



Diffusion MRI Tractography Filtering Techniques Change the Topology of Structural Connectomes

Matteo Frigo, Samuel Deslauriers-Gauthier, Drew Parker, Abdol Aziz Ould
Ismail, Junghoon John Kim, Ragini Verma, Rachid Deriche

► To cite this version:

Matteo Frigo, Samuel Deslauriers-Gauthier, Drew Parker, Abdol Aziz Ould Ismail, Junghoon John Kim, et al.. Diffusion MRI Tractography Filtering Techniques Change the Topology of Structural Connectomes. Journal of Neural Engineering, In press, 10.1088/1741-2552/abc29b . hal-02972198

HAL Id: hal-02972198

<https://hal.science/hal-02972198>

Submitted on 20 Oct 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Diffusion MRI Tractography Filtering Techniques Change the Topology of Structural Connectomes

Matteo Frigo

Athena Project-Team, Inria Sophia Antipolis-Méditerranée, Université Côte D'Azur,
Nice, France

E-mail: `matteo.frigo@inria.fr`

Samuel Deslauriers-Gauthier

Athena Project-Team, Inria Sophia Antipolis-Méditerranée, Université Côte D'Azur,
Nice, France

Drew Parker

Penn Applied Connectomics and Imaging Group, Department of Radiology,
University of Pennsylvania, Philadelphia, PA, USA

Abdol Aziz Ould Ismail

Penn Applied Connectomics and Imaging Group, Department of Radiology,
University of Pennsylvania, Philadelphia, PA, USA

Junghoon John Kim

Department of Molecular, Cellular, and Biomedical Sciences, CUNY School of
Medicine, The City College of New York, NY, USA

Ragini Verma

Penn Applied Connectomics and Imaging Group, Department of Radiology,
University of Pennsylvania, Philadelphia, PA, USA

Rachid Deriche

Athena Project-Team, Inria Sophia Antipolis-Méditerranée, Université Côte D'Azur,
Nice, France

Submitted to: *J. Neural Eng.*

Abstract. *Objective.* The use of non-invasive techniques for the estimation of structural brain networks (i.e. connectomes) opened the door to large-scale investigations on the functioning and the architecture of the brain, unveiling the link between neurological disorders and topological changes of the brain network. This study aims at assessing if and how the topology of structural connectomes estimated non-invasively with diffusion MRI is affected by the employment of tractography filtering techniques in structural connectomic pipelines. Additionally, this work investigates the robustness of topological descriptors of filtered connectomes to the common practice of density-based thresholding. *Approach.* We investigate the changes in global efficiency, characteristic path length, modularity and clustering coefficient on filtered connectomes obtained with the spherical deconvolution informed filtering of tractograms and using the convex optimization modelling for microstructure informed tractography. The analysis is performed on both healthy subjects and patients affected by traumatic brain injury and with an assessment of the robustness of the computed graph-theoretical measures with respect to density-based thresholding of the connectome. *Main Result.* Our results demonstrate that tractography filtering techniques change the topology of brain networks, and thus alter network metrics both in the pathological and the healthy cases. Moreover, the measures are shown to be robust to density-based thresholding. *Significance.* The present work highlights how the inclusion of tractography filtering techniques in connectomic pipelines requires extra caution as they systematically change the network topology both in healthy subjects and patients affected by traumatic brain injury. Finally, the practice of low-to-moderate density-based thresholding of the connectomes is confirmed to have negligible effects on the topological analysis.

1. Introduction

Mapping the brain architecture and functioning is a core health ambition of the 21st century and one of the greatest challenges of modern science [1]. Correspondingly, many important neurological diseases and disorders have been shown to be related to pathological alterations in the connectivity of the brain, calling for specific efforts in research to better understand the network of neural connections composing the human brain. The graph-like representation of brain networks is called connectome [2, 3]. The nodes of the aforementioned graph are defined by the brain regions that are being studied, while the edges represent the connectivity strength between the nodes. These edges can be defined in such a way that they represent either the function or the structure of the brain. In the first case we are defining the *functional connectome* and in the second case we have the *structural connectome*. The latter is the object of interest of this work. The graph-theoretical analysis of connectomes unveiled some previously unknown features of both anatomy and pathology of the human brain [4, 5]. The only non-invasive technique that allows the reconstruction of such networks is diffusion MRI (dMRI)-based tractography, which relies on the estimation of the white matter pathways as a set of streamlines collected in a tractogram. The classical way to build a structural connectome is based on counting the number of streamlines connecting two different regions of the brain and assigning the obtained number to the corresponding edge of

the graph. The methodological aspects related to this streamline-edge matching have recently been analysed [6].

In the last twenty years more and more accurate tractography algorithms have been developed, opening the doors to the exploration of the architecture of the white matter [7]. Each of them aims at solving specific problems related to the estimation of white matter fiber pathways, but the absence of ground truth knowledge on the white matter architecture has prevented the community from building a consensual agreement on a definitive way to estimate structural connectivity from dMRI.

Among the many problematic issues of tractography-based structural connectivity estimation, the non-quantitative nature of tractography [8, 9], the methodological limitations of dMRI-based tractography [10] and the presence of many false positive connections within connectomes [11] are the ones that inspired some of the most recent methodological advances in the field of non-invasive tractography. A simplistic practice that aims at tackling some of the problems of tractography-based connectomics is *density-based thresholding*, which is the act of excluding all the weakest connections in the network until a graph of the wanted density is obtained [12]. Another class of algorithms going under the name of *tractography filtering techniques* (TFTs) have been developed. These algorithms are designed to post-process the tractograms obtained via dMRI-based tractography. Recently, three TFTs have been presented to the community, namely the spherical-deconvolution informed filtering of tractograms (SIFT2) [13], the convex optimization modelling for microstructure informed tractography (COMMIT) [14] and the linear fascicle evaluation (LiFE) [15]. The latter is designed as a particular case of the more general COMMIT framework and will not be studied in this work. The goal of TFTs is to assign to each streamline a quantitative marker of the connectivity strength. Considering all these connections it is possible to obtain brain networks that are also referred to as *filtered connectomes*. The main difference between SIFT2 and COMMIT lies on the specific issue of tractography that is tackled by employing them in the connectomic pipeline. The SIFT2 weights are global representations of the cross-sectional area of the fiber population represented by each streamline. Conversely, the coefficients retrieved by COMMIT and LiFE correspond to the signal fraction associated to the signal profile generated by each streamline. Despite providing a fundamental contribution to the biological interpretability of the obtained connectomes, the SIFT2 model is not designed to detect false positive connections. This task is explicitly tackled by the COMMIT and LiFE models, which are defined in such a way that the resulting streamline coefficients are zero-valued for a relevant part of the input tractogram. Section 2.2 is devoted to the technical description of the mentioned differences.

In order to understand how connectomic studies are affected by the use of TFTs, it is of fundamental importance to explore the changes in network topology generated by the use of filtered connectomes. With this objective in mind, in this study we are going to investigate how four graph-theoretical metrics (GTM) of network integration and segregation change with respect to the employed TFT and to density based thresholding.

The reason why we specifically inspect measures of segregation and integration in the context of connectomes comes from the small-worldness of the brain network topology, which has been assessed by several studies in the last 15 years [16, 17, 18, 4] and its behaviour is characterized by both high segregation and high integration. Section 2.4 will be devoted to a detailed presentation of the graph-theoretical aspects of small-worldness and of the four studied network metrics.

Changes in brain network topology can also be due to the presence of pathology [19, 5] like traumatic brain injury (TBI), which is a network disorder that exhibits (among others) changes in the small-worldness of the brain network [20]. Moreover, subjects affected by TBI show an increase of the mean diffusivity coupled with a decreased fractional anisotropy [21], which are two indicators of dMRI signal changes that can affect the estimation of fiber tracts, and thus structural brain networks. For these reasons it is not possible a priori to generalize the results obtained on healthy subjects to the considered pathological case.

In addition, for each type of filtered connectome we investigate the robustness of the considered graph-theoretical measures with respect to density-based thresholding. A recent study by Civier et al. [12] partially investigated this robustness. The focus of the mentioned analysis was specifically on the density-based thresholding aspect of connectomic pipelines, therefore they factored out any other potentially influencing factor by studying only connectomes obtained on healthy subjects through one specific connectomic pipeline that involved SIFT2 in the tractography filtering step. As our work will focus on the TFT step, we will extend the analysis of Civier et al. [12] to the case where COMMIT is employed for both healthy subjects and TBI patients. The two mentioned analyses will be performed on connectomes obtained on 100 healthy subjects from the Human Connectome Project (HCP) database [22]. The high quality of the data provided by the HCP database is not a realistic example of the data that are acquired clinically. For this reason, we will extend the analysis to a dataset acquired using a clinical protocol, which mandates short acquisitions. This dataset includes both healthy subject and TBI patients, showing that additional care is needed in the employment of TFTs as they have the potential to change the clinical interpretation of the results.

The employed connectomic pipeline was designed to systematically include a tractography filtering step. A recent and related study investigating the effects of tractography filtering on the topology of brain networks is the work of Yeh et al. [23], where the effects of the first version of the SIFT technique on the graph-theoretical analysis of structural brain networks estimated with dMRI are evaluated. Inspired by the novelties in the field of tractography filtering proposed in the last years [13, 14] and the recently growing interest towards the field of patho-connectomics [19], this work represents the first systematic comparison of the effects of the state-of-the-art tractography filtering techniques on the graph-theoretical analysis of structural brain networks estimated in-vivo[‡]. After this introduction, Section 2 and Section 3 will

[‡] The present work has partially been presented at the 25th Meeting of the Organization for Human

present the theoretical and methodological aspects of our work, respectively. The results presented in Section 4 and discussed in Section 5 describe how the graph-theoretical analysis of structural connectomes estimated non-invasively is affected by the employment of tractography filtering techniques.

2. Background

2.1. Weighted Connectomes

The non-invasive estimation of the structural connectivity of the human brain is a complex task that relies on the ability to track the white-matter pathways between different regions of the brain. As already mentioned, dMRI-based tractography is the only non-invasive technique that yields a representation of the axonal pathways connecting different brain regions. These white matter pathways are obtained in the form of streamlines by following the local orientation of the fiber bundles estimated from dMRI data. The definition of *structural connectome associated to a tractogram* that will be used throughout this work reads as follows. Consider each pair of brain regions i and j , then define

$$c_{ij} = \sum_{s \in \mathcal{S}(i,j)} w_s \quad (1)$$

where $\mathcal{S}(i, j)$ is the set of streamlines terminating in regions i and j and w_s is the coefficient that quantifies the connectivity associated with streamline s . The structural connectome is then represented as the symmetric matrix $C = \{c_{ij}\}$. This matrix is a convenient representation of the connectome where each entry c_{ij} encodes the value associated to edge $i \rightarrow j$ of the connectome and i and j are the two regions connected by the edge. Finally, the connectivity matrix C is normalized by dividing each entry of C by the sum of all the entries of the same matrix. Whenever $c_{ij} = 0$, we will assume that the edge $i \rightarrow j$ is not present in the graph. The resulting graph is sparse.

Defining $w_s = 1$ for every $s \in \mathcal{S}$ we obtain the *streamline-count* (SC) connectome. The reliability of these connectomes as estimators of structural connectivity has been shown [8] to be limited by the complexity of the white matter configurations through which tractography algorithms have to find their way [11] and by the way in which each streamline is assigned to a certain edge in the connectome [6]. More sophisticated definitions of the weights w_s allow to correct for some of these biases, in particular through the employment of tractography filtering techniques. The issues related to the streamline-parcel assignment problem will not be solved by these techniques, as TFTs do not change the geometry of the estimated fiber tracks or the shape of the parcels.

2.2. Tractography Filtering Techniques

The limitations of the SC connectome that we mentioned in Section 2.1 influence the sensitivity and the sensibility of the structural connectomes estimated in-vivo [25, 11].

Brain Mapping [24].

Moreover, dMRI-based tractography is not quantitative per-se, hence it needs to be complemented in order to associate to each streamline the corresponding connectivity strength. To solve this task, some *top-down* approaches have been researched and presented in the form of TFTs, which act as post-processing techniques on an existing tractogram. When a tractogram is *filtered*, a coefficient is associated to each streamline. This coefficient quantifies the connectivity that is associated to the specific streamline and corresponds to the w_s defined in equation (1). Recently, three techniques have been presented to the research community. These are the spherical-deconvolution informed filtering of tractograms (SIFT2) [13], the convex optimization modelling for microstructure informed tractography (COMMIT) [14] and the linear fascicle evaluation (LiFE) [15]. In this study we are going to focus only on SIFT2 and COMMIT, as LiFE can be formulated as a particular case of COMMIT.

The SIFT2 model assigns a coefficient to each streamline in such a way that the fiber density computed from the fiber orientation distribution functions (fODF) matches the one obtained from the weighted tractogram. This formulation has the advantage of attaching a biological interpretation to the coefficients associated to the streamlines, as they can be interpreted as the mean cross sectional area along the fiber. The main limitation of this method is its inability to detect false positive connections, whose presence could affect the graph-theoretical analysis of the estimated brain networks, but it can assign very small weights which will give little impact to the corresponding streamlines. This is not the case of the COMMIT model, whose mathematical formulation yields solutions that include many zero-valued streamline weights. This sparsity-promoting formulation was proved to be effective in false positives detection [11]. The COMMIT framework is based on a forward model that transforms streamlines into the expected dMRI signal profile, allowing to define a convex optimization problem that finds the linear combination of streamline signal profiles that best matches the acquired dMRI signal. The coefficients retrieved by COMMIT correspond to the signal fraction associated to the signal profile generated by each streamline, hence their biological interpretability is more limited with respect to one of the SIFT2 weights. The interpretability issue comes from the type of quantity modelled. While SIFT2 is designed to describe the expected fiber density computed from the fODFs, COMMIT relies on the pure dMRI signal, avoiding the potential biases introduced by the transformation of the signal into the fODFs. Table 1 summarizes the differences between the two considered TFTs.

	SIFT2	COMMIT
Sparsity	No	Yes
Interpretability	High	Low
Reference	Fiber density	dMRI signal

Table 1. Summary of the distinctive features of the two employed TFTs, SIFT2 and COMMIT. The lack of sparsity in SIFT2 comes with an increased interpretability of the result.

2.3. Density-based thresholding

As we mentioned in Section 1, a simple technique that aims at eliminating some spurious connections from the connectome is the *density-based thresholding*. Thresholding a connectome means to exclude all the edges whose weight is below a certain threshold, which is equivalent to applying the transformation

$$\tilde{c}_{ij} = c_{ij} \cdot \delta_{>\tau}(c_{ij}) \quad (2)$$

where \tilde{C} is the thresholded connectome and $\delta_{>\tau}$ is the indicator function that takes value 1 if its argument is greater than the threshold τ and zero otherwise. The arguments in favour of employing thresholding are mostly based on the necessity to eliminate the pervasive spurious weak connections produced by probabilistic tractography in order to reduce their influence in the topological analysis [26, 27]. This is related to the interpretation of probabilistic tractography as a Monte Carlo method, where an event (connection) happening with low frequency (having low edge weight) is assumed to be less likely to belong to the ground truth. While being true in many fields of science, this assumption is wrong in the case of tractography, as the strength of the biological connectivity within the brain spans over several order of magnitudes [4], hence a small connectivity value does not necessarily correspond to a false positive connection.

Density-based thresholding is defined as the act of pruning the weakest connections until a pre-determined network density is reached, where this density is defined as the ratio between the number of edges in the network and the number of potential connections of the same network, namely

$$d = \frac{n_e}{n_v \cdot (n_v - 1) / 2} \quad (3)$$

where n_e is the number of edges, n_v is the number of nodes and d is the graph density. Notice that this is not strictly a tractography filtering technique, but we consider it as it is an attempt to remove false positive connections from the connectome.

2.4. Graph-theoretical analysis

The topological properties of connectomes have been of interest since the introduction of the concept of brain network. It has been established that brain networks exhibit the so-called small world [16, 17, 18, 4] behaviour. This concept describes a *hybrid combination of high clustering and short path length* [4] and was first introduced to explain the small world behaviour of social relations [28], where few acquaintance steps are sufficient to connect two people from very distant (both geographically and socially) regions of the world. The classical approach to small-worldness estimation considers the ratio between two global measures of the integration and the segregation of the network respectively [29].

The integration of a network can be measured through the *characteristic path length* (CPL). Let $d(i, j)$ be the length of the shortest path (in Dijkstra's sense [30]) between

two nodes on connectome C . The CPL of C is

$$\ell = \frac{1}{n} \sum_{i \neq j} d(i, j) \quad (4)$$

where n is the number of nodes considered in the graph and the sum is performed on every pair of distinct nodes i and j . The CPL is a reverted marker of the integration of a network, as a high characteristic path length corresponds to a low integration. The quantity defined in equation (4) is not well-defined for disconnected graphs, as the distance between two nodes belonging to distinct connected components of the network cannot be computed. To overtake this limitation, the *global efficiency* (GE) measure was introduced [31]. Its formulation reads as follows

$$e = \frac{1}{n} \sum_{i \neq j} \frac{1}{d(i, j)}. \quad (5)$$

To ensure the well-posedness of equation (5) we extend the definition of path length to the case where there exists no connected path between i and j . In this case we define $d(i, j) = \infty$. Notice that while this extension would break the definition of CPL, it suits the definition of GE, as the corresponding term in equation (5) would be equal to zero. Consequently, the presence of isolated subnetworks in the connectome decreases its efficiency but it does not disrupt the measure. We also highlight that long paths have a bigger influence on CPL than on GE.

To evaluate the segregation of a network it is possible to compute the *global clustering coefficient* (CC) [32]. Let $t = (i, j, k)$ be a triplet, namely a set of three nodes that form a connected partition of the connectome. If the triplet is connected by exactly two edges it is called *open triplet*, otherwise it is a *closed triplet*. For every node i in the graph we can define its corresponding local CC as

$$CC(i) = \frac{\sum_{i \neq j} \sum_{k \neq i \neq j} c_{ij} \cdot c_{jk} \cdot c_{ki}}{\left(\sum_{i \neq j} c_{ij}\right)^2 - \sum_{i \neq j} c_{ij}^2} \quad (6)$$

where c_{ij} is the weight associated to the edge $i \rightarrow j$ encoded in the connectivity matrix and the sums are to be intended for every pair/triplet of nodes satisfying the specified inequality. The global clustering coefficient is then obtained as the arithmetic average of the local CCs across all the nodes

$$CC = \frac{1}{n} \sum_{i=1}^n CC(i) \quad (7)$$

where n is the number of nodes in the graph. The CC describes the average degree of connectivity in the neighbourhood of the nodes by measuring how each node tends to create a cluster around itself [33]. A different point of view on segregation is offered by the concept of *modularity* (MO) [34]. Given a partition of the network (i.e. a subdivision of the nodes in communities), modularity gives a measure of the proportion between the number of direct connections pointing towards nodes belonging to the same

community and the number of direct connections leaving the community. Modularity is mathematically defined as

$$Q = \frac{1}{2m} \sum_{i,j} \left(C_{i,j} - \frac{k_i \cdot k_j}{2m} \right) \delta(\gamma_i, \gamma_j) \quad (8)$$

where m is the sum of all the weights of the edges in the network, $C_{i,j}$ is the connectivity matrix associated to the studied network, k_i is the sum of all the weights of the edges connecting to node i , γ_i is the community to which i belongs and $\delta(\gamma_i, \gamma_j)$ is the Kronecker delta taking value 1 if $\gamma_i = \gamma_j$ and zero otherwise. The shift of paradigm brought by the concept of modularity as opposed to the one of clustering coefficient is shown in Figure 1 with an example on two simple graphs. The first graph is the junction




Figure 1. Both the represented graphs show an evident modular structure, where the communities are two cycles in the first graph and two cliques in the second graph. The modularity of the two graphs is relatively similar (0.37 for the union of cycles and 0.46 for the union of cliques). Conversely, the clustering coefficient of the two graphs is radically different (0.00 for the union of cycles and 0.92 for the union of cliques). The integration measures computed on the two graphs show the expected behaviour: the higher characteristic path length of the community of cycles is reflected in the lower efficiency of the same network with respect to the one computed on the clique community, where the symmetrical phenomenon appears. Each edge is associated to a weight equal to 1.

of two cycles through a single edge. This graph exhibits a non-trivial modularity, which is coherent with the fact that it was defined as the union of two communities. On the contrary, its clustering coefficient is null as the set of triplets does not contain any closed triplet. The second graph is still defined as the junction of two communities, but each community is defined as a clique (i.e. a fully connected graph). The modularity of this graph is higher than the one of the union of cycles. This reflects the presence of more connections among the nodes of each community. The added edges impacted the clustering coefficient even more, taking a value close to the one that would be obtained

for a fully connected graph. The reason why this happens is that the graph actually is almost fully connected. In particular, only the two central nodes have triplets that are not closed. This example showed how the two measures of segregation (MO and CC) are sensitive to different effects of the presence of community and they are not equivalent.

The ensemble of the measures of integration (GE and CPL) and segregation (CC and MO) gives a global picture of the topology of the studied connectomes with a particular attention towards the small-worldness of brain networks.

3. Methods

3.1. Data and preprocessing

HCP subjects. From the HCP database we considered the list of 100 unrelated subjects (U100 group) dataset available at the Connectome Coordination Facility [35, 36, 37]. These data were acquired on a Siemens Magnetom Skyra 3T MRI system and preprocessed with the minimal preprocessing pipeline for the Human Connectome Project, which includes EPI distortion correction via FSL’s *topup* [38, 39] and eddy current and subject motion correction via FSL’s *eddy* [40, 37]. For a detailed discussion on the preprocessing pipeline employed for this dataset, the interested reader can refer to the original paper of Glasser et al. [22]. Aiming at minimizing the influencing factors in the study, we used the preprocessed data that are available at the Connectome Coordination Facility. For each subject we have 288 images subdivided in 18 volumes at $b = 0s/mm^2$ and 90 diffusion-weighted volumes obtained at uniformly distributed directions at $b = 1000s/mm^2$, $b = 2000s/mm^2$ and $b = 3000s/mm^2$ for a total of 3 shells.

Clinical. The clinical dataset consisted of 39 adults with moderate to severe TBI acquired 3, 6, and 12 months after their injuries. The 3-month subset was selected for this study and 35 age-matched healthy controls with similar gender and duration (years) of education were added. This study was approved by the University of Pennsylvania institutional review board. All participants provided written informed consent either directly or by proxy via a legally authorized representative. The TBI patients were recruited from outpatient clinical programs at the Drucker Brain Injury Center at the MossRehab Hospital. They were screened to include only patients with predominantly diffuse TBI [41]. Healthy controls were recruited through local advertising and word of mouth, and underwent a clinical interview to ensure that they had no known history of TBI that resulted in alteration or loss of consciousness. The MRI data were acquired on a Siemens 3T Tim Trio system. Diffusion weighted images were acquired in two runs of 30 directions at a $b = 1000s/mm^2$ with 7 $b = 0s/mm^2$ images dispersed throughout each acquisition. The data were acquired with $TR = 6500ms$ and $TE = 84ms$ and a 90 degree flip angle at a resolution of $2.18 \times 2.18 \times 2.2mm^3$. A structural MPRAGE image was finally acquired with $TR = 1620ms$ and $TE = 3ms$, a 15 degree flip angle,

and an image resolution of $1 \times 1 \times 1\text{mm}^3$. All images were manually inspected for artifacts. If artifacts were present in $< 25\%$ of the volumes of the DWI acquisition, those volumes were removed from the series before processing. All volumes that were flagged for removal contained motion-induced signal drop-out artifacts. If more than 25% of volumes contained artifacts, the scan was rejected and removed from the sample. The final sample size was 35 TBI patients and 34 healthy controls, giving in total 69 subjects. The TBI cohort includes 23 male and 12 female subjects with age in the 19-53 years range (mean \pm standard deviation = 32.71 ± 13.45 years) and the healthy controls cohort includes 25 male and 9 female subjects with age in the 18-65 years range (mean \pm standard deviation = 34.35 ± 9.8 years). The hypothesis that the TBI patients and the healthy controls populations are age-matched is supported by a two-sample two-sided t -test with an alpha equal to 0.05 comparing the average of the age distribution of the two cohorts ($t = 0.58$, $p = 0.57$).

Diffusion MRI data were denoised using a local PCA method [42], followed by brain extraction with FSL's *BET* tool [43] on the first b0 image. The denoised data and brain mask were input to FSL's *eddy* to correct the data for motion and eddy-current distortion [44]. Because reverse phase-encoded data was not acquired, EPI distortion correction was not possible. Finally, the brain was extracted a second time with *BET* on the motion-corrected average b0 image.

3.2. Connectomic pipeline

A five-tissue-type image [45] of each subject was obtained with the Freesurfer pipeline [46] implemented in Mrtrix3 [47]. The estimation of the fiber orientation distribution functions (fODFs) was tailored on each dataset due to the lack of multi-shell data in the clinical cohort.

- HCP subjects: response functions for each tissue were computed using the multi-shell multi-tissue (MSMT) response function estimation algorithm provided by Jeurissen et al. in their work on MSMT constrained spherical deconvolution (CSD) [48] and the fODFs were computed using the MSMT-CSD algorithm [48, 49] with a maximal spherical harmonics order of $\ell = 8$.
- Clinical: the unsupervised algorithm of Dhollander et al. [50] was used for estimating the white matter response function, while fODFs were computed using the CSD algorithm [51] with a maximal spherical harmonics order of $\ell = 6$.

The obtained fODFs were used for probabilistic anatomically-constrained tractography (ACT) with the iFOD2 algorithm [45]. The seeding was performed from the gray matter - white matter interface (GMWMI) and a total of 2 millions of streamlines was obtained.

On both datasets, the cortical parcellation used for evaluating the structural connectivity between regions was extracted with the automated labeling system of Desikan et al. [52] via Freesurfer [46].

Tractography filtering was then performed on both datasets via SIFT2 and COMMIT. The forward model of COMMIT was defined as follows. The diffusivity

within the intra-cellular (IC) compartment was modelled as a stick with parallel diffusivity equal to $1.7 \cdot 10^{-3} mm^2/s$, the diffusivity in the extra-cellular compartment was modelled with a zeppelin under the tortuosity assumption with a fixed intra-cellular volume fraction equal to $f_{IC} = 0.7$ and the isotropic compartment was described as a linear combination of two isotropic balls of radial diffusivity $1.7 \cdot 10^{-3} mm^2/s$ and $3 \cdot 10^{-3} mm^2/s$ respectively. The terms stick, zeppelin and ball are borrowed from the taxonomy of Panagiotaki et al. [53]. To complete the analysis, we also computed the streamline count (SC) connectome associated to each tractogram.

The computation of global efficiency (GE), characteristic path length (CPL), clustering coefficient (CC) and modularity (MO) as they are defined in Section 2.4 was performed with the Python implementation of the brain connectivity toolbox [29]. In particular, modularity was computed by averaging the results of 100 runs of the Louvain algorithm [54]. These metrics were computed on every connectome of every subject. In order to investigate the robustness of these metrics to density-based thresholding, each GTM was computed also on the same connectomes thresholded at unitary intervals from 1% to the base density $d\%$, where d is the density of the non-thresholded connectome.

Figure 2 gives a graphical overview of the connectomic pipeline employed in this work.

3.3. Statistical analysis

In order to understand the differences between the connectomes obtained on distinct datasets or with a particular TFT, statistical analyses were performed with an alpha of 0.05 in all experiments. First, we evaluated the density of the connectomes obtained with each TFT. A separate analysis was performed for each subject cohort. The normality of the distribution of those values was assessed by inspecting the *normal probability plot* of the raw data for each considered TFT. A *two-tailed dependent-samples t-test* was employed to test if the use of TFTs changes the average density of the connectomes within the subject cohort. Furthermore, we analysed the values of the GTMs computed on each connectome. For each considered GTM we analysed each subject cohort independently. The normality of the raw values of the GTM computed on connectomes built with a specific TFT was tested using the Anderson-Darling test and the statistical significance of the differences between the results obtained with distinct TFTs was tested using the Mann-Whitney U test [55] and measured with the rank-biserial correlation (an effect size measure defined in the $[-1, 1]$ range). To account for multiple comparisons, a false discovery rate (FDR) correction was performed with the Benjamini-Hochberg procedure [56]. Finally, we performed an edge-wise comparison between the filtered connectomes obtained from the two clinical subject cohorts (TBI patients and healthy controls). Each TFT was studied independently. For each edge of the connectome, the normality of the distribution of the edge weights for one subject cohort and one TFT was tested using the Anderson-Darling test and the statistical significance of the differences between the edge weight distribution in TBI patients and

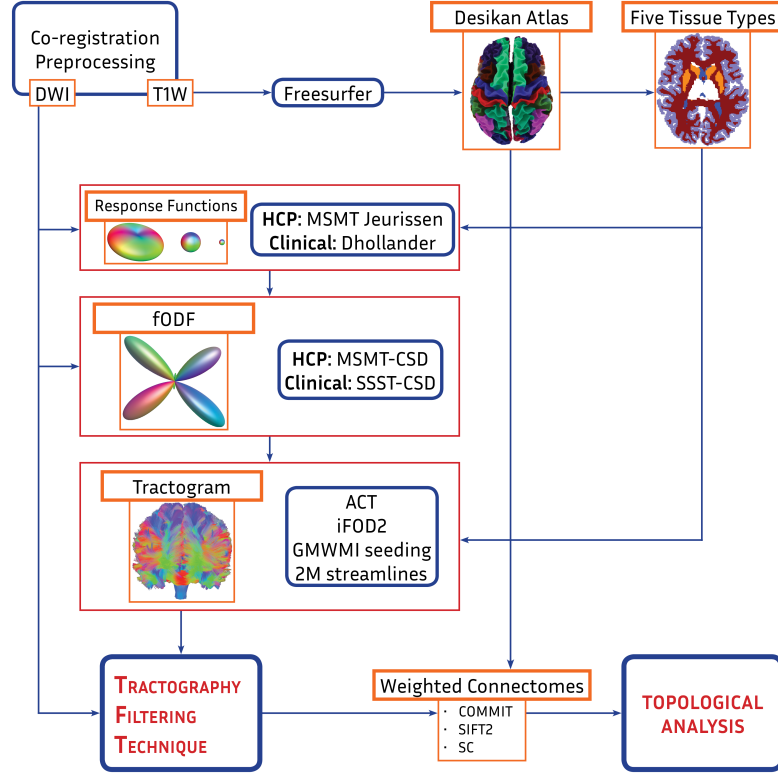


Figure 2. The represented connectomic pipeline is the one employed throughout this work. The first step is the pre-processing and the co-registration of the T1-weighted (T1w) and the diffusion-weighted magnetic resonance images (DWI). The T1w image was used to obtain the Desikan atlas and the five-tissues-type image. The DWI allowed to estimate the response functions and the fODFs (with the multi-shell multi-tissue algorithm for the HCP dataset and the single-shell single-tissue algorithm for the clinical dataset) that were necessary to perform anatomically constrained tractography (ACT) with the second order integration of the FOD (iFOD2) algorithm seeding from the gray matter - white matter interface (GMWMI). Finally, the tractography filtering step allowed to define the filtered connectomes whose topology was the object of interest of this study. Notice that SC connectomes were obtained by skipping the tractography filtering step.

in the healthy controls was tested with the Mann-Whitney U test. The magnitude of this difference was then evaluated using the rank-biserial correlation effect size measure. FDR correction was performed with the Benjamini–Hochberg procedure.

4. Results

We processed the data presented in Section 3.1 and for each subject we built the structural connectivity matrices both by employing one TFT among COMMIT and SIFT2 and by not employing any TFT, hence considering the streamline-count (SC) connectome defined in Section 2.1. Figure 3 shows the effect of the inclusion of TFTs

in connectomic pipelines on the density of the resulting connectomes. Each box-and-whisker plot represents the density of the filtered and non-filtered connectomes obtained from a single subject cohort. We report that SIFT2 has no effect on the density of connectomes. On the contrary, COMMIT has the effect of shifting the mean towards lower values, hence lowering the density of the connectomes with respect to the one of the SC/SIFT2 connectomes. The normal probability plots of the raw values of the density of the connectomes reported in Figure A1 of the supplementary materials support the assumption that these values are normally distributed. The statistical significance of the differences observed between the COMMIT and the SC/SIFT2 connectomes (the latter are equivalent) is supported by the t -test performed on the HCP subjects ($t = 40.69$, $p \leq 0.05$), the healthy controls of the clinical dataset ($t = 39.23$, $p \leq 0.05$) and the TBI patients ($t = 49.11$, $p \leq 0.05$). The mean density of the COMMIT connectomes of the HCP subjects is 2% lower than that of the SC/SIFT2 connectomes, while the shift in the case of both the healthy controls and TBI patients from the clinical dataset is equal to 19%.

In order to understand how the interpretation of patho-connectomic studies can be changed by the use of TFTs, we performed an edge-wise statistical comparison of the connectomes obtained with each TFT on the pathological cohort of subjects affected by TBI against the healthy controls from the clinical dataset. Each TFT was analysed independently. The Anderson-Darling test performed for each edge of the connectome and reported in Figure A2 of the supplementary materials supports the assumption that for a fixed edge, the edge weights from one subject cohort are not normally distributed in general. The results presented in Figure 4 show how the set of edges exhibiting statistically significant differences between TBI patients and healthy controls are different for every TFT.

We analysed the three connectomes per subject from the graph-theoretical point of view by computing the Global Efficiency (GE), the Characteristic Path Length (CPL), the Louvain Modularity (MO) and the Clustering Coefficient (CC) of each connectome. The obtained values are reported in Figures A3 and A4 of the supplementary materials. The effect of TFTs on these metrics is then presented in Figure 5, which shows the results of Mann-Whitney U tests between the values of GTMs calculated under each TFT, displayed as the rank-biserial correlation effect size, within each cohort (HCP, TBI patients or clinical healthy controls). Only significant differences with $p \leq 0.05$ are shown. These plots highlight two distinct phenomena:

- SIFT2 and COMMIT connectomes are not equivalent from the graph-theoretical point of view, as there are significant non-negligible differences between the GTMs computed on connectomes computed with the two TFTs;
- both the SIFT2 and the COMMIT connectomes exhibit topological differences with respect to the unfiltered SC connectomes.

The two phenomena can be observed for every GTM (GE, CPL, CC, MO). We note that the results obtained from the HCP database show the presence of significant differences

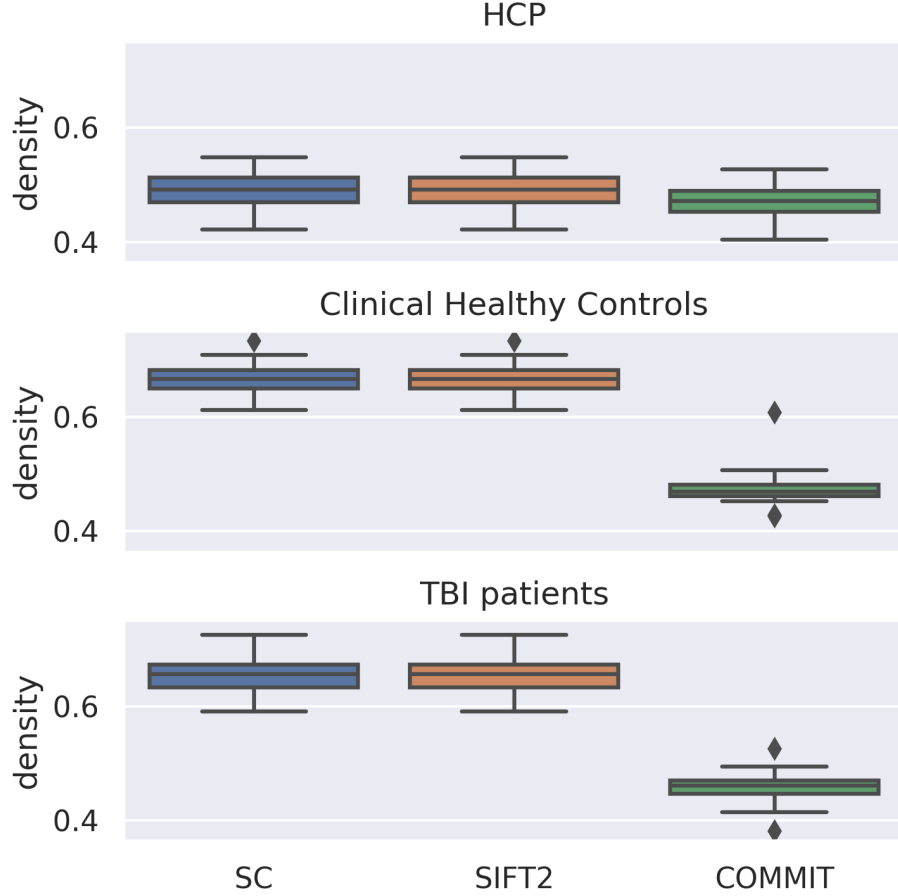


Figure 3. Each box-and-whisker plot represents the density of the connectomes obtained via SC, SIFT2 and COMMIT on subjects in specific cohorts (HCP, healthy controls of the clinical dataset and TBI patients respectively). SC and SIFT2 connectomes systematically have the same density.

between the GTMs computed with each pair of TFTs. On the contrary, the clinical dataset shows less differences, even if they are present in every GTM. In particular, no significant difference was detected between the GE computed with COMMIT and SC. Also, CPL, CC and MO do not exhibit significant differences when computed with SIFT2 and SC.

Figure 6 shows the relative change of the graph-theoretical metrics computed on the filtered connectomes of each subject in the three cohorts pruned with density-based thresholding. In particular, the first row of panels shows the relative change of the GTM computed at specific density with respect to the one computed at base density for the HCP dataset for each considered TFT. The shape of the resulting curves is qualitatively similar to the one reported in [12, Figure 1] not only for the SC and SIFT2 connectomes (as one would expect, since the experimental setup of the present and the cited work are very similar) but also for the COMMIT connectomes, which we recall having a different base density. There is no evident qualitative difference




Figure 4. These panels show the results of the edge-wise statistical comparison of connectomes of patients with brain injury versus the healthy controls of the clinical dataset using the Mann–Whitney U test with FDR multiple comparisons correction on SC, SIFT2 and COMMIT connectomes. The displayed rank-biserial correlations are obtained from the Mann–Whitney U test where the healthy controls are the first group and the TBI patients are the second group. For example, the red dots correspond to connections for which the controls are likely to be stronger than the patients. Only significant differences with $p \leq 0.05$ are shown.

between the deviation curves of the healthy controls from the clinical subjects (second row) and the TBI patients (third row). When compared to the deviation curves of the HCP dataset, it is possible to notice that the latter shows less inter-TFT variability. Tables 2, 3 and 4 report the maximal density across all subjects in the three datasets respectively such that the relative change of the considered graph-theoretical measure is below the 1% and the 5% threshold for each type of connectome. We observe that the robustness of connectomes built via streamline count is similar to the one of connectomes built employing tractography filtering techniques. Table 4 shows that integration measures (GE and CPL) are more robust to density-based thresholding if investigated on connectomes built via COMMIT. Segregation measures show comparable performances among the three filtering techniques.

5. Discussion and Conclusions

In this study, we investigated how the topological analysis of structural brain networks estimated from dMRI is affected by the use of tractography filtering techniques and density-based thresholding in connectomic pipelines. In particular, each considered dataset has been processed separately with SIFT2 and COMMIT in addition to being thresholded.

Analysing data of both clinical quality (shorter acquisition, more noise) and research quality (longer acquisition, less noise) we had the possibility to explore the effects of tractography filtering techniques on the topological analysis of structural brain networks estimated from dMRI data. The clinical cohort involved subjects affected by TBI,

		GE			CPL			CC			MO			
Clinical	HCP	SC		-0.65-0.54		-0.67-0.87		-0.69-0.99		-0.27	0.19			
		SIFT2	0.65		0.23	0.67		-0.58	0.69		-0.97	0.27		0.45
		COMMIT	0.54	-0.23		0.87	0.58		0.99	0.97		-0.19	-0.45	
	Healthy Controls	SC		0.59				0.8				0.73		0.97
		SIFT2	-0.59		-0.75			0.84				0.62		0.96
		COMMIT		0.75		-0.8	-0.84		-0.73	-0.62		-0.97	-0.96	
	TBI patients	SC		0.5				0.5				0.51		0.97
		SIFT2	-0.5		-0.46			0.67				0.5		0.96
		COMMIT		0.46		-0.5	-0.67		-0.51	-0.5		-0.97	-0.96	
		SC	SIFT2	COMMIT	SC	SIFT2	COMMIT	SC	SIFT2	COMMIT	SC	SIFT2	COMMIT	

Figure 5. Each row of 3x3 matrices represents data from a unique subject cohort. Each column made of 3x3 matrices shows the results for a specific GTM. Each 3x3 matrix shows the rank-biserial correlation effect size measure between the GTM computed with the TFTs indexed by row and the column of the observed entry. The displayed rank-biserial correlations are obtained from the Mann-Whitney U test where the first group is indexed by the row and the second group is indexed by the column. For example, the effect size on MO of COMMIT (group 1) with respect to SC (group 2) is reported to be 0.97, which means that whenever one compares the MO of a COMMIT connectome with the one of a SC connectome, it is likely that the first will be higher than the second. Only significant differences with $p \leq 0.05$ are shown.

a connectivity disorder of the brain that changes the topological properties of brain networks and is characterised by high inter-subject heterogeneity. The inclusion of the research quality data represented by the HCP subjects reflected the two necessities of testing the studied state-of-the-art techniques on high quality data for reference and to give a preliminary insight on the effects of these techniques on data that in some years could be available in the daily clinical practice.

Understanding if and how the interpretation of patho-connectomic studies is affected by the use of TFTs is of fundamental importance. In the present work we showed that while performing edge-wise comparisons between cohorts of healthy and TBI-affected subjects, one should take into account that the use of TFTs does change the set of edges showing significant differences. The result can not be straightforwardly generalized to any other pathology. In light of the fact that TFTs have different effects on the HCP dataset and on the clinical dataset and keeping in mind that

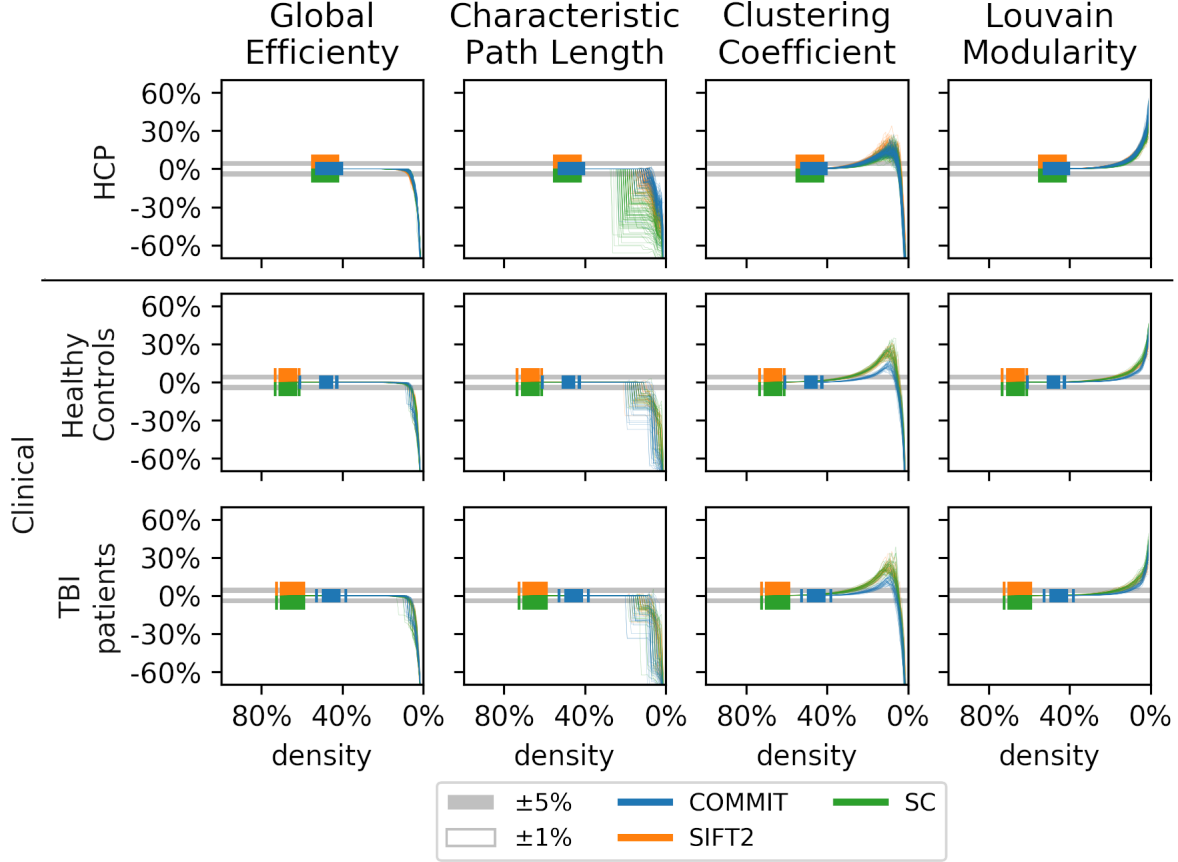


Figure 6. Each panel concerns the analysis of the robustness to density-based thresholding of filtered connectomes obtained from one subject cohort. The plots show the relative change of the graph-theoretical measure computed at a specific density with respect to the one computed at base density. The vertical ticks represent the base density of the connectomes of each subject. The gray bands indicate the 1% range and the 5% range around the no-deviation line.

ground-truth knowledge on the topology of brain networks is not available, we observe that the interpretability of studies of group differences between populations could be unpredictably affected by the use of tractography filtering techniques. Extra caution and further investigations addressing the specific problem of the nature of these effects are required.

The effect on the single edge weights induced by each TFT is reflected by the different density of each type of connectome. Looking at Figure 3 it is possible to notice how the impact of COMMIT on the density of the connectomes of the 100 HCP subjects is relatively small. On the contrary, connectomes obtained from the subjects of the clinical dataset show a remarkable difference between the density of those computed via SC/SIFT2 and those computed via COMMIT. As we reported in Section 4, there is a difference of $\sim 17\%$ between these two densities. Moreover, the average density of the COMMIT connectomes obtained from the clinical dataset is very close (2% difference) to the one obtained for the HCP subjects. This similarity allows us to conjecture the

	SC		COMMIT		SIFT2	
	>1%	>5%	>1%	>5%	>1%	>5%
GE	11%	6%	11%	6%	13%	7%
CPL	26%	26%	23%	23%	19%	19%
MO	24%	13%	30%	16%	30%	16%
CC	35%	22%	37%	25%	39%	26%

Table 2. Maximal density across all subjects in the *HCP* dataset that realized a deviation of at least 1% and 5% from the base value of each studied graph-theoretical metric GE, CPL, MO and CC computed on the three types of connectomes SC, COMMIT and SIFT2. Considering only low-to-moderate pruning, the obtained results are in line with the ones shown in [12].

	SC		COMMIT		SIFT2	
	>1%	>5%	>1%	>5%	>1%	>5%
GE	13%	6%	10%	8%	12%	7%
CPL	19%	19%	19%	19%	20%	20%
MO	36%	18%	31%	12%	36%	18%
CC	48%	30%	42%	26%	49%	30%

Table 3. Maximal density across all subjects in the *healthy clinical* dataset that realized a deviation of at least 1% and 5% from the base value of each studied graph-theoretical metric GE, CPL, MO and CC computed on the three types of connectomes SC, COMMIT and SIFT2. Considering only low-to-moderate pruning, the obtained results are in line with the ones shown in [12].

	SC		COMMIT		SIFT2	
	>1%	>5%	>1%	>5%	>1%	>5%
GE	13%	12%	9%	9%	13%	7%
CPL	18%	18%	19%	19%	18%	17%
MO	36%	18%	28%	14%	36%	18%
CC	49%	30%	39%	23%	48%	29%

Table 4. Maximal density across all subjects in the *TBI patients* dataset that realized a deviation of at least 1% and 5% from the base value of each studied graph-theoretical metric GE, CPL, MO and CC computed on the three types of connectomes SC, COMMIT and SIFT2. Considering only low-to-moderate pruning, the obtained results are in line with the ones shown in [12].

existence of a characteristic density of structural brain networks which may be better captured by sparsity-inducing TFTs. Confirming this conjecture will require further investigations. Nevertheless, the sensibly lower spatial and angular resolution of the clinical data could have affected the sharpness of the fODFs, inducing the creation of many spurious connections that COMMIT detected in the filtering procedure.

In this study we also explored the changes in the global efficiency, the characteristic

path length, the modularity and the clustering coefficient on connectomes obtained through a connectomic pipeline involving a tractography filtering step where SIFT2 or COMMIT were used. Connectomes determined with the streamline count strategy served as a reference since they are computed excluding the tractography filtering step. From the results presented in Figure 5 we confirm that the topology of connectomes is changed by the employment of tractography filtering techniques. These changes appear in both the high resolution data represented by the HCP subjects and the low resolution data acquired in a clinical study. Moreover, clinical data showed changes both in healthy subjects and in TBI patients. On another note, the differences between the results obtained on different datasets by both SIFT2 and COMMIT suggest that the quality of the data is highly influential also in connectomic pipelines involving tractography filtering techniques.

As discussed in the paragraphs above, the changes caused by the employment of TFTs involve every aspect of connectomic studies from edge weights through connectome density to graph-theoretical analyses. In particular, the effect on the distribution of the edge weights within the connectome could be modified by TFTs. Studying this distribution is particularly interesting when selecting the optimal threshold for density-based thresholding of the connectomes, which could be employed in order to remove the weakest edges of the filtered connectomes. This had already been thoroughly studied by Civier et al. [12] in a recent work by considering only SIFT2 connectomes and healthy subjects from both the HCP database and clinical acquisitions. In this work we successfully replicated their experiments and we extended their analysis by considering an additional TFT, i.e. COMMIT, and a supplementary cohort of subjects affected by TBI. In Figure 6 we give evidence to the fact that low-to-moderate density-based thresholding does not affect the analysis of filtered connectomes obtained from data of TBI patients. We confirm that low-to-moderate pruning is not advisable as it would require to justify the choice of the threshold by means of arbitrary or heuristic-based arguments at the price of no evident beneficial effect.

A possible improvement of this work regards the number of employed streamlines, as we considered tractograms made of 2 million streamlines compared to the ones used in [12], where the authors considered tractograms composed of 10 million streamlines. This choice was induced by the limited capability of COMMIT to work on tractograms with more streamlines than the ones used here. The influence of this parameter on the performed analysis will be studied in future works.

Overall, this study highlights that the application of SIFT2 and COMMIT to diffusion MRI-based structural connectomics affects the measurement of global efficiency, characteristic path length, modularity and clustering coefficient of the estimated brain networks. Moreover, the interpretation of group differences in patho-connectomics is altered by the use of TFTs. As such, more research and extra caution are needed prior to incorporating tractography filtering into connectomic analysis pipelines in clinical studies. Finally, the practice of density-based thresholding in the context of graph-theoretical studies of structural brain networks obtained via tractography filtering

is confirmed to have negligible effects both on healthy subjects and patients affected by traumatic brain injury.

Open science

All of the code needed to reproduce this study will be made available under reasonable request to the corresponding author. The connectomes obtained from the subjects of the HCP database are available at <https://osf.io/7cner>.

Acknowledgements

This work received funding from 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) and NIH R01NS065980, PA-Department of Health PACT award and NIH R01NS096606.

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

- [1] Rachid Deriche. “Computational brain connectivity mapping: A core health and scientific challenge”. In: *Medical Image Analysis* 33 (2016), pp. 122–126.
- [2] Olaf Sporns, Giulio Tononi, and Rolf Kötter. “The human connectome: a structural description of the human brain”. In: *PLoS computational biology* 1.4 (2005), e42.
- [3] Patric Hagmann. *From diffusion MRI to brain connectomics*. Tech. rep. EPFL, 2005.
- [4] Danielle S Bassett and Edward T Bullmore. “Small-world brain networks revisited”. In: *The Neuroscientist* 23.5 (2017), pp. 499–516.
- [5] Ragini Verma et al. “Neuroimaging Findings in US Government Personnel With Possible Exposure to Directional Phenomena in Havana, Cuba”. In: *Jama* 322.4 (2019), pp. 336–347.
- [6] Chun-Hung Yeh et al. “Connectomes from streamlines tractography: Assigning streamlines to brain parcellations is not trivial but highly consequential”. In: *NeuroImage* 199 (2019), pp. 160–171. ISSN: 1053-8119. DOI: <https://doi.org/10.1016/j.neuroimage.2019.05.005>.
- [7] Marco Catani and Michel Thiebaut de Schotten. *Atlas of human brain connections*. Oxford University Press, 2012.

- [8] Derek K Jones, Thomas R Knösche, and Robert Turner. “White matter integrity, fiber count, and other fallacies: the do’s and don’ts of diffusion MRI”. In: *Neuroimage* 73 (2013), pp. 239–254.
- [9] Robert E Smith et al. “The effects of SIFT on the reproducibility and biological accuracy of the structural connectome”. In: *Neuroimage* 104 (2015), pp. 253–265.
- [10] Francois Rheault et al. “Common misconceptions, hidden biases and modern challenges of dMRI tractography”. In: *Journal of Neural Engineering* (2020).
- [11] Klaus H Maier-Hein et al. “The challenge of mapping the human connectome based on diffusion tractography”. In: *Nature communications* 8.1 (2017), p. 1349.
- [12] Oren Civier et al. “Is removal of weak connections necessary for graph-theoretical analysis of dense weighted structural connectomes from diffusion MRI?” In: *NeuroImage* 194 (2019), pp. 68–81.
- [13] Robert E Smith et al. “SIFT2: Enabling dense quantitative assessment of brain white matter connectivity using streamlines tractography”. In: *Neuroimage* 119 (2015), pp. 338–351.
- [14] Alessandro Daducci et al. “COMMIT: convex optimization modeling for microstructure informed tractography”. In: *IEEE transactions on medical imaging* 34.1 (2014), pp. 246–257.
- [15] Franco Pestilli et al. “Evaluation and statistical inference for human connectomes”. In: *Nature methods* 11.10 (2014), p. 1058.
- [16] Olaf Sporns et al. “Organization, development and function of complex brain networks”. In: *Trends in cognitive sciences* 8.9 (2004), pp. 418–425.
- [17] Shan Yu et al. “A small world of neuronal synchrony”. In: *Cerebral cortex* 18.12 (2008), pp. 2891–2901.
- [18] Danielle Smith Bassett and ED Bullmore. “Small-world brain networks”. In: *The neuroscientist* 12.6 (2006), pp. 512–523.
- [19] Ragini Verma, Yusuf Osmanlioglu, and Abdol Aziz Ould Ismail. “Multimodal Patho-Connectomics of Brain Injury”. In: *International MICCAI Brainlesion Workshop*. Springer. 2018, pp. 3–14.
- [20] Anand S Pandit et al. “Traumatic brain injury impairs small-world topology”. In: *Neurology* 80.20 (2013), pp. 1826–1833.
- [21] Mehrbod Mohammadian et al. “High angular resolution diffusion-weighted imaging in mild traumatic brain injury”. In: *NeuroImage: Clinical* 13 (2017), pp. 174–180.
- [22] Matthew F Glasser et al. “The minimal preprocessing pipelines for the Human Connectome Project”. In: *Neuroimage* 80 (2013), pp. 105–124.
- [23] Chun-Hung Yeh et al. “Correction for diffusion MRI fibre tracking biases: The consequences for structural connectomic metrics”. In: *Neuroimage* 142 (2016), pp. 150–162.

- [24] Matteo Frigo et al. “Effects of tractography filtering on the topology and interpretability of connectomes”. In: *OHBM Meeting*. 2019.
- [25] Andrew Zalesky et al. “Connectome sensitivity or specificity: which is more important?” In: *Neuroimage* 142 (2016), pp. 407–420.
- [26] Alex Fornito, Andrew Zalesky, and Michael Breakspear. “Graph analysis of the human connectome: promise, progress, and pitfalls”. In: *Neuroimage* 80 (2013), pp. 426–444.
- [27] Stamatios N Sotiropoulos and Andrew Zalesky. “Building connectomes using diffusion MRI: why, how and but”. In: *NMR in Biomedicine* 32.4 (2019), e3752.
- [28] Stanley Milgram. “The small world problem”. In: *Psychology today* 2.1 (1967), pp. 60–67.
- [29] Mikail Rubinov and Olaf Sporns. “Complex network measures of brain connectivity: uses and interpretations”. In: *Neuroimage* 52.3 (2010), pp. 1059–1069.
- [30] Edsger W Dijkstra. “A note on two problems in connexion with graphs”. In: *Numerische mathematik* 1.1 (1959), pp. 269–271.
- [31] Vito Latora and Massimo Marchiori. “Efficient behavior of small-world networks”. In: *Physical review letters* 87.19 (2001), p. 198701.
- [32] Bin Zhang and Steve Horvath. “A general framework for weighted gene co-expression network analysis”. In: *Statistical applications in genetics and molecular biology* 4.1 (2005).
- [33] Tore Opsahl and Pietro Panzarasa. “Clustering in weighted networks”. In: *Social networks* 31.2 (2009), pp. 155–163.
- [34] Mark EJ Newman. “Analysis of weighted networks”. In: *Physical review E* 70.5 (2004), p. 056131.
- [35] David C Van Essen et al. “The Human Connectome Project: a data acquisition perspective”. In: *Neuroimage* 62.4 (2012), pp. 2222–2231.
- [36] 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.
- [37] Stamatios N Sotiropoulos et al. “Advances in diffusion MRI acquisition and processing in the Human Connectome Project”. In: *Neuroimage* 80 (2013), pp. 125–143.
- [38] Jesper LR Andersson, Stefan Skare, and John Ashburner. “How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging”. In: *Neuroimage* 20.2 (2003), pp. 870–888.
- [39] Mark Jenkinson et al. “Fsl”. In: *Neuroimage* 62.2 (2012), pp. 782–790.

- [40] JLR Andersson et al. “A comprehensive Gaussian process framework for correcting distortions and movements in diffusion images”. In: *Proceedings of the 20th Annual Meeting of ISMRM*. Vol. 20. 2012, p. 2426.
- [41] Amanda R Rabinowitz et al. “Neuropsychological recovery trajectories in moderate to severe traumatic brain injury: influence of patient characteristics and diffuse axonal injury”. In: *Journal of the International Neuropsychological Society* 24.3 (2018), pp. 237–246.
- [42] José V Manjón et al. “Diffusion weighted image denoising using overcomplete local PCA”. In: *PloS one* 8.9 (2013), e73021.
- [43] Stephen M Smith. “Fast robust automated brain extraction”. In: *Human brain mapping* 17.3 (2002), pp. 143–155.
- [44] Jesper LR Andersson and Stamatiios N Sotiropoulos. “An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging”. In: *Neuroimage* 125 (2016), pp. 1063–1078.
- [45] Robert E Smith et al. “Anatomically-constrained tractography: improved diffusion MRI streamlines tractography through effective use of anatomical information”. In: *Neuroimage* 62.3 (2012), pp. 1924–1938.
- [46] Bruce Fischl. “FreeSurfer”. In: *Neuroimage* 62.2 (2012), pp. 774–781.
- [47] J-Donald Tournier et al. “MRtrix3: A fast, flexible and open software framework for medical image processing and visualisation”. In: *NeuroImage* (2019), p. 116137.
- [48] Ben Jeurissen et al. “Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data”. In: *NeuroImage* 103 (2014), pp. 411–426.
- [49] J-Donald Tournier et al. “Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution”. In: *NeuroImage* 23.3 (2004), pp. 1176–1185.
- [50] Thijs Dhollander, David Raffelt, and Alan Connelly. “Unsupervised 3-tissue response function estimation from single-shell or multi-shell diffusion MR data without a co-registered T1 image”. In: *ISMRM Workshop on Breaking the Barriers of Diffusion MRI*. Vol. 5. 2016.
- [51] J-Donald Tournier, Fernando Calamante, and Alan Connelly. “Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution”. In: *Neuroimage* 35.4 (2007), pp. 1459–1472.
- [52] Rahul S Desikan et al. “An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest”. In: *Neuroimage* 31.3 (2006), pp. 968–980.
- [53] Eleftheria Panagiotaki et al. “Compartment models of the diffusion MR signal in brain white matter: a taxonomy and comparison”. In: *Neuroimage* 59.3 (2012), pp. 2241–2254.

- [54] Vincent D Blondel et al. “Fast unfolding of communities in large networks”. In: *Journal of statistical mechanics: theory and experiment* 2008.10 (2008), P10008.
- [55] Henry B Mann and Donald R Whitney. “On a test of whether one of two random variables is stochastically larger than the other”. In: *The annals of mathematical statistics* (1947), pp. 50–60.
- [56] Yoav Benjamini and Yosef Hochberg. “Controlling the false discovery rate: a practical and powerful approach to multiple testing”. In: *Journal of the Royal statistical society: series B (Methodological)* 57.1 (1995), pp. 289–300.