark how the concept of *similarity between networks* used throughout this work fits well the standard concept of *matrix similarity* in the particular case where the change of basis matrix is a permutation matrix. In the following, the similarity measures are defined

[1] The existence of the map m is granted whenever $|V_1| \leq |V_2|$.

for equal-sized networks, as typically encountered in connectomics. Classical metrics for this task are based on the comparison of the adjacency matrices of the two graphs by means of Pearson's correlation coefficient, ℓ_p distance, or Frobenius distance (Vogelstein et al., 2015). The norm-based distances estimate the dissimilarity between two graphs G_1 and G_2 by computing the distance between their adjacency matrices as follows

$$d_t(G_1, G_2) = \| \operatorname{Adj}(G_1) - \operatorname{Adj}(G_2) \|_t \tag{3}$$

where t indicates the type of norm (p for ℓ^p norms and F for Frobenius norm). Note that higher distance corresponds to lower similarity. Another similarity measure that has been widely adopted in neuroimaging and brain connectivity is correlation; among the many definitions of correlation, we consider

$$C(G_1, G_2) = \frac{\langle \vec{\operatorname{Adj}}(G_1), \vec{\operatorname{Adj}}(G_2) \rangle}{\| \vec{\operatorname{Adj}}(G_1) \|_2 \cdot \| \vec{\operatorname{Adj}}(G_2) \|_2} \tag{4}$$

where the numerator is the scalar product between the vectorizations of the adjacency matrices of the two graphs and the denominator is the product of their norms. This similarity measure is also known as *cosine similarity*, since it corresponds to the cosine of the angle between the two vectors. More recently, distances based on geometrical (Venkatesh et al., 2020) and homological (Chung et al., 2017) properties of the networks have been introduced. All such measures capture some aspects of the similarity between two graphs, but none of them satisfies the following requirements:

- arising as a natural generalization of other similarity measures for less structured data, e.g., for sets of values without a network structure;
- being applicable to the algorithmic graph isomorphism and induced subgraph isomorphism problems, as fundamental special cases of the problem of measuring the similarity between two graphs;
- being simple enough so that its value can be easily interpreted;
- giving a straightforward notion of metric in the considered space.

We therefore propose a new measure obtained by generalizing the Jaccard similarity index, a similarity metric widely adopted in data mining, so that algorithmic problems such as induced subgraph isomorphism can be retrieved as special cases. Moreover, while our proposed measure assigns a clear meaning to the correspondence between two edges of two given graphs, it also crucially depends on the global network structure.

2.2.1 Weighted graph Jaccard similarity index

The *Jaccard similarity index* was originally proposed in the context of set theory to measure the similarity between two sets S and T. It is computed as the ratio between the size of their intersection and the size of their union, that is

$$J(S,T) = \frac{|S \cap T|}{|S \cup T|}. \tag{5}$$

An example of what is measured by the Jaccard index on sets is given in Figure 1. Notice that $J(S,T)$ is defined in the $[0,1]$ range and the extreme values are attained either when the intersection of the two sets is empty (i.e., $S \cap T = \emptyset \implies J(S,T) = 0$) or when the two sets are equal (i.e., $S = T \implies J(S,T) = 1$).

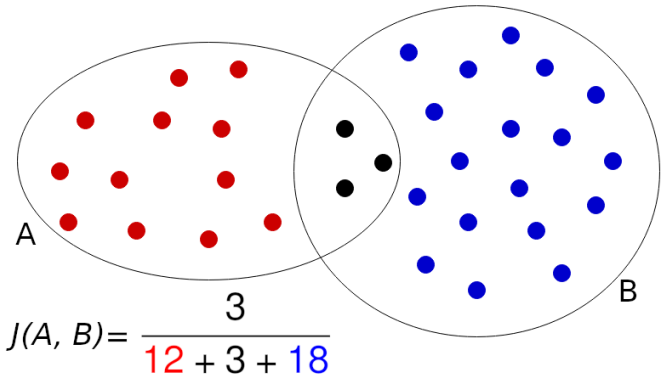

Figure 1: The two sets contoured by the circles have a non-empty intersection marked by the black dots. The Jaccard similarity index between the two sets is the result of the ratio between the number of elements in the intersection and the number of elements in the union of the two sets. The resulting Jaccard index is equal to $J = 3/33 \approx 0.09$.

The Jaccard similarity index has also been generalized to non-negative real vectors and, in this more general setting, is also known as Ruzicka similarity. In detail, given two vectors $x, y \in \mathbb{R}^d$ such that $x_i \geq 0$ and $y_i \geq 0$, their *weighted Jaccard similarity index* can be computed as $J(x,y) = \frac{\sum_{i=1}^{d} \min(x_i, y_i)}{\sum_{i=1}^{d} \max(x_i, y_i)}$. Note that Jaccard similarity index between two sets is a special case whenever the vectors x, y are binary and their dimension d is equal to the size of the union of the two sets. Noticeably, the weighted Jaccard similarity index induces a metric (Charikar, 2002).

Our adaptation of the concept of Jaccard similarity index to weighted graphs is based on the identification of the nodes of the two graphs. Given two brain networks G_1 and G_2 with adjacency matrices $\mathrm{Adj}(G_1) = A$ and $\mathrm{Adj}(G_2) = B$, the *weighted graph Jaccard similarity index* of G_1 and G_2 is

$$J(G_1, G_2) = \frac{\sum_{(i,j) \in \mathcal{E}} \min(A_{i,j}, B_{i,j})}{\sum_{(i,j) \in \mathcal{E}} \max(A_{i,j}, B_{i,j})} \tag{6}$$

where $\mathcal{E}$ is the set of all possible pairs of nodes. For the sake of the present work, we remark that we can

think of B as having been previously aligned to A via Equation 2. Alternatively, the weighted graph Jaccard similarity index is defined as the weighted Jaccard index of the vectorizations of the graphs' adjacency matrices. Notice that $J(G_1, G_2)$ is not well defined when both G_1 and G_2 are empty (i.e., $E_1 = E_2 = \emptyset$). Whenever $\text{Adj}(G_1) = \text{Adj}(G_2)$, the min and the max in Equation (6) coincide and $J(G_1, G_2) = 1$. On the contrary, if G_1 and G_2 do not have any edge in common (i.e., $E_1 \cap E_2 = \emptyset$), the numerator of Equation (6) will be equal to zero and $J(G_1, G_2) = 0$. An example of how the GJI acts on two graphs is given in Figure 2.

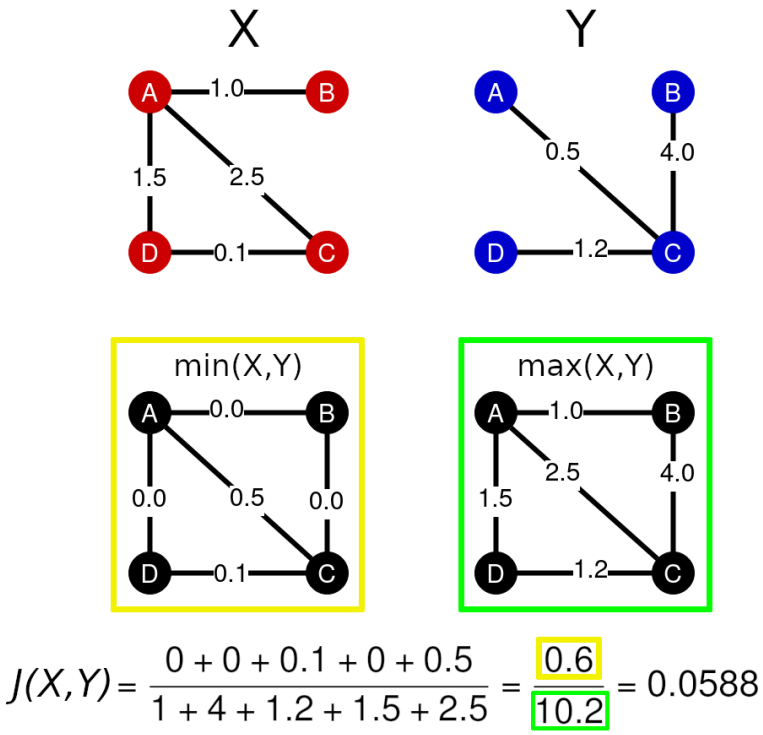

Figure 2: This figure shows an example of how to compute the GJI between two compatible graphs X and Y. For each pair of nodes $i, j \in \{A, B, C, D\}$, one computes the minimum and maximum between $X_{i,j}$ and $Y_{i,j}$. These two quantities will be used to define the numerator and the denominator of the GJI defined in Equation (6). As shown in the min (yellow) and max (green) graphs, edges that are not in a graph are associated to a null weight. The GJI is then computed as the ratio between the sum of the minimal weights and the sum of the maximal weights.

We have so far formally established the notion of network alignment (Equation (1)), and presented the Jaccard index as a principled way to measure the quality of an alignment (Equation (6)). We are thus ready, in the next section, to describe our variant of the Weisfeiler-Lehman heuristic and to show how to employ it to construct a network alignment.

2.3 Weisfeiler-Lehman network alignment

In this work we propose a brain alignment technique that allows to define the graph matching m between two brain networks G_1 and G_2 with a three-step procedure:

1. For each node u in both graphs, define a vector H_u that we call its *signature*.
2. Define a *complete bipartite graph* where on one side there are the nodes of the first graph and on the other side there are the nodes of the second graph; the euclidean distance between two signatures becomes the weight of each edge of the bipartite graph.
3. The graph matching is given by the solution of the *minimum weight bipartite matching problem*, also known as *assignment problem*, on the bipartite graph previously defined.

The novelty element of this brain alignment algorithm is given by the definition of the node signature, defined by an algorithm inspired by the Weisfeiler-Lehman method for graph isomorphism testing (Weisfeiler and Leman, 1968). For this reason, WL-align seems to be a reasonable name for our proposed brain alignment algorithm.

Node signature The signature that we associate to each node of the two graphs describes the local connectivity pattern of the node. It relies on the concept of *volume* of a node, which is defined as the sum of the weights of the edges outgoing from the node itself, namely

$$\mathrm{vol}(v) = \sum_{u \in V} w_{uv} \tag{7}$$

where v is the node of which we compute the volume $\mathrm{vol}(v)$, V is the set of nodes in the graph, and w_{uv} is the weight of the edge connecting nodes u and v. The algorithm that defines the signature of node u considers the subnetwork H induced by the nodes that are reachable from u in at most ℓ hops. At each of these hops, H retains only the k nodes which are connected by the edges with the highest weights. This subnetwork is a complete k-ary tree of depth ℓ which can be obtained from a breadth-first search (BFS) starting from u, and has a total of $d = \sum_{i=0}^{\ell} k^i$ nodes. For this reason the parameters k and ℓ are respectively called *width* and *depth*. The entries of the signature $H_u \in \mathbb{R}^d$ are then computed starting from u and following the BFS by recursively estimating the contribution of each edge to the volume of the considered node. A graphical intuition of how H_u is defined is illustrated in Figure 3 and a rigorous presentation of the whole WL-align algorithm is provided in the Supplementary Materials.

Bipartite graph Once a signature is computed for each node of the two graphs, we define a weighted complete bipartite graph where the nodes on the left represent the nodes of the first graph, while the

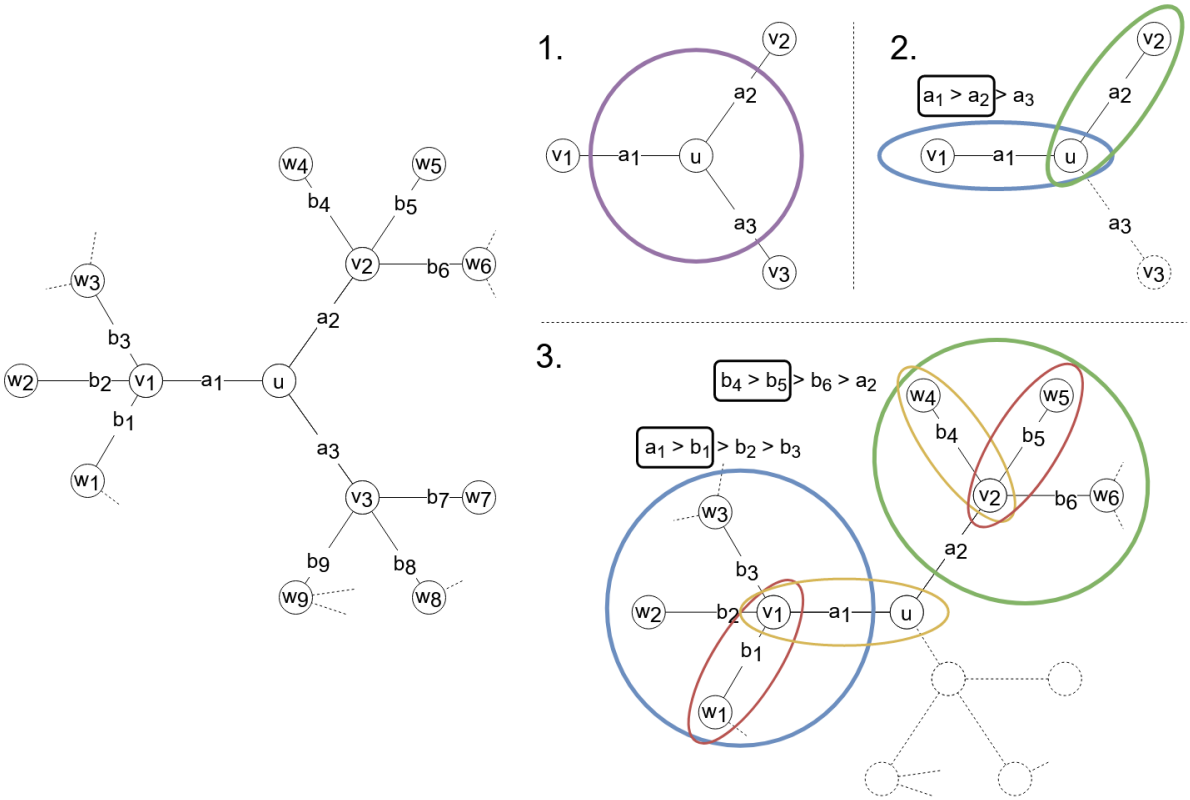

Figure 3: The graph on the left is the one that serves as an example for explaining the algorithm for computing the signature H_u of node u with $k = 2$ and $\ell = 2$. The first entry of H_u is $H_u[1] = \mathrm{vol}(u)$, which is obtained by considering all the edges touching node u contoured by the purple circle in panel 1. The focus moves then to the two edges with highest weights (a_1 and a_2) which are marked by the blue and green circles in panel 2. They are considered in decreasing order and the corresponding entries in H_u are computed from Equation (1) of the Supplementary Material. For instance, the second entry of H_u is equal to $H_u[2] = H_u[1] \cdot a_1/\,\mathrm{vol}(u)$. The third entry is computed in an analogous way as $H_u[3] = H_u[1] \cdot a_2/\,\mathrm{vol}(u)$. This concludes the definition of the first $1 + k$ entries of H_u. The next entries are defined by considering first the blue and then the green subnetworks in panel 3. These are defined by rescaling the previous entries of H_u w.r.t. the weights of the nodes that are visited next. The fourth entry is equal to $H_u[4] = H_u[2] \cdot a_1/\,\mathrm{vol}(v_1)$ and the fifth is $H_u[5] = H_u[2] \cdot b_1/\,\mathrm{vol}(v_1)$. Analogously, the sixth and the last entry will be $H_u[6] = H_u[3] \cdot b_4/\,\mathrm{vol}(v_2)$ and $H_u[7] = H_u[3] \cdot b_5/\,\mathrm{vol}(v_2)$.

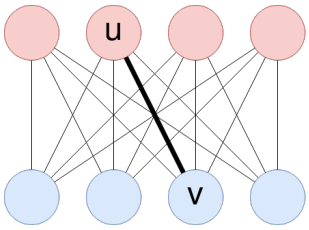

Figure 4: The red and blue nodes in the two rows represent the two graphs $G_1 = (V_1, E_1)$ $G_2 = (V_2, E_2)$ being aligned. The displayed complete bipartite graph is the one constructed in the second step of the WL-align algorithm. Each edge has weight equal to the euclidean distance between the signatures of the nodes that it connects. For instance, the weight associated to the edge connecting nodes $u \in V_1$ and $v \in V_2$ is $\|H(u) - H(v)\|_2$, where $H(\cdot)$ is the node signature defined in the first step of the WL-align algorithm.

nodes on the right represent the nodes of the second graph. The edge-weights encode the distance between the signatures of pairs of nodes belonging to different graphs. Figure 4 shows a simple example of the defined bipartite graph.

Assignment problem The final step towards finding the wanted matching is the resolution of the assignment problem corresponding to the bipartite graph defined in the previous paragraph. The matching can be found by selecting a minimum-weight graph matching, namely a subset of edges of the bipartite graph such that every node has degree 1 and the sum of the weights of all edges of the subset is minimal. This assignment problem is efficiently solved by the Hungarian algorithm (Jacobi, 1890; Kuhn, 1955).

3 Methods

We processed the data of 100 unrelated subjects from the HCP database and obtained the structural brain networks via dMRI-based tractography as described in Section 3.1. For each of the 100 subject we considered 23 parcellations (Desikan, Glasser, Gallardo at 11 different resolutions, Schaefer at 10 different resolutions) described in Section 3.1.1, obtaining a total of 2300 weighted graphs. For each parcellation, we retrieved a network alignment between each of the 5050 pairs of subjects using WL-align, which is the novel technique introduced in this work, and the state-of-the-art competitor FAQ, as described in Section 3.2, for a total of 232300 alignments. The quality of the obtained alignments was then assessed using four network similarity measures described in Section 3.3.

3.1 Data and preprocessing

To build the structural brain networks, we considered the preprocessed data of the Human Connectome Project (HCP) database (U100 subject group) (Van Essen et al., 2012; WU-Minn Human Connectome Project consortium, 2017; Glasser et al., 2013). For each subject, a five-tissue-type image (Smith et al., 2012) was obtained using the Freesurfer pipeline (Fischl, 2012) invoked through Mrtrix3 (Tournier et al., 2019). A response function was estimated for the white matter, gray matter, and cerebrospinal fluid using a maximal spherial harmonic order of 8 for all tissues (Jeurissen et al., 2014). The fiber orientation distribution functions (fODFs) were then computed using the multi-shell multi-tissue constrained spherical deconvolution algorithm (Jeurissen et al., 2014). Finally, the fODFs were used as input for probabilistic anatomically-constrained tractography performed with the iFOD2 algorithm (Smith et al., 2012) seeding from the gray matter - white matter interface and obtaining a total of five million streamlines per subject.

3.1.1 Parcellations

The four parcellations considered in this work subdivide the cerebral cortex following different characteristics of the brain. The Desikan (Desikan et al., 2006) parcellation is based on the manual segmentation of a template of the brain cortex that takes into account the morphological consistencies of healthy human brains. For each subject, it was obtained directly from the Human Connectome Project database (*aparc+aseg.nii.gz*) together with the cortical surface in fslr32k space. The Glasser parcellation Glasser et al. (2016) follows a multi-modal approach that considers cortical architecture, function, connectivity, and topography. Its projection onto the fslr32k space was obtained from the BALSA repository (of Medicine, 2020). The Gallardo parcellation Gallardo et al. (2018b) is based on the segmentation of the structural connectivity profiles associated to each point of the cortical surface and the Schaefer parcellation Schaefer et al. (2017) is based on the analysis of the co-activation patterns of the brain by means of the analysis of resting-state functional connectivity. The Gallardo and the Schaefer parcellations were computed with a granularity of 100, 200, 300, 400, 500, 600, 700, 800, 900, and 1000 parcels. The Gallardo parcellation was computed also with a granularity of 50 parcels. We extracted the 11 Gallardo atlases from the extrinsic connectivity parcellation of Gallardo et al. (Gallardo et al., 2018b). The used Schaefer atlas Schaefer et al. (2017) was downloaded from the repository of the CBIG laboratory (Yeo, 2020) for the *seven-networks* parcellation (Yeo et al., 2011). The use of multi-resolution parcellations reflects the multi-scale nature of the brain network and allows to inspect how the atlas resolution affects the similarity and the alignment of brain networks.

3.1.2 Connectomes

For each subject and parcellation an in-house software was used for counting the number of streamlines connecting each pair of regions. The obtained quantity was encoded as the weight of the edge connecting the two parcels in the brain network. All the edge weights were then divided by the sum of all the weights in the graph. A total of 23 connectomes of different sizes was obtained for each subject. Given the limitations of dMRI-based tractography, self-connections were excluded from the connectomes, i.e., the diagonal of the adjacency matrix is set to zero. Because of the high resolution of some parcellations, some regions turned out to be isolated (i.e., not connected to any other region). In order to have a connected graph, which is a requirement of the WL-align algorithm, we artificially connected these isolated (i.e., zero-volume) nodes to the others by adding small-weighted edges connecting each of these nodes to all the other nodes in the graph. This weight was set to 1 (before normalization), which from the point of view of tractography is equivalent to the existence of one single streamline connecting the region to the others. The obtained graphs are undirected and weighted.

3.2 Network alignments

In order to assess the ability of WL-align to retrieve the wanted brain alignment map, we prepared the dataset in a way that allows to test the quality of the alignment against a known ground truth. In practice, for each parcellation p, we randomly permuted the brains of all subjects and kept track of the permutation maps. By using the permutation maps, we computed the ground truth matching m^* for each pair of brains of each parcellation.

For the same set of brains, we also computed two graph matchings. The first is m_{WL}, which is computed with the proposed WL-align technique. The width and depth parameters of the WL-align algorithm were fixed to $k = \lfloor \log_2 n \rfloor$, where n is the number of nodes in the considered network (i.e., one hemisphere), and $\ell = 2$.

The second is m_{FAQ}, which is computed with the Fast Approximate Quadratic Programming for Graph Matching (FAQ) algorithm (Vogelstein et al., 2015), which is the state-of-the-art technique for network alignment. Given two graphs, FAQ first finds a bi-stochastic matrix which acts as a relaxed permutation matrix that transforms the adjacency matrix of the first graph into one with minimal Frobenius distance from the adjacency matrix of the second graph, then projects the obtained solution onto the proper space of permutation matrices. Notice that optimality with respect to the Frobenius distance might not correspond to absolute optimality. The chosen termination criteria for FAQ are a maximum number of 5000 iterations and an absolute tolerance on the objective function of 10^{-7}. Since an implementation of FAQ was not available, we provide our own implementation that also corrects some errors that we identified in the description provided in the original paper (Vogelstein et al., 2015) (see Supplementary Materials for a description of the error).

Both algorithms were run separately on each hemisphere of the brain and the two resulting partial alignments were then combined into a single one.

3.3 Quality of alignments

Given two networks $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ defined on the same parcellation and given a matching m between them, we consider the following metrics to evaluate the quality of the matching m.

- Node Matching ratio (NMr): the fraction of nodes that have been correctly matched by m with respect to the ground truth matching m^* (known a priori), namely

$$\mathrm{NMr}(m) = \frac{|\{u \in V_1 : m(u) = m^*(u)\}|}{|V_1|}. \tag{8}$$

The NMr metric is defined in the $[0, 1]$ range and higher values correspond to better alignments.

- Graph Jaccard index J: as defined in Equation (6), namely

$$\mathrm{J}(m) = J(m(G_1), G_2) \tag{9}$$

where, with an abuse of notation, we write $m(G_1)$ to denote the relabeling of the nodes obtained by applying the matching m on the nodes of G_1. Recall that the graph Jaccard index is defined in the $[0, 1]$ range and higher values correspond to better alignment.

- J-ratio (Jr): the ratio between the graph Jaccard index $J(m)$ obtained by m and the graph Jaccard index J_p^* obtained by the ground truth matching m^*, namely

$$\mathrm{Jr}(m) = \frac{J(m)}{J(m^*)}. \tag{10}$$

When the ground truth matching m^* is also an optimal matching, the denominator $J(m^*)$ acts as a normalization factor, which takes into account how complex it is to retrieve the matching m^* in terms of Jaccard similarity; under such assumption of ground-truth optimality, the Jr metric takes value in the $[0, 1]$ range, where higher values correspond to better alignment.

- Frobenius norm (FRO): the Frobenius norm of the difference between the adjacency matrices of $m(G_1)$ and G_2, namely

$$\mathrm{FRO}(m) = \|\mathrm{Adj}(m(G_1)) - P_m^{\intercal} \mathrm{Adj}(G_2) P_m\|_F \tag{11}$$

where, as also done for J, we write $m(G_1)$ to denote the relabeling of the nodes obtained by applying the matching m on the nodes of G_1 and where P_m is the permutation matrix associated to the matching m, as defined in Equation (1). The FRO metric is defined in the $[0, 2]$ range (since the adjacency matrices both have norm 1) and lower values correspond to better alignment.

For each considered parcellation p and for each network alignment algorithm of interest x (either WL-align or FAQ), we report the average quality metric, computed among all pairs of brains in the parcellation. For example, considering NMr as quality metric, we compute

$$\mathrm{NMr}_p^x = \frac{1}{|\mathcal{P}|} \sum_{(G_1, G_2) \in \mathcal{P}} \mathrm{NMr}(m)$$

where $\mathcal{P}$ is the set of all pairs of brains with parcellation p and m is the matching found by algorithm x for the input pair of graphs G_1, G_2. Analogously, this is done for all quality metrics.

4 Results

4.1 Experiments

We processed the data of 100 unrelated subjects from the HCP database obtaining the structural brain networks as detailed in Section 3. For each of the 100 subjects we considered 23 parcellations (Desikan, Glasser, Gallardo x 11, Schaefer x 10), obtaining 2300 weighted graphs. For each parcellation, we retrieved a network alignment between each pair of subjects using WL-align and FAQ. The ability of WL-align to retrieve the correct brain-alignment map was quantitatively evaluated by means of four similarity measures. First, a novel measure of similarity between brain networks called graph Jaccard Index was introduced in Section 2 as an adaptation of the concept of Jaccard index between sets. While behaving in a way which is similar to the commonly used correlation index defined in Equation (4), the graph Jaccard index has the property of defining a metric in the space of connectomes. This is a remarkable property in the context of modern data science, as many standard machine learning techniques can be applied only in metric spaces. The second considered similarity measure is the aforementioned correlation index defined in Equation (4), also known as cosine similarity. The third similarity measure is the Frobenius distance defined in Equation (3), which actually is a dissimilarity measure, therefore connectomes showing higher Frobenius distance are less similar and vice-versa. The node matching ratio defined in Equation (8) is the last considered similarity measure.

4.2 Comparison between similarity measures

Each employed similarity metric answers a specific question. The node matching ratio corresponds to what the expression suggests, namely it counts how many nodes were correctly matched and normalizes the result by the number of nodes in the graph. The other similarity measures have less intuitive definitions. For this reason, we measured how much the subjects in the considered datasets are similar to each other with respect to each metric and each parcellation. We recall that the dataset contains only healthy unrelated subjects which do not exhibit any family structure (WU-Minn Human Connectome Project consortium, 2017). This allows to compare how the within-group similarity reacts to the change in resolution and type of the used parcellation.

Figure 5 shows the average similarity between all the subjects in the cohort evaluated with the graph Jaccard index, the Frobenius norm, and the correlation. The used alignment is the one defined by the ground truth, which for our experiments is known a priori. The most noticeable fact is that the graph Jaccard index and the correlation show an inverted trend with respect to the one of the Frobenius norm. A higher number of parcels gives both lower Jaccard/correlation index and lower Frobenius distance,

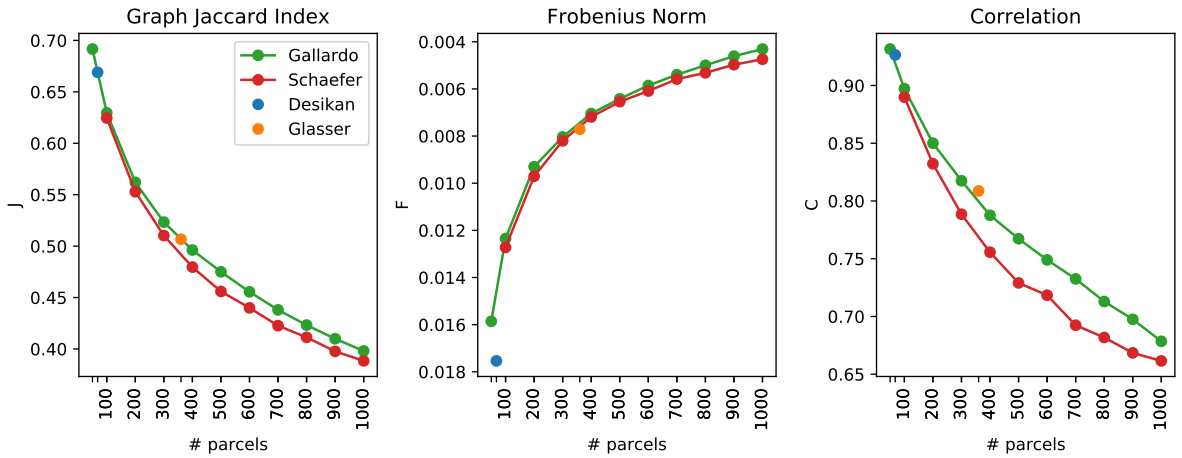

Figure 5: Each point shows the average similarity between every pair of subjects in the considered cohort measured on connectomes obtained with a specific parcellation. The used alignment is the one defined by the ground truth, which in our experiments is known a-priori. All panels show the similarity measure as a function of the number of parcels of the considered atlas. A higher graph Jaccard index and correlation corresponds to higher similarity. On the contrary, a higher Frobenius norm corresponds to lower similarity. In order to keep the intuition that *higher is better*, the y axis of the Frobenius norm is flipped.

which a priori is counter-intuitive. This phenomenon is due to the fact that the Frobenius norm is incapable of capturing the relative difference between edge weights and instead considers only the absolute difference between them. As a matter of fact, parcellations with a higher number of parcels will create brain networks with lower edge weights, since the same amount of connectivity (i.e., the same number of streamlines) is distributed among a number of edges that grows quadratically with the number of regions. For this reason, the absolute value of the edge weights will be lower, giving also a lower absolute difference. On the contrary, the graph Jaccard index and the correlation, which are able to capture the relative difference between edge weights, show lower similarity values between brain networks obtained with a higher number of parcels compared to brain networks obtained with a lower number of parcels. This difference suggests that the graph Jaccard index and the correlation mitigate the influence of the number of parcels in the estimation of the similarity between the compared brain networks. Another observation can be done on the singular nature of the Desikan and Glasser parcellations. While the results obtained on the Glasser parcellation are in line with those obtained with the Gallardo and Schaefer parcellations, connectomes obtained with the Desikan parcellation show a similarity in line with the Gallardo parcellation only when this is measured with the graph Jaccard index and the correlation.

4.3 Computing brain alignments with WL-align

In this work, the concept of *similarity between networks* was used as a proxy for the quality of a brain alignment, since a good graph matching is expected to correspond to a higher similarity between the

aligned graph and the ground truth. A separate analysis was performed for each of the 23 considered parcellations. First, an alignment was computed between each pair of subjects with the proposed technique WL-align and the state-of-the-art algorithm FAQ, then the similarity between the aligned network and the ground truth network was computed with the similarity measures listed in Section 3.3. The node matching ratio (NMr) tells the proportion of nodes that were correctly matched by the alignment. This measure does not give any information about the topological differences between the original and the aligned graph, but it gives an important insight on how many nodes are correctly labeled, which may be of fundamental importance in connectomic studies where the regions are associated to a specific function of the brain. The second used metric is the Jaccard similarity index introduced and described in this paper, while the third employed metric is the Jaccard index ratio. The latter measures how the Jaccard index performed with respect to the Jaccard index of the ground-truth matching shown in Figure 5, which is known a priori from the design of the experiment. It differs from the raw Jaccard index in the sense that it takes into account the complexity of the alignment problem, which we showed in the previous section to be more difficult when the number of parcels is higher. A final comparison was made using the Frobenius distance, which is what the FAQ algorithm is designed to minimize. This makes it particularly interesting since we expect FAQ to give Frobenius distance which is less or equal to the one obtained with WL-align.

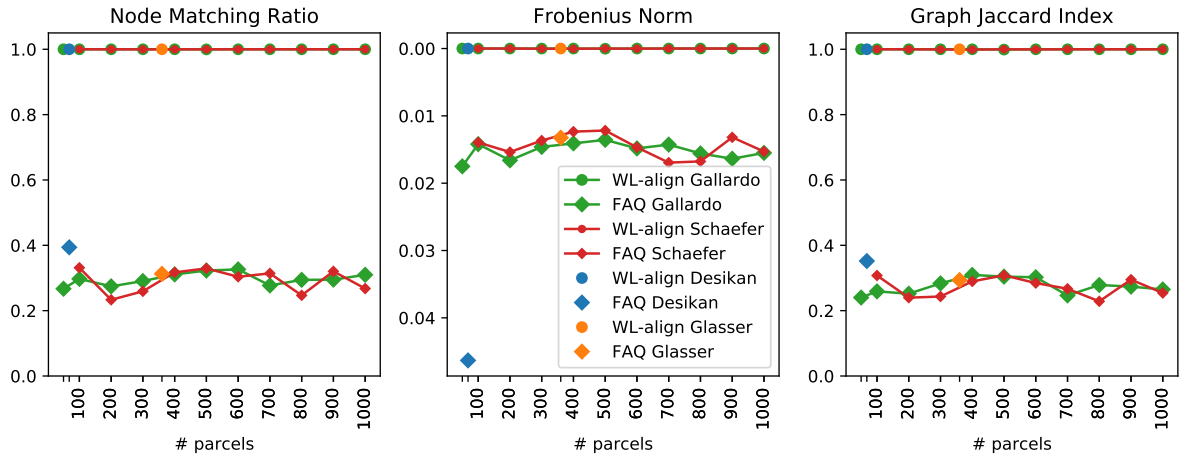

Figure 6: The displayed results concern the alignment between the structural brain network of one subject and its randomly-permuted version. Each panel shows one type of similarity between the aligned networks. Higher values of NMr, Jaccard index and Jaccard index ratio correspond to higher similarity, whereas the Frobenius norm is higher when similarity is lower. In order to keep the intuition that *higher is better*, the y axis of the Frobenius norm is flipped. In each panel, one point corresponds to the average (among subjects) similarity computed between brain networks obtained on a specific parcellation and aligned with one technique among WL-align and FAQ. All the four plots show the similarity as a function of the number of parcels in the considered atlas.

Subject-wise analysis In the context of this work, the simplest non-trivial alignment to be retrieved is the one between the brain network of a subject and its randomly-permuted version. In this case, a good alignment algorithm is expected to always retrieve the ground truth alignment. In Figure 6 we report the average similarity between the ground truth and the obtained alignment. We notice that WL-align consistently achieves the best possible performance with respect to all the considered metrics. In particular, the naive metric of the node matching ratio always gives similarity equal to 1, meaning that WL-align correctly labels all the nodes whenever a structural brain network is aligned against a randomly-permuted version of itself. These considerations are true for every parcellation. On the contrary, FAQ does not solve the self-alignment problem exactly. All the considered metrics highlight a poor performance of FAQ both in absolute terms and compared to WL-align. As a matter of fact, FAQ on average yields at most 40% of correctly matched nodes, while WL-align consistently gives 100% of correctly matched nodes. Also, different parcellations behave differently when FAQ is employed; for instance, the Desikan parcellation gives lower Frobenius similarity with respect to the other parcellations but shows higher Jaccard index and node matching ratio.

Full cohort analysis When all the subject are aligned with the permuted version of each other, the problem is more complicated. Even though we considered healthy subjects whose acquisition followed the same protocol and that have been processed in an identical way, the subject-specific differences and the intrinsic noise of the data yield estimated structural brain networks that are in practice different among each other, despite being substantially coherent. In order to assess the ability of the proposed alignment technique to overcome these differences and yield an alignment as close as possible to the ground truth, we considered all the alignments between each pair of subjects, including the ones between a subject and a randomly-permuted version of itself. The brain alignments obtained with WL-align are compared to the one computed with FAQ and presented in Figure 7, which reports the average similarity between the obtained alignment and the ground truth alignment among all the possible pairs of subjects. In terms of Frobenius norm, the alignments obtained with WL-align and FAQ are very similar, with WL-align systematically showing slightly higher Frobenius similarity. The performance of the Gallardo parcellation is indistinguishable from the one of the Schaefer parcellation. Also, the Desikan and the Glasser parcellations have results in line with the Schaefer and Gallardo parcellation when the alignment is obtained with WL-align. This is not true for the Desikan parcellation when FAQ is employed. Recalling that FAQ is a technique that is inherently based on the Frobenius norm and WL-align is not, we can notice that WL-align gives a brain alignment that does satisfies also the optimality criteria of FAQ, additionally to its own. A second thing that we can notice about the Frobenius norm is that it exhibits the same phenomenon as in Figure 5, where the Frobenius similarity increases with the number of parcels.

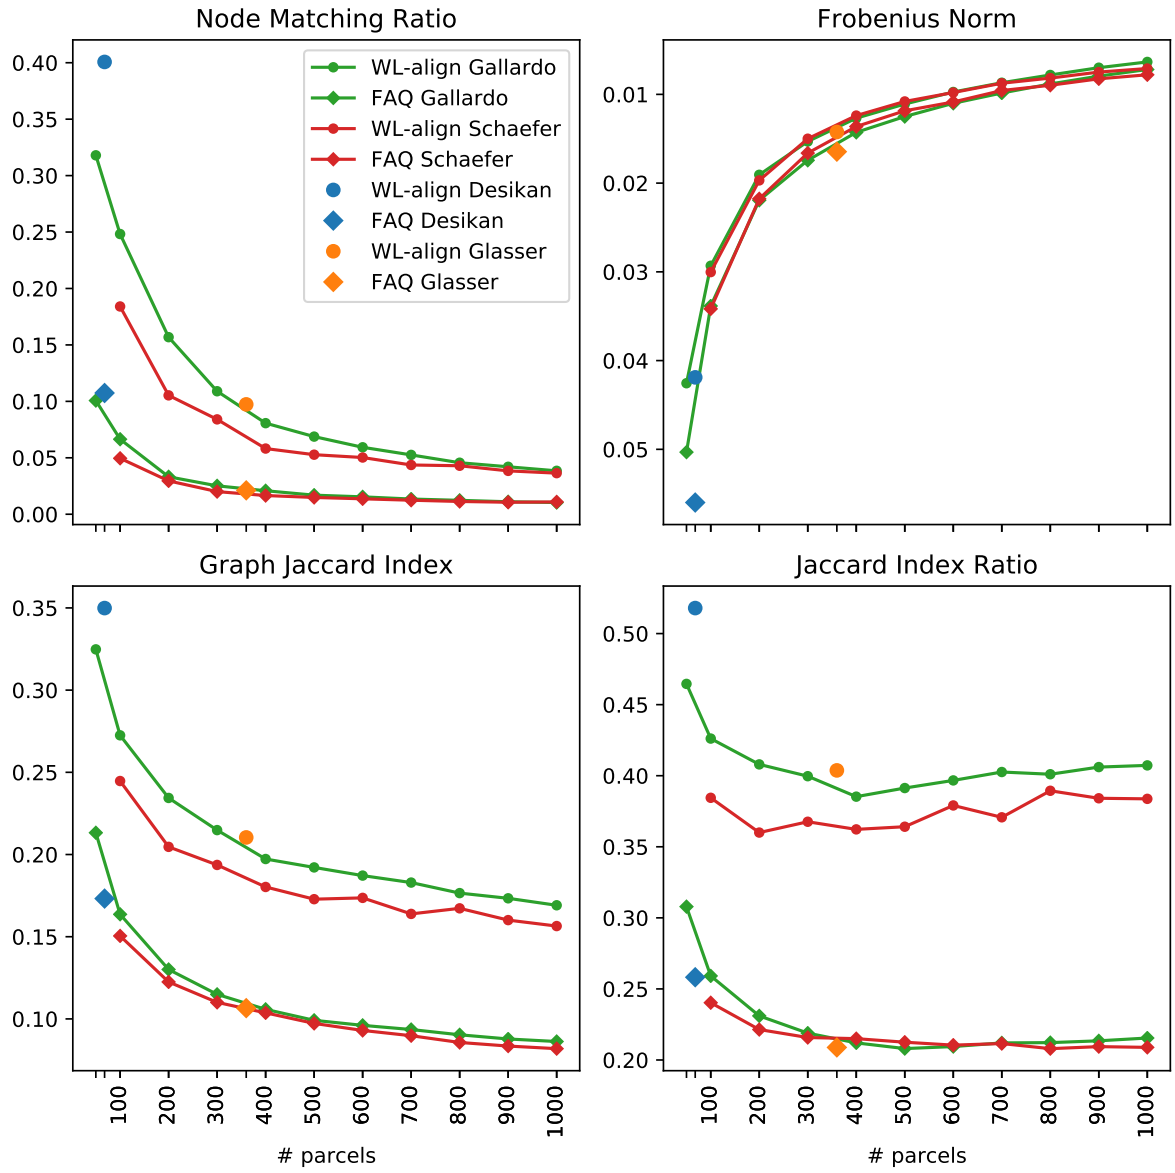

Figure 7: The displayed results concern the alignment between the structural brain networks of each pair of subjects including the self-comparisons. Each panel shows one type of similarity between the aligned brain networks. Higher values of NMr, Jaccard index and Jaccard index ratio correspond to higher similarity, whereas the Frobenius norm is higher when similarity is lower. In order to keep the intuition that *higher is better*, the y axis of the Frobenius norm is flipped. In each panel, one point corresponds to the average (among subjects) similarity computed between brain networks obtained on a specific parcellation and aligned with one technique among WL-align and FAQ. All the four plots show the similarity as a function of the number of parcels in the considered atlas.

This phenomenon appears for the same reason as before, namely the Frobenius norm does not capture the relative difference between the edge weights in the compared networks. All the other employed similarity metrics suggest that WL-align has superior performance with respect to FAQ. While FAQ has almost identical performances when applied on the Gallardo and the Schaefer parcellations, WL-align shows relevant and previously unobserved differences between the performances of the two. In particular the Gallardo parcellation allows to retrieve better alignments with respect to the Schaefer parcellation. This may be due to the fact that we are studying structural connectivity, therefore the use of a function-based parcellation like the one of Schaefer may affect the quality of the alignment when compared to the structural connectivity computed on a structure-based parcellation like the one of Gallardo. Looking at the behavior of the Desikan and the Glasser parcellation, we notice two different scenarios. The Glasser parcellation shows Jaccard similarity slightly higher than the one of the Gallardo parcellation, suggesting that the multi-modal nature of the atlas well captures the structural connectivity features that we are looking at. This is true both when we consider WL-align and FAQ. The Desikan parcellation behaves quite differently. Its performance does not follow the pattern of the other parcellations and is not aligned with neither the Gallardo nor the Schaefer parcellation. It has a sensibly superior performance compared to any other parcellation when the WL-align algorithm is considered and this is visible through every employed similarity metric. The same difference does not emerge so clearly when FAQ is employed. We finally notice that atlases with > 400 parcels all behave very similarly, namely they reach a plateau in terms of Jaccard index, Jaccard index ratio and node matching ratio. This is true both when WL-align and FAQ are employed. The performance in this range is lower than the one in the $50 - 400$ parcels range.

4.4 Self matching rate

Figure 8 illustrates the self matching rate for each region of 9 example atlases, i.e., the fraction of times regions were correctly matched when aligning different brains represented using the same atlas. It is clear that, as the number of parcels is increased, the matching rate is reduced. This can be explained by the increased difficulty of the alignment problem, but also by a decrease in the signal-to-noise ratio of the connectomes driven by the reduction in parcel size. It is also interesting to note that the matching rate does not appear to be symmetric across hemispheres. For example, the right inferior parietal region of the Desikan atlas obtains relatively high matching rate of roughly 0.8, whereas the contralateral region only obtains roughly 0.4.

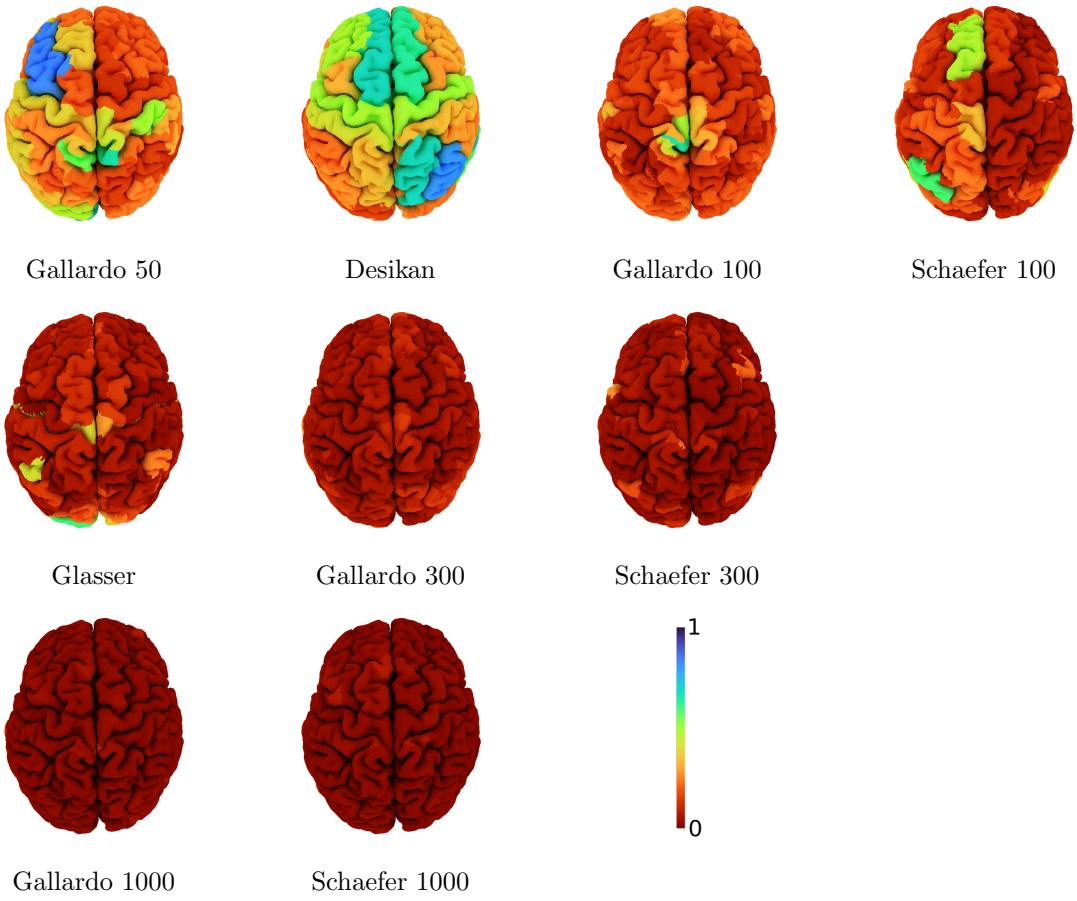

Figure 8: Self matching rate of the labeling per region for different atlases using WL-align. Atlases with 100 regions or less are illustrated in the first row. The second row illustrates atlases with approximately 300 regions and the third row those with 1000 regions.

5 Discussion and conclusions

Among the fundamental problems of network neuroscience at the scale of whole-brain structural connectivity, finding correspondences between brain regions and quantitatively assessing the similarity between brain networks are particularly important when it comes to considering massive heterogeneous datasets and modern data science techniques. In this work we considered these two problems in relation with the unresolved question concerning the choice of the parcellation for structural connectivity studies.

We proposed and analyzed a similarity index between brain networks, inspired by the Jaccard index between sets, that behaves in a way similar to the classical correlation index. Additionally, it enjoys the remarkable property of defining a metric in the space of connectomes, which is interesting both from the theoretical point of view and for data science applications. The proposed graph Jaccard index showed to be less affected by the number of regions in the chosen parcellation than the Frobenius distance, which is one example of (dis)similarity index from the class of norm-based distances.

The second object introduced in this paper is WL-align, a novel algorithm that allows to find the

graph alignment between two brain networks. It relies solely on topological features of the brain network, which makes it particularly suitable for being applied also outside the domain of network neuroscience. When WL-align is used in our experiments in order to retrieve the alignment between a network and a permuted version of itself, it gives the exact solution. This does not happen when the main competitor FAQ is employed. The superior performance of WL-align is evident also when brain networks of different subjects are aligned. In this case, the WL-align algorithm was shown to retrieve brain alignments that are closer to the ground truth with respect to the alignments obtained by FAQ. Notice that as it is designed, the WL-align algorithm builds on the construction of a feature vector for each node of the graph, which is then used as an edge weight in an assignment problem on a bipartite graph. This does not include any prior knowledge other than the topological similarity between the two networks to be aligned. The analysis provided in this work was intentionally constrained to the pure topological comparison of networks. Nevertheless, it would be possible to extend the feature vector defined in WL-align with any prior of geometrical, spatial, anatomical or connectomic nature or to add any constraints in the assignment problem on the bipartite graph. Future works will be devoted to the design of these constraints and features.

Some remarkable conclusions concerning the parcellations to be used in structural brain connectivity studies can be drawn from the ability of WL-align to find the correct alignment between two brain networks. First, the function-based parcellation of Schaefer is a poorer choice than the structure-based parcellation of Gallardo, the multimodal parcellation of Glasser and the morphological parcellation of Desikan. This was expected as the whole study is centered on measuring *structural* connectivity, hence the choice of a function-based parcellation was never expected to be optimal from any point of view. Allowing to express this concept quantitatively is one of the merits of WL-align. A second remarkable aspect is the performance of the Desikan atlas, which gave much better results in terms of alignability than any other parcellation of any granularity. For this reason, whenever the number of parcels considered in the study is low (less than 200), one should consider using the Desikan atlas as a first choice. Not only it would be a highly reliable choice that has been consistently used throughout time in the community, but with this study we showed that it would also allow to define brain networks with more consistent topological features, in particular those captured by WL-align. As far as brain atlases with a higher number of parcels (more than 200) is concerned, we showed that parcellations with a number of parcels in the > 400 range have lower performance in terms of alignment. For this reason, we suggest to use atlases with a lower number of parcels, if the designed experiment allows it. This difference in performance may be due to the number of streamlines used for defining the structural brain network, which may be lower than what would be necessary for a reliable definition of the brain networks. The lack of a consensus on the number of streamlines to be used in structural connectivity studies leaves this possibility open.

As highlighted throughout the paper, this work analyzes the problems of parcellation selection and brain alignment in the context of *structural* connectivity. Any conclusion we made should not be straightforwardly generalized to functional connectivity studies, which would require a separate analysis which was out of the scope of this work.

6 Open science

The data and code used in this work are all available at open repositories, as indicated in the text. We uploaded the used code and the obtained connectomes and alignments on the Open Science Framework. They can be found at this link: https://osf.io/depux/.

7 Acknowledgements

The authors would like to thank Dr. Guillermo Gallardo for the help in computing the Gallardo parcellation. Also, we are grateful to the OPAL infrastructure from Université Côte d'Azur and Inria Sophia Antipolis - Méditerranée "NEF" computation platform for providing resources and support. This work was funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (ERC Advanced Grant agreement No 694665: CoBCoM - Computational Brain Connectivity Mapping). Data were provided in part by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

References

N. Ayache and B. Faverjon. Efficient registration of stereo images by matching graph descriptions of edge segments. *International Journal of Computer Vision*, 1(2):107–131, 1987.

B. Barak, C.-N. Chou, Z. Lei, T. Schramm, and Y. Sheng. (Nearly) Efficient Algorithms for the Graph Matching Problem on Correlated Random Graphs. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d. Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 9190–9198. Curran Associates, Inc., 2019. URL http://papers.nips.cc/paper/9118-nearly-efficient-algorithms-for-the-graph-matching-problem-on-correlated-random-graphs.pdf.

M. Bayati, D. F. Gleich, A. Saberi, and Y. Wang. Message-Passing Algorithms for Sparse Network Alignment. *ACM Trans. Knowl. Discov. Data*, 7(1):3:1–3:31, Mar. 2013. ISSN 1556-4681. doi: 10.1145/2435209.2435212.

C. O. Becker, S. Pequito, G. J. Pappas, M. B. Miller, S. T. Grafton, D. S. Bassett, and V. M. Preciado. Spectral mapping of brain functional connectivity from diffusion imaging. *Scientific Reports*, 8(1411), 2018.

K. Brodmann. *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Barth, Leipzig, 1909.

M. S. Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing*, STOC '02, page 380–388, New York, NY, USA, 2002. Association for Computing Machinery. ISBN 1581134959. doi: 10.1145/509907.509965.

M. K. Chung, H. Lee, V. Solo, R. J. Davidson, and S. D. Pollak. Topological distances between brain networks. In *International Workshop on Connectomics in Neuroimaging*, pages 161–170. Springer, 2017.

D. Conte, P. Foggia, C. Sansone, and M. Vento. Thirty years of graph matching in pattern recognition. *International Journal of Pattern Recognition and Artificial Intelligence*, 18(03):265–298, May 2004. ISSN 0218-0014. doi: 10.1142/S0218001404003228.

R. S. Desikan, F. Ségonne, B. Fischl, B. T. Quinn, B. C. Dickerson, D. Blacker, R. L. Buckner, A. M. Dale, R. P. Maguire, B. T. Hyman, M. S. Albert, and R. J. Killiany. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage*, 31(3):968–980, 2006. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2006.01.021.

S. Deslauriers-Gauthier, M. Zucchelli, M. Frigo, and R. Deriche. A unified framework for multimodal structure–function mapping based on eigenmodes. *Medical Image Analysis*, 66, 2020.

S. Feizi, G. Quon, M. Mendoza, M. Medard, M. Kellis, and A. Jadbabaie. Spectral Alignment of Graphs. *IEEE Transactions on Network Science and Engineering*, pages 1–1, 2019. ISSN 2327-4697. doi: 10.1109/TNSE.2019.2913233.

B. Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012.

G. Gallardo, N. Gayraud, R. Deriche, M. Clerc, S. Deslauriers-Gauthier, and D. Wassermann. Solving the Cross-Subject Parcel Matching Problem using Optimal Transport. In *International Conference*

on Medical Image Computing and Computer-Assisted Intervention 2018, Granada, Spain, Sept. 2018a. URL https://hal.archives-ouvertes.fr/hal-01935684.

G. Gallardo, W. Wells, R. Deriche, and D. Wassermann. Groupwise structural parcellation of the whole cortex: A logistic random effects model based approach. *NeuroImage*, 170:307–320, 2018b. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2017.01.070.

M. F. Glasser, S. N. Sotiropoulos, J. A. Wilson, T. S. Coalson, B. Fischl, J. L. Andersson, J. Xu, S. Jbabdi, M. Webster, J. R. Polimeni, et al. The minimal preprocessing pipelines for the human connectome project. *Neuroimage*, 80:105–124, 2013.

M. F. Glasser, T. S. Coalson, E. C. Robinson, C. D. Hacker, J. Harwell, E. Yacoub, K. Ugurbil, J. Andersson, C. F. Beckmann, M. Jenkinson, S. M. Smith, and D. C. V. Essen. A multi-modal parcellation of human cerebral cortex. *Nature*, 536, 2016. doi: 10.1038/nature18933.

P. Hagmann. From diffusion mri to brain connectomics. Technical report, EPFL, 2005.

M. Hayhoe, F. Barreras, H. Hassani, and V. M. Preciado. SPECTRE: Seedless Network Alignment via Spectral Centralities. *arXiv:1811.01056 [cs, math]*, May 2019. URL http://arxiv.org/abs/1811.01056. arXiv: 1811.01056.

M. Heimann, H. Shen, T. Safavi, and D. Koutra. REGAL: Representation Learning-based Graph Alignment. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management - CIKM '18*, pages 117–126, Torino, Italy, 2018. ACM Press. ISBN 978-1-4503-6014-2. doi: 10.1145/3269206.3271788.

C. G. J. Jacobi. The reduction to normal form of a non-normal system of differential equations. de aequationum differentialium systemate non normali ad formam normalem revocando. *III. C. G. J. Jacobi manuscriptis posthumis in medium protulit A. Clebsch*, 1890.

B. Jeurissen, J.-D. Tournier, T. Dhollander, A. Connelly, and J. Sijbers. Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion mri data. *NeuroImage*, 103:411–426, 2014.

N. Korula and S. Lattanzi. An efficient reconciliation algorithm for social networks. *Proceedings of the VLDB Endowment*, 7(5):377–388, 2014. doi: 10.14778/2732269.2732274.

H. W. Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2 (1-2):83–97, 1955. doi: 10.1002/nav.3800020109.

C. Li, S. Wang, P. S. Yu, L. Zheng, X. Zhang, Z. Li, and Y. Liang. Distribution Distance Minimization for Unsupervised User Identity Linkage. In *Proceedings of the 27th ACM International Conference*

559 *on Information and Knowledge Management*, CIKM '18, pages 447–456, Torino, Italy, Oct. 2018.
560 Association for Computing Machinery. ISBN 978-1-4503-6014-2. doi: 10.1145/3269206.3271675.

561 L. Liu, W. K. Cheung, X. Li, and L. Liao. Aligning Users across Social Networks Using Network Embedding.
562 In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, IJCAI'16,
563 pages 1774–1780, New York, New York, USA, 2016. AAAI Press. ISBN 978-1-57735-770-4. URL
564 https://www.ijcai.org/Proceedings/16/Papers/254.pdf.

565 R. Mars, L. Verhagen, T. Gladwin, F. Neubert, J. Sallet, and M. Rushworth. Comparing brains by
566 matching connectivity profiles. *Neurosci. Biobehav.*, 60:90–97, 2016.

567 H. Nassar, N. Veldt, S. Mohammadi, A. Grama, and D. F. Gleich. Low Rank Spectral Network Alignment.
568 In *Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW '18*, pages
569 619–628, Lyon, France, 2018. ACM Press. ISBN 978-1-4503-5639-8. doi: 10.1145/3178876.3186128.

570 W. U. S. of Medicine. The brain analysis library of spatial maps and atlases (balsa), sep 2020. URL
571 https://balsa.wustl.edu/WN56.

572 Y. Osmanlıoğlu, B. Tunç, D. Parker, M. A. Elliott, G. L. Baum, R. Ciric, T. D. Satterthwaite, R. E.
573 Gur, R. C. Gur, and R. Verma. System-level matching of structural and functional connectomes in the
574 human brain. *NeuroImage*, 199:93–104, 2019.

575 S. Parisot, S. Arslan, J. Passerat-Palmbach, W. Wells, and D. Rueckert. Tractography-Driven Groupwise
576 Multi-scale Parcellation of the Cortex. *Inf. Process. Med. Imaging*, 24:600–612, 2015.

577 A. Schaefer, R. Kong, E. M. Gordon, T. O. Laumann, X.-N. Zuo, A. J. Holmes, S. B. Eickhoff, and B. T. T.
578 Yeo. Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity
579 MRI. *Cerebral Cortex*, 28(9):3095–3114, 07 2017. ISSN 1047-3211. doi: 10.1093/cercor/bhx179.

580 R. Singh, J. Xu, and B. Berger. Global alignment of multiple protein interaction networks with application
581 to functional orthology detection. *Proceedings of the National Academy of Sciences*, 105(35):12763–12768,
582 Sept. 2008. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.0806627105.

583 R. E. Smith, J.-D. Tournier, F. Calamante, and A. Connelly. Anatomically-constrained tractography:
584 improved diffusion mri streamlines tractography through effective use of anatomical information.
585 *Neuroimage*, 62(3):1924–1938, 2012.

586 O. Sporns, G. Tonomi, and R. Kötter. The human connectome: A structural description of the human
587 brain. *PLoS Computational Biology*, 1:245–251, 2005.

J.-D. Tournier, R. Smith, D. Raffelt, R. Tabbara, T. Dhollander, M. Pietsch, D. Christiaens, B. Jeurissen, C.-H. Yeh, and A. Connelly. Mrtrix3: A fast, flexible and open software framework for medical image processing and visualisation. *NeuroImage*, page 116137, 2019.

D. C. Van Essen, K. Ugurbil, E. Auerbach, D. Barch, T. Behrens, R. Bucholz, A. Chang, L. Chen, M. Corbetta, S. W. Curtiss, et al. The human connectome project: a data acquisition perspective. *Neuroimage*, 62(4):2222–2231, 2012.

M. Venkatesh, J. Jaja, and L. Pessoa. Comparing functional connectivity matrices: A geometry-aware approach applied to participant identification. *NeuroImage*, 207:116398, 2020.

J. L. Villareal-Haro, A. Ramirez-Manzanares, and J. A. Pichardo-Corpus. A community-based topological distance for brain-connectome classification. *Journal of Complex Networks*, 8(4):cnaa034, 2020.

J. T. Vogelstein, J. M. Conroy, V. Lyzinski, L. J. Podrazik, S. G. Kratzer, E. T. Harley, D. E. Fishkind, R. J. Vogelstein, and C. E. Priebe. Fast Approximate Quadratic Programming for Graph Matching. *PLoS ONE*, 10(4), Apr. 2015. ISSN 1932-6203. doi: 10.1371/journal.pone.0121002.

B. Y. Weisfeiler and A. A. Leman. The reduction of a graph to canonical form and the algebra which appears therein. English translation of the original paper published in Russian., 1968. URL https://www.iti.zcu.cz/wl2018/pdf/wl_paper_translation.pdf.

R. C. Wilson and P. Zhu. A study of graph spectra for comparing graphs and trees. *Pattern Recognition*, 41(9):2833–2841, Sept. 2008. ISSN 0031-3203. doi: 10.1016/j.patcog.2008.03.011.

WU-Minn Human Connectome Project consortium. 1200 subjects data release reference manual, 2017. URL https://www.humanconnectome.org/storage/app/media/documentation/s1200/HCP_S1200_Release_Reference_Manual.pdf.

B. T. Yeo, F. M. Krienen, J. Sepulcre, M. R. Sabuncu, D. Lashkari, M. Hollinshead, J. L. Roffman, J. W. Smoller, L. Zöllei, J. R. Polimeni, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of neurophysiology*, 2011.

T. Yeo. Computational brain imaging group (cbig) repository, sep 2020. URL https://github.com/ThomasYeoLab/CBIG/tree/master/stable_projects/brain_parcellation/Schaefer2018_LocalGlobal/Parcellations/HCP.