



HAL
open science

Computational Brain Connectivity Mapping: From Multi-Compartment Modeling To Network Topology Via Tractography Filtering

Matteo Frigo

► **To cite this version:**

Matteo Frigo. Computational Brain Connectivity Mapping: From Multi-Compartment Modeling To Network Topology Via Tractography Filtering. Signal and Image processing. Université Côte d'Azur, 2021. English. NNT: 2021COAZ4003 . tel-03273337

HAL Id: tel-03273337

<https://theses.hal.science/tel-03273337>

Submitted on 29 Jun 2021

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.

THÈSE DE DOCTORAT

Cartographie computationnelle de la connectivité cérébrale

De la modélisation multi-compartiments
à la topologie des réseaux
via la tractographie par filtrage

Matteo FRIGO

Inria Sophia Antipolis – Méditerranée, Équipe-Projet Athéna

**Présentée en vue de l'obtention
du grade de docteur en**
Automatique et Traitement du Signal
et des Images

d'Université Côte d'Azur

Dirigée par : Rachid Deriche

Co-encadrée par : Samuel
Deslauriers-Gauthier

Date de soutenance :
22 Février 2021

Devant le jury, composé de :

Alfred Anwander, Chargé de Recherche, Max Planck
Institute for Human Cognitive and Brain Sciences,
Germany – Examineur

Rachid Deriche, Directeur de Recherche, Inria Sophia
Antipolis Méditerranée – Directeur de thèse

Maxime Descoteaux, Professeur titulaire, Université
de Sherbrooke – Rapporteur

Samuel Deslauriers-Gauthier, Chargé de Recherche
SRP ERC, Inria Sophia Antipolis Méditerranée – Co-
Encadrant

Théodore Papadopoulos, Directeur de Recherche, Inria
Sophia Antipolis Méditerranée – Examineur

Jean-Philippe Thiran, Professeur, École polytechnique
fédérale de Lausanne – Rapporteur

Ragini Verma, Professeur, Section of Biomedical
Image Analysis, Department of Radiology, University
of Pennsylvania – Examinatrice

PhD Thesis

Computational Brain Connectivity Mapping

**From Multi-Compartment Modeling To Network Topology Via
Tractography Filtering**

Matteo Frigo

Date of defense: February 22nd 2021

Athena Project Team
Inria Sophia Antipolis - Méditerranée
Université Côte d'Azur

To my parents
for their love and care.

To my grandparents
for whom emigrating was a necessity
not a privilege.

Funding

This doctoral project 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](#)).



European Research Council
Established by the European Commission

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.

The authors are grateful to Inria Sophia Antipolis - Méditerranée "[Nef](#)" computation cluster for providing resources and support.

Abstract (english)

Mapping the human brain is one of the complex challenges of contemporary science. It is a task that concatenates several problems from acquisition design to preprocessing, modelling, analysis, visualisation and assessment of the coherence with the state-of-the-art knowledge on the architecture and functioning of the human brain. For each of these steps a plethora of solutions has been and is being developed. It is of fundamental importance that the assumptions made in each step align with each other, demanding extra care in the verification of the theoretical requirements of the employed tools. In this thesis we focus on three specific parts of the chain of problems that leads to a comprehensive view of the brain architecture, highlighting the theoretical aspects that characterise the posed challenges and providing experimental evidence of the soundness of the proposed solutions. We present four contributions on three topical research areas of diffusion MRI methods for human brain mapping: brain tissue microstructure, tractography filtering and topological analysis of brain networks. First, we propose a new method for the estimation of tissue-specific volume fractions by means of multi-compartment models of the single-TE diffusion MRI signal. Then, we review the state of the art of tractography filtering and unveil its effects on the graph-theoretical analysis of the structural connectomes of both healthy subjects and patients affected by traumatic brain injury. In addition, we propose a novel filtering technique that integrates structural and functional information in the process. Finally, we propose a new similarity measure between brain networks and a new graph alignment techniques, allowing to obtain original insights into the problem of selecting the suitable parcellation for brain connectivity studies.

Keywords: Diffusion MRI, Brain Connectivity, Microstructure, Tractography Filtering, Graph Alignment, Spherical Harmonics, Brain Network, Topology, Connectome.

Abstract (français)

La cartographie du cerveau humain est l'un des défis complexes de la science contemporaine. Il s'agit d'une tâche qui concatène plusieurs problèmes allant du design de l'acquisition au prétraitement, en passant par la modélisation, l'analyse, la visualisation et l'évaluation de la cohérence avec l'état de l'art sur l'architecture et le fonctionnement du cerveau humain. Pour chacune de ces étapes, une pléthore de solutions a été et est en cours de développement. Il est d'une importance fondamentale que les hypothèses faites dans chaque étape soient alignées les unes avec les autres, ce qui exige une attention particulière dans la vérification des aspects théoriques des outils employés. Dans cette thèse, nous nous concentrons sur trois parties spécifiques de la chaîne de problèmes qui mène à une vision globale de l'architecture du cerveau, en soulignant les aspects théoriques qui caractérisent les défis posés et en fournissant des preuves expérimentales de la solidité des solutions proposées. Nous présentons quatre contributions sur trois domaines de recherche actuels des méthodes d'IRM de diffusion pour la cartographie du cerveau humain : la microstructure du tissu cérébral, le filtrage de la tractographie et l'analyse topologique des réseaux cérébraux. Tout d'abord, nous proposons une nouvelle méthode pour l'estimation des fractions volumiques spécifiques aux tissus au moyen de modèles multi-compartiments du signal d'IRM de diffusion à un seul TE. Ensuite, nous passons en revue l'état de l'art du filtrage de la tractographie et dévoilons ses effets sur l'analyse grapho-théorique des connectomes structurels de sujets sains et de patients affectés par traumatisme cranio-cérébral. En outre, nous proposons une nouvelle technique de filtrage qui intègre les informations structurelles et fonctionnelles dans le processus. Enfin, nous proposons une nouvelle mesure de similarité entre les réseaux cérébraux et une nouvelle technique d'alignement de graphes, ce qui permet d'obtenir des informations originales sur le problème de la sélection de l'atlas approprié pour les études de connectivité cérébrale.

Keywords: IRM de Diffusion, Connectivité Cérébrale, Microstructure, Tractographie par Filtrage, Alignement de Graphe, Harmoniques Sphériques, Réseaux Cérébraux, Topologie, Connectome.

Acknowledgements

A doctoral thesis has one single author, but its conception, growth and finalisation all require a collective effort. The two people that have been with me in all these stages, and that I will never thank enough for this, are my advisor Rachid Deriche and my co-advisor Samuel Deslauriers-Gauthier. Their trust, patience, wisdom and mentorship have been my companions in this three-year-long pursuit. I will forever be thankful for all the dedication that they have put into this project.

I would also like to thank the members of my jury for taking the time and responsibility to review and judge this dissertation. Thank you Alfred Anwander, Maxime Descoteaux, Théodore Papadopoulo, Jean-Philippe Thiran and Ragini Verma. Your comments and advice helped me put in perspective what I learnt in these years. Addressing each of your concerns undoubtedly increased the scientific value of this thesis.

Among all the people that I need to thank for helping me get through these years, I will start from Mauro Zucchelli, Ragini Verma, Emanuele Natale and Rutger Fick. They have been by my side in many steps of this thesis. Their support has been irreplaceable and I hope that one day I'll be able to help someone in the same way they helped me. May their brilliant minds continue to inspire many young researchers and students.

I also want to thank all the past, present and temporary members of the ATHENA team. Thank you Abib, Amandine, Antonia, Côme, Demian, Federica (both of you), Gloria, Guillermo, Imogen, Isa, Ivana, Joan, Kostiantyn, Lars, Lavinia, Maureen, Max, Nathalie, Oceane, Patryk, Paul-Emmanuel, Pierre, Rebecca, Romain, Sandra, Sara, Sofiane and Theo. It has been my greatest honour and pleasure to share these years with you and I hope that we will cross our paths again. In particular I would like to thank Kostia and Guillermo for constantly raising the bar of the weirdness that can be found on the internet. Having someone with which we can share a good laugh is the greatest of all blessings. Thank you for being that blessing. Merci aussi à Claire, qui m'a beaucoup aidé chaque fois que j'ai dû batailler avec la bureaucratie française. Nous pouvons avoir perdu quelque bataille, mais nous avons gagné la guerre. Merci pour tous !

During these years I've also had invaluable collaborators that helped me understand and value the plurality of challenges that populates science. Their contribution to this thesis goes well beyond the co-authorship of some of the works that it contains. Thank you Abdol Aziz Ould Ismail, David Coudert, Drew Parker, Emilio Cruciani, Julio Villalon Reina, Junghoon John Kim, Ragini Verma and Talia Nir. To me you represent the pleasure of looking at science as a mindful social exercise.

A special thank goes to the LTS5 team at EPFL, where I first got in touch with the world of diffusion MRI. Thank you Jean-Philippe Thiran and Alessandro Daducci for welcoming me into this universe and giving me the opportunity to move my first steps among yours and the ones of Alonso Ramirez-Manzanares, Anna Auria Rasclosa, David Romascano, Erick Jorge Canales-Rodríguez, Gabriel Girard, Giorgio Innocenti, Jonathan Rafael Patino Lopez, Marco Pizzolato and Muhamed Barakovic. I keep a clear memory of the curiosity that animated all our discussions. Thank you also to Alessandro, Alexis, Gianluca, Jacopo, Monsieur and Peggy for making my life in Lausanne as happy as I remember it.

I also want to dedicate a few words to those that first ignited my passion for mathematics and its applications. It all started at highschool with my teacher Emiliana Tonin, who made me realise that math was my thing. Then I found my mentors in Verona in Giandomenico Orlandi, Marco Caliari, Luca di Persio and Simone Zuccher, who helped me shape my way of thinking and directed me towards my first experiences in signal processing and imaging sciences with Manuele Bicego and Roberto Foroni. Speaking about the Verona times, I have to mention my classmate and partner in crime Alessandro Mella, who owns an exceptionally brilliant mind and that I will never thank enough for everything that we shared in those years. Last but not least, thank you Andrea, Anna, Melissa, Michele and Franco. The world that we built for ourselves changed our way of sharing part of our lives, but I always think about the great times we shared.

These years have certainly been made easier by my old friends Marco, Alberto, Ilaria, Davide (all of you), Claudio, Greta and Marelsa. Grassie tusi par tuto. Spero tanto che se vedimo presto a Carmignan (a rima ve a metì da soi, che fa brutto scriverla su na tesi...).

The happiness that characterised my french times is also due to my friends in Nice, starting from Giovanni and Giulia, who have been there since the very beginning, and continuing with Alberto, Aldo, Andi, Anne, Chiara, Felice, Giandonato, the other Giulia, Isabel, Isabella, Lorenzo, Luca, Marco, Maria Giulia, Marie, Marika, Massimo, Matteo, Riccardo, Saveria, Sofia and Vittorio. Thank you for constantly reminding me to never take myself too seriously.

I want to thank my parents and grandparents, to which this thesis is dedicated. You are the pillars on which everything I said before is based. Nothing would have been possible without you. Thank you from the bottom of my heart.

And thank you Ottavia. You have been the most beautiful surprise of these years. Being at your side fills me with joy every single day. May this never end.

Contents

Abstract (english)	vi
Abstract (français)	vii
Acknowledgements	viii
Contents	x
1 Introduction	1
2 Background	7
2.1 Essential Anatomy of the Human Brain	7
2.1.0.1 Gray Matter	8
2.1.0.2 White Matter	8
2.2 Brain Imaging	9
2.2.1 Magnetic Resonance Imaging	9
2.2.2 Diffusion MRI	12
2.2.2.1 The Tensor Representation	14
2.2.2.2 HARDI	16
2.2.2.3 The Spherical Harmonics Representation	17
2.2.2.4 Tractography	18
2.3 Connectomics	20
2.3.1 Atlases of the Human Brain	21
2.3.2 Structural Connectivity	23
2.3.3 Functional Connectivity	24
2.3.4 Effective Connectivity	25
2.3.5 Topology of Brain Networks	26
3 MT-MC models of dMRI	29
3.1 Introduction	29
3.2 Theory	31
3.2.1 Multi-Compartment models	31
3.2.2 The Standard Model of White Matter	32
3.2.3 MC models do not account for T2 differences	34
3.2.4 Multi-TE MC models of dMRI signal	35
3.2.5 Multi-Tissue MC models	37
3.2.5.1 Fitting MT-MC models	39
3.2.6 Generalised MT-MC Modelling	40
3.2.7 The MT Standard Model of dMRI in White Matter	41
3.3 Experiments	41
3.3.1 Synthetic phantom	42
3.3.1.1 Dataset	42
3.3.1.2 Results	43
3.3.2 Real Data	44
3.3.2.1 Dataset	44
3.3.2.2 Results	44
3.4 Discussion	46
3.5 Conclusion	48

4	Filtering dMRI-based tractograms	49
4.1	Introduction	49
4.2	Theory	51
4.2.1	Tractography Filtering Techniques	51
4.2.1.1	SIFT2	52
4.2.1.2	COMMIT2	53
4.2.2	Unified Tractography Filtering Framework	57
4.2.2.1	Computational aspects	58
4.2.3	Functionally Informed Tractography Filtering	60
4.3	Experiments	62
4.3.1	TFTs Change the Topology of Structural Connectomes	62
4.3.1.1	Data and Methods	63
4.3.1.2	Results	67
4.3.1.3	Discussion	70
4.3.2	Functionally-informed TFT	72
4.3.2.1	Data and methods	72
4.3.2.2	Results	75
4.3.2.3	Discussion	78
4.4	Conclusion	80
5	Brain Alignment and Similarity	81
5.1	Introduction	81
5.2	Theory	83
5.2.1	Brain Alignment	83
5.2.2	Quality of Brain Alignments	84
5.2.2.1	Graph Jaccard Similarity Index	85
5.2.3	Weisfeiler-Lehman Network Alignment	87
5.3	Experiments	90
5.3.1	Data and Methods	90
5.3.1.1	Parcellations	91
5.3.1.2	Connectomes	91
5.3.1.3	Intra-cohort variability	92
5.3.1.4	Network alignments	92
5.3.1.5	Quality of alignments	93
5.3.2	Statistical analysis	94
5.3.3	Results	94
5.3.3.1	Comparison between similarity measures	95
5.3.3.2	Computing brain alignments with WL-align	96
5.3.3.3	Region matching rate	99
5.4	Discussion and conclusion	100
6	Conclusion	104
A	Software contributions	110
A.1	TALON: Tractograms As Linear Operators in Neuroimaging	110
A.2	WL-align	112
A.3	COMMIT	112
A.4	PyUNLocBoX: Optimization by Proximal Splitting	112
A.5	Dmipy: Diffusion Microstructure Imaging in Python	113
A.6	Mrtrix3	113

B Other works	114
B.1 Structure-function mapping	114
B.2 dMRI-PLI resolution gap	114
B.3 Multi-Compartment models of Gray Matter	115
B.4 MEMENTO Challenge	115
Bibliography	117

Chapter 1

Introduction

Mapping the human brain is a core health ambition of the 21st century and one of the greatest collective efforts of modern science (Deriche 2016). In particular, investigating the architecture and the functioning of the brain provides fundamental insights into the evolutionary, developmental, anatomical, pathological and behavioural aspects of the human brain. The immense complexity of the challenge of mapping the brain network is induced by the multiplicity of spatial and temporal scales that need to be considered at the same time when studying the brain. The first obstacle concerns the number of neurons that compose the brain, which has been estimated to be in the order of a hundred billions, generating $\sim 160km$ of myelinated axons (Marner et al. 2003). Such numerosness makes it impossible represent the brain as a whole without simplifying assumptions at spatial scale. A second encumbrance is due to the temporal changes to which the brain is subject. Evolution, development and pathology all influence the brain plasticity, and the brain's functioning is certainly a dynamic phenomenon, given that we think, act and react in a variety of different ways over time. A third layer of complexity is brought by the strong multi-scale nature of the brain, as both the architecture and the functioning of the neuronal network are exhibit multiple levels of organisation (Bassett et al. 2017). This implies that in modelling the brain one has to take into account the interdependence between the different representations of anatomy and phenomena that are consequently used at each different scale.

Focusing on the *microscopic* level ($\sim 10^{-6}/^{-5}m$), one can observe the cytoarchitecture of the neurons, the diameter of their axons and the variety of bodies that populate the brain tissues. An analysis at what we will call *mesoscopic* level ($\sim 10^{-4}/^{-3}m$) provides descriptors of the brain tissues such as the fraction of volume occupied by a certain type of body or the local orientation of the white matter fibers. Finally, at *macroscopic* scale ($\sim 10^{-2}/\dots/2m$) it is possible to inspect how the brain can be subdivided into regions or how these regions are connected with each other. Encoding brain regions as nodes of a network whose edges are the aforementioned connections yields a graph-like representation of brain networks called connectome (Hagmann 2005; O. Sporns et al. 2005).

Diffusion Magnetic Resonance Imaging (dMRI) is an in-vivo imaging technique that allows to explore each of the three mentioned scales. In this sense, an illuminating example is the one of tractography, which is the art of tracing the trajectory followed by the axonal pathways in the white matter. With dMRI we are able to estimate how easily water molecules diffuse in a certain direction, hence we sense properties of microscopic phenomena. With this information we are able to gain insights on the local geometry of the white matter configurations and on the local orientation of the underlying axonal pathways, which are characteristics

of the tissues that are expressed at mesoscopic level. Following these orientations it is possible to obtain the trajectories followed by axonal pathways, obtaining a description of the macroscopic connectivity between grey matter regions. The network whose nodes are cortical grey matter regions and whose edges encode the axonal connectivity estimated with tractography is called *structural* connectome. Analogously, a *functional* connectome can be obtained by defining edges whose connection strength represents the degree of co-activation between the functional activity recorded in the linked brain regions via functional MRI (fMRI) or magneto/electroencephalography (M/EEG).

Representing the variety of information listed in the previous paragraphs requires the use of several types of map. The three-dimensional space where the brain lies is tessellated with voxels, which are the three-dimensional correspondent of one-dimensional intervals. For each voxel it is possible to encode scalar values (e.g., the fraction of volume populated by cerebro-spinal fluid) or vector values. The latter allows to describe more complex data like the distribution of axon diameters in the voxel, the signal obtained for different acquisition parameters of dMRI, the distribution of fiber orientations or the time series of the functional activity in the voxel. A different type of data is the one that describes the streamlines representing the trajectories of the axonal pathways estimated via tractography, which are encoded as three-dimensional curves. Also, the subdivision of the cerebral cortex that provides the description of the brain regions can be given on a three-dimensional mesh that describes the surface on which the cortex lies. In this case, different regions are encoded as different colours associated to the vertices of the mesh. Finally, connectomes are defined as graphs, hence they do not have a geometric interpretation in the three-dimensional space. However, it is possible to associate to each node the position that the corresponding region occupies in the brain. All together, the five types of brain maps listed above allow to give a comprehensive view on the architecture and functioning of the brain.

In this thesis we explored three different fields of brain mapping: brain tissue microstructure, tractography filtering and brain network topology. The first aims at estimating microstructural features of brain tissues by leveraging their water diffusion properties via dMRI. Tractography filtering encompasses some recent developments in the field of tractography that aim at overcoming limitations of tractography algorithms that have recently been identified, such as the systematic generation of a relevant portion of false positive connections and the limited quantitative interpretability of the obtained tractograms (i.e., sets of streamlines) (Jeurissen et al. 2019). Finally, the study of brain network topology allows to inspect connectomes from a purely graph-theoretical point of view. This approach has been showed to provide unique insights into the functioning and the architecture of the brain, such as the simultaneous insurgence of high integration and high segregation in the connectivity patterns, which in turn highlight how the brain is organised as a small-world network (Bassett et al. 2006, 2017).

Organisation of the manuscript

This document is structured in six chapters and two appendices, including this introduction, a background chapter, three main chapters devoted to the three research lines followed in this thesis and a conclusive chapter. Each reference is given in author-year format and the reader will find the complete bibliography in alphabetical order at the end of the manuscript.

In Chapter 2 we are going to introduce the theoretical background and the domain knowledge on which this thesis is built by reviewing the basic concepts of neuroanatomy, dMRI image formation and processing and the fundamental aspects of brain connectivity. Thereafter, in Chapter 3 we will propose a novel multi-compartment model of tissue microstructure that allows to estimate the volume fraction of different tissues from single-TE dMRI data. In Chapter 4 we will review the state-of-the-art tractography filtering techniques and show how their use has non-negligible effects on the network topology. Also, we will present a novel filtering technique that integrates both structural and functional information in the process. Chapter 5 will be devoted to addressing the question on whether the choice of the brain atlas has an impact on the robustness of the topological properties of the generated brain network. In order to do this, we will introduce a novel similarity measure for graphs and a new graph alignment heuristic. Finally, in Chapter 6 we will conclude by discussing our contributions and posing new questions for the future development of the field.

Chapter 3: Brain Tissue Microstructure

Recent studies highlighted how all the available multi-compartment (MC) models of white matter tissue microstructure via dMRI are defined in such a way that they are transparent to differences between the T_2 times of the modelled tissues (Lampinen et al. 2019; Veraart et al. 2018). As a consequence, they implicitly assume that all the considered tissues have the same S_0 response. We showed how the concept of *signal fraction* is more appropriate to describe what have always been referred to as *volume fraction* (Frigo et al. 2020c). Standard methods for overcoming such problem are based on the acquisition of multi-TE multi-shell dMRI data instead of the more classical single-TE multi-shell dMRI data (Gong et al. 2020; Lampinen et al. 2019, 2020; Veraart et al. 2018). This allows to define a multi-TE version of standard MC models that simultaneously fits the volume fraction and the T_2 time of each modelled tissue.

In this chapter we will show how it is possible to define multi-tissue (MT)-MC models that describe *volume* fractions using only single-TE multi-shell dMRI data. This is made possible by the inclusion of prior knowledge on the S_0 response of each modelled tissue obtained from third-part heuristics such as the one of Dhollander et al. 2016a. In Chapter 3 we will provide a complete theoretical presentation of the motivations and formalisation of our MT-MC model, exposing the limitation of the state-of-the-art multi-TE models and proposing a generalised framework for modelling multiple tissues with MC models. Moreover, we present how this method is implemented in the Dmipy framework Fick et al. 2019.

Finally, we discuss the applicability of the MT-MC model following the results we obtained on both simulated data prepared in-house and real data from the Human Connectome Project (HCP) database. This work was done in collaboration with Rutger Fick*. Partial results have been presented at the International Symposium on Biomedical Imaging (ISBI) of 2020 Frigo et al. 2020c and at the 26th meeting of the Organization for Human Brain Mapping Frigo et al. 2020d.

- ▶ *Matteo Frigo, Mauro Zucchelli, Rutger Fick, Samuel Deslauriers-Gauthier, Rachid Deriche. Multi-compartment modelling of diffusion MRI signal shows TE-based volume fraction bias. OHBM 2020.*
- ▶ *Matteo Frigo, Rutger Fick, Mauro Zucchelli, Samuel Deslauriers-Gauthier, Rachid Deriche. Multi Tissue Modelling of Diffusion MRI Signal Reveals Volume Fraction Bias. ISBI 2020.*
- ▶ *Matteo Frigo, Rutger Fick, Mauro Zucchelli, Samuel Deslauriers-Gauthier, Rachid Deriche. Multi Tissue Multi Compartment Modelling of Single-TE Diffusion MRI. Journal paper submitted to Neuroimage.*

Chapter 4: Tractography Filtering Techniques

Several works analysed the limitations of dMRI-based tractography (Jbabdi et al. 2011; Jones et al. 2013), in particular in the context of structural connectivity analysis, where a non-trivial quantity of false positive connections have been shown to be present (Maier-Hein et al. 2017) and detrimental (Zalesky et al. 2016). Some methods that address this issue go under the name of *tractography filtering* techniques (TFTs) (Daducci et al. 2013, 2014; Pestilli et al. 2014; Schiavi et al. 2020; R. E. Smith et al. 2013, 2015a). These techniques act as a post-processing of pre-computed tractograms by assigning a coefficient to each streamline. This coefficient represents the amount of signal explained by the streamline or the “*connectivity strength*” associated to it, depending on the employed technique, and it is used for the definition of the so-called weighted structural connectome. This is a type of structural network encoded as a weighted graph whose nodes are brain regions and whose weighted edges encode the strength of the axonal connection that they represent. Standard methods associate to each edge the number of streamlines connecting the two corresponding regions. Weighted connectomes extend this definition by allowing each streamline to contribute differently, e.g., with the coefficients retrieved by TFTs. In this way the edge weights can be computed as the sum of the weights associated to the streamlines that connect two regions.

In this chapter we present a review of the most common TFTs, focusing on how they can be framed into a more general formulation. Additionally, we assess if and how the state-of-the-art TFTs have an effect on the topology of the connectomes that they produce. We show that including TFTs in connectomic pipelines changes the topology of structural brain networks, and thus alters network metrics both in the pathological (traumatic brain injury) and the healthy cases (Frigo et al. 2020b). Also, we propose a novel TFT that integrates structural and functional criteria in the filtering process. In particular, our proposed method extends

* Former PhD student in the team, now at TRIBVN Healthcare, Paris, France

the approaches of the state-of-the-art TFTs by simultaneously fitting the dMRI signal (or some transformation of it) and penalising the magnitude of coefficients associated to streamlines that are not coherent with the co-activation patterns measured with resting-state functional MRI. This novel formulation is tested on both simulated and real data (Frigo et al. 2018a,b) and the results of the corresponding experiments are discussed. Finally, we present *Tractograms As Linear Operators in Neuroimaging* (TALON), a Python package that allows a generalised definition of TFTs which will be publicly released together with this thesis. The work presented in Chapter 4 was done in collaboration with Ragini Verma[†] and Junghoon John Kim[‡]. The analysis of the topological effects of the use of TFTs has been presented at the 24th meeting of the Organization for Human Brain Mapping (Frigo et al. 2019) and published in the Journal of Neural Engineering (Frigo et al. 2020b) together with the data, which are made publicly available in an open repository. Partial results on the functional extension of standard TFT approaches have been presented at the 24th meeting of the Organization for Human Brain Mapping (Frigo et al. 2018a) and at the MICCAI 2018 International Workshop on Computational Diffusion MRI (Frigo et al. 2018b). A journal article is in preparation.

- ▶ *Matteo Frigo*, Samuel Deslauriers-Gauthier, Drew Parker, Abdol Aziz Ould Ismail, Junghoon John Kim, Ragini Verma and Rachid Deriche. Diffusion MRI Tractography Filtering Techniques Change the Topology of Structural Connectomes. *Journal of Neural Engineering* 17.6 (2020): 065002.
- ▶ *Matteo Frigo*, Samuel Deslauriers-Gauthier, Drew Parker, Abdol Aziz Ould Ismail, Junghoon John Kim, Ragini Verma and Rachid Deriche. Effects of tractography filtering on the topology and interpretability of connectomes. OHBM 2019.
- ▶ *Matteo Frigo*, Guillermo Gallardo, Isa Costantini, Alessandro Dadducci, Demian Wassermann, Rachid Deriche and Samuel Deslauriers-Gauthier. Reducing false positive connection in tractograms using joint structure-function filtering. OHBM 2018.
- ▶ *Matteo Frigo*, Isa Costantini, Rachid Deriche, Samuel Deslauriers-Gauthier. Resolving the crossing/kissing fiber ambiguity using Functionally Informed COMMIT. MICCAI 2018 International Workshop on Computational Diffusion MRI.
- ▶ *Matteo Frigo*, Samuel Deslauriers-Gauthier, Rachid Deriche. Diffusion MRI Tractography Filtering with Functional Priors. Paper in preparation.
- ▶ *Matteo Frigo*, Rachid Deriche, Samuel Deslauriers-Gauthier. TALON: Tractograms As Linear Operators in Neuroimaging. Python package. <https://pypi.org/project/cobcom-talon/>

Chapter 5: Brain Alignment and Similarity

Selecting a suitable parcellation for structural connectomic studies is a problem whose solution is most of the times driven by non-quantitative

[†] Diffusion and Connectomics In Precision Healthcare Research Lab, Department of Radiology, University of Pennsylvania, Philadelphia, PA, United States of America

[‡] Department of Molecular, Cellular, and Biomedical Sciences, CUNY School of Medicine, The City College of New York, NY, United States of America

factors such as what the usual choice in the field is or the straightforward applicability of some software for computing the atlas. In this chapter we present a novel and quantitative perspective on this problem based on the assumption that subjects from a homogeneous cohort should exhibit brain networks with highly similar topology. We formulate the problem as the following question: given an atlas, how robustly does it capture the network topology across different subjects? To answer the question, we define two mathematical objects that provide insights on the similarity between the topology of two networks. First, we define 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 design WL-align, a new technique for aligning connectomes obtained by adapting the Weisfeiler-Lehman (WL) graph-isomorphism test. We validated both the GJI and WL-align on data from the Human Connectome Project (HCP) database (Van Essen et al. 2012), inferring a strategy for choosing a suitable parcellation for structural connectivity studies. This work was developed in collaboration with our colleagues from the Combinatorics, Optimization, and Algorithms for Telecommunications (COATI) team at Inria Sophia Antipolis - Méditerranée. An article on this study has been submitted to a scientific journal and the code and data are published in open repositories.

- ▶ *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. In press, *Network Neuroscience*.
- ▶ Emilio Cruciani, *Matteo Frigo*, David Coudert, Rachid Deriche, Emanuele Natale, Samuel Deslauriers-Gauthier. WL-align. Python package. <https://osf.io/depux/>

Appendix A: Software Contributions

In Appendix A we will present the software that has been produced or to which we contributed during the development of the works presented in this thesis.

Appendix B: Other works

In the second appendix we present other works where we have partly contributed with some colleagues, but not at the heart of this thesis.

Chapter 2

Background

Overview In this chapter we are going to introduce the reader to the pillars that constitute the backbone of the present thesis. It is by no means intended to give exhaustive presentations of each included topic, as it aims at introducing the reader to the terminology used in this work and at preparing the presentation of the research topics of the following chapters. After presenting the essential elements of human brain anatomy in Section 2.1, some elements of Magnetic Resonance Imaging (MRI), the diffusion phenomenon and the diffusion MRI (dMRI) techniques are introduced in Section 2.2. Finally, Section 2.3 is devoted to the presentation of the topic of brain connectivity analysis.

2.1 Essential Anatomy of the Human Brain	7
2.2 Brain Imaging	9
Magnetic Resonance Imaging	9
Diffusion MRI	12
2.3 Connectomics	20
Atlases of the Human Brain	21
Structural Connectivity . . .	23
Functional Connectivity . . .	24
Effective Connectivity	25
Topology of Brain Networks	26

2.1 Essential Anatomy of the Human Brain

The nervous system is the part of the human body deputed to the regulation of the analysis and integration of internal and external stimuli. It is subdivided into *central* nervous system (CNS) and *peripheral* nervous system (PNS), where the latter consists of more simple conductors connecting the peripherally situated receptor end effector organs to each other in a biomechanism mediated by the more complex organs of the CNS, i.e. the brain and the spinal chord (Gray et al. 1973). The human brain, which is the object of interest of this thesis, is composed of ~ 100 billion specialised cells called neurons. These are independent entities interacting with each other through electro-chemical processes called synapses. A vast literature on the classification of neurons has built the foundations for the neuroscientific analysis of development, evolution and pathology (Zeng et al. 2017). Neurons consist of a cell body (soma), dendrites, and a single axon as depicted in Figure 2.1. The soma contains




Figure 2.1: Schematic representation of a neuron with its soma, dendrites and axon. The myelin sheath surrounding the axon isolates it from the external bodies, allowing a rapid conduction of the electric impulses.




Figure 2.2: The left panel illustrates the different composition of the GM and the WM (adapted from Fick 2017). The nuclei, the dendrites and the somata are confined in the GM region, while the axonal tracts stem from the GM, traverse the GM-WM interface and develop in the WM. The right panel is adapted from (Kay et al. 2013) and (Liewald et al. 2014) and displays the electron micrograph images of a GM and a WM region. The GM shows less structure with respect to the WM. In particular, the picture of the WM cross section allows to distinguish the different axons having different diameters and non-uniform packing.

the nucleus of the cell, while dendrites and axons are elongated structures that are dedicated to the reception and transport of the electric stimuli respectively. The anatomy of the human brain is highly consistent across healthy adult subjects. This opened the door to the definition of *maps* of the brain that describe specific aspects of its structure and functioning.

The human brain is composed of the cerebrum, the cerebellum and the brainstem. In this thesis we focus on the cerebrum (from now on also *brain* by synecdoche), which is the largest organ of the human brain and is subdivided in two hemispheres that occupy the left and right part of the cranium respectively. These hemispheres are made of Gray Matter (GM) and White Matter (WM). The former is mostly composed of cell bodies and is located in the external boundary of the brain, while the inner space is occupied by the WM, which is made of axons, and the Cerebrospinal Fluid (CSF). The two hemispheres are connected by the Corpus Callosum (CC), which is a region of the WM composed of densely packed axons.

2.1.0.1 Gray Matter

The GM can be subdivided into *cortical* and *subcortical* GM. The cortical GM is folded into gyri and sulci whose shape is highly consistent across subjects and exhibits an approximate symmetry between the two hemispheres. Since the beginning of medicine, anatomists have defined atlases of the cortex that subdivide it into contiguous regions having coherent properties. A gross subdivision of the cortex can be obtained by separating it into lobes as shown in Figure 2.3. More refined atlases of the cortex have been defined following anatomical, cytoarchitectonic, functional or structural criteria or a combination of the aforementioned (Eickhoff et al. 2018).

2.1.0.2 White Matter

The WM is formed of axonal tracts that connect different GM regions. These tracts are subdivided in projection, association and commissural tracts. Projection tracts form two types of connection. The first links




Figure 2.3: The cortex can be subdivided into frontal, occipital, parietal and temporal lobes. Image adapted from Gray et al. 1973.

the brain and the spinal chord, while the second connects the subcortical GM structures to the cerebral cortex. Association tracts define intra-hemispheric connections. If the connection is established between regions in the same lobe, they are called *U*-fibers because of their shape. Commissural tracts are inter-hemispheric connections, hence they are the ones that form the corpus callosum.

2.2 Brain Imaging

Being able to *see* the structure of the brain in-vivo has a strong impact in terms of diagnosis and research. The history of medicine has been profoundly changed by in-vivo imaging, which is a class of imaging techniques that provide information for research, diagnosis and pre-surgical planning with a non-invasive approach. Since the discovery of x-rays in 1895 (Röntgen 1896), several other techniques have been developed. Among these we mention ultrasound imaging and its first medical uses in the early 1940's (Dussik 1942), computed tomography, whose starting point dates back to the works of Pollak 1953 and Hounsfield 1973, and Magnetic Resonance Imaging (MRI), which was officially invented in 1971 by Paul Lauterbur (Lauterbur 1973) but whose foundations trace back to the equations of nuclear magnetization first defined in the works of Bloch 1946. Diffusion MRI (dMRI) is the technique on which most of this thesis is based and it is an evolution of MRI that takes into account the effects of diffusion phenomena in the nuclear magnetization process thanks to the generalization of the Bloch equations given by Torrey 1956. In the past three decades several reviews have been published on the topic of dMRI, covering the image formation process from the physics of nuclear magnetisation through the different aspects of the diffusion phenomenon up to the post-processing of the obtained signal. The interested reader should refer to the handbook of Derek Jones (Jones 2010) for a detailed presentation of the technical aspects of dMRI, and to the review of Hagmann et al. 2006 for an overview on why the dMRI signal plays a fundamental role in the in-vivo imaging of the white matter.

2.2.1 Magnetic Resonance Imaging

Magnetic Resonance Imaging (MRI) is an imaging technique that senses the effects of magnetization in atomic nuclei and forms images that show the anatomy or the physiology of the studied tissues. The theory behind MRI builds on top of the quantum mechanical properties of spins that interact with an external magnetic field. The MRI scanners used in clinical practice (like the one in Figure 2.4) are able to induce magnetic fields of 1.5 to 3 Tesla of magnitude and in research context they can reach a strength of 10 to 15 Tesla, which corresponds to some million times the strength of the terrestrial magnetic field. The static magnetic field generated by an MRI scanner is denoted as $\mathbf{B}_0 = B_0 \cdot \mathbf{v}$ where $B_0[T]$ is the strength of the field and $\mathbf{v} \in S^2$ is the direction of the field. When a body is placed in the scanner and the static field is induced, the magnetic moment of the hydrogen nuclei (1H isotope) aligns with the magnetic field, hence with \mathbf{v} . The spins also precess with frequency equal to the product of the strength of the magnetic field B_0 and the gyromagnetic ratio of the



Figure 2.4: Siemens Healthcare GmbH[®] MAGNETOM Skyra 3T scanner.

spinning particle γ . This product is called Larmor frequency [Hz] and is equal to

$$\omega = B_0 \cdot \gamma \quad (2.1)$$

where the gyromagnetic ratio of the proton is $\gamma = 42.58 \text{ MHz/T}$. The total magnetic field \mathbf{M} to which the sample is subject is the sum of \mathbf{B}_0 and of the magnetic field induced by the spinning protons. This magnetic field can be described as the sum of a longitudinal component \mathbf{M}_z (aligned with \mathbf{v} and with modulus $|\mathbf{M}_z| = M_z$) and a transversal component \mathbf{M}_{xy} (perpendicular to \mathbf{v} and with modulus $|\mathbf{M}_{xy}| = M_{xy}$). When all the spins are in equilibrium, the longitudinal component \mathbf{M}_z is the only one present and its magnitude is directly proportional to the density of protons and to B_0 , obtaining

$$|\mathbf{M}| = |\mathbf{M}_z| \sim [H] \cdot B_0 \quad (2.2)$$

where $[H]$ is the density of protons. No signal is emitted by the nuclei up to now, as they are only subject to the static magnetic field \mathbf{B}_0 . In order to move away from this equilibrium, an ad-hoc radio-frequency (RF) pulse that resonates with irradiation frequency matching the Larmor frequency is applied. In this way, the alignment of \mathbf{M} moves away from the z axis and, after the RF pulse is applied, the total magnetic field \mathbf{M} will be composed only of the transversal component \mathbf{M}_{xy} . The RF pulse is then turned off and the spins will start to realign with the external magnetic field \mathbf{B}_0 , hence with the longitudinal component \mathbf{M}_z . The realignment is complete once both the longitudinal and the transversal components of \mathbf{M} are back to \mathbf{M}_z and zero respectively. This happens at two different times which are called T_1 relaxation time and T_2 relaxation time respectively, with $T_2 < T_1$. These two quantities are characteristic to the studied tissue. After the application of the RF pulse, the longitudinal and transversal magnetization will see their moduli exponentially decaying. In particular, from the Bloch equations (Bloch 1946) one can show that the longitudinal relaxation $\mathbf{M}_z(t)$ has modulus

$$M_z(t) = M_z(eq) - (M_z(eq) - M_z(0)) \cdot e^{-t/T_1}, \quad (2.3)$$

therefore it depends on the longitudinal magnetization at equilibrium $M_z(eq)$ and the longitudinal magnetization right after the RF pulse have been applied $M_z(0)$. At the same time, the relaxation of the transversal magnetization $\mathbf{M}_{xy}(t)$ happens with the law

$$M_{xy}(t) = M_{xy}(0) \cdot e^{-t/T_2} \quad (2.4)$$

where $M_{xy}(0)$ is the transversal magnetization right after the RF pulse has been applied. At time t one can acquire the signal $S(t)$ which is the effect of the magnetization of all the nuclei present in the xy plane and can then be described as

$$S(t) = \int_{\mathbf{y} \in \mathbb{R}^3} \rho(\mathbf{y}) \cdot M_{xy}(\mathbf{y}, t) \cdot e^{-i\omega t} d\mathbf{y} \quad (2.5)$$

where $\rho(\mathbf{y})$ is the proton density at location \mathbf{y} , $M_{xy}(\mathbf{y}, t)$ is the transversal magnetization at time t in location \mathbf{y} and the exponential decay $e^{-i\omega t}$ describes the spins oscillating in the xy plane with Larmor frequency ω . The signal resulting from Equation (2.5) can be expressed in terms of T_1

The Resonance in MRI comes from the RF pulse resonating with the precession of the spinning nuclei.

The T_1 and T_2 relaxation times define how long the longitudinal and the transversal components of the magnetic field take to be restored to the equilibrium position.

and T_2 times of the sample as follows:

$$S \sim [H] \cdot (1 - e^{-TR/T_1}) \cdot e^{-TE/T_2} \quad (2.6)$$

where $[H]$ is the density of protons in the sample, TR is the repetition time of the sequence and TE is the echo time of the acquisition. Equation (2.6) will play a central role in the following sections, as the relationship between the acquisition parameters TR and TE , the proton density (PD) and the T_1 and T_2 times of the gives fundamental insights on the type of contrast that will be observed in the MR image.

The expression given in Equation (2.5) aggregates the contribution of the whole sample in a single signal and is in fact the core of the nuclear magnetic resonance spectroscopy technique, which is undoubtedly interesting for biomolecular application, but when the scale of the studied sample is at the level of brain tissues (as in our case) it is not sufficient. An additional element must be included in the MRI system in order to transform the presented procedure into an imaging technique. The original intuition behind this third element, the first being the static field B_0 and the second being the RF pulse, is due to Lauterbur and Mansfield, who earned the Nobel Prize in Physiology or Medicine “for their discoveries concerning magnetic resonance imaging” in 2003. The idea is to introduce a third controlled magnetic field that localizes the behavior of the nuclei in the z axis, in such a way that an image of a slice of the sample can be generated. This field is called *gradient field* and is aligned with the z axis. Its strength depends on the location $p \in \mathbb{R}^3$ within the sample according to the following law:

$$B_G(t, \mathbf{p}) = [0, 0, \mathbf{K}(t, \mathbf{p}) \cdot \mathbf{p}] \quad (2.7)$$

with $\mathbf{K}(t, \mathbf{p}) = [K_x(t, \mathbf{p}), K_y(t, \mathbf{p}), K_z(t, \mathbf{p})]$ at time t . Each component of $\mathbf{K} \in \mathbb{R}^3$ plays a specific role in the image reconstruction process: the longitudinal component K_z encodes the slice in the xy plane, the K_x component encodes the frequency of the signal and the K_y component encodes the phase of the signal. This allows to sense the signal in a slice of fixed thickness perpendicular to the z axis. The resulting fluctuation of the transversal components M_x and M_y of the total magnetic field is then composed of the static field B_0 , the magnetic field produced by the RF pulse and the gradient field. The total magnetic field induces a current in a coil surrounding the sample. This current could be expressed in terms of differential equations by Faraday’s law of induction, but it is more convenient to encode the two components in a single complex-valued function

$$M^*(t, p) = M_x(t, p) + i \cdot M_y(t, p) \quad (2.8)$$

which allows to treat the current in terms of frequency and phase for the application of a first bank of filters for noise reduction. In practice, this is done directly in the MRI machine. Acquiring the signal with fixed K_z and varying K_x and K_y yields the so-called k -space image of a single slice of the sample, of which an example is presented in Figure 2.5a. The obtained k -space image is the Fourier transform of M^* weighted by the density of nuclei and is defined by the following *imaging* equation:

$$S(t) \sim \int \rho(\mathbf{p}) \cdot M^*(t, \mathbf{p}) \cdot e^{-i\omega t} d\mathbf{p} \quad (2.9)$$

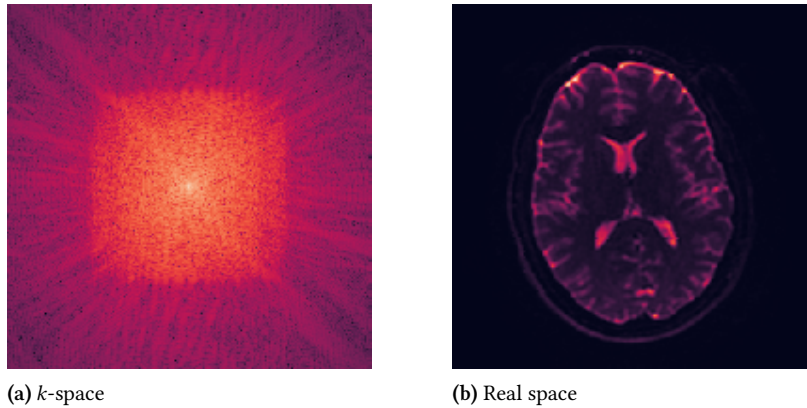


Figure 2.5: These panels show the same axial slice in two different spaces. Figure 2.5a displays the acquired *k*-space version of the image, while Figure 2.5b shows the inverse Fourier transform of the *k*-space image, namely the image in real space. The field of view of the image is 256x256mm with a voxel size of 1mm. Different tissues can be distinguished and localized. Also, potential artifacts and noisy measurements can be spotted and corrected later.

where $\rho : \mathbb{R}^3 \rightarrow \mathbb{R}^+$ is the aforementioned density. An inverse Fourier transform of the *k*-space image $S(t)$ yields the real-space image that can be observed in Figure 2.5b. A three-dimensional volumetric image of the sample can then be obtained by stacking several images acquired by varying the slice encoding parameter $K_z(t, \mathbf{p})$.

Several types of images can be sensed using the procedure described above. The type of tissue contrast imaged by the sequence is mostly determined by the chosen TR and TE . The relationship between these parameters (which depend on the acquisition scheme) and the proton density (PD), the T_1 time and the T_2 time of the studied sample was made explicit a few lines above in Equation (2.6). The proportion between TR and TE determines the so-called T_1 -, T_2 - or PD -weighting of the image. As a rule of thumb, a T_1 -weighted image is obtained with both short TE and TR , a T_2 -weighted image is obtained from both long TE and long TR and a PD -weighted image is obtained from a long TR and a short TE . A common misconception is that X -weighting of the acquisition yields a map of the values of X , be X the proton density, the T_1 or the T_2 of the tissue. This is not true, as Equation (2.6) shows how the contribution of the PD can not be avoided in the image formation and the two terms containing the TR/T_1 ratio and the TE/T_2 ratio can at most approximately be considered irrelevant (i.e. equal to 1 in the multiplication). Nevertheless, the reader should keep in mind the acquisition of T_2 -weighted images is strongly related to the TE .

The procedure presented in this section allows to take static images of the brain tissues, possibly highlighting the insurgence of pathology or anomalies like edema and cancer, or the presence of a contrast agent. These information are specific to each voxel and depend solely on the composition of the tissue in each volume element.

2.2.2 Diffusion MRI

Brain tissues are vastly heterogeneous, since neurons are highly specialized cells that have different shape depending on their function and they are surrounded by several types of non-neuronal cells that allow the neurons to operate as expected. This heterogeneity is reflected in the local anisotropy of the diffusivity of water molecules in the brain, which will tend to be more free to diffuse along directions where they do not encounter barriers like cell membranes or other constraints. As

showed by Einstein 1905 the displacement of water molecules in a free medium due to thermal excitation follows random trajectories that can be formally described by a Brownian motion. The probability that a water molecule will travel at distance $\mathbf{r} \in \mathbb{R}^3$ during the time interval $[0, \tau]$ is described by the ensemble average propagator (EAP) $P(\mathbf{r}, \tau)$, whose probability density function follows the Gaussian law and reads as follows:

$$P(\mathbf{r}, \tau) = \frac{1}{(4\pi D\tau)^{3/2}} e^{-\frac{|\mathbf{r}|^2}{4D\tau}} \quad (2.10)$$

where D is the diffusion coefficient. From Equation (2.10) we can notice that the diameter of the Brownian motion increases when the diffusion time τ or the diffusion coefficient D increase, hence when the particle can move faster or for longer times. As we showed in Figures 2.1 and 2.2, the WM is an agglomerate of strongly anisotropic structures. As a consequence, the diffusion of water molecules within these structures and the corresponding EAP exhibit the same anisotropy. This phenomenon can be captured via diffusion MRI (dMRI), which is an MRI technique that senses the diffusivity along a certain direction. The technique originates from the work of Stejskal and Tanner (Stejskal et al. 1965), who designed the Pulsed Gradient Spin Echo (PGSE) sequence for the acquisition of what was later called the dMRI signal. A schematic depiction of the PGSE protocol is presented in Figure 2.6. Several parameters are




Figure 2.6: Visual representation of the Pulsed Gradient Spin Echo acquisition sequence.

involved in the design of such scheme, namely the echo time $TE[s]$, the strength of the pulsating gradient $G[T/m]$ with direction \mathbf{n} , the pulse length $\delta[s]$ and the separation time $\Delta[s]$ between the two gradients. The obtained signal is weighted by the diffusivity along the direction of G . For this reason, the technique is also called *diffusion-weighted MRI (DW-MRI)* or *diffusion-weighted imaging (DWI)*. The combination of the aforementioned parameters controls how much the diffusivity weights the measured signal. To quantify in a synthetic way the extent to which diffusivity weights the signal, Le Bihan and Breton (Le Bihan et al. 1985) proposed what is now known as b -value, which is measured in s/mm^2 and summarizes the acquisition scheme as follows:

$$b = \gamma^2 G^2 \delta^2 \left(\Delta - \frac{\delta}{3} \right) \quad (2.11)$$

where γ is the gyromagnetic ratio and G is the gradient strength. Equation (2.11) is often reported also as

$$b = 4\pi q^2 \tau \quad (2.12)$$

where $\tau = \Delta - \delta/3$ and $q = \gamma \delta G / 2\pi$. The expression of the b -value given in Equation (2.12) allows to describe the acquisition parameters in terms of q , which is also known as the wave vector. This lexical choice is due to the fact that the signal attenuation can be expressed as the Fourier transform of the EAP:

$$E(q, \tau) = \frac{S(q, \tau)}{S_0} = \int_{\mathbb{R}^3} P(\mathbf{r}, \tau) \cdot e^{2\pi i(\mathbf{q}\mathbf{r})} d\mathbf{r} \quad (2.13)$$

where $P(\mathbf{r}, \tau)$ is the EAP, \mathbf{q} is the vector that has the same direction of the gradient and strength equal to q and S_0 is the signal recorded without diffusion attenuation, which in fact is a T_2 -weighted image. If we multiply both sides of Equation (2.13) by S_0 , we can see that the dMRI signal can be represented as the product of the non-diffusion-weighted component S_0 , which will also be referred to as the *amplitude* of the signal S , and the diffusion-weighted component E , also called *signal shape*. This product of amplitude and shape can be expressed as follows:

$$S(G, \mathbf{n}, \Delta, \delta, TE) = S_0(TE) \cdot E(G, \mathbf{n}, \Delta, \delta) \quad (2.14)$$

where one should notice that the signal amplitude S_0 depends solely on TE, while the signal shape E depends on all the other acquisition parameters, among which the gradient direction \mathbf{n} makes E anisotropic.

2.2.2.1 The Tensor Representation

A well-known representation of the dMRI signal $S(b)$ is the Taylor expansion of its natural logarithm (Novikov et al. 2018), which reads as follows:

$$\ln S(b) = \ln S_0 - D \cdot b + K \cdot b^2 - \dots \quad (2.15)$$

where S_0 is the signal amplitude, D represents the diffusivity and K the kurtosis. When the concept of diffusion *coefficient* is substituted with the concept of *tensor*, which captures the three dimensional anisotropy that we're trying to describe, one obtains the so-called Diffusion Tensor Imaging (DTI) signal representation cutting the Taylor expansion after the first order, while considering all the three terms yields the Diffusion Kurtosis Imaging (DKI) representation (Novikov et al. 2018). Cutting the Taylor expansion after the first order term gives the so-called Diffusion Tensor Imaging (DTI) signal representation, while considering all the three terms yields the Diffusion Kurtosis Imaging (DKI) representation (Novikov et al. 2018). The DTI equation of the signal shape $E(b)$ reads as follows:

$$E(b) = \frac{S(b)}{S_0} = e^{-b\mathbf{G}^T \mathbf{D} \mathbf{G}} \quad (2.16)$$

where \mathbf{G} is the gradient direction and it suggest that higher diffusivity (i.e. higher D) correspond to a lower signal. The particular case where \mathbf{D} is the three-dimensional identity matrix yields the diffusion within a ball, hence a configuration of isotropic diffusivity. This justifies the use of the expression *signal decay* to describe the signal response to the diffusion

phenomenon. Such shape is uniquely determined by a 3D symmetric tensor (the diffusion tensor) whose shape is defined by its eigenvectors \mathbf{v}_i and eigenvalues λ_i and is mathematically expressed as follows:

$$\mathbf{D} = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{bmatrix} = \lambda_1 \mathbf{v}_1 \times \mathbf{v}_1 + \lambda_2 \mathbf{v}_2 \times \mathbf{v}_2 + \lambda_3 \mathbf{v}_3 \times \mathbf{v}_3 \quad (2.17)$$

where \times denotes the outer product between two vectors. The fact that \mathbf{D} is symmetric means that $D_{ij} = D_{ji}$, therefore it is uniquely determined by 6 coefficients in total and it requires exactly six acquired samples to be computed. If more measurements are available, ordinary least squares can be used for increasing the robustness of the estimate.

The eigenspace decomposition of the diffusion tensor allows to define an algebraic framework for the assessment of the geometry of the tensor (O'Donnell et al. 2011). This formalism provides scalar descriptors of the anisotropy and the magnitude of the tensors. The simplest of these measures is the mean diffusivity (MD), which is defined as the arithmetic mean of the three eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$:

$$MD = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3} \quad (2.18)$$

and provides a description of the magnitude of the signal represented by the tensor. Also, one can define the fractional anisotropy (FA), which measures of how far the tensor is from being spherical.

$$FA = \sqrt{\frac{1}{2} \frac{(\lambda_1 - \lambda_2)^2 + (\lambda_1 - \lambda_3)^2 + (\lambda_2 - \lambda_3)^2}{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}} \quad (2.19)$$

Voxels with high FA (e.g., $FA > 0.95$) are associated to strongly anisotropic WM configurations, which usually correspond to uniformly parallel fibers. For this reason, FA has been mistaken for a measure of white matter integrity. The tensors showed in Figure 2.7 highlight how this misconception is proved wrong in the presence of crossing fibers. These are legitimate WM configurations without integrity problems that are only guilty of not allowing a meaningful tensor representation of their corresponding diffusion signal. More sophisticated scalar measures have been




Figure 2.7: Four different fiber configurations are displayed in the first row of plots. The corresponding tensors obtained with the DTI model are shown immediately below. Fiber populations with a single anisotropic component can be properly described with a tensor. On the contrary, the DTI representation of crossing fibers is not accurate. The FA of the two configurations on the left is much higher than the one of the two configurations on the right.

defined for describing the planarity, sphericity and other properties of the tensor. The interested reader should refer to Batchelor et al. 2005, Ar-

signy et al. 2006, Westin et al. 1997 and O'Donnell et al. 2011 for further details.

The DTI and DKI representations are appropriate only in certain ranges of b -values, notably in the range that ensure the higher order terms to be in the little- o of the lower order terms. In the case of DTI, this range can be determined by looking at when $K \cdot b^2 \ll D \cdot b$, which is typically observed when $b < 1000s/mm^2$. In absolute terms, these are relatively low b -values, which are known to yield images with superior signal-to-noise ratio (SNR) (Xie et al. 2015). To obtain an appropriate description of the signal with the DKI representation in the WM, higher b -values in the $2500 - 3000s/mm^2$ range are necessary (Chuhutin et al. 2017).

DTI has been extensively used in clinical applications (Tae et al. 2018) thanks to its simple formalism and its relatively short acquisition time. However, its capability to describe complex white matter configurations has been proved to be limited, as shown in Figure 2.7.

2.2.2.2 High Angular Resolution Diffusion Imaging

Describing complex fiber architectures like the kissing, fanning and crossing configurations requires the use of the so-called *high angular resolution diffusion imaging* (HARDI) techniques (Descoteaux et al. 2008; Yo et al. 2009). Methods in this class are able to sense the complexity of the complex WM configurations by using diffusion-weighted images acquired with several different gradient direction (and possibly different intensity).

The gradient directions are described in terms of q -space, where q is the wave vector that defines the signal attenuation as the Fourier transform of the diffusion propagator. Techniques such as Q-ball imaging (Descoteaux et al. 2007), Constrained Spherical Deconvolution (CSD) (Tournier et al. 2007) and Diffusion Spectrum Imaging (DSI) (Wedeen et al. 2008) rely on such acquisitions. In particular, if the q -points are sampled on the surface of one or multiple isocentric spheres, the acquisition is called single- or multi-shell. These shells are determined by the employed b -value and are called b -shells. Conversely, when the q -space is discretised on a Cartesian grid it yields the DSI scenario. The main advantage of using shell-based protocols with respect to DSI is the lower number of samples required to achieve the wanted angular resolution (Descoteaux et al. 2011). In this thesis we are going to consider only single-shell and




Figure 2.8: Three different HARDI schemes are displayed. The DSI scheme is made of q -space points sampled on a three-dimensional Cartesian grid. Conversely, the single shell and multi shell schemes have q -space points uniformly distributed on spheres, which in this context are called b -shells in analogy with the b -value used for describing the ensemble of the acquisition parameters.

multi-shell schemes. Among these, single-shell acquisitions are usually faster than multi-shell acquisitions.

2.2.2.3 The Spherical Harmonics Representation

In a shell-based acquisition, the signal can be represented in spherical coordinates as a function of the direction of the gradient. These functions can be expressed via spherical-harmonics (SH), which are the S^2 equivalent of the Fourier basis in \mathbb{R}^n . The SH functions are denoted as $Y_\ell^m(\theta, \phi)$, where ℓ denotes the order, m the phase, $\theta \in [0, \pi]$ and $\phi \in [-\pi, \pi]$, where Y is π -periodic in θ and 2π -periodic in ϕ . Their definition relies on the notion of Legendre polynomial P_ℓ^m of degree ℓ and order m and reads as follows:

$$Y_\ell^m(\theta, \phi) = \sqrt{\frac{2\ell + 1}{4\pi} \frac{(\ell - m)!}{(\ell + m)!}} P_\ell^m(\cos \theta) e^{im\phi}. \quad (2.20)$$

Any square-integrable function $f : S^2 \rightarrow \mathbb{C}$ can be represented as a linear combination of spherical harmonics as

$$f(\theta, \phi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} c_{\ell m} Y_\ell^m(\theta, \phi) \quad (2.21)$$

where each $c_{\ell m}$ is a SH coefficient. In the context of brain imaging via dMRI, SHs are typically used for representing the orientation distribution function (ODF) of diffusion (dODF) or fibers (fODF) within a voxel. The dODF describes the amount of diffusion in a certain direction as a function on the sphere (i.e. the EAP), while fODFs aim at representing the dispersion and the multiplicity of orientation of the fibers in the considered voxel. The two concepts can be seen as a model-free representation of the dMRI signal (the dODF) and a model-based representation of the underlying white matter structure (the fODF). The interested reader should refer to (Descoteaux et al. 2008) and (Dell'Acqua et al. 2019) for a detailed comparison of the two. Diffusion ODFs and fODFs share two mathematical properties that allow to significantly simplify their SH representation.

- ▶ The orientation distributions formally defined as dODFs and fODFs are real functions.
- ▶ The data they describe are symmetric with respect to the origin, i.e., $f(\theta, \phi) = f(\pi - \theta, \pi + \phi)$. The translations of θ and ϕ are to be intended in the respective periodic domains.

The SH basis of the space of square-integrable functions $\{f : S^2 \rightarrow \mathbb{R}\}$ that satisfies the two aforementioned properties reads as follows:

$$Y_{\ell m}(\theta, \phi) = \begin{cases} 0 & \text{if } \ell \text{ is odd,} \\ \sqrt{2} \Im [Y_\ell^{-m}(\theta, \phi)] & \text{if } m < 0, \\ Y_\ell^0(\theta, \phi) & \text{if } m = 0, \\ \sqrt{2} \Re [Y_\ell^m(\theta, \phi)] & \text{if } m > 0, \end{cases} \quad (2.22)$$

where $Y_{\ell m}$ are the elements of the basis, Y_ℓ^m are the SH functions and $\sqrt{2}$ is a normalisation factor that makes the basis orthonormal. This basis is the one employed by default in the main software packages dedicated to dMRI data processing, i.e., Dipy (Garyfallidis et al. 2014) and




Figure 2.9: The two upper panels show the diffusion tensors (left) and fODFs (right) fitted on a healthy subject of the HCP database. The fODFs are obtained with the multi-shell multi-tissue constrained spherical deconvolution technique of Jeurissen et al. (Jeurissen et al. 2014). The image shows how the SH representation of fODFs gives a much more detailed description of the fiber bundles traversing a voxel with respect to DTI. In particular, fODFs contain information about *multiple* orientations of the fiber bundles, while DTI gives an estimate of the *average* orientation which might not correspond to any of the directions of the peaks of fODFs. The lower panels show the diffusion tensor and the fODF corresponding to the circled voxel in the upper panels.

Mrtrix3 (Tournier et al. 2019). For historical reasons, the dMRI community refers to this one as the Descoteaux et al. 2007 basis, while the same basis without the $\sqrt{2}$ factor is known as the Tournier et al. 2007 basis.

An example of what information is carried by fODFs is given in Figure 2.9, where fODFs are shown to yield a high angular resolution description of the local direction of white matter fibers. Figure 2.9 also shows how the angular resolution of the fODFs allow to achieve the detail level that Figure 2.7 was missing in the DTI representation.

2.2.2.4 Tractography

The information about the fiber orientation carried by both diffusion tensors and fODFs can be leveraged in the task of tracking the trajectories of the WM pathways. These can be obtained with the so-called tractography algorithms, which yield a set of streamlines collected in a tractogram. The rationale behind such techniques is that axonal pathways are tangent to the main direction described by the local diffusivity pattern. Designing tractography algorithms yielding realistic subject-tailored tractograms has been one of the main technical challenges tackled by the dMRI research community in the last twenty years. As a result, a wide variety of tractography techniques have been developed and systematically compared in international challenges (Maier-Hein et al. 2017; Schilling et al. 2019a). These techniques can be subdivided into *local* and *global* tractography algorithms (Anastasopoulos et al. 2014). The former determine the trajectory of streamlines by propagating a line from a seed following the orientations estimated with some local model (Basser et al. 2000; Descoteaux et al. 2008; Girard et al. 2014, 2017; Jones 2008; Mori et al. 1999; Tournier et al. 2012), while the latter retrieve the whole tractogram by solving a global problem aiming at fitting the measured dMRI data (Christiaens et al. 2015; Konopleva et al. 2018; Kreher et al. 2008; Reisert et al. 2011; Schreiber et al. 2014). The differences between these two approaches have recently been reviewed by Jeurissen et al. 2019. In this thesis we will consider only local tracking algorithms.

Streamlines are parametrised curves $\mathbf{r}(s) \in \mathbb{R}^3$ obtained by integrating the field $\mathbf{v}(x)$ defined by the local directions yielded by the employed local model, be it a tensor or an fODF. If the local directions are uniquely

determined, the tractography algorithm will be called deterministic. Conversely, if the local direction is sampled each time from a probability distribution like a normalised ODF or a diffusion tensor, the algorithm is called probabilistic (Sarwar et al. 2019).

The purpose of tractography is two-fold:

1. Describe the anatomy of the white matter fascicles (Johansen-Berg et al. 2006; Francois Rheault et al. 2020b; Sarubbo et al. 2019; Schilling et al. 2020a,b; Schotten et al. 2011).
2. Determine how two gray matter regions are connected (Maier-Hein et al. 2017; Sotiropoulos et al. 2019; C.-H. Yeh et al. 2019).

In this work we will focus only on tractography techniques that are designed with the second purpose in mind, hence for the description of the structural connectivity between brain regions.

Several works have been published providing a comprehensive view of the advances made in the last twenty years with respect to each step of the tractography pipeline (Jbabdi et al. 2011; Jeurissen et al. 2019; Jones et al. 2013; Maier-Hein et al. 2017; Francois Rheault et al. 2020a; Schilling et al. 2019b). In spite of the notable improvements granted by these advances, tractography remains an ill-posed problem. The specific terms in which this ill-posedness is declined cover all the aspects of the design of tractography algorithms, which can be summarised as:

- ▶ Estimation of the local geometry of the fiber orientation.
- ▶ Integration scheme used for propagating the trajectory of the streamline.
- ▶ Interpolation of the discrete grid where the local orientations are defined.
- ▶ Placement of the seeds from which the streamlines are propagated.
- ▶ Assessment of the appropriate termination and acceptance criteria for the drawn streamlines.

The tractography pipeline that will be used in this thesis can be completely obtained with the Mrtrix3 (Tournier et al. 2019) suite and, despite the lack of consensus in the design of tractography experiments, aims at including any technical aid and biological prior that may improve the reliability of the obtained tractograms while balancing the applicability of the pipeline in the context of large cohort studies. The pipeline reads as follows:

- ▶ The fODFs describing the local model of the white matter are obtained with the multi-shell multi-tissue constrained spherical deconvolution (MSMT-CSD) pipeline of Jeurissen et al. 2014 when multi-shell data are available, or with the heuristic algorithm of Dhollander et al. 2016a coupled with the CSD algorithm (Tournier et al. 2007) when single-shell data are considered.
- ▶ A Runge-Kutta method of order 2 is used for the integration of the fODFs (Jeurissen et al. 2019).
- ▶ Trilinear interpolation of the fODFs is employed.
- ▶ Streamlines are seeded from the gray-matter white-matter interface (GMWMI) (Girard et al. 2014; R. E. Smith et al. 2012).
- ▶ Anatomically Constrained Tractography (ACT) (R. E. Smith et al. 2012) is employed, therefore streamlines have endpoints in the GMWMI.

This pipeline is still exposed to limitations posed by the spatial resolution of the acquired data, the complexity of the fiber configurations through which the algorithm is asked to track and the variety of tissues that populate the brain (Francois Rheault et al. 2020a). Moreover, it does not include a strategy for the quantitative evaluation of the tracked streamlines, which are by definition curves in a three-dimensional space, hence they do not carry any information about the shape or the volume of the occupied space. This issue has been highlighted in several articles published in the last ten years (Jbabdi et al. 2011; Jeurissen et al. 2019; Jones et al. 2013; Maier-Hein et al. 2017; Francois Rheault et al. 2020a; Schilling et al. 2019b), bringing to the conclusion that tractography per-se is not quantitative. A significant part of this thesis (Chapter 4) is dedicated to methods that aim at re-establishing this quantitiveness. In particular, we will focus on the so-called Tractography Filtering Techniques (TFTs), which are designed to process an input tractogram and return a coefficient for each streamline that evaluates how much the estimated track is necessary to explain the dMRI signal or some transformation of it.

2.3 Connectomics

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. Also, the graph-theoretical analysis of connectomes unveiled some previously unknown features of both anatomy and pathology of the human brain (Bassett et al. 2017; Verma et al. 2019). As a consequence, brain networks have been proved to be fundamental objects in modern neuroscience. They allow to look at the ensemble of all the neurons as a single integrated entity that shapes the behavioural and functional aspects of the brain. In this sense, the neural network has a structure that shapes its functioning (Honey et al. 2007). To define such networks, the concept of *connectome* was formulated independently by Hagmann 2005 and O. Sporns et al. 2005. Connectomes are graph-like representations of the human brain where nodes represent brain regions and edges encode the existence and possibly the strength of the connection between two nodes. The aforementioned structure-function duality is present also in the definition of connectomes, which have been defined both in the context of the structure and the functioning of the brain. Functional connectomes encode the co-activation patterns of the activity in cortical regions. This is measured with some form of correlation analysis between the time series of the activation of neurons that can be obtained from resting-state functional MRI (fMRI) or from magneto/electro encephalography (M/EEG). These connectomes are usually fully connected, meaning that they encode information about the correlation between the activity in *every* pair of regions. On the contrary, structural connectomes describe the architecture of the axonal connections and they are sparse. The information they encode is representative of the strength of the connection established by the axons that connect the two regions represented by the nodes and are estimated with dMRI-based tractography.

From the mathematical point of view, a connectome is 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 matrix that encodes in position (i, j) the weight of the edge w_{ij} between nodes i and j is called *adjacency* or *connectivity* matrix of G and is denoted as $\text{Adj}(G)$. Edges that are not present in the graph are associated to a null weight.

The following sections are devoted to the nodes and the edges of connectomes. Section 2.3.1 contains the presentation of the segmentation of the brain cortex that define the nodes of the connectomes. Finally, in Sections 2.3.2 and 2.3.3 we will define *structural* and *functional* connectomes respectively.

2.3.1 Atlases of the Human Brain

Ideally, one would like to be able to describe how each neuron is connected with the others. This is perfectly doable for simple organism like the *C. Elegans* (Cook et al. 2019), whose whole nervous system (neurons, muscles and on-muscle end organs) contains less than 600 nodes. The feasibility in the case of the human brain is challenged by its immense complexity, which is well represented by the approximate number of neurons in a healthy brain, being it in the scale of the hundreds of billions (10^{11}) (Herculano-Houzel 2009). This is beyond the current computational capability. A simplifying assumption that preserves the ability to represent the whole brain which has been widely used in neuroscience, is that the brain, and in particular the cortical surface, can be divided into distinct and homogeneous areas. Neurons that show homogeneous characteristics are grouped into *parcels* which are in turn collected in *brain parcellations*. A variety of brain atlases has been proposed in the past years and each of them subdivides the cortex based on assumptions of anatomical, cytoarchitectonic, functional and structural criteria. Some atlases cover the whole cortex, while some other provide parcellations of specific cortical region*.

Anatomical parcellations characterise each parcel by its localisation or shape in the brain. Thanks to the high inter-subject similarity between the shape of different subjects, it is possible to define a certain region as “the k -th gyrus from the n -th sulcus”. Notable examples of anatomical atlas are the AAL atlas (Tzourio-Mazoyer et al. 2002), the Desikan atlas (Desikan et al. 2006), the parcellation of Destrieux et al. 2010 and the MarsAtlas (Auzias et al. 2016). Each of the mentioned parcellations subdivides the cortex in a different number of regions, but in general this number is lower than 100, as the delineation of the boundaries of each region is performed manually. In this thesis we’re going to make use of the *Desikan* atlas, which is based on the manual segmentation of a template of the brain cortex that takes into account the morphological consistencies of healthy human brains. This parcellation is composed of 68 regions and is one of the most appreciated in brain imaging thanks to the high accuracy when defined on new subjects. Its popularity has also been established thanks to the accessibility granted by Freesurfer (Fischl 2012), a free software that (among other things) allows to easily project the Desikan atlas onto an arbitrary brain.

* The interested reader should refer to Gallardo Diez 2018 for a detailed presentation of the topic.

Cytoarchitectonic parcellations subdivide the cortex based on its cellular composition, following criteria like the thickness of the cortex or the presence of a certain type of neuron. Atlases of this type have been introduced starting from the early studies of Campbell 1905 and Brodmann 1909 through the atlas of von Economo and Koskinas (Economo et al. 1925) and the more recent works of Eickhoff et al. 2008, Ding et al. 2016 and Rosenke et al. 2018.

Functional atlases are designed to separate regions involved in different functions of the brain. In modern times, these atlases are defined using imaging techniques such as functional MRI (fMRI) which are able to characterise the cortical activity while specific tasks are being performed. Examples of functional atlases are the parcellation of B. T. Yeo et al. 2011, the atlas of Craddock et al. 2012, the ones of Deen et al. 2011 and Huth et al. 2016 and the parcellation of Schaefer et al. 2017. The latter is among the ones used in this thesis and can be constructed with an arbitrary number of parcels. It is based on the analysis of the co-activation patterns of the brain by means of the analysis of resting-state functional connectivity. The *Schaefer* parcellation technique condensates the two major approaches to cortical parcellations based on resting-state functional MRI, namely the local gradient and the global similarity approaches¹.

Structural parcellations are based on the axonal connectivity patterns exhibited by different cortical regions. The use of in-vivo techniques such as dMRI-based tractography for tracking the axonal pathways allowed to define several of these atlases. Structural atlases of specific regions have been obtained by Anwander et al. 2007, Bajada et al. 2017 and Thiebaut de Schotten et al. 2017, while Moreno-Dominguez et al. 2014, Fan et al. 2016 and Gallardo et al. 2018a proposed structural parcellations of the whole cortex. The *Gallardo* parcellation is based on the segmentation of the structural connectivity profiles associated to each point of the cortical surface. The segmentation is obtained by defining a local coherence criterion that defines areas that show similar connectivity fingerprints within the region against a known reference. As for the Schaefer parcellation, the atlas of Gallardo has the advantage of not being limited to a pre-defined number of parcels, which can be set arbitrarily. This parameter controls the granularity of the parcellation, hence the average extension of the parcels.

A modern parcellation paradigm involves the concurrent use of multiple methodologies, obtaining atlases that are coherent with more than one of the aforementioned criteria. Examples of multi-modal atlas are the one of Schubotz et al. 2010, Diez et al. 2015, Sarah Parisot et al. 2017 and Matthew F. Glasser et al. 2016. In Chapter 5 we will make use of the *Glasser* parcellation, which was designed delineating 360 cortical areas “*bounded by sharp changes in cortical architecture, function, connectivity, and/or topography*” (Matthew F. Glasser et al. 2016).

The analysis of parcellation-based connectomes has been shown to yield high qualitative reproducibility across parcellations, but the quantification of this reproducibility may be strongly influenced by the choice of the parcellation (Reus et al. 2013). Some attempts at defining a parcellation-free connectome have been presented in the last years (Gutman et al. 2014; Moyer et al. 2016). These approaches are based on the relaxation

1: See (Schaefer et al. 2017) for a detailed discussion of these two approaches and the methodological aspects of their unification.

of the concept of connectivity matrix, which is transformed into a symmetric real function $K : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{R}$ defined on the product space of the cortex surface \mathcal{C} with itself. The cortex \mathcal{C} is here treated as a continuous surface embedded in \mathbb{R}^3 and each point of the surface is potentially a “node” of the graph. The obtained network is also called *continuous connectome* in contrast with the *discrete connectome* defined on a finite set of nodes. The work of Moyer et al. 2016 showed how every parcellation can be reduced to a special case of this continuous representation. Notice that the “*continuous*” label has sometimes been inappropriately attached to connectomes that are defined using as nodes the vertices of a high resolution three-dimensional mesh that mimics the shape of the cortex. These meshes can be obtained for instance with Freesurfer (Fischl 2012) and are usually defined on a set of thousands of points (a frequently used representation involves $\sim 32k$ nodes). Even if the high spatial resolution of the mesh suggests that the representation of the network is close to the one of a continuous graph, the connectome is still parcellation-based and the cortical atlas is given by the Voronoi diagram of the mesh (Kang 2008).

2.3.2 Structural Connectivity

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 in Section 2.2.2, 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. Streamlines are then associated to parcels in a cortical atlas where they terminate in order to gain insights on the connectivity between the two regions. 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 in an atlas, then define

$$c_{ij} = \sum_{s \in \mathcal{S}(i,j)} w_s \quad (2.23)$$

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 through its connectivity matrix $C = \{c_{ij}\}$. Since tractography is not able to distinguish afferent and efferent connections Jbabdi et al. 2011, matrix C will be symmetric. 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 normalised by dividing each entry of C by the sum of all the entries of the same matrix. The normalisation is not strictly necessary, but it mitigates the effects of comparing networks obtained from tractograms containing different number of streamlines. 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

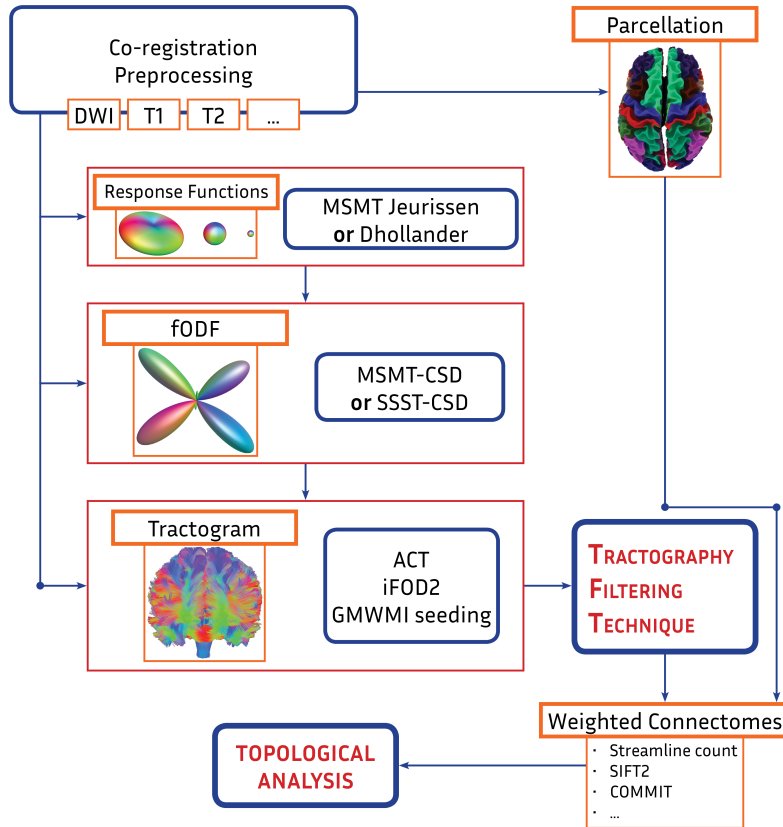


Figure 2.10: The diagram gives a schematic representation of the general connectomic pipeline employed in this thesis. From co-registered and pre-processed data (dMRI, T_1 , T_2 , ...) we obtain a cortical parcellation for defining the nodes of the connectome. The dMRI data are then used to obtain the fODFs that are integrated with a second order method (iFOD2) seeding from the GMWMI and following the ACT paradigm. The resulting streamlines can be filtered before defining a weighted connectome. The network can then undergo a topological analysis.

connectivity has been shown to be limited by the complexity of the white matter configurations and the fallacies of tracking algorithms (Girard et al. 2014; Jeurissen et al. 2019; Jones et al. 2013; Maier-Hein et al. 2017) and by the way in which each streamline is assigned to a certain edge in the connectome (C.-H. Yeh et al. 2019). 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 (TFTs). This topic will be thoroughly analysed in Chapter 4.

The connectomic pipeline employed in this thesis uses the tractography pipeline summarised in Section 2.2.2.4. The obtained streamlines are then assigned to a node by performing a voxel-based local search around the endpoints with $2mm$ radius (C.-H. Yeh et al. 2019). If a TFT is employed, the weights associated to each streamline connecting two regions are *summed* and encoded in the corresponding edge weight (Frigo et al. 2020b). The whole pipeline is summarised in Figure 2.10.

2.3.3 Functional Connectivity

The estimation of functional connectivity can be obtained by processing the data provided by functional MRI (fMRI) or magneto/electro encephalography (M/EEG). These modalities measure distinct effects of the brain activity at different time scales. In particular, fMRI records time series with frequency $\sim 1Hz$, while M/EEG time series are sampled with a frequency of $\sim 1000Hz$. In the absence of task, the recorded time series is said to be in *resting-state* regime. The only type of functional information that will be considered in this thesis is the resting-state functional

connectivity obtained from fMRI. In this technique, the subject is not conducting any task and a time series is recorded for each voxel in the cortex.

Once the time series are recorded, a cortical atlas is used for clustering them and extracting a single pre-processed time series per region. These time series can then be processed to get a static or a dynamic description of the synchronisation patterns of distinct regions of the brain (Preti et al. 2017; Van Den Heuvel et al. 2010). The literature on the methods for extracting measures of the functional synchronisation between the activity in different regions is vast and the reader should refer to S. M. Smith et al. 2013 for a review on the topic. The concept of *functional connectivity between two regions* used throughout this thesis relies on the correlation-based approach that allows to quantify the synchronisation between two regions as the correlation between the time series associated to each region. The resulting connectome is a complete weighted graph whose edge weights take values in the $[-1, 1]$ range.

Static functional connectomes are networks that summarise the dynamics of the functional activity in a single static graph that is build looking at the correlation between the whole time series recorded in each region, which can be as long as 10 minutes or more. Despite representing a dynamic phenomenon through a static mathematical object, static functional connectomes have been proven extremely useful in the exploration of the functioning of the brain (S. M. Smith et al. 2013). A less explored paradigm is the one of the dynamic functional connectomics, which restores the time-dependence of the functional phenomenon by considering a dynamic graph, i.e. a network whose edge weights are a function of time. This formulation allows to capture temporary co-activation patterns that in the static paradigm would be more difficult to appreciate as they would get absorbed in the macroscopic comparison between the time series.

2.3.4 Effective Connectivity

Describing the brain network as a structure-function dichotomy is an efficient representation exercise that gives a realistic picture of what is the architecture (i.e., structural connectivity) that supports the co-activation of neurons in different regions (i.e., functional connectivity). Structural connectivity describes the existence of a WM structure that connects two GM regions, while functional connectivity tells us the extent at which the activity in two GM regions is coherent. What is missing from the picture is temporal causality. The study of *effective connectivity* comes in aid answering the question: what is the causal relation that generates the influence that the activity in one neural population has over the activity in another one? In this sense, the knowledge carried by effective connectivity is the missing link between structural connectivity and functional connectivity. Effective connectivity can be described as a directed network, i.e., with a symmetric adjacency matrix, estimated from fMRI (Stephan et al. 2010; Zarghami et al. 2020), M/EEG (Deslauriers-Gauthier et al. 2019) or in general from any functional imaging technique (Horwitz 2003). Observing this causality is a complex task that finds the most solid mathematical tools to be investigated with in Granger causality (Seth et al. 2015) and dynamical systems theory within the broader framework of

Bayesian modelling (Deslauriers-Gauthier et al. 2019; Friston 2011; Zarghami et al. 2020; Zou et al. 2009).

2.3.5 Topology of Brain Networks

Topology is the discipline of mathematics that studies the properties that are invariant under homeomorphism. In the context of the graph-theoretical analysis of brain networks, this corresponds to studying those properties that characterise connectomes by abstracting the neuroscientific meaning attached to edges and nodes.

The exploration of the topology of a network usually starts from the assessment of the presence of two properties: segregation and integration. The former highlights the presence of densely connected local clusters of nodes and the latter indicates the presence of few connections that link these clusters and make the long-distance communication very efficient. For the point of view of brain networks, this corresponds to having both highly localised and highly connected clusters of neurons (Friston 2011). These two properties together give a global picture of the organisation of the brain network. Highly integrated networks allow efficient communication between cortical regions, while highly segregated networks come at low wiring cost (Bullmore et al. 2009; Tononi et al. 1994; Watts et al. 1998). The simultaneous presence of high segregation and high integration has been referred to as *small-worldness* (Watts et al. 1998), a term which was first introduced to explain the small world behaviour of social relations Milgram 1967, where few acquaintance steps are sufficient to connect two people from very distant (both geographically and socially) regions of the world. The tendency of the brain network to exhibit small-worldness has been assessed by several studies in the last 20 years Bassett et al. 2006, 2017; Olaf Sporns et al. 2004; S. Yu et al. 2008.

Small-worldness is usually quantified as the ratio between two global measures of the integration and the segregation of the network respectively, but, as highlighted by Papo et al. 2016, this approach poses several limitations in the context of brain connectivity. An alternative approach is to independently assess the integration and segregation properties of the network may be preferable. A detailed review of the variety of measures of integration and segregation that have been proposed in the past years can be found in the paper of Rubinov et al. 2010, where the authors also the *Brain Connectivity Toolbox*, which is the reference software used in this work for the evaluation of these graph-theoretical measures. In the following paragraphs we are going to give a formal description of the integration and segregation measures that are used in this thesis.

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 Dijkstra 1959) between two nodes on connectome C . The CPL of C is

$$\ell = \frac{1}{n} \sum_{i \neq j} d(i, j) \quad (2.24)$$

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 (2.24) 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 (Latora et al. 2001). Its formulation reads as follows

$$e = \frac{1}{n} \sum_{i \neq j} \frac{1}{d(i, j)}. \quad (2.25)$$

To ensure the well-posedness of equation (2.25) 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 (2.25) 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) B. Zhang et al. 2005. 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

$$CC(i) = \frac{\sum_{i \neq j} \sum_{k \neq i \neq j} c_{ij} \cdot c_{jk} \cdot c_{ki}}{\left(\sum_{i \neq j} c_{ij}\right)^2 - \sum_{i \neq j} c_{ij}^2} \quad (2.26)$$

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

$$CC = \frac{1}{n} \sum_{i=1}^n CC(i) \quad (2.27)$$

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 (Opsahl et al. 2009). A different point of view on segregation is offered by the concept of *modularity* (MO) (Newman 2004). 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

$$Q = \frac{1}{2m} \sum_{i,j} \left(C_{i,j} - \frac{k_i \cdot k_j}{2m} \right) \delta(\gamma_i, \gamma_j) \quad (2.28)$$

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 2.11 with an example on two simple graphs. The first graph




Figure 2.11: 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.

is the junction 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.

Chapter 3

Multi-Tissue Multi-Compartment models of diffusion MRI

Overview State-of-the-art multi-compartment microstructure models of diffusion MRI in the human brain have limited capability to model multiple tissues at the same time. In particular, the available techniques that allow this multi-tissue modelling are based on multi-TE acquisitions. In this chapter we propose a novel multi-tissue formulation of classical multi-compartment models that relies on more common single-TE acquisitions and can therefore be employed in the analysis of previously acquired datasets. We show how modelling multiple tissues provides a new interpretation of the concepts of signal fraction and volume fraction in the context of multi-compartment modelling. The software that allows to inspect single-TE diffusion MRI data with multi-tissue multi-compartment models is included in the publicly available Dmipy Python package.

3.1 Introduction

Diffusion MRI (dMRI) is an imaging technique that allows to inspect the brain tissue microstructure in-vivo non-invasively. One of the most commonly studied microstructural feature is the *volume fraction* of a tissue in a sample. In particular, the intra-axonal (a.k.a. intra-cellular - IC), extra-axonal (a.k.a. extra-cellular - EC) and isotropic (ISO) or cerebrospinal fluid (CSF) volume fractions have been investigated in the past literature with several models. Among all, we mention the neurite orientation dispersion and density imaging (NODDI) (H. Zhang et al. 2012) and NODDI-X (Farooq et al. 2016), ActiveAx (Alexander et al. 2010), the multi-compartment microscopic diffusion imaging framework (Kaden et al. 2016), the CHARMED model (Assaf et al. 2005), the intravoxel incoherent motion model (Le Bihan et al. 1988), the Stanisz model (Stanisz et al. 1997), the AxCaliber model (Assaf et al. 2008), the ball and stick model (T. E. Behrens et al. 2003), the Bingham-NODDI model (Tariq et al. 2016), FERNET (Parker et al. 2020), CODIVIDE (Lampinen et al. 2017), COMMIT (Daducci et al. 2014), VERDICT (Panagiotaki et al. 2014) and the DIAMOND model (Scherrer et al. 2016). The differences between these models lie on the used signal representation for specific tissues (e.g. the intra-axonal diffusion can be modelled as the diffusion within a stick, a cylinder, ...) or specific assumptions on the model parameters motivated by histology or limitations of the used technique (e.g. fixed

3.1 Introduction	29
3.2 Theory	31
Multi-Compartment models	31
The Standard Model of White Matter	32
MC models do not account for T2 differences	34
Multi-TE MC models of dMRI signal	35
Multi-Tissue MC models . .	37
Generalised MT-MC Mod- elling	40
The MT Standard Model of dMRI in White Matter . . .	41
3.3 Experiments	41
Synthetic phantom	42
Real Data	44
3.4 Discussion	46
3.5 Conclusion	48

diffusivity or tortuosity of parameters). A unifying aspect that characterizes most of the brain microstructure models is the *building-blocks* concept behind their formalisation. In other words, models are defined in a multi-compartment (MC) fashion, where the dMRI signal is described as a linear combination of single-tissue models. The resulting models are called MC models and they rely on multi-shell acquisitions of the dMRI signal (Scherrer et al. 2010). Thorough reviews have been dedicated to the design and validation of such models (Jelescu et al. 2017), to the sensitivity of MC models to experimental factors and microstructural properties of the described tissues (Afzali et al. 2020) and to the abstraction of these models that allows to obtain a unified theory (Fick et al. 2019).

Recent studies highlighted how all the available MC models of dMRI are defined in such a way that they are transparent to differences between the T_2 times of the modelled tissues (Alexander et al. 2019; Lampinen et al. 2019; Veraart et al. 2018). As a consequence, they implicitly assume that all the considered tissues have the same S_0 response. While this is a reasonable assumption in some contexts, it can not be considered true in general. As a consequence, tissue fractions obtained with this assumptions are called *signal* fractions, in contrast with the *volume* fractions that are obtained with models that account for different S_0 responses of the modelled tissues. The first measures the linear relation between the signal generated by a single tissue compartment and the acquired signal, while the second measures the volume of single tissue compartment that is present in the voxel. In this sense, signal fractions will be shown to be biased estimators of the volume fractions. Given the known interdependence between the T_2 times of tissues and the TE of the acquisition, some attempts at addressing this issue have been formulated making use of multi-TE multi-shell dMRI acquisitions (Gong et al. 2020; Lampinen et al. 2020; Veraart et al. 2018). Despite allowing to increase the signal-to-noise ratio (SNR) (Eichner et al. 2020), these techniques require a complete re-design of the experiments from acquisition to post-processing, posing severe limitations in terms of usability of already acquired data. This aspect is crucial in modern neuroimaging, where large studies like the Human Connectome Project (HCP) (Van Essen et al. 2012), the UK Biobank (Sudlow et al. 2015) and the Alzheimer Disease Neuroimaging Initiative (ADNI) (Mueller et al. 2005) invest significant amounts of time and financial resources to acquire data of large cohorts with standardised protocols that need to be carefully designed a priori.

With the objective to allow the estimation of *volume* fractions with more common single-TE multi-shell dMRI data, we formulated the Multi-Tissue MC (MT-MC) model of the dMRI signal. This novel formulation is inspired by the technique of Jeurissen et al. 2014 for the estimation of tissue-specific orientation distribution functions. The use of the MT-MC formulation solves some limitations of the previously mentioned multi-TE approaches and opens the door to the multi-tissue investigation of brain microstructure with data acquired with standard single-TE multi-shell dMRI protocols. Two algorithms for fitting the MT-MC model are proposed, one of which is designed to build on top of data already processed with standard MC models. We included it in the Diffusion Microstructure Imaging in Python (Dmipy) (Fick et al. 2019) framework, which is an open source tool designed for the abstraction, simulation and fitting of MC models of dMRI. The ability of the MT-MC model to retrieve

the unbiased volume fractions is tested on both synthetic data generated with Dmipy and real data obtained from the HCP database. Results obtained from these data confirm that the concepts of *signal fraction* and *volume fraction* are not interchangeable. Moreover, all the studies involving white matter tissue microstructure models should be reinterpreted taking into account the difference between the concepts of signal fraction and volume fraction.

The MT-MC model have already been presented at the International Symposium on Biomedical Imaging of 2020 Frigo et al. 2020c and at the 26th meeting of the Organization for Human Brain Mapping Frigo et al. 2020d. A manuscript in the form of article is being prepared to be submitted to a scientific journal (Frigo et al. 2021a).

Section 3.2 is devoted to the presentation of the theoretical bases of MC modelling, to the highlighting of why *signal* and *volume* fractions are not equivalent in general, and to the formalization of the MT-MC model that we proposed. Section 3.3 will include the design of the experiments and the presentation of their results, which are finally discussed in Section 3.4.

3.2 Theory

3.2.1 Multi-Compartment models

Complex microstructural configurations can be modelled as a linear combination of few elementary compartments. For example, the diffusion within axons can be described as the motion of water molecules along a stick or within a cylinder, while diffusion in free water like the one that can be observed in the CSF can be modelled as the diffusion within a ball. A vast portion of the dMRI literature of the last twenty years is devoted to the definition of compartmental models for anisotropic intra-axonal and extra-axonal diffusivity and for the isotropic diffusivity. These are known as Multi Compartment (MC) models and they all describe the shape of the dMRI signal by means of the following linear combination of compartment-specific shapes:

$$E(b, \mathbf{G}) = \frac{S(b, \mathbf{G}, TE)}{S_0(TE)} = \sum_{i=1}^{N_c} \phi_i \cdot E_i(b, \mathbf{G}) \quad (3.1)$$

where TE is the echo time, b is the b -value, \mathbf{G} is the gradient direction, N_c is the number of considered compartments, E_i is the signal attenuation (a.k.a. response function) of compartment i and ϕ_i is the portion of E explained by compartment i , i.e. the *signal fraction* of the compartment. The derivation of analytical expressions for the compartment-specific response functions has been researched broadly and deeply in the past literature. See the work of Panagiotaki et al. 2014 for a thorough review of the topic.

(Frigo et al. 2020c): Frigo et al. (2020), 'Multi Tissue Modelling of Diffusion MRI Signal Reveals Volume Fraction Bias'

(Frigo et al. 2020d): Frigo et al. (2020), 'Multi-compartment modelling of diffusion MRI signal shows TE-based volume fraction bias'

3.2.2 The Standard Model of White Matter

Several MC models have been proposed to describe the microstructural composition of the brain, among which we mention the stick-and-ball model of T. E. Behrens et al. 2003, the ActiveAx model of Alexander et al. 2010 and the neurite orientation dispersion and density imaging (NODDI) model of H. Zhang et al. 2012. The common factor of these formulations is the presence of an anisotropic component that describes the WM and an isotropic component accounting for the CSF. This coherence is reflected in what Novikov et al. called the *standard model* of dMRI in the brain (Novikov et al. 2019). The anisotropic component of the standard model is described as a bundle of dispersed fibers composed of a stick and a zeppelin compartment (accounting for the intra- and extra-axonal diffusion respectively), while the isotropic component is modelled with a ball.

Borrowing the taxonomy from (Panagiotaki et al. 2014), the three compartments are mathematically defined as follows:

- ▶ The IC compartment is modelled as a stick whose free parameters are the parallel diffusivity λ_{\parallel} and the spherical coordinates of the direction of the fiber population $\mathbf{n} = (\theta, \varphi)$. The corresponding signal shape along gradient direction $\mathbf{G} \in \mathbb{S}^2$ with b -value b is:

$$E_{IC}(b, \mathbf{G}, \lambda_{\parallel}, \mathbf{n}) = e^{-b\lambda_{\parallel}\langle \mathbf{n}, \mathbf{G} \rangle} \quad (3.2)$$

where $\langle \mathbf{n}, \mathbf{G} \rangle$ denotes the usual scalar product in \mathbb{R}^2 .

- ▶ The EC component is described by a Zeppelin, which can be defined as a diffusion tensor that depends on the parallel diffusivity λ_{\parallel} , the perpendicular diffusivity λ_{\perp} and the direction of the fiber population $\mathbf{n} = (\theta, \varphi)$. The signal shape is given by the tensor model of Equation (2.16), which is here reported with the customised notation:

$$E_{EC}(b, \mathbf{G}, \lambda_{\parallel}, \lambda_{\perp}, \mathbf{n}) = e^{-b\mathbf{G}^T D \mathbf{G}} \quad (3.3)$$

where the diffusion tensor defined as $D = (\lambda_{\parallel} - \lambda_{\perp}) \mathbf{n}\mathbf{n}^T + \lambda_{\perp} I$ and I is the 3-by-3 identity matrix.

- ▶ The CSF compartment is modelled an isotropic ball, which is defined in the same way as the zeppelin for $\lambda_{\parallel} = \lambda_{\perp} = \lambda_r$ where λ_r is the radial diffusivity. Notice that the first term in the definition of the diffusion tensor disappears, hence the model does not depend on the principal direction \mathbf{n} , making the compartment *isotropic* as wanted. The expression for the signal shape reads as follows:

$$E_{CSF}(b, \lambda_r) = e^{-b\lambda_r} \quad (3.4)$$

where one can notice that the signal decay depends solely on the b -value, making it independent from the gradient direction.

Additionally, fiber dispersion is formalized as the convolution of the stick and zeppelin compartments with an ODF denoted by \mathcal{P} . An example of such orientation function is the Watson distribution $W(\mathbf{n}, \kappa)$ (Mardia et al. 1990), which assumes axial symmetry of the dispersion around the main direction of the bundle $\mathbf{n} \in \mathbb{S}^2$ with concentration κ . The corresponding orientation dispersion index (ODI) can be computed as $ODI = 2/\pi \cdot \arctan(1/\kappa)$ (H. Zhang et al. 2012).

Given the elements described in the previous lines, the MC formulation of the standard model is defined as follows:

$$E(\mathbf{n}, \kappa, \lambda_{\parallel}, \lambda_{\perp}, \lambda_r, \phi_{IC}, \phi_{EC}, \phi_{CSF}) = \mathcal{P}(\mathbf{n}) * [\phi_{IC} \cdot E_{IC}(\lambda_{\parallel}, \mathbf{n}) + \phi_{EC} \cdot E_{EC}(\lambda_{\parallel}, \lambda_{\perp}, \mathbf{n})] + \phi_{CSF} \cdot E_{CSF}(\lambda_r) \quad (3.5)$$

where the dependence on the acquisition parameters b and \mathbf{G} has been omitted for the sake of readability.

Several constraints can be applied to the model given in Equation (3.5), among which the most common are hereafter reported.

1. The sum of the signal fractions is unitary: $\phi_{IC} + \phi_{EC} + \phi_{CSF} = 1$;
2. The perpendicular diffusivity of the EC compartment is tortuous, which in mathematical terms means the following:

$$\lambda_{\perp} = \frac{\phi_{EC}}{\phi_{IC} + \phi_{EC}} \cdot \lambda_{\parallel}; \quad (3.6)$$

3. The parallel diffusivity of the IC and EC compartments is fixed (e.g. $\lambda_{\parallel} = 1.7 \cdot 10^{-9} m^2/s$ as in (H. Zhang et al. 2012));
4. The radial diffusivity of the CSF compartment is fixed (e.g. $\lambda_r = 3.0 \cdot 10^{-9} m^2/s$ as in (H. Zhang et al. 2012)).

Several recent studies questioned the validity of these constraints (Dell'Acqua et al. 2019; Jelescu et al. 2016; Lampinen et al. 2017).

In presence of two distinct compartments for the intra- and extra-axonal components of the white matter, the so-called *tortuosity* constraint (Szafer et al. 1995; Aaron Szafer et al. 1995) can be employed to reduce the number of independent parameters. This means that the perpendicular diffusivity of the EC compartment is defined as a function of the parallel diffusivity and the proportion between the intra-axonal and the extra-axonal volume fractions (Jelescu et al. 2016). To employ tortuosity, first the fraction of the axonal compartments explained by the extra-axonal component is computed:

$$v_e = 1 - \frac{f_i}{f_i + f_e}$$

where v_e is the extra-axonal portion and f_i and f_e are the volume fractions of the intra- and extra-axonal compartments respectively. The perpendicular diffusivity λ_{\perp} can then be obtained by multiplying the parallel diffusivity λ_{\parallel} by the extra-axonal fraction:

$$\lambda_{\perp} = v_e \cdot \lambda_{\parallel}. \quad (3.7)$$

As highlighted by the left hand side of Equation (3.5), the model depends on eight parameters, where \mathbf{n} is two-dimensional, yielding 9 degrees of freedom, to which one has to subtract the degrees of freedom covered by the constraints. The remaining parameters can be estimated via standard least squares by defining the following minimization problem:

$$p^* = \operatorname{argmin}_p \frac{1}{2} \left\| \frac{\hat{S}}{\hat{S}_0} - E(p) \right\|_2^2 \quad (3.8)$$

where p is the parameter vector, \hat{S} is the acquired dMRI signal, \hat{S}_0 is the mean $b = 0$ image and $E(p)$ is the realization of the forward model given in Equation (3.5). Fitting such parameters requires the acquisition of multi-shell data with at least one shell per compartment (Scherrer et al. 2010). The obtained parameters p^* are the microstructural parameters that can finally be analysed for clinical or research purposes. In practice, the fitted signal fractions ϕ_i will likely not sum to 1, as they absorb any discrepancies between the normalised signal in the left hand side and the signal shapes in the left hand sides of equation (3.5), in particular when more than one image is acquired at $b = 0$.

A thorough review on the variety of MC models of WM that can be defined with the current state-of-the-art tools is the one of Fick et al. 2019, where also the Dmipy package is presented. This software is the reference tool used throughout this work for the study of microstructure. More recently, MC models have been used to assess also the microstructural composition of the gray matter (GM) (Fukutomi et al. 2019; Ganepola et al. 2018; Villalon-reina et al. 2020), but the literature is still sparse and there is a lack of agreement on how to model the GM with MC models.

3.2.3 MC models do not account for T2 differences

As shown in Equation (3.1), MC models aim at fitting the signal shape E as the ratio of the PGSE signal S and the S_0 amplitude. The implicit assumption that lies behind this formulation is that the S_0 by which the acquired signal is divided is the same for all the modelled compartments. In particular, as the S_0 image corresponds to the signal coming from the non-diffusion-weighted spin-echo sequence, we know that its amplitude depends on the TE and TR of the acquisition and on the T_1 and T_2 times of the sample. The mathematical relationship between these quantities is the one already presented in Equation (2.6), which read as follows:

$$S_0 \sim [H] \cdot (1 - e^{-TR/T_1}) \cdot e^{-TE/T_2} \quad (3.9)$$

where $[H]$ is the proton density in the sample. Before going into the interpretation of the two exponential terms of the equation, we must remember that different tissues have different T_1 and T_2 times. While in the formation of the S_0 image the different T_1 times of the tissues are negligible thanks to the length of TR (which is usually one order of magnitude longer than TE), tissues with different T_2 will generate sensibly different contrast in the image (Just et al. 1988; Plewes 1994; Veraart et al. 2018). Figure 3.1 illustrates how this difference is visible in the S_0 response of the WM and the CSF. These differences are the result of the contrast between the compartments in T_2 -weighted images. In order to understand how this difference in the T_2 impacts the signal-fraction estimation, consider the following example. Let a voxel in the WM containing some partial volume of CSF, which is common in the corpus-callosum near the ventricles. In particular, let's assume that the volume fractions are $f_{WM} = 0.9$ and $f_{CSF} = 0.1$. The corresponding signal equation will be

$$S = 0.9 \cdot S_0^{WM} \cdot E_{WM} + 0.1 \cdot S_0^{CSF} \cdot E_{CSF}. \quad (3.10)$$

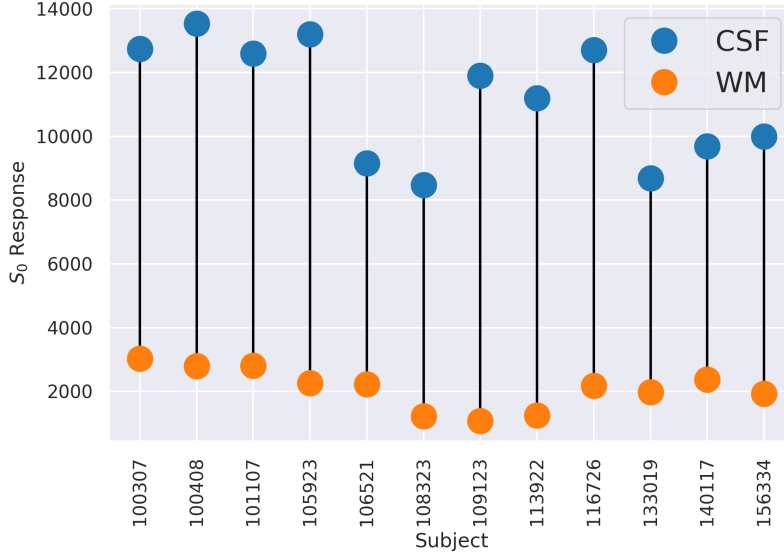


Figure 3.1: The figure shows the S_0 response of the WM and of the CSF for twelve randomly picked subjects from the HCP database. Values are obtained with the heuristic technique of Dhollander et al. 2016b via Mrtrix3 (Tournier et al. 2019).

As highlighted by Figure 3.1, the value of S_0^{CSF} can be up to six times the one of S_0^{WM} . Including this into our toy model, hence defining $S_0^{CSF} = 6 \cdot S_0^{WM}$, Equation (3.10) becomes

$$S = 0.9 \cdot S_0^{WM} \cdot E_{WM} + 0.6 \cdot S_0^{WM} \cdot E_{CSF} \quad (3.11)$$

which after dividing both sides of the equation by the composite $S_0 = f_{WM} \cdot S_0^{WM} + f_{CSF} \cdot S_0^{CSF}$ becomes

$$\frac{S}{S_0} = 0.6 \cdot E_{WM} + 0.4 \cdot E_{CSF} \quad (3.12)$$

yielding the signal fractions $\phi_{WM} = 0.6$ and $\phi_{CSF} = 0.4$. This example showed how signal fractions and volume fractions are not interchangeable concepts when it comes to modelling multiple tissues having different S_0 responses. Not taking into account these differences can lead to significant misrepresentations of the tissue composition, as shown in the previous example and in the results reported in Section 3.3.

Recent works (Gong et al. 2020; Lampinen et al. 2019, 2020; Veraart et al. 2018) addressed the possibility of using multi-TE acquisitions to disentangle the contribution of tissue with different T_2 times in MC modelling. The next section is devoted to the presentation of the advantages and limitations of such formulations.

3.2.4 Leveraging multi-TE sequences in Multi-Compartment modelling of the dMRI signal

If the problem of MC models is that they do not distinguish the S_0 of different tissues because of the limitations of single-TE acquisition sequences like the one considered in the previous sections, the solution could simply be to use multi-TE (MTE) acquisitions, despite the required longer acquisition time. This idea has been investigated in recent works of Veraart et al. 2018, Lampinen et al. 2019, 2020 and Gong et al. 2020. The common ground of these works is the following consideration: the

analysis of volume fractions as we intend them is not doable with conventional multi-shell dMRI acquired with a single echo time. In this section we are going to present how the MTE information is included in MC modelling following the flow of the definition of the TE-dependent Diffusion Imaging (TEdDI) technique proposed by Veraart et al. 2018. The same principles are followed in the more recent works of Lampinen et al. 2019, 2020 and of Gong et al. 2020. For the sake of coherence, we adapted the original notation used in the articles to the one used in this thesis.

The TEdDI technique considers a rewriting of the MC equation that directly includes the contribution of the T_2 time of the tissue modelled by the compartment and the TE of the acquisition into the volume fraction of each compartment. The contribution of the compartment-specific S_0^i considered in Equation (3.16) is modelled explicitly as follows:

$$S(b, TE, T_2^i, \mathbf{p}_i) = S_0 \cdot \sum_{i=1}^{N_c} \phi_i \cdot e^{-TE/T_2^i} \cdot E_i(b, \mathbf{p}_i) \quad (3.13)$$

where e^{-TE/T_2^i} plays the role of the compartment-specific contribution of the T_2 time and S_0 is the proton density- and T_1 -weighted image, which corresponds to Equation (2.6) for $TE = 0$. Notice that the T_2 time of each compartment is an independent variable of the model, hence it must be estimated in the fitting process. This requires the acquisition of multi-shell (to allow the use of multiple compartments) and multi-TE (to avoid degeneracy in the joint fitting of ϕ_i and T_2^i) dMRI data. The *volume fraction* of each compartment is defined by Veraart et al. 2018 and Gong et al. 2020 as follows:

$$f_i(TE) = \frac{\phi_i \cdot e^{-TE/T_2^i}}{\sum_j \phi_j \cdot e^{-TE/T_2^j}} \quad (3.14)$$

where one should notice how the volume fraction f_i depends on the echo time TE . Conversely, Lampinen et al. 2019, 2020 opted for defining the volume fractions as

$$f_i = \frac{\phi_i}{\sum_j \phi_j}, \quad (3.15)$$

which corresponds to the normalisation of the ϕ_i retrieved from fitting the model given in Equation (3.13). The formulation provided in Equation (3.13) can be regarded as the multi-TE standard model of the dMRI signal in the human brain, in analogy with what reported by Novikov et al. 2019 and in Section 3.2.2 of this thesis.

The MTE standard model has already been used in the previously cited works of Veraart et al. 2018, Lampinen et al. 2019, 2020 and Gong et al. 2020 to investigate the microstructure of the white matter of the brain. They showed that particular instances of the MTE standard model allow to assess how the T_2 time of the acquired sample is formed by the different compartments. Also, with MTE-MC models they showed that the concept of *volume fraction* should not just be abandoned in favor of the concept of *signal fraction*. Its straightforward interpretability is of much appeal in brain pathology research (Hara et al. 2018; Suzuki et al. 2017; Vestergaard-Poulsen et al. 2007), where biomarkers are not only quantified but also contextualized, related to other non-microstructural information and interpreted.

Some limitations come with the use of such formulation. First, the vol-

ume fractions defined in Equation (3.14) are TE -dependent. This poses several limitations in terms of usability and prevents from having a single index for the volume fraction of a compartment, which intuitively should be a characteristic of the sample, not of the acquisition. This ambiguity adds itself to the second limitation of the MTE formulation, which concerns *how* the T_2 time of the tissue modelled by each compartment is included in the model. As shown in Equation (3.13), the MTE framework corrects the classical MC model by multiplying each term in the sum by the T_2 -dependent factor Equation (2.6) of the S_0 image, which is exactly e^{-TE/T_2} . This contribution to the signal S is counted twice, as it is implicitly present also in the S_0 term that multiplies the sum on the right hand side of Equation (3.13). Notice that relaxing the $\sum_{i=1}^{N_c} \phi_i = 1$ constraint would solve the issue, as the double contribution would be corrected by an identical scaling of each ϕ_i , which is then absorbed in the normalization performed in Equation (3.14). Assessing if and how the repeated presence of the term affects the fitting of the T_2 times is out of the scope of this work. The third limitation of MTE-MC modelling that we highlight is of methodological nature. Classical MC models are representations of the dMRI signal that rely on standard multi-shell acquisitions designed in a HARDI fashion which have been used in the last 15 years for the study of both microstructure and tractography-based structural connectivity. The MTE framework does have the merit to correct the signal/volume fraction ambiguity, but this is achieved by increasing the complexity of the acquisition, which requires multiple TE to be considered. For this reason, the MTE framework is not to be considered an alternative to the MC formulation but rather a new method for the estimation of microstructural parameters that spans the whole range from acquisition design to post-processing, preventing from correcting the estimation of volume fractions on datasets acquired in the past.

In Section 3.2.5 we are going to propose a generalization of the MC models that takes into account the differences in terms of S_0 between the tissues modelled by each compartment without requiring multi- TE acquisitions to be considered. This will allow us to restore the concept of *volume fraction* in the context of MC modelling of the dMRI signal acquired with a single TE .

3.2.5 Multi-Tissue Multi-Compartment models

The standard formulation of MC models includes a normalization of the dMRI signal S_0 by its non-diffusion-weighted component S_0 . This operation is performed in order to retrieve the shape E of the acquired signal. The shape is then modelled as a linear combination of signal shapes of different compartments. In Section 3.2.3 we showed how this formulation hides the assumption that all the tissues modelled by the compartments have the same T_2 time (hence S_0), highlighting how this is not true a-priori. The solutions to the multi-tissue problem proposed in the literature and summarised in the previous section have the remarkable limitation of requiring the acquisition of multi- TE data to be used.

In the context of multi-tissue modelling, a solution has been proposed by Jeurissen et al. 2014 for the estimation of fODFs from multi-shell data, where the shell- and tissue-specific signal amplitude is leveraged in order

to rescale the fODF that describes the signal shape of each considered tissue. This includes the response of each tissue in the $b = 0$ shell, hence the S_0 of the tissues. The technique we are proposing in the following paragraphs builds on top of this idea. We highlight how similar solutions have been exploited also in the estimation of single-shell single-tissue response functions for the estimation of fODFs (Descoteaux et al. 2007; Tournier et al. 2007).

Let N_c be the number of compartments included in the model we want to design and let S_0^i be the S_0 response of compartment i . We define the *Multi-Tissue Multi-Compartment* (MT-MC) model as follows:

$$S(b, TE) = \sum_{i=1}^{N_c} f_i \cdot S_0^i(TE) \cdot E_i(b, p_i) \quad (3.16)$$

where f_i is the *volume fraction* of compartment i and $S_0^i(TE)$ is the S_0 response of the tissue modelled by compartment i . The dependence on TE is added to the notation only to remark that this formulation restores the factorization of the PGSE signal given in Equation (2.14), where the signal attenuation is defined as the *amplitude* \times *shape* product, with the *amplitude* being the TE-dependent component and the *shape* being the diffusion-dependent component. This mechanism is here replicated at compartmental level, implying that we model the *whole* signal instead of its shape only. Recalling that TE and T_2 are linked by the exponential term $S_0 \sim \exp(-TE/T_2)$ in Equation (2.6), we highlight how the MT-MC formulation does indeed take into account the differences between the S_0 responses in each tissue.

Notice that ϕ_i and f_i are not equivalent in general, as the former describes the *signal fraction* of compartment i and the latter describes the *volume fraction* of the tissue modelled by the compartment. The only case in which they are equivalent is when all the tissues described by the MT-MC model have equal S_0 responses. In that case, Equation (3.16) is equivalent to (3.1) after multiplying both the sides of the equation by S_0 . For this reason we say that ϕ_i is a *biased estimator* of f_i . One could argue that the relationship between the signal fractions ϕ_i and the volume fractions f_i is just a rescaling, in which case the volume fractions could be retrieved with a simple correction that takes into account the S_0 signal and the S_0^i response of the compartment. This is always true except when the volume fraction of the compartment is an independent variable in some other compartment. Using the tortuosity constraint to define the perpendicular diffusivity of the extra-axonal compartment, we establish a non-linear dependence between the compartmental fraction of the intra- and extra-cellular compartments. In this way, the diffusivity of the extra-axonal compartment, which has a non-linear relationship with the model, is defined as in Equation (3.7), forcing the intra- and extra-cellular signal/volume fractions to be non-linear parameters of the model. For instance, if the intra- and extra-axonal compartments have different S_0 , the perpendicular diffusivity of the EC computed with the tortuosity constraint defined on the signal fractions will be different from the one obtained from the volume fractions. As a consequence, two models defined with the two possible tortuosity constraints are not interchangeable and rescaling one's volume fractions does not yield the other's signal fractions.

3.2.5.1 Fitting MT-MC models

The fitting of a MT-MC model is designed in a fashion similar to the one of MC models, but two different approaches can be employed. The first provides only the volume fractions (VF), while the second yields both the signal and the volume fractions (SVF). Given the acquired dMRI signal \hat{S} , the corresponding S_0 , the number of compartments N_c , the signal shape $E_i(\mathbf{p}_i)$ of compartment i depending on the parameter vector \mathbf{p}_i and the compartment-specific signal amplitude S_0^i , the fitting can be performed in the two following ways.

VF Solve the following least squares problem with respect to the microstructural parameters f_i and \mathbf{p}_i :

$$f^*, \mathbf{p}^* = \underset{f_i, \mathbf{p}_i}{\operatorname{argmin}} \left\| S - \sum_{i=1}^{N_c} f_i \cdot S_0^i \cdot E_i(\mathbf{p}_i) \right\|_2^2 \quad (3.17)$$

which can be solved with standard inverse-problem solvers where the forward model is the one given in Equation (3.16). The procedure yields the volume fractions (VF) of the compartments.

SVF First fit the signal fractions ϕ_i and the microstructural parameters \mathbf{p}_i from the MC formulation of Equation (3.5), then consider the MT-MC formulation of Equation (3.16) and fix the microstructural parameters \mathbf{p}_i to the ones fitted in the previous step. The volume fractions can then be retrieved as a rescale of the signal fractions. In mathematical terms, the described procedure reads as follows:

1. Solve the associated MC problem:

$$\phi^*, \mathbf{p}^* = \underset{\phi_i, \mathbf{p}_i}{\operatorname{argmin}} \left\| \frac{S}{S_0} - \sum_{i=1}^{N_c} \phi_i \cdot E_i(\mathbf{p}_i) \right\|_2^2 \quad (3.18)$$

where the product of the minimization problem is the signal fraction ϕ_i and the parameter vector \mathbf{p}_i of each compartment i ;

2. Fix the fitted non-signal-fraction parameters in the MT-MC model. At this point the volume fractions are not related to each other (or to other compartments in general) and it is therefore possible to estimate them by rescaling the signal fractions. The rescale is the one suggested by the comparison of the coefficients that multiply the signal shapes in Equations (3.1) and (3.16) and reads as follows:

$$\begin{aligned} f_i \cdot S_0^i &= \phi_i \cdot S_0 \\ f_i &= \phi_i \cdot \frac{S_0}{S_0^i} \end{aligned} \quad (3.19)$$

yielding a simple operation that allows to retrieve volume fractions from signal fractions once the S_0 of each compartment is known. Both the signal and volume fractions of each compartment are returned.

To employ the SVF strategy, extra caution must be taken towards the use of the tortuosity constraint. The intra- and extra- axonal fractions used for the definition of the perpendicular diffusivity can be either the signal fractions or the volume fractions of the compartments. The choice influences the whole model design and should be taken with extra care. In an effort to keep the notation coherent with the previous literature, we will say that whenever the tortuosity constraint is defined using the volume fractions we will have a *MT-corrected* tortuosity constraint.

The major difference between the two approaches is that VF fits directly the MT-MC model, while SVF first fits the associated MC model, then does a rescale of the obtained signal fractions. The reason why the two-step procedure SVF was designed is two-fold:

- ▶ to be able to access both the signal and the volume fractions of the compartments at the same time;
- ▶ to allow to re-process results that had previously been obtained on standard MC models in a MT fashion with a simple operation.

The SVF strategy is the one implemented in Dmipy (Fick et al. 2019), which at our knowledge is the only available framework for generalised MC modelling that includes the definition of MT-MC models. As far as specific instances of MT-MC models are concerned, the multi-shell multi-tissue CSD technique of (Jeurissen et al. 2014) is implemented in Mrtrix3 (Tournier et al. 2019).

3.2.6 Generalised MT-MC Modelling

The proposed MT-MC framework is in fact a general formulation that contains both MC and MTE-MC models as particular cases. The MC case has already been discussed in Section 3.2.5.

Multi-TE MC models are a particular instance of MT-MC models in the sense that for a specific choice of S_0^i in Equation (3.16) one recovers the formulation of MTE-MC models given in Equation (3.13). Comparing the coefficients that multiply the (non-TE-dependent) signal shapes in the two equations, one obtains

$$\begin{aligned} f_i \cdot S_0^i &= S_0 \cdot (\phi_i \cdot e^{-TE/T_2^i}) \\ f_i \cdot S_0^i &= \phi_i \cdot S_0 \cdot e^{-TE/T_2^i} \end{aligned} \quad (3.20)$$

which suggests that the MTE-MC model can be obtained by setting $S_0^i = S_0 e^{-TE/T_2^i}$. The fact that we do not include the $\sum_{i=1}^{N_c} f_i = 1$ constraint makes our formulation immune to the sub-optimality of this choice discussed in Section 3.2.4. In a different context, the same problem is solved in an identical way in the COMMIT framework (Daducci et al. 2014), where the signal fractions are not subject to the unitary sum constraint.

The same idea used for including the T_2 effects in the MC modelling could be exploited to include any PD or T_1 correction in a fashion similar to the one of (3.13). Acquisitions obtained with state of the art protocols do not allow to sense such differences between tissues all together, but future developments could open this opportunity. The model is ready.

3.2.7 The MT Standard Model of dMRI in White Matter

Inspired by the standard model of dMRI in WM as described by Novikov et al. 2019, in this section we define the MT generalization of such model. We recall that the model includes a stick and a zeppelin compartment for the intra- and extra-cellular diffusivity respectively and a ball that accounts for the diffusivity in the CSF and other isotropic structures. Let S_0^i be the S_0 response of the tissue modelled by each compartment i and $\mathcal{P} : \mathbb{S}^2 \rightarrow \mathbb{R}^+$ the orientation distribution. The MT standard model of dMRI in WM reads as follows:

$$S(\mathbf{n}, \kappa, \lambda_{\parallel}, \lambda_{\perp}, \lambda_r, f_{IC}, f_{EC}, f_{CSF}) = P(\mathbf{n}) * \left(\underbrace{f_{IC} \cdot S_0^{IC} \cdot E_{IC}(\lambda_{\parallel}, \mathbf{n})}_{\text{intra-axonal}} + \underbrace{f_{EC} \cdot S_0^{EC} \cdot E_{EC}(\lambda_{\parallel}, \lambda_{\perp}, \mathbf{n})}_{\text{extra-axonal}} \right) + \underbrace{f_{CSF} \cdot S_0^{CSF} \cdot E_{CSF}(\lambda_r)}_{\text{CSF}} \quad (3.21)$$

where the compartment specific parameters are defined as in Section 3.2.2. Three scenarios can be described with the model given in Equation (3.21):

- ▶ *3-tissue model* - The three compartments describe tissues with distinct S_0 responses. This corresponds to the explicit case of Equation (3.21).
- ▶ *2-tissue model* - The two anisotropic compartments (IC and EC) model tissues whose S_0 is equal. Typically, it is the S_0 of the WM, so we say that $S_0^{IC} = S_0^{EC} = S_0^{WM}$ and $S_0^{WM} \neq S_0^{CSF}$.
- ▶ *1-tissue model* - In absence of any prior knowledge on the S_0 of the three tissues, they are considered all equal. We denote this as $S_0^{IC} = S_0^{EC} = S_0^{CSF} = S_0$ where S_0 is the average across the images acquired at $b = 0 \text{ smm}^{-2}$.

Notice that the 1-tissue scenario is mathematically equivalent to the single-tissue (ST) standard model of Section 3.2.2, hence the associated volume fractions are equivalent to the signal fractions.

In the following section we are going to study the impact on the estimation of compartment-specific volume fractions of the assumptions used for defining the three models. In particular, a global picture on these differences will be obtained on a simulated dataset, while a more specific insight on the differences between the 2-tissue model and the 1-tissue model will be addressed on real data.

3.3 Experiments

The model considered in the performed experiments is the MT standard model defined in the previous section with fixed S_0^i and three additional constraints:

- ▶ The fiber orientation distribution is modelled with a Watson distribution of axis \mathbf{n} and fixed ODI.

- ▶ The perpendicular diffusivity is subject to the tortuosity constraint, hence

$$\lambda_{\perp} = \left(1 - \frac{f_{IC}}{f_{IC} + f_{EC}}\right) \cdot \lambda_{\parallel}. \quad (3.22)$$

- ▶ The parallel and radial diffusivity are fixed to $\lambda_{\parallel} = 1.7 \cdot 10^{-9} m^2 s^{-1}$ and $\lambda_r = 3.0 \cdot 10^{-9} m^2 s^{-1}$ respectively.

The free parameters that are left are $[f_{IC}, f_{EC}, f_{ISO}, \mathbf{n}, \kappa]$, where we recall that the normalised direction $\mathbf{n} = [\theta, \varphi]$ is two-dimensional. Fitting is done with the SVF procedure described in Section 3.2.5.1 in order to retrieve both the signal fractions and the volume fractions to be compared.

3.3.1 Synthetic phantom

3.3.1.1 Dataset

The simulated dataset is obtained from the forward model given by Equation (3.21) and generated with Dmipy (Fick et al. 2019). A total of 10'000 voxels was simulated on a multi-shell acquisition scheme identical to the one that will be considered on the real dataset, which includes a TE of 0.0895s and is composed of 288 samples subdivided in 18 points at $b = 0s/mm^2$ and 90 diffusion-weighted samples obtained with uniformly distributed directions at $b = 1000s/mm^2$, $b = 2000s/mm^2$ and $b = 3000s/mm^2$ for a total of 3 non- $b = 0$ shells plus the $b = 0$ shell. The direction of the IC and EC compartments was set to $\mathbf{n} = [0, 0] \in S^2$ for all the voxels. The T_2 time of each tissue was randomly sampled from a uniform distribution in the range specified in Table 3.1. The corresponding S_0 was then computed as $S_0 = c \cdot e^{-TE/T_2}$ where c is a scaling parameter that positions the value of S_0 in a realistic range and we tuned to $c = 1400$. The ODI of the Watson distribution was sampled from a uniform distribution in the range specified in Table 3.1. Finally, the volume fractions of each compartment were randomly generated from a uniform distribution in the range specified in Table 3.1, then normalized in such a way that their sum was equal to 1. The choice of each range was tuned to mimic the single-bundle configuration in the WM that one expects to be able to model with the considered formulation. An additive rician noise

Parameter	Min	Max
f_{IC}	0.5	0.8
f_{EC}	0.3	0.5
f_{CSF}	0.3	0.7
T_2^{IC}	0.080s	0.100s
T_2^{EC}	0.050s	0.070s
T_2^{CSF}	0.900s	1.100s
ODI	0.02	0.99

Table 3.1: For each parameter used in the definition of the forward model of the synthetic dataset we report the minimum and maximum value of the uniform distribution from which it was drawn.

was added to the simulated data to obtain a signal-to-noise ratio equal to 30.

3.3.1.2 Results

We fitted the volume fractions \hat{f}_i of each compartment with the SVF procedure explained in Section 3.2.5.1 with the 1-tissue (1T), 2-tissue (2T) and 3-tissue (3T) model, both with standard and MT-corrected (3T-MTT) tortuosity. The latter is actually different from the standard tortuosity only when the 3-tissue model is considered. When the 2-tissue model is employed, the S_0^{WM} response is computed as the weighted average of the S_0 response of the IC and EC compartments, hence

$$S_0^{WM} = \frac{f_{IC} \cdot S_0^{IC} + f_{EC} \cdot S_0^{EC}}{f_{IC} + f_{EC}}. \quad (3.23)$$

Once each \hat{f}_i was estimated, we computed the absolute fitting error $|f_i - \hat{f}_i|$. The difference between the absolute errors obtained with each model is tested with a Wilcoxon signed-rank test (Wilcoxon 1945) with $\alpha = 0.05$ and measured with the rank biserial correlation, i.e., an effect size measure in the $[0, 1]$ range. In Figure 3.2 we report the boxplot of the distribution of the absolute fitting error across 10000 simulations (top row) and the rank biserial correlation (bottom row). As we expected, the signal




Figure 3.2: The first row shows the boxplot of the absolute error of the estimated volume fraction of each compartment computed on the synthetic dataset. The 1T categorical variable corresponds to the 1-tissue model, 2T to the 2-tissue, 3T to the 3-tissue with standard tortuosity and 3T MTT to the 3-tissue model with MT-corrected Tortuosity (MTT). The second row shows, for each compartment, the rank biserial correlations measuring the difference between the absolute errors of each model obtained from a Wilcoxon test with $\alpha = 0.05$: the lower the correlation value, the higher the difference between the underlying distributions. Only statistically significant results are reported.

fractions retrieved by the 1-tissue model are biased estimates of the volume fractions retrieved with the 2-tissue and 3-tissue model. The bias in the estimation of the volume fraction of the CSF compartment is four times bigger than the one of the IC and EC compartment. This is coherent with the fact that the S_0 of the CSF compartment is much higher than the one of the IC and EC compartments. The error decreases importantly

when the 2-tissue model is used. Here, the IC volume fraction has absolute error comparable to the one of the 3-tissue models. The first factor that could induce such phenomenon is the definition of S_0^{WM} , which by design of the experiment will be closer to the S_0 of the IC than to the one of the EC compartment ($f_{IC} > f_{EC}$ as reported in Table 3.1). This induces the estimated EC volume fraction to be farther from the ground truth than the one of the IC compartment. This difference is reflected in the absolute error of the estimated volume fraction of the CSF compartment, which is affected by the presence of the non-zero perpendicular diffusivity of the EC compartment. Nevertheless, the estimation error of the CSF volume fraction is much lower than in the 1-tissue model thanks to the inclusion of the specific S_0^{CSF} in the formulation. Finally, the 3-tissue model retrieves volume fractions that are in line with the ground truth ones. A notable aspect concerns the inclusion of the MTT correction, which is shown to significantly improve the estimation of the IC volume fraction. This improvement corresponds to a deterioration (albite lower in scale) of the estimation of the EC volume fraction. The estimation of the CSF volume fraction does not change significantly when the MTT correction is employed.

3.3.2 Real Data

3.3.2.1 Dataset

From the Human Connectome Project (HCP) database we considered three randomly picked subjects* available at the Connectome Coordination Facility (Sotiropoulos et al. 2013; Van Essen et al. 2012). For each subject a total of 288 images is acquired, 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.

To our knowledge, current state-of-the-art techniques do not allow to estimate subject-specific S_0 responses of the IC and EC compartments, while the S_0 of the CSF compartment can be estimated together with S_0^{WM} via techniques such as the heuristic approach of Dhollander et al. 2016a. For this reason, we analysed the aforementioned data with a 1-tissue and a 2-tissue model where S_0^{WM} and S_0^{CSF} have been estimated with the Dhollander technique. The obtained values of S_0 are displayed in Figure 3.1 and reported in Table 3.2.

The difference between the absolute errors obtained with each instance of the model is tested with a Wilcoxon signed-rank test (Wilcoxon 1945) with $\alpha = 0.05$ and measured with the rank biserial correlation, i.e., an effect size measure in the $[0, 1]$ range.

3.3.2.2 Results

For each model, we fitted the signal and volume fractions with the SVF technique. Figure 3.3 shows the distribution of the signal fraction and volume fraction of each compartment in the WM for three HCP subjects. The WM mask was computed with FSL *fast* from the T_1 -weighted image

Table 3.2: S_0 response of the WM and of the CSF of the three studied HCP subjects. Values are obtained with the heuristic technique of Dhollander et al. 2016b via Mrtrix3 (Tournier et al. 2019).

Subject	S_0^{WM}	S_0^{CSF}
#1	3024	12740
#2	2794	13531
#3	2811	12598

* ID subject 1: 100307, ID subject 2: 100408, ID subject 3: 101107.

Subject	IC		EC		CSF	
	r	p	r	p	r	p
#1	0.243	0.0	0.036	0.0	0.002	0.0
#2	0.221	0.0	0.142	0.0	0.002	0.0
#3	0.197	0.0	0.100	0.0	0.001	0.0

Table 3.3: For each subject (1, 2, 3) and compartment (IC, EC, CSF), the table displays the value of the rank biserial correlation r and the corresponding p-value p computed obtained from a Wilcoxon signed rank test (Wilcoxon 1945) with $\alpha = 0.05$. The showed values are computed with Scipy (`scipy`) and all the performed comparisons exhibit statistically significant differences.

with 1.25mm voxel size available at the HCP database, then dilated by one voxel with Mrtrix3's (Tournier et al. 2019) `maskfilter` command to smooth the boundary. For each subject, the difference between the distribution of the signal and the volume fractions is tested with a Wilcoxon signed-rank test (Wilcoxon 1945) with $\alpha = 0.05$ and measured with the rank biserial correlation, i.e., an effect size measure in the $[0, 1]$ range. The results of this test are reported in Table 3.3.




Figure 3.3: The displayed data are obtained three subjects of the HCP database (solid lines, dashed lines, and dotted lines). Each panel shows the distribution of the signal fraction and the volume fraction of the IC, EC and CSF compartments respectively. The blue lines correspond to *signal* fractions and the orange lines to *volume* fractions.

We recall that we considered a 2-tissue model by compressing the IC and EC compartments in a unique block that describes the WM tissue. The distribution of the volume fractions of the IC and EC compartments showed in Figure 3.3 is right-shifted with respect to the distribution of the corresponding signal fractions. On the contrary, the distribution of CSF volume fractions in the WM mask is shifted towards lower values with respect to the corresponding signal fractions. This means that the signal fraction *underestimates* the presence of the intracellular compartment in favour of the CSF compartment. This behaviour is consisted in all the tested subjects.

This is coherent with the proportion between S_0^{WM} and S_0^{CSF} , as the former is typically lower than the acquired S_0 and the latter is higher. The results displayed in Figure 3.4 show how, within the WM mask, the WM volume fraction is globally higher than the WM signal fraction. Also, the absolute difference between the two exhibits some uniformity within the considered sample. The macroscopic differences between the left and right hemispheres present in all the three subjects may be due to some bias field effect that we did not include in the model and survived the minimal preprocessing of the data (Matthew F Glasser et al. 2013).




Figure 3.4: Signal fractions (first column), volume fractions (second column), and their absolute difference (third column) for three subjects of the HCP database. Brighter colors correspond to higher fractions (in the first two columns) and errors (in the third column). Voxels shown in orange/red and black correspond to decreasingly lower values of the same fractions and errors.

3.4 Discussion

In this chapter we analysed how multi-compartment models of brain tissue microstructure can be adapted to account for the presence of tissues having different T_2 relaxation times. In particular, we focused on the capability of such models to estimate the volume fraction of each tissue in the WM. We proposed a solution based on single-TE dMRI data, in contrast with the state-of-the-art techniques that require multi-TE dMRI data. Our results on both synthetic and in-vivo data show that *signal fraction* and *volume fraction* are not interchangeable concepts in the context of MC microstructure modelling. The shift of paradigm from signal fractions to volume fractions has already been shown to improve the estimation of fODFs (Jeurissen et al. 2014) and in this work we transferred the same approach to the field of MC models of brain tissue microstructure, leveraging the differences between the S_0 responses of each modelled tissue.

Overall, the presented results yielded an empirical confirm of the theoretical considerations made in this work. In particular, the following aspects are highlighted:

- With *single-TE dMRI data* it is possible to retrieve tissue-specific volume fractions. Under the assumption that the IC and EC compartments have equal S_0 response, techniques like the one of Dholander et al. 2016b allow to define the 2-tissue model used in this section, opening the door to a better estimation of the compartment-specific volume fractions. This is made possible by the MT-version of the standard model of dMRI in the WM that we presented in this work. It models multiple tissues in a MC fashion *without* requiring multi-TE acquisition, which are conversely necessary in order to use other state-of-the-art models.

- *Signal fractions and volume fractions are not equivalent in general.* This fact has considerable implications in clinical context. Previous studies that drew conclusions based on the idea of inspecting volume fractions with single-TE dMRI need to be re-interpreted in light of the fact that what they are based on is the *signal* fraction of the tissues and not their volume fraction. How those differences are expressed in the presence of pathology or group differences remains unexplored and needs to be assessed in future studies.

We designed a multi-tissue version of the standard model of dMRI in the WM, which allows to separate the contribution of the intra-axonal, the extra-axonal and the CSF compartments and estimate the corresponding three volume fractions. The results reported in Figure 3.2 suggest that 2-tissue and 3-tissue models are always preferable to the 1-tissue model. A bigger improvement is obtained by considering two tissues instead of one, compared to the shift from the 2-tissue to the 3-tissue model. This is due to the proportion between the S_0 s of each tissue, which sees $S_0^{CSF} \gg S_0^{WM}$, with $S_0^{IC} > S_0^{EC}$ but the latter difference is lower than the former (Jeurissen et al. 2014).

A remarkable property of the proposed MT-MC model is that not only it can be straightforwardly fitted on single-TE dMRI data (VF strategy), but it can also re-use the results obtained with the MC version of the model (which in principle would have returned only the signal fraction of each compartment) and yield the volume fractions by means of an elementary rescaling operation (SVF strategy). While employing the SVF solution, extra care must be devoted to the use of the tortuosity constraint. Rescaling signal fractions obtained using the non-MT-corrected tortuosity constraint yields the volume fractions of a model where the perpendicular diffusivity of the EC compartment has been obtained using signal fractions, configuring an ambiguous (if not degenerate) solution.

The proposed model strongly relies on the external estimation of the T_2 or the S_0 of the modelled tissues. Our experiments on real data leveraged the heuristic of Dhollander et al. 2016b to retrieve the S_0 of the WM and the CSF. Understanding how this choice affects the estimation of volume fractions is out of the scope of this work, but the raised question suggests that further efforts should be devoted to researching techniques that estimate tissue-specific S_0 responses using single-TE data. Additionally, analysing the proportion between the S_0 of each tissue in a large cohort of subjects could highlight patterns that could be exploited. If hypothetically the T_2 of extra-axonal was showed to be a constant fraction of the T_2 of the intra-axonal compartment, this could straightforwardly be encoded in the model.

The difference between signal fractions and volume fractions has implications also in the field of tractography filtering (Frigo et al. 2020b), where a coefficient is assigned to each streamline in a tractogram weighing its contribution to the formation of the dMRI signal. In the COMMIT framework (Daducci et al. 2014) these coefficients are the signal fractions associated to each streamline. The model can be easily adapted to obtain the volume fraction associated to each streamline, in particular in the context of the recent work of Barakovic et al. 2020, where streamlines are associated to bundle-specific T_2 times.

3.5 Conclusion

In this chapter we presented our work in the field of brain tissue microstructure estimation via multi-compartment models of dMRI. A known limitation of these models is their inability to model multiple tissues at the same time. We showed how what have always be considered to be the volume fraction of a certain tissue is shown to actually be the signal fraction of the same tissue. State-of-the-art techniques for overtaking such limitation rely on multi-TE dMRI data. In this chapter we proposed a multi-tissue the Multi-Tissue Multi-Compartment models of dMRI, which allow to model multiple tissues at the same time using single-TE dMRI data. Moreover, we formulated a generalised multi-tissue modelling framework that encompasses both single-TE and multi-TE multi-tissue models. Our results indicate that with single-TE dMRI data alone one can model multiple tissues with multi-compartment models.

This work was done in collaboration with Rutger Fick [†]. This chapter is based in part on the following published works.

- ▶ *Matteo Frigo*, Mauro Zucchelli, Rutger Fick, Samuel Deslauriers-Gauthier, Rachid Deriche. Multi-compartment modelling of diffusion MRI signal shows TE-based volume fraction bias. OHBM 2020.
- ▶ *Matteo Frigo*, Rutger Fick, Mauro Zucchelli, Samuel Deslauriers-Gauthier, Rachid Deriche. Multi Tissue Modelling of Diffusion MRI Signal Reveals Volume Fraction Bias. ISBI 2020.
- ▶ *Matteo Frigo*, Rutger Fick, Mauro Zucchelli, Samuel Deslauriers-Gauthier, Rachid Deriche. Multi Tissue Multi Compartment Modelling of Single-TE Diffusion MRI. bioRxiv 2021.01.29.428843.

[†] Former PhD student in the team, now at TRIBVN Healthcare, Paris, France

Chapter 4

Filtering dMRI-based tractograms

Overview Several works exposed the limitations of dMRI-based tractography, in particular in relation to the lack of quantitiveness in the information carried by the estimated fiber pathways. A class of methods that has been designed for addressing this issue goes under the name of *tractography filtering* techniques (TFTs). These techniques act as post-processing of a pre-computed tractogram by assign to each streamline a coefficient representing how much the streamline is necessary for explaining the acquired dMRI signal or some transformation of it. The first contribution of this chapter is a review of the state of the art TFTs and definition of a unified framework that generalises the previous approaches. The *Tractograms As Linear Operators in Neuroimaging* (TALON) Python package allows this generalised approach to TFTs and has been publicly released together with this thesis. We show how the inclusion of TFTs in connectomic pipelines changes the topology of the obtained structural connectomes on both healthy subjects and patients affected by traumatic brain injury. This work was done in collaboration with Ragini Verma* and Junghoon John Kim†. The second contribution here presented concerns the presentation of a novel TFT that extends the previous approaches by including not only the structural information represented by diffusion MRI, but also the functional information that is encoded in functional connectomes. Our preliminary results show that the proposed method is sound with respect to the criteria adopted for the evaluation of related techniques.

4.1 Introduction	49
4.2 Theory	51
Tractography Filtering Techniques	51
Unified Tractography Filtering Framework	57
Functionally Informed Tractography Filtering	60
4.3 Experiments	62
TFTs Change the Topology of Structural Connectomes	62
Functionally-informed TFT	72
4.4 Conclusion	80

4.1 Introduction

The increasing interest in connectomics shed light on the intrinsic limitations of tractography based connectomes (Daducci et al. 2016; Jbabdi et al. 2011; Jones et al. 2013; Maier-Hein et al. 2017; Francois Rheault et al. 2020a). Tractography itself is an ill-posed problem that aims at reconstructing the white matter fiber tracts from dMRI data. For example, straight bundles are easier to recover than curved or fanning bundles (Schilling et al. 2019b). Moreover, as exemplified in Figure 4.1, crossing and kissing fiber configurations require specific prior knowledge to be distinguished (Francois Rheault et al. 2020a). Considering that crossing and kissing configurations are vastly present in the human white

* Penn Applied Connectomics and Imaging Group, Department of Radiology, University of Pennsylvania, Philadelphia, PA, United States of America

† Department of Molecular, Cellular, and Biomedical Sciences, CUNY School of Medicine, The City College of New York, NY, United States of America




Figure 4.1: *Left panel:* Assuming that the gray region is the portion of WM that is being studied and that the anatomically correct bundles that form such geometry are the ones that follow the red and the blue lines, the figure gives an example of a crossing bundles configuration. The limited resolution of dMRI does not allow to distinguish the red and blue bundles from the yellow and azure ones. *Right panel:* Without additional priors, tractography yields all the four bundles independently of what the ground truth configuration is.

matter, the number of bottlenecks that are encountered in the tracking process grows quickly, drastically reducing the accuracy of the recovered streamlines. As a result of the tractography challenge organised within ISMRM 2015 focused on the structural connectomics implications of tractography, Maier-Hein et al. 2017 reported that state-of-the-art tractography algorithms retrieve most of the valid bundles but the retrieved invalid bundles are more than the valid ones. Notice that these results rely on the mere presence of such streamline bundles in the tractogram. When a quantitative approach is followed in the evaluation of connectomes, such as the *streamline count* measure defined in Chapter 2, the problem inherits the uncertainty carried by the lack of volume that affects the streamlines. As a matter of fact, streamlines are defined as curves in the three-dimensional space, hence they do not occupy a certain portion of volume per-se. This makes the density of streamlines much more affected by co-factors such as the seeding strategy, the curvature of the bundle that is being tracked, the degree of branching of the fiber configuration and any other configuration that deviates from the simple case of the straight bundle not traversed by any other fiber (Jones et al. 2013; Francois Rheault et al. 2020a). Finally, several studies (Jbabdi et al. 2011; Jeurissen et al. 2019; Jones et al. 2013) pointed the finger towards the fundamental sin of dMRI-based tractography: tracking algorithms assume that axons are aligned with the direction of minimum water hindrance. While this is a reasonable approximation at axonal resolution, it loses validity when brought to the dMRI resolution, where water encounters a variety of barriers and restrictions limiting the acceptability of the simplification that describes the anisotropy of water diffusion as a direct consequence of the microstructure of the sole axons. Finally, dMRI is not able to distinguish afferent and efferent connections. For these and other reasons (Jbabdi et al. 2011; Jeurissen et al. 2019; Jones et al. 2013; Maier-Hein et al. 2017; Francois Rheault et al. 2020a; Schilling et al. 2019a,b), the universally accepted truth is that *tractography per-se is not quantitative*. This feels like a waste of information, since dMRI is a complex quantitative technique that measures with incredible precision the average displacement of water molecules, which is a phenomenon at a much more refined scale than the one described with tractograms.

Among the many problematic issues of tractography-based structural connectivity estimation, the non-quantitative nature of tractography (Jones et al. 2013; R. E. Smith et al. 2015b), the methodological limitations of dMRI-based tractography (Francois Rheault et al. 2020a) and the presence of many false positive connections within connectomes (Maier-Hein et al. 2017) are the ones that inspired some of the most recent methodological advances in the field of non-invasive tractography. With the aim of re-establishing the quantitative properties that one expects from tractography, i.e. the strength of the structural connection being associated to quantitative aspects of the fiber population represented by each streamline, several evolution of the standard tracking paradigm have been developed. One such evolution is the so-called *global tractography* approach (Christiaens et al. 2015; Kreher et al. 2008; Schreiber et al. 2014), which makes use of a simulated annealing process and has been proven to mitigate the potential bias induced by the geometry of the streamlines to be tracked (bending, fanning, crossing, kissing, ...). Another technique called AxTract (Girard et al. 2015, 2017) simultaneously tracks the streamlines and associates an axonal diameter to the tracked bundles. The global and AxTract methods are *generative* (a.k.a. bottom-up) techniques, as they build streamlines according to criteria that are more restrictive and better informed than the ones of ordinary tractography algorithms.

In contrast with the bottom-up approach, a series of top-down techniques have been proposed, namely the Spherical-deconvolution Informed Filtering of Tractograms (SIFT) technique (R. E. Smith et al. 2013) and its evolution (R. E. Smith et al. 2015a), the Convex Optimization Modelling for Microstructure Informed Tractography (COMMIT) (Daducci et al. 2014) and its evolution COMMIT2 (Schiavi et al. 2020) and the Linear Fascicle Evaluation (LiFE) model (Pestilli et al. 2014). These techniques go under the name of Tractography Filtering Techniques (TFTs). They take a predefined tractogram and associate to each streamline a single weight that represents the amount of signal explained by or the connectivity strength associated to the streamline depending on the used technique. The following section is devoted to a review of these methods, to their unification under a more general formulation and to the presentation of the *Functionally Informed TFT*, which is a novel filtering technique that extends the previous ones by integrating functional information in the model. Section 4.3 will be dedicated to two sets of experiments. In the first we will assess how the use of TFTs affects the graph-theoretical analysis of structural connectomes estimated with dMRI. This work has been published in the Journal of Neural Engineering (Frigo et al. 2020b). The second set of experiments is dedicated to the preliminary analysis of the proposed FIT framework on both synthetic and real data. The results are then discussed and a conclusion is given in Section 4.4

4.2 Theory

4.2.1 Tractography Filtering Techniques

As we discussed in the previous section, tractography is intrinsically ill-posed and several of its limitations have been unveiled in the last fifteen

years, inspiring the development of novel tractography algorithms. Tractography filtering techniques are top-down approaches that *augment* a tractogram by associating a coefficient w_s to each streamline s according to how much of the dMRI signal or of a certain map is explained by the streamline. This coefficient quantifies the connectivity that is associated to the specific streamline and corresponds to the w_s defined in Equation (2.23). In the context of tractometry (Côté et al. 2013), streamlines associated to null weights are marked as false positives, hence discarded from the tractogram in the construction of the connectome. In this study we are going to focus only on SIFT2, and COMMIT2, as LiFE can be formulated as a particular case of COMMIT and SIFT2 and COMMIT2 substitute SIFT and COMMIT respectively.

More recently, deep learning techniques for labelling false positive streamlines have been proposed (Astolfi et al. 2020; Legarreta et al. 2020). The reason why these techniques are not included in our analysis is two-fold. First, they follow a completely different paradigm. These deep learning approaches rely solely on the geometry of the streamlines, while the TFTs studied in this thesis are based on the idea that streamlines should be re-projected on some reliable map in order to assess their role in the construction of a weighted structural connectome. Secondly, these techniques were proposed when this thesis started to see its end.

In the next sections we are going to define the mathematical foundations of the TFTs studied in this work. The used notation is as close as possible to the one of the original works. In some occasions, further simplifications have been included in order to ease the comprehension of the core elements that characterise each TFT.

4.2.1.1 SIFT2

The SIFT2 model assigns a coefficient to each streamline in such a way that the fiber density computed from the fODFs matches the one obtained from the weighted tractogram. It is implemented as part of the Mrtrix3 suite (Tournier et al. 2019).

The technique is formulated as a regularized least squares problem that fits the reference vector whose entries are the integrals of each lobe of the fODFs as they are showed in Figure 4.2. The fiber density associated to lobe ℓ is denoted FD_ℓ and can be computed beforehand. This will not change during the optimization process. What changes is the *streamline* density, which is denoted as TD_ℓ and defined as follows. Given the L fODF lobes that are present in the image, let $s \in \ell$ denote the streamlines that are associated to the direction of lobe ℓ , then:

$$TD_\ell = \sum_{s \in \ell} |s_\ell| \cdot e^{w_s} \quad (4.1)$$

where $|s_\ell|$ is the length of the segment of the streamline associated to lobe ℓ traversing the voxel and the exponential is employed in order to impose a soft non-negativity constraint on the weighting of each streamline contribution in the formation of the track density. The minimization




Figure 4.2: This fODF has four lobes, two of which follow the superior-inferior direction (blue) while the other two follow the left-right direction (red). The volume of each lobe is what in Equation (4.2) is denoted as FD_ℓ , where ℓ is the lobe index.

problem is then defined as follows:

$$x^* = \underset{x}{\operatorname{argmin}} \sum_{\ell=1}^L \left(FD_{\ell} - \sum_{s \in \ell} |s_{\ell}| \cdot e^x \right)^2 + \Omega(x) \quad (4.2)$$

where Ω is a regularization function. The original paper (R. E. Smith et al. 2015a) proposes the use of regularization terms such as the standard Tikonov regularization and the total variation penalty.

The minimisation problem written in Equation (4.2) can be re-written to match the more common dictionary-based paradigm. Let $FD \in \mathbb{R}^L$ be the vector whose entries are the volumes of each fODF lobe and n the number of streamlines in the tractogram, then define the matrix $S \in \mathbb{R}^{L \times n}$ having one row per lobe and one column per tracked streamline. Each entry $S_{\ell,s}$ encodes the length of the segment $|s_{\ell}|$ that was used in the definition of TD_{ℓ} . If a streamline s does not traverse lobe ℓ , the corresponding entry will be null $S_{\ell,s} = 0$. In this way, the optimisation problem given in Equation (4.2) can be written as follows:

$$x^* = \underset{0 < x \in \mathbb{R}^n}{\operatorname{argmin}} \|Sx - FD\|_2^2 + \Omega(x) \quad (4.3)$$

where Sx denotes the standard matrix-vector product and $x > 0$ accounts for the positivity constraint inherited from the exponential in Equation (4.2).

This SIFT2 model has the advantage of attaching a biological interpretation to the coefficients associated to the streamlines. Since the data that are fitted are the volume integrals (mm^3) of the lobes and the track density is defined as the product of a length (mm) and some weighting factor, such weighting factor has the dimension of an area (mm^2). For this reason, the weighting factors e^{w_s} associated to each streamline can be interpreted as the mean cross sectional area along the streamline. The main limitation of this method is its inability to associate a null weight to a streamline. In the context of connectomics, this means that the method does not mark any streamline as a false positive connection. Nevertheless, the influence of such streamlines can be mitigated by assigning a very small weights to it, limiting its impact in the construction of the weighted connectome.

4.2.1.2 COMMIT2

The COMMIT2 framework is based on a forward model that transforms each streamline into its profile in some dMRI-based map, allowing to define a convex optimization problem that finds the linear combination of streamline profiles that best matches the acquired dMRI signal. The implementation of COMMIT2 is delivered as a Python package (Daducci et al. 2020).

As for SIFT2, the technique is formulated as a regularised least squares problem. The differences lie in what data are used as reference, in the design of the forward model and in the specific choices made in the regularization. As far as the data and the forward model are concerned, there are two main choices that have been made by the authors of the original papers of COMMIT (Daducci et al. 2014) and COMMIT2 (Schiavi et

al. 2020). The first can be categorised under the *microstructure-informed tractography* (MIT) umbrella, while the second follows the trend of SIFT2, hence models the fiber density in each voxel. In both cases, the dMRI signal $y \in \mathbb{R}^{n_d \cdot n_v}$ (n_d is the number of acquired diffusion-weighted images and n_v is the number of voxels) is modelled as $y = Ax + \eta$, where x is the vector of coefficients associated to each streamline, A is a linear operator that maps those coefficients to the space of the reference volume and η accounts for the noise in the measurements. To retrieve the streamline-specific coefficients, the procedure is analogous to the one of SIFT2, hence:

$$x^* = \underset{x \geq 0}{\operatorname{argmin}} \|Ax - y\|_2^2 + \lambda \Omega(x) \quad (4.4)$$

where $\lambda \geq 0$ is a regularisation parameter and Ω is a regularisation term. What is solved is the regularised *non-negative* least squares (rNNLS) problem, which reduces to the NNLS case when $\lambda = 0$. The hard non-negativity constraint is the main difference with the SIFT2 model, as it allows to retrieve streamline coefficients that are actually equal to zero. This marks the corresponding streamlines as candidate false positives without the need to threshold the coefficients, as it was the case for SIFT2. Whenever Ω is a convex lower semi-continuous proper function whose proximal operator can be computed, the problem in Equation (4.4) can be solved efficiently with routines like the Fast Iterative Shrinkage Thresholding Algorithm (FISTA) routine of Beck and Teboulle (Beck et al. 2009).

The definition of the forward operator A and the regularisation term Ω is discussed in the following paragraphs.

Microstructure-Informed Tractography The MIT approach to the definition of the forward operator of COMMIT2 aims at finding the linear combination of streamline profiles that best fits the pure dMRI signal. This idea was first mentioned by Smith et al. in the paper where the first version of SIFT was presented (R. E. Smith et al. 2013), then Daducci et al. (Daducci et al. 2014) developed and implemented it in the first version of COMMIT. The forward model is inspired by the standard MC model of dMRI in WM discussed in Chapter 3. The dMRI signal y is modelled as the combination of three compartments that account for the diffusion in the intra-cellular (IC), the extra-cellular (EC) and CSF compartments. The IC compartment is modelled by the streamlines in the tractogram. For each streamline, the corresponding signal profile is computed by defining a vector $S_s \in \mathbb{R}^{n_d \cdot n_v}$ whose voxel-wise blocks are the response function of corresponding to the diffusivity within a stick aligned with the streamline s with parallel diffusivity equal to $\lambda_{\parallel} = 1.7 \cdot 10^{-9} m^2 s^{-1}$. The taxonomy is borrowed from (Panagiotaki et al. 2012). The response function in each voxel is weighted by the length of the streamline segment traversing the voxel, hence voxels that are not traversed by the streamline are associated to a null response function. All these vectors define a dictionary $A_{IC} \in \mathbb{R}^{(n_d \cdot n_v) \times n_s}$ that generates the dMRI signal as a linear combination of the expected signal profiles of the streamlines in the given tractogram, where n_s is the number of streamlines in the tractogram. A dictionary for the EC compartment is build from the peaks of the fODFs in each voxel. With the same principle of the streamline-specific dictionary, we define $A_{EC} \in \mathbb{R}^{(n_d \cdot n_v) \times n_p}$ as the set of n_p vectors $Z_p \in \mathbb{R}^{n_d \cdot n_v}$ such that the block of coordinates corresponding to the voxel v where peak p is present en-

codes the response function of a zeppelin aligned with the peak with parallel diffusivity equal to $\lambda_{\parallel} = 1.7 \cdot 10^{-9} m^2 s^{-1}$ and perpendicular diffusivity equal to $\lambda_{\perp} = 0.5 \cdot 10^{-9} m^2 s^{-1}$. Finally, a dictionary A_{CSF} for the CSF compartment is defined by including one vector per voxel that encodes the response function of a ball with radial diffusivity $\lambda_r = 3.0 \cdot 10^{-9} m^2 s^{-1}$ within the voxel. Stacking the three dictionaries horizontally yields the linear operator $A = [A_{IC}, A_{EC}, A_{CSF}]$ that characterises the MIT paradigm. The linear operator has dimension $(n_d \cdot n_v) \times (n_s + n_p + n_v)$ as it predicts the full dMRI signal ($n_d \cdot n_v$ rows) by combining linearly the signal profiles of each streamline (n_s in the tractogram), peak (n_p in the volume) and per-voxel ball (n_v voxels). Equation (4.5) gives a global view on the definition of the forward operator A in the MIT paradigm.

$$A = \left[\begin{array}{ccc|ccc|ccc} S_1^1 & \dots & \dots & S_{n_s}^1 & Z_1 & 0 & \dots & 0 & B & 0 & \dots & 0 \\ \vdots & & & \vdots & 0 & \ddots & \ddots & \vdots & 0 & B & \ddots & \vdots \\ \vdots & & & \vdots & \vdots & \ddots & \ddots & 0 & \vdots & \ddots & \ddots & 0 \\ S_1^{n_v} & \dots & \dots & S_{n_s}^{n_v} & 0 & \dots & 0 & Z_{n_p} & 0 & \dots & 0 & B \end{array} \right] \quad (4.5)$$

where S_i^j is the expected signal profile of the segment of streamline i in voxel j , Z_p is the response function of the p -th peak in the volume and B is the response function of the ball. The IC block of matrix A has n_s columns, the EC block has n_p columns and the CSF block has n_v columns. The entries of the matrix are a-dimensional, as they correspond to an attenuation of the dMRI signal.

Ideally every possible MC microstructure model can be encapsulated in the MIT paradigm. The one presented in the previous lines is intended to be a general example from which one can tailor the needed specific instance while respecting the MC paradigm and the block-matrix definition. For instance, the LiFE framework is designed to include only the IC component and fit the demeaned version of y . An attempt to use the multi-TE multi compartment models described in Chapter 3 into the MIT framework has been recently presented by Barakovic et al. 2020.

Fiber Density A different approach to the definition of the forward operator of COMMIT2 involves the use of a *fiber density* model. From the methodological point of view, this approach is a simplification of the forward model of SIFT and SIFT2. Both of them aim at fitting the fiber density in each voxel. In SIFT2 this is further refined by considering separately the fiber density in each fODF lobe, while in COMMIT2 the fitted vector is the $y \in \mathbb{R}^{n_v}$ encoding the IC volume fraction in each voxel. This can be obtained with some third-part technique such as the MT-MC model presented in Chapter 3, NODDI (H. Zhang et al. 2012) or SMT (Kaden et al. 2016). In dictionary A each streamline is encoded as a column vector whose entry corresponding to voxel v is the length of the streamline segment traversing that voxel. Using the same notation as for the SIFT2 model we define $A \in \mathbb{R}^{n_v \times n_s}$ as the matrix whose entry (v, s) encodes the length of the segment $|s_v|$ of streamline s in voxel v . That this quantity is measured in mm , meaning that the streamline coefficients x_s with which the matrix is multiplied have to be measured in mm^2 in order to balance the dimensional equation against the volumetric (mm^3) data that are fitted. As for the SIFT2 model, these coefficients can be interpreted as the mean cross-sectional area of the fiber population

represented by the corresponding streamline.

Regularisation The regularised NNLS problem refined in Equation (4.4) includes the use of a regularisation term Ω . The design of this function shapes the solution of the minimisation problem. Notice that the hard non-negativity constraint can be seen as a regularisation per-se, as the constrained optimisation problem of Equation (4.4) can be written as a minimisation problem as follows:

$$x^* = \underset{x}{\operatorname{argmin}} \|Ax - y\|_2^2 + \lambda\Omega(x) + \iota_{\geq 0}(x) \quad (4.6)$$

where $\iota_{\geq 0}(x) : \mathbb{R}^n \rightarrow \{0, +\infty\}$ is the indicator function (in the sense of convex analysis) that takes value $\iota_{\geq 0}(x) = 0$ if all the entries of x respect the condition $x \geq 0$, otherwise it takes value $+\infty$. While in SIFT2 Ω is set to account for a regularisation of the solution that reduces the ill-posedness of the minimisation problem, COMMIT2 uses sparsity-inducing regularisation functions. In the paper where COMMIT was introduced (Daducci et al. 2014), the used penalty term was $\Omega(x) = \|x\|_{\ell_1}$, which is the ℓ_1 norm of the coefficients. This regularisation promotes sparsity in the space of streamlines, as it is designed to retain the minimum number of streamlines while explaining the data by means of the employed forward model. In this case, Equation (4.6) is called the *least absolute shrinkage and selection operator* (LASSO) problem.

To actually account for the minimal number of retained variables in a regularised least squares problem, one would have to employ the ℓ_0 norm as a penalty term:

$$\ell_0(x) = |x_0|^0 + |x_1|^0 + \dots + |x_n|^0 \quad (4.7)$$

which under the assumption that $0^0 = 1$ returns the number of entries of x that have non-zero value. This function is not convex. The ℓ_1 norm is a more suitable choice since it acts as a convex relaxation of the ℓ_0 norm, allowing to keep the optimisation problem of Equation (4.6) convex.

More recent works (Ocampo-Pineda et al. 2021; Schiavi et al. 2020) aimed at imposing sparsity in a structured fashion, hence to promote solutions where the least number of groups (a.k.a. bundles) of streamlines is used for explaining the data through the forward model. This is motivated by the observation that streamlines represent populations of neuronal fibers which are in turn organised in fascicles that connect the same pair of cortical regions (Schiavi et al. 2020). The mathematical formulation of such regularisation term was originally proposed by Jenatton et al. (Jenatton et al. 2010, 2011). Given a group structure \mathcal{G} and a real positive weight w_g associated to each group $g \in \mathcal{G}$, the structured sparsity penalty term reads as follows:

$$\Omega(x) = \sum_{g \in \mathcal{G}} w_g \|x_g\|_2 \quad (4.8)$$

where x_g is the restriction of vector x to those entries belonging to group g . When the groups in \mathcal{G} are disjoint, Ω defines the so-called *group sparsity* regularisation (obtaining the *group LASSO* problem), while when the groups overlap in a hierarchical fashion, the regularisation is called *hierarchical sparsity*. When there is no structure in the overlapping of the groups, the penalty is then called *overlapping group*. The coefficients w_g

associated to each streamline bundle can be designed to encapsulate any sort of normalisation and prior information. The simplest definition of w_g normalises for the different number of streamlines in each bundle and is defined as $w_g = 1/\sqrt{|g|}$, where $|g|$ is the number of streamlines in group g . This normalisation prevents from promoting a bundle to be zeroed-out just because it involves a larger number of streamlines. A more sophisticated choice gives higher weights to bundles whose streamlines have been associated to low weights in the NNLS formulation of the same problem (hence for $\lambda = 0$). This choice is inspired by the work of Candés et al. (Candés et al. 2008), where the reweighted ℓ_1 approach was first introduced. This is achieved by $w_g = 1/\|\hat{x}_g\|_2$, where \hat{x} is the solution of the NNLS problem. In COMMIT2, the groups weights are designed as $w_g = \sqrt{|g|}/\|\hat{x}_g\|_2$ (Schiavi et al. 2020).

The sparsity paradigm have been explored also in the LiFE framework in the specific form of the ℓ_1 regularisation (Caiafa et al. 2017).

4.2.2 Unified Tractography Filtering Framework

As shown in the previous sections, the presented TFTs show a common structure that reduces tractography filtering to four fundamental elements.

Data y to be fitted:

- ▶ SIFT2: fODF lobe-specific fiber density.
- ▶ COMMIT2: voxel-wise fiber density or raw dMRI signal.

Forward model mapping streamlines onto data through a forward operator A :

- ▶ SIFT2: length of the streamline segment associated to the fODF lobes that it traverses.
- ▶ COMMIT2: length of the streamline segment traversing each voxel or MC microstructure model.

Regularisation term promoting solutions having a property that minimises $\lambda\Omega$:

- ▶ SIFT2: Tikonov and total variation.
- ▶ COMMIT2: sparsity-promoting regularisation possibly organised in group or hierarchical structure.

Constraints encoded through the use of indicator functions ι (in the sense of convex analysis):

- ▶ SIFT2: soft non-negativity constraint on the coefficients that does never associate null weights to streamlines.
- ▶ COMMIT2: hard non-negativity constraint that allows streamlines to be associated to null weights, hence to be marked as candidate false positives.

These four elements are all part of the convex optimisation problem already stated in Equation (4.6) and hereafter reported.

$$x^* = \underset{x}{\operatorname{argmin}} \underbrace{\|Ax - y\|_2^2}_{\text{data fitting}} + \underbrace{\lambda\Omega(x)}_{\text{regularisation}} + \underbrace{\iota(x)}_{\text{constraint}}. \quad (4.9)$$

This problem is the core operation that retrieves the streamline coefficients x_s which are then used for defining the corresponding weighted structural connectome as prescribed in Chapter 2. We call the connectomes obtained using TFTs *filtered connectomes*. The topology of these networks is supposed to be different from the one of the streamline-count connectomes, but the literature lacked a systematic evaluation of the effects of TFTs on the graph-theoretical analysis of structural connectomes. As we are going to show in Section 4.3.1, we presented the aforementioned analysis in a recent article (Frigo et al. 2020b), showing how TFTs indeed change the topology of structural connectomes.

Not all the combinations of the four fundamental factors that define the TFTs are implemented in publicly available software (Daducci et al. 2020; Pestilli et al. 2014; Tournier et al. 2019). One of the outcomes of this thesis is a general implementation of all these models in the *Tractograms As Linear Operators in Neuroimaging* (TALON) Python package, whose description is present in Appendix A.

The presented TFTs rely uniquely on the dMRI signal and its transformations, hence to data that carry only information about the *structure* of the brain network. In the next sections we are going to present how *functional* information can be included in the tractography filtering process.

4.2.2.1 Computational aspects

As we showed in the previous paragraphs, the design of a TFT experiment reduces to four choices: fitted data, forward model, regularisation term and constraints. From the neuroscientific point of view, these choices affect the interpretability of the obtained streamline coefficients and the assumptions and limitations carried by the model. One side that has been less discussed in previous works concerns the computational aspects of tractography filtering. A typical connectomic pipeline involves the tracking of several million streamlines. This means that direct methods for the solution of (4.9) are not usable, as the inversion of matrix A is not feasible on reasonable hardware. Fortunately, SIFT2, COMMIT and TALON are all capable of treating the forward operator A in a sparse fashion, hence significantly reducing the necessary resources. Nevertheless, the gap between sparse methods and direct methods for the solution of minimisation problems such as the one in Equation (4.9) remains important.

The standard algorithm used for solving the problem stated in Equation (4.9) is FISTA, which is an accelerated proximal gradient descent. This method is a generalisation of the classic gradient descent that allows to consider non-smooth regularisation terms. At each iteration, it requires evaluations of the fitting and regularisation terms, of the gradient of the fitting term, of the proximal operator of the regularisation term (i.e., the non-smooth generalisation of its gradient) and of the projection onto the constraining set. Convergence is granted any time the regularisation term is convex, proper and lower semi-continuous and the constraining set is convex, which is the case for the non-negativity constraint (Beck et al. 2009).

The backtracking strategy used for adapting the step size (a.k.a. learning rate) during the optimisation process requires multiple evaluations of the

fitting term $1/2 \|Ax - y\|_2^2$ and its gradient $A^T(Ax - y)$ between two iterations. From the computational point of view, the large part of the burden is carried by the Ax and $A^T Ay$ operations, which are matrix-vector products involving the forward operator A . While this can be optimised in terms of memory consumption, two main concerns on the computational complexity have to be considered.

- ▶ A higher number of streamlines n increases the complexity of computing the matrix-vector product Ax , as the dimension of the domain where x lies increases together with n .
- ▶ The forward operator can map streamline coefficients to voxel-specific scalar or vector data. This choice changes the size of the codomain of A , which affects any matrix-vector product $A^T y$, which is necessary for the computation of the gradient of the data fitting term. In practice, a MIT forward model that directly describes the dMRI signal (vector valued in each voxel) implies a much heavier computational complexity than e.g. the volume fraction model, which models one scalar per voxel. On the other side, modelling maps other than the pure dMRI signal requires to explicitly compute those maps, which could itself require heavy computations or further acquisitions.

Two further factors act as obstacles in the optimisation process, namely the design of the regularisation term and the constraints to be applied to the streamline coefficients. As far as the regularisation term is concerned, we are going to limit our considerations to the case where it depends solely on the streamline coefficients x . What we are going to say may be not true in general, but it will be true in most of the cases to which machine learning have accustomed us, like the sparsity inducing norms, the Tikonov regularisation and the total variation regularisation (Sra et al. 2011). Except few notable smooth exceptions like the aforementioned Tikonov regularisation ($\Omega(x) = \|x\|_2^2$), in most cases the regularisation term is non-smooth, convex, proper and lower semi-continuous. This means that the computation of its gradient is not possible, hence it requires the use of the concepts of subgradient and proximal operator. The former is the non-smooth equivalent of the gradient, while the latter provides the mathematical tool that grants the ability to perform the proximal gradient descent. The definition of the proximal operator reads as follows[‡]:

$$\text{prox}_\Omega(v) = \underset{x}{\operatorname{argmin}} \frac{1}{2} \Omega(x) + \|x - v\|_2^2 \quad (4.10)$$

which makes straightforwardly clear that an additional layer of complexity is added to the standard gradient descent. As a matter of fact, the gradient descent is based solely on the evaluation of the gradient, while the proximal iteration requires the evaluation of the proximal operator, which in turn requires to solve a minimisation problem. Whenever the regularisation term is an ℓ_p norm, such operation reduces to the projection onto an appropriate convex set, which in the common cases when p is equal to $p = 1$ or $p = 2$ is a $\mathcal{O}(n)$ operation (Sra et al. 2011), where n is the number of streamlines. The degree of sophistication induced by definitions like the group lasso or the hierarchical sparsity terms mentioned in Section 4.2.1.2 strongly depends on the geometry of the encoded prior. For instance, the evaluation of the proximal operator of the group LASSO

[‡] The interested reader should refer to the thorough presentation of Sra et al. 2011.

regularisation has complexity linear in the number of groups, while for the hierarchical sparsity it becomes $\mathcal{O}(d \cdot n)$, where n is the number of coefficients (i.e., streamlines) and d is the depth of the considered hierarchical structure. In brief, the design of the regularisation term brings an additional layer of complexity that can not be omitted in the global economy of TFT methods.

A final note is devoted to the use of constraints defined through indicator functions $\iota_C(x)$, where C is the convex set of the acceptable solutions. Again, the computational burden reduces to the projection onto C at each iteration of the optimisation process. In the case of the hard non-negativity constraint employed by COMMIT, B is the first n -dimensional orthant and the corresponding projection is the coordinate-wise projection onto the positive half line. More sophisticated constraint may not be treatable as simply.

With the previous paragraphs we gave a simplified picture of the computational bottlenecks that affect the solution of the optimisation problem of Equation (4.9). The combination of choices that yields the simplest (from the computational point of view) optimisation problem is the one employed by SIFT2, which avoids the use of hard constraints and maps streamline coefficients in a voxel-wise vector-valued map of relatively low dimension, as highlighted in the previous section. We highlight how these choices are very similar to the ones that can be employed in COMMIT with the volume fraction model, modulo the hard non-negativity constraint. In practice, the difference in terms of performance should not be important. Nevertheless, the SIFT2 framework is known to be the fastest among the available TFTs. This reputation is partially due to its implementation, which not only is in C++ (against the Python implementations of the competitors) but is also highly engineered[§]. This difference becomes relevant when large cohorts are studied.

Finding the appropriate balance between a sufficient number of streamline, an informative forward model, meaningful regularisation terms and reasonable constraints is a challenging task whose solution is sometimes driven by exogenous constraints such as the available hardware or the reliability of the data.

4.2.3 Functionally Informed Tractography Filtering

As highlighted in the previous section, SIFT2 and COMMIT2 are purely based on the structural information obtained from dMRI. One of our main contributions presented in this thesis is the development of a TFT that uses functional information to drive the fitting process of structure-based TFTs like the ones summarised in Section 4.2.2. The procedure we are proposing is called *Functionally Informed TFT* (FIT) and it is a general framework that builds on top of any forward model among the ones employed by SIFT2 or COMMIT2, as the functional information is included in the design of the regularisation term. We presented preliminary works concerning this formulation at the 24th meeting of the Organisation for Human Brain Mapping (Frigo et al. 2018a) and at the Computational Diffusion MRI workshop of MICCAI 2018 (Frigo et al. 2018b).

[§] The interested reader should refer to the sources available at the Github repository of Mrtrix3 (Tournier et al. 2019).

The prior knowledge we want to include in the formulation of FIT is comes from two observations:

1. Axons are organised in coherent fascicles that connect pairs of cortical regions, so should be streamline bundles.
2. Significant correlation has been highlighted between axonal connectivity and functional interaction (Hagmann et al. 2008), suggesting that streamline bundles connecting brain regions that show higher functional connectivity are more likely to represent true axonal connections with respect to bundles connecting regions with low functional connectivity.

As far as the first statement is concerned, the concept of non-overlapping group sparsity defined in Section 4.2.1.2 provides the appropriate tools. The fact that it drives the optimisation to retain the fewest groups of variables that explain the data, perfectly fits the concept of *sparsity in the space of streamline bundles* that we described in the first listed point. The design of such regularisation term requires the definition of a weight associated to each streamline bundle. The higher the weight is, the more the optimisation process will be driven towards solutions where the streamlines belonging to the group have null weights. This comes in handy for including the second prior, which concerned the establishment of a lower penalisation for streamline bundles connecting regions with high functional correlation. Let \mathcal{G} be the group structure defined in Section 4.2.1.2 and let $g \in \mathcal{G}$ be a streamline bundle that connects cortical regions i and j . Let then $c_{ij} = c_g \in [0, 1]$ be the functional connectivity between regions i and j , then the wanted regularisation term can be defined as follows:

$$\Omega(x) = \sum_{g \in \mathcal{G}} \frac{1}{\sqrt{|g|} (1 + c_g)} \|x_g\|_2 \quad (4.11)$$

which corresponds to the particular case of Equation (4.8) where $w_g = 1/(\sqrt{|g|} (1 + c_g))$. The regularisation term defined in Equation (4.11) includes both a normalisation by the number of streamlines in each group and a function-dependent component. The $1/(1 + c_g)$ factor that multiplies each term in the sum takes value in the $[1/2, 1]$ interval. In particular, the weight associated to streamline bundles connecting pairs of regions with low functional connectivity ($c_g \sim 0$) is up to twice the weight of bundles of streamlines connecting regions with high functional connectivity ($c_g \sim 1$). This proves that the designed penalisation term respects the second prior information that we wanted to include.

The definition of the coefficient c_g that encodes the functional connectivity between the regions connected by streamline bundle g can be undertaken in different ways. Among the techniques that allow to estimate the functional connectivity between regions, resting-state functional MRI (rs-fMRI) has showed unprecedented robustness for the study of the functional activity in the human brain (Van Den Heuvel et al. 2010). Nothing prevents from considering other techniques such as M/EEG to define the c_g coefficient, but in this study we are going to restrict the analysis to rs-fMRI. We remark that c_g must indicate if and how much the streamlines connecting regions $(i, j) = g$ represent a connection that is observed also from the functional point of view. One possible solution to the problem of estimating c_g is to employ a *maximum sliding window correlation* strategy, which corresponds to taking the maximal correlation observed between

the two regions in the dynamic functional connectome (dFC) obtained through sliding window correlation (Preti et al. 2017). This choice provides a highly sensitive estimator of the functional connectivity. Another solution comes from the use of static (as opposed to dynamic) functional connectivity (sFC) index (Van Den Heuvel et al. 2010). The first possibility requires the design of a strategy for the computation of the dFC, increases the complexity of the problem in a way that the second strategy simply avoids at the price of a lower sensitivity. Notice that if a streamline bundle connects two regions that show low functional correlation it will not be automatically marked as a false positive. In fact, the presence of the fitting term relying on the dMRI data in the optimisation problem of Equation (4.6) guarantees that streamline bundles that are required for the explanation of the reference data (be it the dMRI signal or the IC volume fractions) are retained. For this reason, the higher complexity of the dFC strategy might not be worth its higher specificity. In this work we are going to consider the definition of c_g coming obtained from the sFC strategy, which corresponds to the one given in Chapter 2.

In Section 4.2.3 we are going to present some unpublished preliminary results obtained with FIT on both synthetic and real data, showing how the inclusion of functional priors in the filtering process drives the optimisation.

4.3 Experiments

4.3.1 TFTs Change the Topology of Structural Connectomes

With the aim of assessing how connectomic studies are affected by the use of TFTs, in this experiment we investigate how four graph-theoretical metrics (GTM) of network integration and segregation change with respect to the employed TFT and to density based thresholding. In particular we consider the Characteristic Path Length (CPL), the Global Efficiency (GE), the global Clustering Coefficient (CC), and the Louvain Modularity (MO). These metrics are calculated on connectomes obtained via streamline count (SC), SIFT2 and COMMIT¹ and a detailed presentation of why they are representative of the topology of the network is given in Chapter 2. For each type of filtered connectome we also investigate the robustness to density-based thresholding of the considered graph-theoretical measures, which is the act of excluding all the weakest connections in the network until a structural connectome of the wanted density is obtained. A recent study by Civier et al. 2019 partially investigated the latter aspect, but the focus was specifically on the density-based thresholding aspect of connectomic pipelines. Density-based thresholding is the act of excluding all the weakest connections in the network until a structural connectome of the wanted density is obtained (Civier et al. 2019). This practice was proved to reduce the number of false positive connections in structural networks (Sarwar et al. 2019), but its use was showed to yield no detectable effect on the topology of the connectome when coupled with the use of SIFT2 (Civier et al. 2019). As our experiment focuses on the TFT step, our analysis is an extension of the one of (Civier et al. 2019) to the case where COMMIT is employed.

1: At the time when this work was prepared, COMMIT2 had not been published yet.

When identical connectomic pipelines are employed, changes in brain network topology can be due to the presence of pathology (Verma et al. 2018, 2019) like traumatic brain injury (TBI), which is a network disorder that exhibits (among others) changes in the small-worldness of the brain network (Pandit et al. 2013). Moreover, subjects affected by TBI show an increase of the mean diffusivity coupled with a decreased fractional anisotropy (Mohammadian et al. 2017), 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 cases.

The analyses of the effects of TFTs and density-based thresholding are performed on connectomes obtained on the 100 unrelated healthy subjects from the Human Connectome Project (HCP) database (Matthew F Glasser et al. 2013). 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 extend our 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. (C.-H. Yeh et al. 2016), 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 and the recently growing interest towards the field of patho-connectomics (Verma et al. 2018), this experiment represents the first systematic comparison of the effects of TFTs on the graph-theoretical analysis of structural brain networks estimated in-vivo.

This analysis was published in the Journal of Neural Engineering (Frigo et al. 2020b) and 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), the NIH R01NS065980 grant, the PA-Department of Health PACT award and the NIH R01NS096606 grant.

4.3.1.1 Data and Methods

Data: HCP subjects From the HCP database we considered the list of 100 unrelated subjects (U100 group) dataset available at the Connectome Coordination Facility (WU-Minn Human Connectome Project consortium 2017; Sotiropoulos et al. 2013; Van Essen et al. 2012). 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* (Jesper LR Andersson et al. 2003; Jenkinson et al. 2012) and eddy current and subject motion correction via FSL's *eddy* (J. Andersson et al. 2012; Sotiropoulos et al. 2013). 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. (Matthew F Glasser et al. 2013). 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.

Data: 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 (Rabinowitz et al. 2018). 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, both on the same 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 $1 \times 1 \times 1mm^3$. 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 (Manjón et al. 2013), followed by brain extraction with FSL's *BET* tool (S. M. Smith 2002) on the first b_0 image. The denoised data and brain mask were input to FSL's *eddy* to correct the data for motion and eddy-current distortion (Jesper LR Andersson et al. 2016). 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 b_0 image.

Connectomic pipeline A five-tissue-type image (R. E. Smith et al. 2012) of each subject was obtained with the Freesurfer pipeline (Fischl 2012) implemented in Mrtrix3 (Tournier et al. 2019). 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) (Jeurissen et al. 2014) and the fODFs were computed using the MSMT-CSD algorithm (Jeurissen et al. 2014; Tournier et al. 2004) with a maximal spherical harmonics order of $\ell = 8$.
- ▶ Clinical: the unsupervised algorithm of Dhollander et al. 2016b was used for estimating the white matter response function, while fODFs were computed using the CSD algorithm (Tournier et al. 2007) with a maximal spherical harmonics order of $\ell = 6$.

The obtained fODFs were used for probabilistic anatomically-constrained tractography (ACT) with the iFOD2 algorithm (R. E. Smith et al. 2012). 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. 2006 via Freesurfer (Fischl 2012).

Tractography filtering was then performed on both datasets via SIFT2 and COMMIT. The forward model of COMMIT was defined with the MIT paradigm. 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 with radial diffusivity $1.7 \cdot 10^{-3} mm^2/s$ and $3 \cdot 10^{-3} mm^2/s$ respectively. The use of a multi-compartment model with the single-shell data of the clinical dataset is allowed by the fact that the only parameters that are not fixed are the signal fractions. A weighted connectome was computed for each employed TFT and we also computed the streamline count (SC) connectome associated to each tractogram. Figure 4.3 gives a graphical overview of the connectomic pipeline employed in this work.

The GE, CPL, CC and MO are defined as presented in Chapter 2 and computed for each obtained connectome with the Python implementation of the brain connectivity toolbox (Rubinov et al. 2010). In particular, modularity was computed by averaging the results of 100 runs of the Louvain algorithm (Blondel et al. 2008). 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.

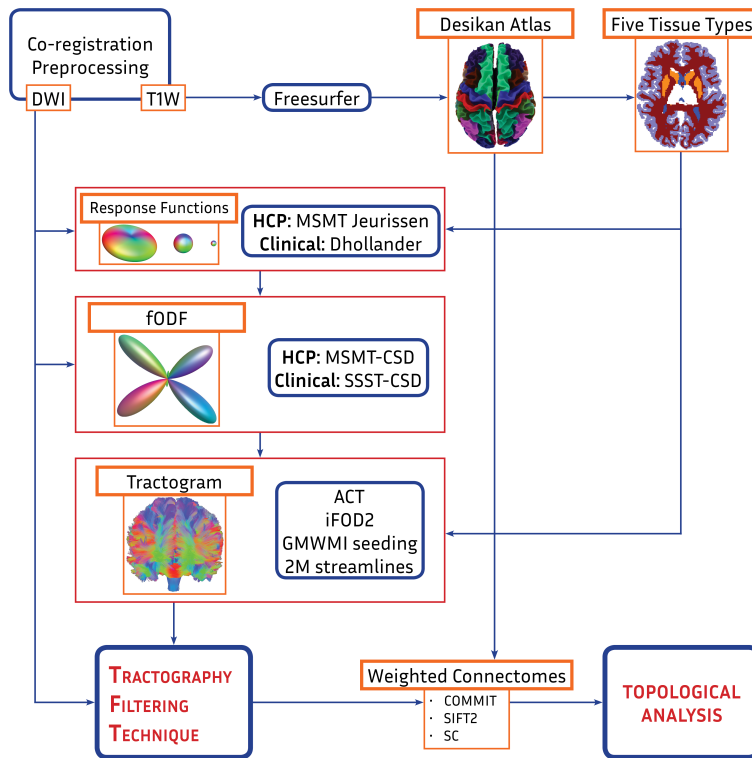


Figure 4.3: The represented connectomic pipeline is the one employed throughout this work. The first step is the pre-processing and the co-registration of the T1-weighted (T1w) and the diffusion-weighted magnetic resonance images (DWI). The T1w image was used to obtain the Desikan atlas and the five-tissues-type 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.

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 (Mann et al. 1947) 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 (Benjamini et al. 1995). 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 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 ef-

fect size measure. FDR correction was performed with the Benjamini-Hochberg procedure.

4.3.1.2 Results

We processed the data presented in the previous section 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. Figure 4.4 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 ($t = 64.02$, $p \leq 0.05$), 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 8% 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




Figure 4.4: 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.

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 of (Frigo et al. 2020b) 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 4.5 show how the set of edges exhibiting statistically significant differences between TBI patients and healthy controls are different for every TFT.




Figure 4.5: 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.

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 of (Frigo et al. 2020b). The effect of TFTs on these metrics is then presented in Figure 4.6, 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




Figure 4.6: 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) for HCP subjects is reported to be 0.87, 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.

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 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 4.7 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 (Civier et al. 2019) 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 between the de-

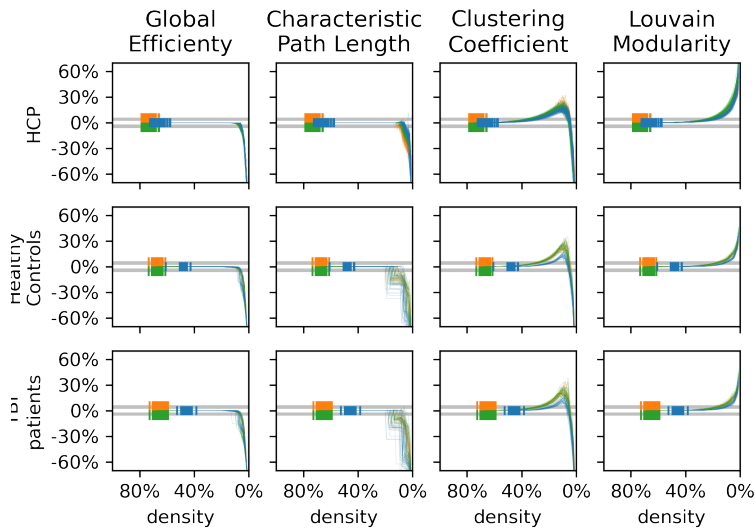


Figure 4.7: Each panel concerns the analysis of the robustness to density-based thresholding of filtered connectomes obtained from one subject cohort. The plots show the relative change of the graph-theoretical measure computed at a specific density with respect to the one computed at base density. The 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.

viation 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. Table 4.1 reports 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 4.1 shows that integration measures (GE and CPL) are more robust to density-based

		SC		COMMIT		SIFT2	
		>1%	>5%	>1%	>5%	>1%	>5%
HCP	GE	9%	7%	9%	6%	11%	7%
	CPL	11%	9%	11%	8%	13%	13%
	MO	43%	25%	38%	21%	41%	22%
	CC	57%	36%	51%	32%	54%	35%
C-HC	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%
TBI	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 4.1: Maximal density across all subjects in the HCP, clinical healthy controls (C-HC) and TBI datasets 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 Civier et al. 2019.

thresholding if investigated on connectomes built via COMMIT. Segregation measures show comparable performances among the three filtering techniques.

4.3.1.3 Discussion

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, 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 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 4.4 it is possible to notice how COMMIT lowers the density of the connectomes in both the high resolution and the clinical resolution datasets. Moreover, the shift of density is different in the two datasets.

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 4.6 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. (Civier et al. 2019) 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 4.7 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 (Civier et al. 2019), 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.

4.3.2 Functionally-informed TFT

In this experiment we tested the validity of the FIT formulation proposed in Section 4.2.3. The lack of a golden standard of the structural connectivity of the human brain poses a severe challenge to the possibility to provide an unequivocal answer to the question whether a TFT² is valid or not. Moreover, Zalesky et al. 2016 recently showed how the sensitivity and specificity of structural connectomes affect the graph-theoretical analysis of brain networks, highlighting how, in specific settings, the presence of a false positive connection can be twice as detrimental as having a false negative. This consideration acquires an alarming tone when put next to the observations of Thomas et al. 2014 and Maier-Hein et al. 2017, who verified the presence of a non-trivial portion of false positive connections in dMRI tractography-based structural connectomes.

2: The same challenges are posed by the validation of tractography algorithms.

To assess if and how FIT improves the quality of the retrieved structural connectomes, we tested it on synthetic datasets that allow to observe its effects on the sensitivity and the specificity of the retrieved networks in a tractometric fashion (Côté et al. 2013). Moreover, preliminary insights are obtained on the real data of a healthy subject from the HCP database

4.3.2.1 Data and methods

Data Two synthetic datasets have been employed in the experiment. Both of them aim at assessing if FIT is able to distinguish false positive (FP) and true positive (TP) connections. The employed synthetic datasets are designed in such a way that they allow to analyse the filtered tractograms in terms of true positive (TP) and false positive (FP) connections (a.k.a. in a *tractometric* fashion, see Côté et al. 2013).

In particular, the focus is on how the presence of a prior knowledge on the connectivity between the connected regions affects the resulting connectome. Preliminary results are obtained also on a healthy subject from the HCP database.

Synthetic Dataset – Crossing/Kissing

The first employed phantom describes crossing and kissing bundles configurations and mimics the one employed in the work of Girard et al. 2017. As shown in Figure 4.1, the tractograms obtained from the two configurations have all the four possible streamline bundles, when they should have only two. The dMRI signal was simulated using Phantomas (Caruyer et al. 2014) for a geometry of 94x94x94 voxels of (2.0mm)³ size. A single-shell acquisition scheme was employed including one $b = 0s/mm^2$ image and 64 diffusion weighted images uniformly distributed on a b -shell Caruyer et al. 2013 at $b = 1500s/mm^2$. Fiber ODFs and probabilistic anatomically constrained tractography were computed with Dipy (Garyfallidis et al.

2014), obtaining a total of 50'000 streamlines per phantom (one for the crossing and one for the kissing configuration). This dataset is the one employed in our preliminary work (Frigo et al. 2018b).

Synthetic Dataset – ISBI 2013

The second phantom is the well known three-dimensional dataset employed in the second HARDI reconstruction challenge (ISBI2013). This phantom represents a step forward with respect to the simplistic phantom presented in the previous paragraph in the sense that it includes the complexity of several major white matter tracts of the human brain that go beyond the simple crossing/kissing ambiguity. The dMRI signal was simulated using Phantomas (Caruyer et al. 2014) for a geometry of $88 \times 88 \times 88$ voxels of $(1.25\text{mm})^3$ size with $SNR = \infty$. The same acquisition scheme of the HCP dataset was employed, which includes 288 diffusion-weighted images subdivided in 18 volumes at $b = 0\text{s/mm}^2$ and 90 diffusion-weighted volumes obtained at uniformly distributed directions at $b = 1000\text{s/mm}^2$, $b = 2000\text{s/mm}^2$ and $b = 3000\text{s/mm}^2$ for a total of 3 shells. A tailored five-tissue-type image (Jeurissen et al. 2014) was build in-house and used for running the standard tractography pipeline presented in Chapter 2, which includes an estimation of the fODFs via MSMT-CSD (Jeurissen et al. 2014) and a second order integration of the fODFs with a probabilistic anatomically constrained tractography algorithm (R. E. Smith et al. 2012) performed via Mrtrix3 (Tournier et al. 2019) seeding from the white matter mask.

Both the synthetic datasets do not have native functional prior knowledge associated to the dMRI data. Simulating a functional network using neural mass models (Sanz Leon et al. 2013) or structure-function maps (Deslauriers-Gauthier et al. 2020) would mean including an additional modeling step that could affect the analysis. As the use of functional connectivity information is the core of FIT, we opted for using the ground-truth structural connectivity matrix as a functional prior, given the strong relationship between the structural and the functional brain network (Honey et al. 2007).

HCP subject

The real dataset used for the analysis of the proposed FIT technique is represented by subject 100307 of the HCP database. For this subject, both dMRI and resting-state fMRI data are available. In both cases, we employed the pre-processed data available at the Connectome Coordination Facility (Matthew F Glasser et al. 2013).

The employed cortical atlas is the Desikan parcellation (Desikan et al. 2006), which is directly available on the HCP database.

The dMRI data include 288 images subdivided in 18 volumes at $b = 0\text{s/mm}^2$ and 90 diffusion-weighted volumes obtained at uniformly distributed directions at $b = 1000\text{s/mm}^2$, $b = 2000\text{s/mm}^2$ and $b = 3000\text{s/mm}^2$ for a total of 3 shells. A five-tissue-type image (R. E. Smith et al. 2012) of each subject was obtained with the Freesurfer pipeline (Fischl 2012) implemented in Mrtrix3 (Tournier et al. 2019). 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) (Jeurissen et al. 2014) and the fODFs were computed using the MSMT-CSD algorithm (Jeurissen et al. 2014; Tournier et al. 2004) with a maximal spherical harmonics order

of $\ell = 8$. The obtained fODFs were used for probabilistic anatomically-constrained tractography (ACT) with the iFOD2 algorithm (R. E. Smith et al. 2012). The seeding was performed from the gray matter - white matter interface (GMWMI) and a total of 2 millions of streamlines was obtained.

The resting state fMRI data used in this experiment are the ones from the REST1 session, which includes two acquisitions. The time series were filtered using a Butterworth bandpass filter with critical frequencies of 0.01 and 0.1 Hz. We discarded the first 20 volumes of each acquisition and we detrended linearly the remaining voxels. The movement parameters and their derivatives were regressed out and the data was motion scrubbed (Power et al. 2012). At this point both acquisitions were concatenated to produce a single dataset. The mean signal for each parcel was computed and used to build the functional connectome. Each entry (i, j) of its adjacency matrix was computed as the Pearson correlation between the time series associated to regions i and j . Finally all negative entries of the matrix, corresponding to anticorrelations between cortical regions, were set to zero.

Filtering Pipeline As explained in Section 4.2.2, the design of a tractography filtering pipeline requires four choices: the fitted data, the forward model, the regularisation term and the constraints. The choices employed in our experiments are motivated mostly by the limitations posed by the considered datasets and the computational complexity of the resulting optimisation problem.

For all dataset, a hard non-negativity constraint $t_{\geq 0}(x)$ was included. Also, the regularisation term was designed according to the FIT formulation proposed in Section 4.2.3 and hereafter reported:

$$\Omega(x) = \lambda \sum_{g \in \mathcal{G}} w_g \|x_g\|_2, \quad w_g = \frac{1}{\sqrt{|g|} (1 + \alpha \cdot c_g)} \quad (4.12)$$

where λ is the regularisation parameter and $\alpha \in \{0, 1\}$ controls the use of the functional information encoded in c_g . For $\alpha = 0$, c_g does not have any effect on the regularisation and one retrieves a group sparsity (GS) formulation that is a close variation of the one used in COMMIT2. For $\alpha = 1$ the functional information is included and gives less penalty to solutions that contain bundles connecting regions with higher functional correlation. The optimal λ is selected by drawing the L-curve and selecting the value that yields the solution closest to the origin after having normalised the two axes (P. C. Hansen 1992).

The two elements that are still to be determined are the forward model and the fitted data. In the following paragraphs they are going to be specified for each dataset.

Synthetic Dataset: crossing/kissing

The forward model was designed to fit the raw dMRI data in a MIT fashion by means of a stick-zeppelin-ball MC model implemented in COMMIT (Daducci et al. 2014). Sticks model the intra-cellular diffusivity of streamlines. Finally, one zeppelin per fODF peak and one ball per voxel are included. The diffusion parameters are set as specified in Section 4.2.1.2.

As we reported in Frigo et al. 2018b, a slight variation of the bundle-specific weight w_g is employed, namely

$$w_g = \frac{1}{\sqrt{|g|}} \left(\frac{1}{c_g^2} \exp\left(\frac{1 - c_g^2}{2s^2}\right) - 1 \right). \quad (4.13)$$

where s is a shape parameter. This formulation was an attempt to introduce a non-linearity in the control of the influence of c_g in the model.

Synthetic Dataset: ISBI 2013

The forward model was designed to fit the raw dMRI data in a MIT fashion by means of a stick-zeppelin-ball MC model implemented in TALON. Sticks model the intra-cellular diffusivity of streamlines. Finally, one zeppelin per fODF peak and one ball per voxel are included. The diffusion parameters are set as specified in Section 4.2.1.2.

HCP subject

The fitted data is the intra cellular volume fraction computed with the MT-MC model presented in Chapter 3. The employed volume fraction map is the same as the one that was analysed in the corresponding results presented in Section 3.3.2. We refer to that section for further details on the computation of such map. The forward operator is defined with the fiber density model described in Section 4.2.

4.3.2.2 Results

Synthetic Dataset – crossing/kissing We report here the same results presented in (Frigo et al. 2018b). The optimisation was performed with the solver provided in the COMMIT package. The objective of this experiment was to assess the ability of FIT to distinguish the crossing and kissing configurations by using the prior knowledge on the functional connectivity between the connected regions. Figures 4.8 and 4.9 show how the crossing/kissing fibers ambiguity is solved by FIT in the context of this simplistic phantom dataset.

Synthetic Dataset – ISBI2013 The ISBI 2013 phantom dataset mimics several complex fiber configurations that are present in the human brain. The optimisation was performed using the solver provided by the TALON package. The regularisation parameter was defined as follows. In an attempt to balance the weights of the fitting term and the regularisation term, we set

$$\lambda = \sigma \cdot \|A^T y\|_2 \max_{g \in \mathcal{G}} \left(\frac{1}{w_g} \right) \quad (4.14)$$

where σ is a scale parameter, A is the forward operator, y is the fitted data and w_g is the weight defined in Equation (4.12). The $\|A^T y\|_2$ factor normalises the cost function by the energy of the gradient of the fitting term evaluated at $x = 0$, hence for the solution that would be obtained for $\lambda \rightarrow \infty$. The normalisation by the maximum of the reciprocal of w_g is required in order to fix the interval spanned by the set of bundle-specific weights. Notice that, differently from COMMIT2, this definition does not rely on the solution of the associated NNLS problem, i.e., the problem obtained from $\lambda = 0$. We performed the optimisation for 25 logarithmically

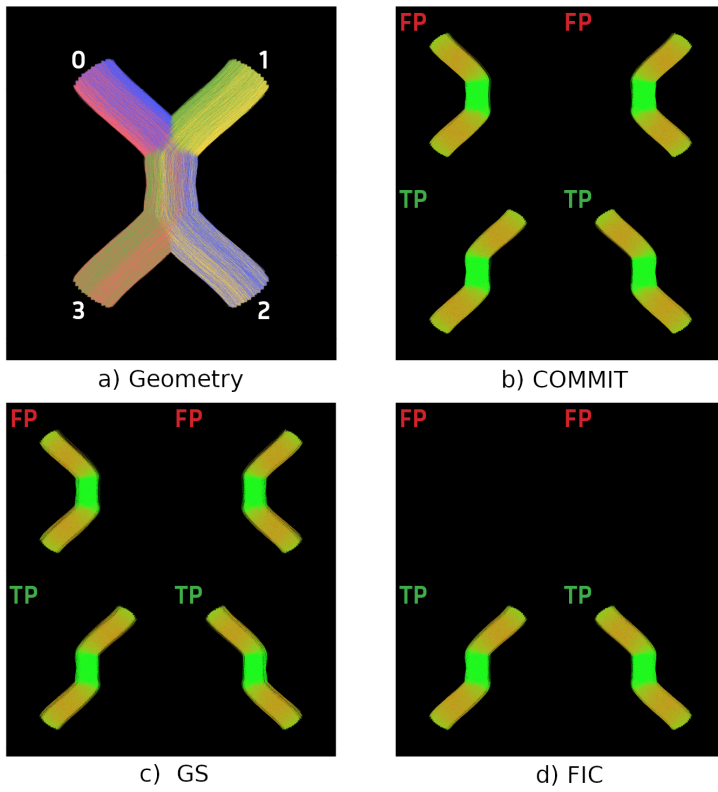


Figure 4.8: Crossing. The displayed results were computed for $\lambda = 2.5$ and $s = 1$. Subplot *a* shows the ground truth geometry of the phantom, where each bundle has a distinctive color. Subplots *b*, *c*, and *d* show every bundle belonging to the filtered tractogram. Those marked as FP should not be present after the tractography optimisation while those marked as TP should be.




Figure 4.9: Kissing. The displayed results were computed for $\lambda = 2.5$ and $s = 1$. Subplot *a* shows the ground truth geometry of the phantom, where each bundle has a distinctive color. Subplots *b*, *c*, and *d* show every bundle belonging to the filtered tractogram. Those marked as FP should not be present after the tractography optimisation while those marked as TP should be.




Figure 4.11: Results of ISBI 2013 phantom. The top-left panel displays the L-curve and the corresponding optimal regularisation parameters for the FIT and GS experiments. The top-right panel shows how the Youden's J index changes with respect to the employed regularisation parameter for each technique. The two bottom panels are devoted to the presentation of how the number of TPs and FPs is affected by the choice of the technique (FIT or GS) and of the regularisation parameter.

spaced values of σ spanning the $[10^{-16}, 1]$ range as shown in Figure 4.10.

The results for FIT AND gs are reported in Figure 4.11, where the L-curve displayed in the top left panel shows the asymptotic behaviour on the left, a plateau in the middle and the convergence towards the $x = 0$ solution on the right. The behaviour is exhibited both by FIT and GS. Notice how the optimal regularisation parameter of FIT and GS is obtained for different values of σ .

As mentioned before, the dataset is designed in such a way to allow to label each bundle as a True Positive (TP) or a False Positive (FP). The analysis of these aspects is reported in the other three panels of Figure 4.11. The results highlight how the same σ yields solutions with systematically less FPs when FIT is employed with respect to GS. For the respective optimal values of σ the number of retrieved FP bundles is 23 for FIT and 21 for GS. This difference is present also in the analysis of TPs. For the same value of σ , FIT retrieves more TPs than GS. At the respective optimal values of σ the FIT solution has 20 TPs, while the GS solution has 17 TPs. Notice how the highest value of the Youden's J statistic (Youden 1950) is not attained for the optimal value of σ neither for the FIT nor for the GS problem. The value of Youden's J at the optimal σ is equal to $J_{FIT} = 0.72$ for FIT and to $J_{GS} = 0.61$ for GS.

HCP subject The studied HCP subject[¶] is a healthy young adult. The optimisation was performed using the solver provided by the TALON package. The choice of the regularisation parameter was done as for the ISBI 2013 phantom, hence with an L-curve strategy based on the 25-points logarithmic grid. The results of the optimisation process are

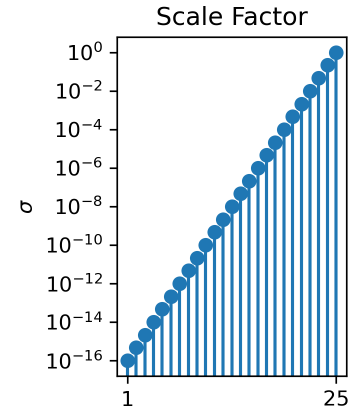


Figure 4.10: Scale factors employed in the definition of the L-curve for the ISBI 2013 and the HCP datasets. The considered values of σ are distributed logarithmically in the $[10^{-16}, 1]$ range.

Youden's J : statistic that encompasses the effects of both TPs and FPs. It's defined as $J = TPr + FPr - 1$, where TPr and FPr are the TP and FP rate respectively.

[¶]Subject id: 100307.

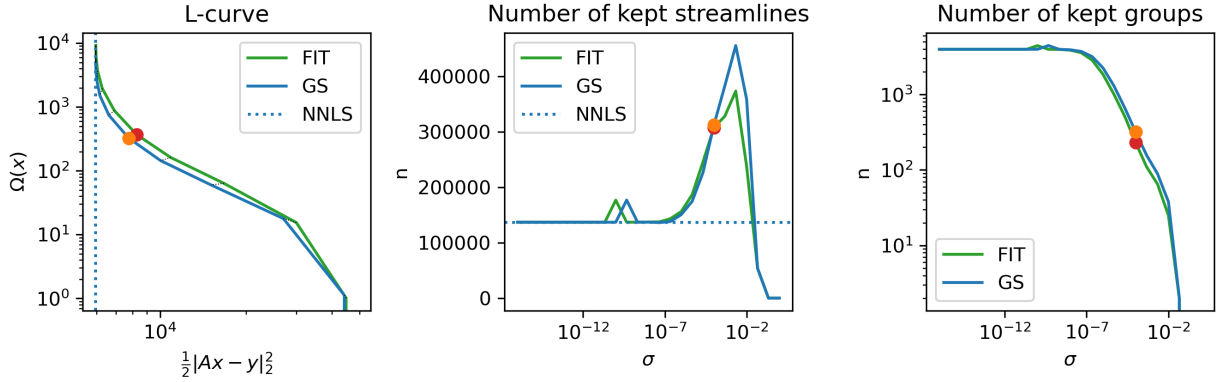


Figure 4.12: Results for the HCP subject. The left panel shows the L-curve for the FIT and GS experiments. The central panel reports the number of kept streamlines for each technique and σ . The right panel shows the number of groups (i.e., streamline bundles) retained for each technique and σ . The optimal parameter for each technique is marked in the corresponding curves (red for FIT, orange for GS). Optimality is attained for both experiments at $\sigma = 10^{-4}$.

presented in Figure 4.12. The left panel shows the L-curve for both FIT and GS and the respective optimal points, which are both attained for $\sigma = 10^{-4}$. The shape of the L-curve indicates that explored range of σ covers the whole spectrum of meaningful cases. The central panel of the figure reports the number of streamlines retained after the optimisation process for each technique and for each σ . A streamline s is considered to be retained if the optimisation process associates to it a coefficient $x_s^* > 0$. Despite using a sparsity-inducing penalty term, the number of retained streamlines does not monotonically decrease as σ grows. This is due to the design of the employed sparsity-inducing penalty, which promotes sparsity in the space of bundles, instead of promoting it in the space of streamlines. The sparsity in the space of bundles is displayed in the right panel of the same figure, where it is possible to see how the number of retained bundles decreases as expected. Considering the results obtained with the optimal σ , the FIT solution has 1% less streamlines and 27% less bundles than the GS solution. This indicates a superior sparsity of the solution with respect to not only the criterion explicitly encoded in the regularisation term (group sparsity) but also to the sparsity in the space of streamlines, which demonstrates a higher global parsimony of the obtained solution. All this at the price of a slight loss in terms of data fitting, as the fitting term evaluated at the GS solution is 6% lower than the one evaluated at the FIT solution.

4.3.2.3 Discussion

The results reported in the previous section show how filtering with FIT as a different effect with respect to filtering GS. This comparison is meaningful in the sense that FIT can be formulated as an extension of GS that includes additional priors on the functional connectivity between the regions connected by each streamline bundle.

The simple experiment involving the crossing-kissing fibers ambiguity presented in Figures 4.8 and 4.9 highlighted how the employed regularisation term does indeed include the desired type of information. The geometry of the simulated phantom and the topology of the ground truth network are too simplistic to draw any neuro-scientific conclusion from

it, but it confirms that the prior can be included in the process and solves simple problems and the designed regularisation terms acts in this sense as expected.

Looking at the results obtained on the ISBI 2013 dataset, we can have a more realistic picture of the effects of FIT on the sensitivity and specificity of the filtered tractograms. Ideally, one would want to have less FPs and more TPs. This would imply an increase in terms of Youden's J statistic. In Figure 4.11 we showed how in a wide neighbourhood of the optimal regularisation parameter the FIT solution has systematically more TPs and less FPs, hence a higher J , than GS. This result was expected since FIT includes the prior designed explicitly on the ground truth connectivity, as specified in Section 4.3.2.1. The solutions of FIT and GS obtained with the respective optimal regularisation parameters highlight a fundamental aspect: the highest Youden's J is attained for regularisation parameters different from the optimal ones in both GS and FIT. This consideration indicates that the sole use of Youden's J as a criterion for the selection of the regularisation parameter³ is not an optimal strategy, as it could be transparent to relevant differences in terms of data fit. The potential bias introduced by the use of Youden's J as the only criterion for choosing the regularisation parameter is particularly evident when real data are employed, for which no ground truth knowledge is available. In particular, the ground truth connectivity between the considered regions is not known, hence it is not possible to perform the the TP/FP analysis.

The experiments on the real data highlights one difference in the selection of the optimal regularisation parameter, which in this case, contrary to the experiments on the synthetic datasets, is the same for FIT and GS. This suggests that the assessment of the optimal strategy for the selection of the regularisation parameter is far from being concluded. Further studies are necessary in order to understand if the choice can be made at subject level, at cohort level or at TFT level.

Another relevant aspect highlighted by the real data is the different sparsity yielded by FIT with respect to GS. At the price of a slight reduction of the performance in terms of data fitting, the sparsity increases significantly when FIT is applied with respect to GS. This increases our confidence towards the idea that the higher parsimony observed when the functional prior knowledge is included does indeed come from an improved soundness of the whole framework.

One limitation faced in these experiments concerns the nature itself of the bundle sparsity framework, which is what associates COMMIT2, GS and FIT. For the way in which we formulated it, streamline bundles are defined with a *cortical* criterion, hence by clustering together streamlines connecting different pairs of cortical regions defined with a cortical atlas. Keeping in mind that the fitting term of the defined TFTs models the maps generated by the streamlines along their trajectories (i.e., in the WM), there is a WM-GM duality in the definition of the fitting and the regularisation terms. Designing the streamline clustering in a way that takes into account also the geometry of the streamlines (e.g., with QuickBundles (Garyfallidis et al. 2012)) could allow a better modelling of the prior. A drawback of such formulation would be that the definition of the functional connectivity associated to a bundle would need

3: In (Schiavi et al. 2020) the authors selected the regularisation parameter of COMMIT2 (GS-like model) based uniquely on the maximisation of the J statistic on a synthetic phantom.

to be rethought to match the new definition of *regions connected by the streamline bundle*.

Another limitation is posed by the use of Phantomas (Caruyer et al. 2014) in the generation of the synthetic phantoms. This software does not allow a reliable description of the intra-cellular volume fraction that is required for the use of the fiber density model as a forward model for the TFT. This weakens the possibility to infer properties of the solutions on real data from solutions on synthetic data. On the other side, the use of the fiber density model on the real data is mandated by the superior complexity of the MIT modelling, whose use would make the problem intractable at large scale.

4.4 Conclusion

In this chapter we analysed the expanding field of tractography filtering by reviewing the principal tools available in the literature, i.e., SIFT2 and COMMIT2. We proposed a more abstract unified formulation of those techniques discussing the modelling and the computational aspects. We presented a systematic comparison of the effects of TFTs on the graph-theoretical analysis of structural connectomes of the human brain estimated with dMRI. This phenomenon is shown to be present in both healthy subjects and patients affected by traumatic brain injury. The second contribution presented in this chapter is the *Functionally Informed TFT* (FIT), a novel filtering technique that extends the previous ones by including prior information on the functional connectivity between the regions connected by the streamlines. The preliminary results presented in this chapter show that connectomes obtained with FIT satisfy the expected properties, opening the door towards the exploration of functionally-informed structural connectomes.

The discussed results have partially been included in the following works.

- ▶ *Matteo Frigo*, Samuel Deslauriers-Gauthier, Drew Parker, Abdol Aziz Ould Ismail, Junghoon John Kim, Ragini Verma, Rachid Deriche. "Diffusion MRI tractography filtering techniques change the topology of structural connectomes." *Journal of Neural Engineering* 17.6 (2020): 065002.
- ▶ *Matteo Frigo*, Guillermo Gallardo, Isa Costantini, Alessandro Dadducci, Demian Wassermann, Rachid Deriche, Samuel Deslauriers-Gauthier. "Reducing false positive connection in tractograms using joint structure-function filtering." 24th Meeting of the Organization for Human Brain Mapping. Singapore, 2018.
- ▶ *Matteo Frigo*, Isa Costantini, Rachid Deriche, Samuel Deslauriers-Gauthier. "Resolving the crossing/kissing fiber ambiguity using Functionally Informed COMMIT." *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, Cham, 2019. In: Elisenda Bonet-Carne et al., (Eds.) *Computational Diffusion MRI: International MICCAI Workshop*, Granada, Spain, September 2018. Springer, 2019.
- ▶ *Matteo Frigo*, Rachid Deriche, Samuel Deslauriers-Gauthier. *TALON: Tractograms As Linear Operators in Neuroimaging*. Python package. 2020. gitlab.inria.fr/cobcom/talon

Chapter 5

Brain Alignment and Similarity

Overview 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 problem of choosing a suitable atlas for structural connectivity studies by assessing how robustly an atlas captures the network topology across different subjects in a homogeneous cohort. We measure this robustness by assessing the alignability of the connectomes, namely the possibility to retrieve graph matchings that provide highly similar graphs. To answer such a question, in Section 5.2 we introduce two novel concepts arising as natural generalisations 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 a new technique for aligning connectomes obtained by adapting the Weisfeiler-Lehman (WL) graph-isomorphism test called WL-align. In Section 5.3.3 we present results obtained on 100 unrelated subjects from the Human Connectome Project database, which allowed us to validate the GJI and WL-align, inferring a strategy for choosing a suitable parcellation for structural connectivity studies. Finally, a discussion on the potential developments and limitations of our work is presented in Section 5.4.

5.1 Introduction	81
5.2 Theory	83
Brain Alignment	83
Quality of Brain Alignments	84
Weisfeiler-Lehman Network	
Alignment	87
5.3 Experiments	90
Data and Methods	90
Statistical analysis	94
Results	94
5.4 Discussion and conclusion .	100

5.1 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. As showed in Chapter 2, the cortical surface has been subdivided based on its cytoarchitecture, shape, functional organisation, axonal connectivity, and combinations of these and other features. 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, the field of connectomics studies their relationship and interac-

tions. In this context, the focus is shifted from understanding how information is segregated in the brain to how it is integrated. Through the use of dMRI-based tractography, structural connections between brain areas can be recovered, as presented in Chapters 2 and 4. 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. Alternatively, one could determine the functional connectome of a brain, hence the network that describes the similarity between the patterns of activation of brain regions. Given structural or functional connectomes, their features can be compared across subjects and populations to link network changes to pathology or to further increase our understanding of its organisation. An underlying assumption is that a correspondence between nodes of the network exists across subjects, a condition which is usually satisfied by using a group (Gallardo et al. 2018a; S. Parisot et al. 2015) or template-based parcellation (see Chapter 2). The drawback of this strategy is that it ignores any subject specific changes in cortical organisation and reduces the specificity of the results. The use of subject-specific atlases and the subsequent comparison of the resulting connectomes requires the a-posteriori definition of a mapping between the nodes of the two networks. This is also known as *brain atlas correspondence* and *parcel matching* problem (Gallardo et al. 2018b; Mars et al. 2016).

The construction of a mapping between network nodes corresponds to what is known in various fields as *network alignment* or *graph matching* (Ayache et al. 1987; Barak et al. 2019; Conte et al. 2004; Korula et al. 2014; Singh et al. 2008). Graph alignment solutions (called *alignments*) correspond to a permutation of the labels of the nodes of a graph which maximises 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 (Becker et al. 2018; M. K. Chung et al. 2017; Deslauriers-Gauthier et al. 2020; Osmanlıoğlu et al. 2019; Shen 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 maximises 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 parcel matching problem. The main difference is that in those problems the permutation acts only on the rows of the connectivity matrix as 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 a naïve brute force strategy 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 (Feizi et al. 2019; Hayhoe et al. 2019; Nassar et al. 2018), despite being subject to limitations (R. C. Wilson et al. 2008). Modern machine learning paradigms exploit deep learning techniques for finding an alignment (Heimann et al. 2018; C. Li et al. 2018;

Liu et al. 2016), however they make use of partially available information about the alignment itself (Liu et al. 2016), or lack explainability and interpretability.

We first introduce in Section 5.2 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 penalising mismatches, overcoming limitations of previous approaches (Feizi et al. 2019). We then propose in Section 5.2.3 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 et al. 1968). WL-align is amenable to concrete interpretability in terms of local network structure around each node (Figure 5.3) and can be integrated with other heuristics. We compare WL-align against the Fast Approximate Quadratic Programming for Graph Matching (FAQ) (Vogelstein et al. 2015), another efficient brain-alignment heuristic which is solely based on network structure.

5.2 Theory

A brain network is characterised 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 an edge $(i, j) \notin G$ 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 $\text{Adj}(G)$. In the context of connectomics (Hagmann 2005; O. Sporns et al. 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.

5.2.1 Brain Alignment

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 $G_2 = (V_2, E_2)$ of n_1 and n_2 nodes respectively, it is possible to define an injective map¹ $m : V_1 \rightarrow V_2$ 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

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

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

$$(P_m)_{ij} = \begin{cases} 1 & \text{if } m(j) = i, \\ 0 & \text{otherwise.} \end{cases} \quad (5.1)$$

In the special case where $n_1 = n_2$, 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 characterised by the matching matrix and is given by

$$\text{Adj}(G_1) = P_m^T \text{Adj}(G_2) P_m \quad (5.2)$$

where T denotes the transpose of the matrix.

5.2.2 Quality of Brain Alignments

Once a brain alignment is identified, its quality can be assessed by evaluating the (dis)similarity of the resulting two resulting networks. In the following, the similarity measures are defined 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

$$d_t(G_1, G_2) = \|\text{Adj}(G_1) - \text{Adj}(G_2)\|_t \quad (5.3)$$

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

$$c(G_1, G_2) = \frac{\langle \vec{\text{Adj}}(G_1), \vec{\text{Adj}}(G_2) \rangle}{\|\vec{\text{Adj}}(G_1)\|_2 \cdot \|\vec{\text{Adj}}(G_2)\|_2} \quad (5.4)$$

where the numerator is the scalar product between the vectorisation 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. Other distances based on geometrical (Venkatesh et al. 2020) and homological (M. K. Chung et al. 2017) properties of the networks have been proposed.

We remark how the concept of *similarity between networks* used in this work fits well the standard concept of *matrix similarity* in the particular case where the change of basis matrix is a permutation matrix. Given two n -by- n matrices A and B , they are said to be *similar* if there exists an invertible matrix P such that

$$B = P^{-1}AP$$

where P can be interpreted as the change of basis matrix. In the context of graph alignment, the expression given in the last equation appears also in the mapping between two graphs of Equation (5.2). In particular, the change of basis matrix P corresponds to the matrix representation of the graph matching.

All such measures capture some aspects of the similarity between two graphs, but none of them satisfies all the following requirements:

- ▶ arising as a natural generalisation 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 generalising 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 in two given graphs, it also depends on the global network structure.

5.2.2.1 Graph Jaccard Similarity Index

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

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}. \quad (5.5)$$

An example of what is measured by the Jaccard index on sets is given in Figure 5.1. Notice that $J(A, B)$ 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., $A \cap B = \emptyset \implies J(A, B) = 0$) or when the two sets are equal (i.e., $A = B \implies J(A, B) = 1$). Both the sets need to be non-empty.




Figure 5.1: 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 generalised 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

$$J(x, y) = \frac{\sum_{i=1}^d \min(x_i, y_i)}{\sum_{i=1}^d \max(x_i, y_i)}. \quad (5.6)$$

Note that the Jaccard similarity index between two sets follows as 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. A remarkable property of the weighted Jaccard similarity index is that it induces a metric in the space where it is defined (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 $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ with adjacency matrices $\text{Adj}(G_1) = A$ and $\text{Adj}(G_2) = B$, the *weighted graph Jaccard similarity index* (GJI) of G_1 and G_2 is

$$J(G_1, G_2) = \frac{\sum_{(i,j) \in \mathcal{E}} \min(A_{i,j}, B_{i,j})}{\sum_{(i,j) \in \mathcal{E}} \max(A_{i,j}, B_{i,j})} \quad (5.7)$$

where \mathcal{E} is the set of all possible pairs of nodes. In the context of the present work, we remark that we can think of B as having been previously aligned to A via Equation 5.2. Alternatively, the weighted graph Jaccard similarity index is defined as the weighted Jaccard index of the vectorisation 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 (5.7) 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 (5.7) will be equal to zero and $J(G_1, G_2) = 0$. A remarkable property of the GJI is that it induces a metric in the space where it is defined. As a matter of fact, the function $d_J(x, y) = 1 - J(x, y) \in [0, 1]$ respects the three properties of metrics:

1. Identity: $d_J(x, y) = 0$ if and only if $x = y$;
2. Symmetry: $d_J(x, y) = d_J(y, x)$;
3. Triangle inequality: $d_J(x, y) \leq d_J(x, z) + d_J(z, y)$.

The first two properties trivially follow from the definition of J , while the third follows as a particular case of what is presented in Charikar 2002, Lemma 1. An example of how the GJI acts on two graphs is given in Figure 5.2.

We have so far formally established the notion of network alignment (Equation (5.1)), and presented the Jaccard index as a principled way to measure the quality of an alignment (Equation (5.7)). 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.

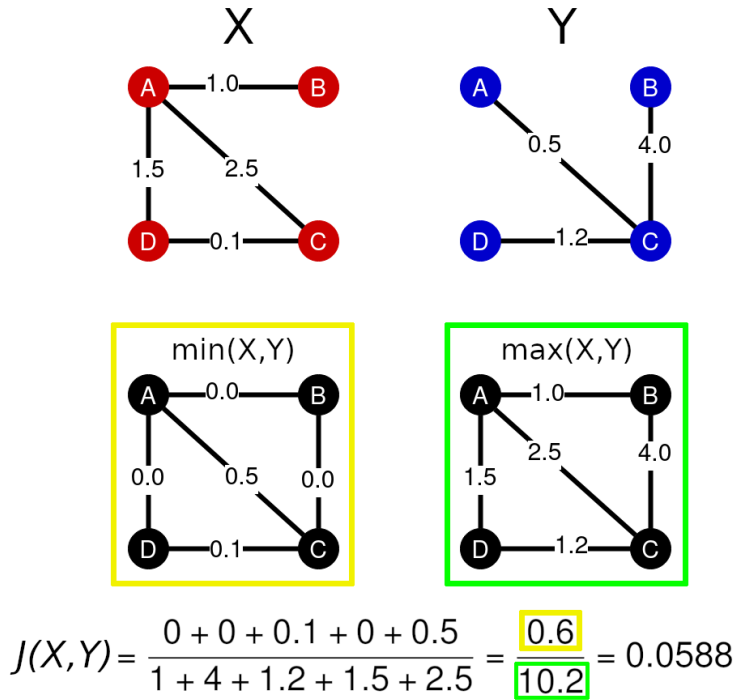


Figure 5.2: This figure shows an example of how to compute the GJI between two compatible graphs X and Y . For each pair of nodes $i, j \in \{A, B, C, D\}$, one computes the minimum and maximum between $X_{i,j}$ and $Y_{i,j}$. These two quantities will be used to define the numerator and the denominator of the GJI defined in Equation (5.7). 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. The displayed example shows a very low GJI, indicating poor similarity between the compared graphs.

5.2.3 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 called *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 as 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, which is determined with an algorithm inspired by the Weisfeiler-Lehman (WL) method for graph isomorphism testing (Weisfeiler et al. 1968). For this reason, *WL-align* is the name we propose for our brain alignment algorithm.

The WL Graph Isomorphism Heuristic We first recall the classical WL algorithm (Färer 2017) for graph isomorphism testing. Consider an undirected unweighted graph $G = (V, E)$. At the onset, we color all the nodes of V with the same color. We then refine the coloring in consecutive rounds: at round $i + 1$, nodes u and v receive new and different colors whenever they already had different colors at round i , or if the multi-sets² of colors of their neighbors at round i were different. For example, if two nodes have different degrees, the multi-set of neighbor colors are

2: A multi-set is a set whose elements can appear multiple times.

initially different since their cardinality differs, hence they will be assigned a different color after the first iteration of the algorithm.

At each round, the colors assigned by the algorithm define a partition of the nodes, which stabilizes after at most n rounds. Since WL is solely based on the topology of the graph, if the partitions that it constructs for two graphs G_1 and G_2 differ, then the two graphs are not isomorphic. If, instead, the two partitions are the same, the generated partitions can be exploited to search for an isomorphism between the two graphs (McKay et al. 2014).

The core operation of WL is the computation of the multi-set of neighbouring colors of a node, which is some sort of *signature* of the node. In the next paragraphs we're going to see how this can be generalised to the case of weighted graphs. This will allow us to define the signature that will be exploited in the WL-align technique.

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 incident to the node itself, namely

$$\text{vol}(v) = \sum_{u \in V} w_{uv} \quad (5.8)$$

where V is the set of nodes in the graph, v is the node of which we compute the volume $\text{vol}(v)$ 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 G' induced by the nodes that are reachable from u in at most ℓ hops. At each of these hops, G' retains only the k nodes with highest contribution, weighted according to a function of the path that connects them to u . In detail, such a contribution is computed via the following function

$$f(v_0, \dots, v_h) = \begin{cases} \text{vol}(v_0) & \text{if } \pi = (v_0), \\ \frac{w(v_0, v_1)}{\text{vol}(v_0)} \cdot f(v_1, \dots, v_h) & \text{otherwise} \end{cases} \quad (5.9)$$

where $w(u, v)$ is just a more verbose notation for the edge weight w_{uv} . The subnetwork G' 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 via Equation (5.9). A formal description of the algorithm for computing the signature H_u is given in Algorithm 1, while a graphical intuition is illustrated in Figure 5.3.

Bipartite graph A bipartite graph is a network whose nodes can be divided in two distinct and non-overlapping sets, such that there are no edges connecting nodes in the same set. Once a signature is computed for each node of the two graphs, we define a weighted complete bipartite graph $G_m = ((V_1 \cup V_2), (V_1 \times V_2))$. The nodes on the left, i.e., V_1 , represent the nodes of the first graph, while the nodes on the right, i.e., V_2 , represent




Figure 5.4: 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.

Algorithm 1 WL-align signature

Input: graph G ; node u ; width k ; depth ℓ **Output:** signature H_u

```

1:  $H_u \leftarrow$  empty list ▷ append right
2:  $Q \leftarrow$  empty queue ▷ FIFO data-structure; pop left (get and remove);
   append right
3:  $Q \leftarrow$  append  $\pi = (u)$  ▷  $\pi$  is the zero-length, single-node path
4: while  $Q$  is not empty do
5:    $\pi = (u, \dots, v_h) \leftarrow$  pop path from  $Q$ 
6:    $H_u \leftarrow$  append  $f(\pi)$ 
7:   if  $h < \ell$  then ▷  $h$  is the length of  $\pi$ 
8:      $\pi_z \leftarrow (u, \dots, v_h, z), \forall z \in V$  ▷ if  $(v_h, z) \notin E$  then
    $w(v_h, z) = 0 \implies f(\pi_z) = 0$ 
9:      $z_1, \dots, z_n \leftarrow$  nodes s.t.  $f(\pi_{z_1}) \geq \dots \geq f(\pi_{z_n})$  ▷ ties broken
   uniformly at random
10:     $Q \leftarrow$  append the  $k$  paths  $\pi_{z_1}, \dots, \pi_{z_k}$ .
11: return  $H_u$ 

```

the nodes of the second graph. The edge-weights encode the distance between the signatures of pairs of nodes belonging to different graphs, i.e., each edge (u, v) with $u \in V_1$ and $v \in V_2$ is weighted according to a function $b : V_1 \times V_2 \rightarrow \mathbb{R}$ defined as the euclidean distance between the signatures of the two endpoints $b(u, v) = \|H_u - H_v\|_2$. Figure 5.4 shows a simple example of the defined bipartite graph.

Assignment problem The final step towards finding the wanted matching with WL-align is the resolution of the assignment problem corresponding to the bipartite graph G_m 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. In formal terms, given the two sets V_1 and V_2 and the weighting function b that define G_m , the problem asks to find a bijection $m : V_1 \rightarrow V_2$, i.e., the matching, that minimizes the function $\sum_{u \in V_1} b(u, m(u))$. This assignment problem is efficiently solved by the Hungarian algorithm (Jacobi 1890; Kuhn 1955), which is a combinatorial optimization routine that finds the wanted matching with a minimum total cost. The Hungarian algorithm has $\mathcal{O}(n^3)$ complexity and it is implemented in Scipy (`scipy`).




Figure 5.3: 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 $H_u[1] = \text{vol}(u)$, 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 neighbors that create a path with highest f , namely v_1 and v_2 , which are marked by the blue and green circles in panel 2. They are considered in decreasing order (w.r.t. the volume) and the corresponding entries are computed with Equation (5.9). For instance, the second entry of H_u is equal to $H_u[2] = \text{vol}(v_1) \cdot a_1 / \text{vol}(u)$. The third entry is computed in an analogous way as $H_u[3] = \text{vol}(v_2) \cdot a_2 / \text{vol}(u)$. This concludes the definition of the first $1 + k$ entries of H_u . The following entries are defined by considering first the blue and then the green subnetwork in panel 3. The fourth entry is equal to $H_u[4] = \text{vol}(u) \cdot (a_1 / \text{vol}(v_1)) \cdot (a_1 / \text{vol}(u))$ and the fifth is $H_u[5] = \text{vol}(w_1) \cdot (b_1 / \text{vol}(v_1)) \cdot (a_1 / \text{vol}(u))$. Analogously, the sixth and the last entry will be $H_u[6] = \text{vol}(w_4) \cdot (b_4 / \text{vol}(v_2)) \cdot (a_2 / \text{vol}(u))$ and $H_u[7] = \text{vol}(w_5) \cdot (b_5 / \text{vol}(v_2)) \cdot (a_2 / \text{vol}(u))$.

5.3 Experiments

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 5.3.1. For each of the 100 subjects we considered 23 parcellations (Desikan, Glasser, Gallardo at 11 different resolutions, Schaefer at 10 different resolutions) described in Section 5.3.1.1, 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 5.3.1.4, for a total of 232300 alignments. The quality of the obtained alignments was then assessed using four network similarity measures described in Section 5.3.1.5. All the code and data that are necessary for reproducing the presented experiments are available on the Open Science Framework Frigo et al. 2020a.

5.3.1 Data and Methods

To build the structural brain networks, we considered the preprocessed data of the Human Connectome Project (HCP) database (U100 subject group) (Matthew F Glasser et al. 2013; WU-Minn Human Connectome Project consortium 2017; Van Essen et al. 2012) available at the Connectome Coordination Facility*. For each subject, a five-tissue-type image (R. E. 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

(Frigo et al. 2020a): Frigo et al. (2020), *Code and data for "Network alignment and similarity reveal atlas-based topological differences in structural connectomes"*

* <https://www.humanconnectome.org/>

fluid using a maximal spherical 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 (R. E. Smith et al. 2012) seeding from the gray matter - white matter interface and obtaining a total of five million streamlines per subject.

5.3.1.1 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 together with the cortical surface in fsLR32k space. The Glasser parcellation (Matthew F. 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 (Washington University School of Medicine 2020). The Gallardo parcellation (Gallardo et al. 2018a) 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. 2018a). The used Schaefer atlas (Schaefer et al. 2017) was downloaded from the repository of the CBIG laboratory (T. Yeo 2020) for the *seven-networks* parcellation (B. T. 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.

5.3.1.2 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. We removed the self-connections from each connectome to force the comparison of connections *between* regions, allowing to work with the sheer topological properties of the network (Stratos 2020). 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 normalisation), 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.

5.3.1.3 Intra-cohort variability

In order to assess the variability between the brain networks of the subjects in the studied cohort, for each subject we measured the similarity between the connectomes of each pair of subjects with three different similarity metrics: the weighted graph Jaccard index (Equation (5.7)), the Frobenius norm (Equation (5.3)) and the correlation (Equation (5.4)).

5.3.1.4 Network alignments

In order to assess the ability of WL-align to retrieve the wanted 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 node labels of the connectomes of all subjects keeping track of the permutation maps. These permutation maps allow to compute the ground truth matching m^* between each pair of brain networks computed with the same 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 $\ell = 2$. We limited the width for efficiency reasons (the size of the signature is higher than k^ℓ , as described in the previous section) and the depth since further increasing it does not lead to substantial gain w.r.t. the quality of the alignments (the deeper the nodes in the search, the smaller the contribution of the nodes to the signature, as described in Equation (5.9)).

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. FAQ works in three main steps: i) arbitrarily choose a starting bistochastic matrix, which acts as a relaxed permutation matrix that aligns the two networks; ii) find a local solution to the Relaxed Quadratic Assignment Problem (rQAP), a dual version of the graph matching problem; iii) project back onto the set of permutation matrices. The solution found by FAQ transforms the adjacency matrix of the first graph into one with approximately minimal Frobenius distance from the adjacency matrix of the second graph. Notice that optimality with respect to the Frobenius distance might not correspond to absolute optimality. We used the implementation of FAQ available in the *graspologic* package (J. Chung et al. 2019) (<https://graspologic.readthedocs.io/>), setting the number of random initializations to 30.

Both WL-align and FAQ were run separately on each hemisphere of the brain and the two resulting partial alignments were then combined into a

single one. The motivation for this choice is that the correct hemisphere can always be assigned to a cortical region, and this property is independent from any influence potentially caused by the registration of the template atlas onto the subject-specific cortical mesh, while other properties, e.g., the location of a region, would be. Moreover, by studying single-hemisphere alignments we bypass the issue concerning the high degree of left-right similarity that characterizes the brain, which could drive the solution towards sub-optimal alignments that are hardly distinguishable without external criteria such as the localization or geometry of the brain regions. Notice that this choice concerns the design of the experiment, not the setup of the graph matching algorithm, which could still be obtained using the full brain network, hence including the inter-hemispheric connections.

5.3.1.5 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

$$\text{NMr}(m) = \frac{|\{u \in V_1 : m(u) = m^*(u)\}|}{|V_1|}. \quad (5.10)$$

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 (5.7), namely

$$J(m) = J(m(G_1), G_2) \quad (5.11)$$

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

$$\text{Jr}(m) = \frac{J(m)}{J(m^*)}. \quad (5.12)$$

When the ground truth matching m^* is also an optimal matching, the denominator $J(m^*)$ acts as a normalisation 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

$$\text{FRO}(m) = \|\text{Adj}(m(G_1)) - \text{Adj}(G_2)\|_F \quad (5.13)$$

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 . 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

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

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.

A further qualitative assessment of the accuracy of the alignments obtained with WL-align was performed by projecting the matching ratio of each node onto the cortical surface of a randomly picked subject, obtaining a visual indication of the localization of the regions that have been more or less frequently correctly matched. Projecting this information directly on the cortical surface provides insights into the spatial organization of the errors and of the correct matches.

5.3.2 Statistical analysis

In order to understand the differences between the alignments obtained with WL-align and FAQ, statistical analyses were performed with an alpha of 0.05 in all experiments. A separate analysis was performed for each of the four similarity metrics presented in the previous section. First, for each atlas and pair of subjects we computed an alignment with WL-align and FAQ. For each atlas, we compared the distributions of the values of the similarity metric computed on the alignments obtained with the two techniques using the non-parametric paired-samples Wilcoxon signed-rank test (Wilcoxon 1945).

5.3.3 Results

We processed the data of 100 unrelated subjects from the HCP database obtaining the structural brain networks as detailed in Section 5.3. 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 5.2 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 (5.4), 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 (5.4), also known as cosine similarity. The third similarity measure is the Frobenius distance defined in Equation (5.3), 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 (5.10) is the last considered similarity measure.

5.3.3.1 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, and in order to assess the intra-cohort similarity of the connectomes, 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.

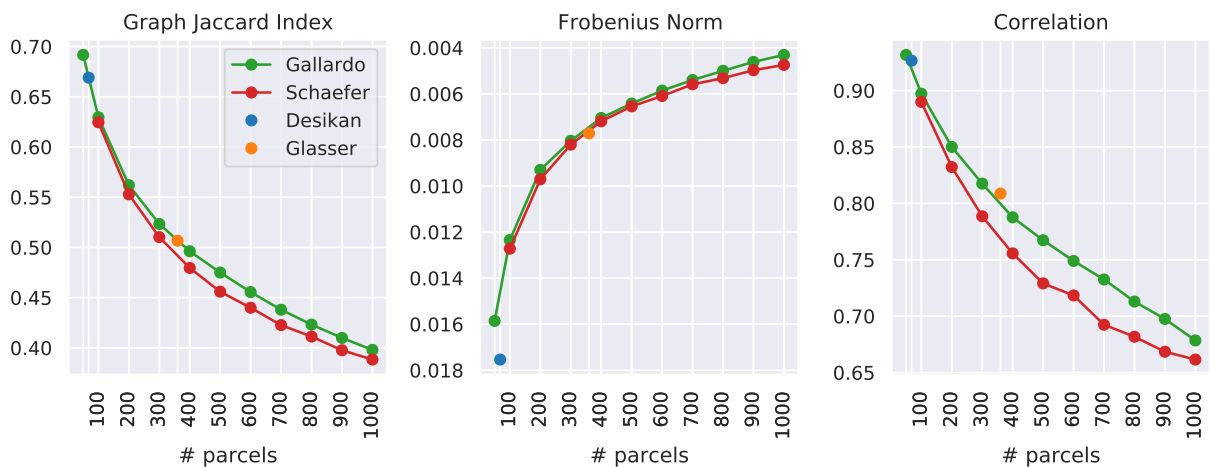


Figure 5.5: Each point shows the average similarity between every pair of subjects in the considered cohort measured on connectomes obtained with a specific parcellation. The used alignment is the one defined by the ground truth, which in our experiments is known a-priori. All panels show the similarity 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.

For each parcellation, Figure 5.5 shows how similar the subjects are with respect to the graph Jaccard index, the Frobenius norm, and correlation. In particular, the figure reports for each parcellation the average similarity across each pair of subjects, which can be computed from the ground truth matching whose existence is granted by the fact that each network is defined on a known set of nodes. Despite using the ground truth matching, the graphs are not expected to exhibit perfect similarity (i.e., $J = 1$, $FRO = 0$ or $C = 1$), as their edge weights are subject-specific. This specificity is what determines the intra-cohort variability that is taken into

account by the J ratio similarity metric defined in Equation (5.12). Figure 5.5 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, 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. When measured with the GJI and the correlation, both these parcellations exhibit an intra-cohort similarity in line with the one of the Gallardo parcellation at the corresponding resolutions.

5.3.3.2 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 5.3.1.5. 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 5.5, 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 5.6: 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. We do not report the results for the J -ratio since in this experiment its denominator $J(m^*) = 1$, making the plot identical to the one of the graph Jaccard index. 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 5.6 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




Figure 5.7: 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.

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 ones computed with FAQ and presented in Figure 5.7, which reports the average similarity between the obtained alignment and the ground truth alignment among all the possible pairs of subjects. The statistical significance of the differences between results obtained with WL-align and FAQ is assessed using the non-parametric paired-samples Wilcoxon signed-rank test (Wilcoxon 1945). For the studied cohort, statistically significant differences are observed for each atlas and each employed similarity metric, as shown in Section B of the supplementary material. 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 Glasser parcellation is in line with the Schaefer and Gallardo parcellations when the alignment is obtained with WL-align, while this is not true for the Desikan parcellation. 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 5.5, where the Frobenius similarity increases with the number of

parcels. 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 exhibits Jaccard similarity slightly lower than the one of the Gallardo parcellation but still higher than Schaefer's, suggesting that the multi-modal nature of the atlas allows to capture, at least in part, the structural connectivity features that we are looking at. This contrast is evident only when WL-align is employed. The Desikan parcellation behaves differently. While exhibiting lower performance with FAQ, when the WL-align is employed it emerges as a slightly superior parcellation with respect to the NMr, the GJI and the J ratio. 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.

5.3.3.3 Region matching rate

Figure 5.8 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. This analysis gives important insights into the type of errors that are made by WL-align. In particular, it shows that the incorrect matchings do not have a particular structure that can be related to the geometry and morphology of the brain, be it some regional concentration of errors or some consistent symmetry with respect to the hemispheres.




Figure 5.8: 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.

5.4 Discussion and conclusion

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 analysed 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.

The proposed WL-align algorithm can be further adapted to work with types of network other than the structural networks studied in this work, which are undirected and have non-negative edge weights. The most intuitive way to adapt WL-align is to change the way in which the node signatures are defined, then set up the bipartite graph and find the matching with the Hungarian algorithm in the canonical way. A first interesting case is represented by weighted networks having both positive and negative weights. This is the typical case of *functional* connectivity studies, where the connectivity between regions is evaluated as the correlation (i.e., $w_{ij} \in [-1, 1]$) between the activation in different regions (Van Den Heuvel et al. 2010). As we defined it in this work, WL-align would select the most relevant d nodes in an unpredictable way due to the presence of negative-valued entries in Equation (5.9). A possible adaptation of it would be to select the relevant edges performing the breadth-first search ignoring the sign of the weights, then evaluating the corresponding entries of the WL signature using the signed edge weights in Equation (5.9). Another possible adaptation would require the decomposition of the adjacency matrix of the network as $\text{Adj}(A) = A_p - A_n$, where A_p is the positive part of the matrix and A_n is the negative part of the matrix. Notice that the graphs corresponding of both A_p and A_n will have non-negative edge weights. For each node, the WL signatures obtained from A_p and A_n can be concatenated, then used in the canonical way. Another interesting case is represented by directed networks, which in the context of brain imaging represent the concept of *effective* connectivity (Friston 2011). Here, the only further adaptation that would be required is a careful definition of the breadth-first search that gives the selection and the order of nodes that are used for defining the WL signature. For directed positive-weighted network, the algorithm works as it is, while for directed networks with signed weights it would require the adaptations mentioned for the case of functional connectivity. Finally, we discuss the adaptation of WL-align to temporal networks. This type of graph has gained much interest in the context of brain imaging since the concept of *dynamic functional connectome* (Prete et al. 2017) has been introduced and the consequent definition of specific tools for the graph-theoretical analysis of these time-dependent networks (Sizemore et al. 2018). In this case, at least two options can be explored: first, one could concatenate the WL signatures of each node obtained at each time point, then run WL-align in the canonical way. Alternatively, it would be possible to perform the breadth-first search by taking into account the temporal component, hence traversing the graph both in space and time.

An important instance of the graph matching problem which we did not consider in this work corresponds to when the two networks that are being aligned have different numbers of nodes. Being a generalization of a graph isomorphism test, WL-align does not appear to be trivially adaptable to this case. A possible solution would be to employ some dimensionality reduction technique (e.g., clustering via community detection) in the larger graph to reduce the number of nodes to the one of the smaller graph, then use WL-align to retrieve the wanted alignment.

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 from the fact that 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 better results in terms of alignability than any other parcellation of any granularity. For this reason, whenever a study is designed using a coarse parcellation of the cortex (in the 50-200 parcels range), 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 are concerned, we showed that parcellations with a number of parcels in the > 400 range have lower performance in terms of GJI and NMr. However, when the inter-subject variability is taken into account in the evaluation of the similarity, as for the case of the Jaccard index ratio, we see that the performance is nearly constant for atlases with > 300 parcels.

The change in performance that we observe with the growing resolution of the atlas could also be due to the number of streamlines employed in the tractography pipeline, which could be adapted to the used atlas, but in practice is the same for every atlas at each resolution. On the other hand, the standardized tractography pipeline (including the identical number of streamlines in each tractogram) is what allowed us to present a comprehensive analysis and comparison of the performance across resolutions. In order to disambiguate this point, it would first be necessary to analyze how the strongest connections in the network (hence those considered by WL-align) are affected by the number of tracked streamlines. An alternative solution could be to employ a tractography filtering technique such as SIFT2 (R. E. Smith et al. 2015a), COMMIT (Daducci et al. 2014) or LiFE (Pestilli et al. 2014) in order to mitigate the limited reliability of streamline count as a proxy of axonal connectivity (Jbabdi et al. 2011). Given that tractography filtering techniques have non-negligible effects on the topology of structural connectomes (Frigo et al. 2020b), an independent analysis is due in order to assess how their use affects the alignability of connectomes.

Notice that in our analysis we used the defined similarity metrics to as-

sess which atlas yields connectomes with higher or lower robustness *in a certain resolution range*. This means that we could not have used the similarity argument to claim that, for instance, the Desikan atlas (68 parcels) should in general be preferred to the Gallardo 1000 atlas. In this sense, we highlight how the considered similarity metrics (GJI, Jr, NMr and Fro) should not be used for selecting the appropriate resolution at which structural connectivity studies should be designed, but they provide a well grounded tool for assessing which of the available atlases at the wanted resolution is most suitable for the considered type of study.

As highlighted throughout the chapter, this work analyses 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 or effective connectivity studies, which would require a separate analysis which was out of the scope of this work.

This work was developed in collaboration with our colleagues from the Combinatorics, Optimization, and Algorithms for Telecommunications (COATI) team at Inria Sophia Antipolis - Méditerranée and has been accepted for publication at Network Neuroscience (Frigo et al. 2021b). The Python code necessary to reproduce the result has been included together with the data in a repository stored at the OSF[†].

- ▶ *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. 2021. In press. Network Neuroscience.
- ▶ Emilio Cruciani, *Matteo Frigo*, David Coudert, Rachid Deriche, Emanuele Natale, Samuel Deslauriers-Gauthier. WL-align. Python package. <https://osf.io/depux/>

[†] <https://osf.io/depux/>

Chapter 6

Conclusion

Mapping the human brain is one of the complex challenges of contemporary science. It is a task that concatenates several problems from acquisition design to preprocessing, modelling, analysis, visualisation and assessment of the coherence with the state-of-the-art knowledge on the architecture and functioning of the human brain. For each of these steps a plethora of solutions has been and is being developed. It is of fundamental importance that the assumptions made in each step align with each other, demanding extra care in the verification of the theoretical requirements of the employed tools. In this thesis we focused on three specific parts of the chain of problems that leads to a comprehensive view of the brain architecture, highlighting the theoretical aspects that characterise the posed challenges and providing experimental evidence of the soundness of the proposed solutions.

We presented four contributions on three topical research areas of diffusion MRI methods for human brain mapping: brain tissue microstructure, tractography filtering and brain network topology. We proposed a new method for the estimation of tissue-specific volume fractions by means of multi-compartment models of the single-TE diffusion MRI signal. Then, we reviewed the state of the art of tractography filtering and unveiled its effects on the graph-theoretical analysis of the structural connectomes of both healthy subjects and patients affected by traumatic brain injury. In addition, we proposed a novel filtering technique that integrates structural and functional information in the process. Finally, we proposed a new similarity measure between brain networks and a new graph alignment techniques, which allowed to obtain original insights into the problem of selecting the suitable parcellation for brain connectivity studies.

Multi-Tissue Multi-Compartment models of diffusion MRI

Our first contribution is a novel model that allows to estimate tissue-specific volume fractions from single-TE dMRI data. Recent works highlighted how multi-compartment models of dMRI are not in principle designed to take into account the T_2 differences between the modelled tissues. As a result, what has always been referred to as *volume fraction* is in fact the *signal fraction* of a compartment. The latter is shown to be a biased estimator of the former. Solutions involving the acquisition of multi-TE dMRI data have been proposed by Veraart et al. 2018, Lampinen et al. 2017, 2019 and Gong et al. 2020. These techniques have the drawback of requiring a re-design of the whole pipeline from acquisition to post processing, as standard dMRI datasets are acquired with a single TE. For this reason, large datasets as the one published by the Human Connectome

Contribution #1

Project or the UK Biobank can not be treated with these multi-TE models. In Chapter 3 we proposed a Multi-Tissue Multi-Compartment (MT-MC) model of single-TE dMRI inspired by the aforementioned multi-TE models and by the multi-tissue formulation of the CSD algorithm proposed by Jeurissen et al. 2014. With our MT-MC model it is possible to retrieve an unbiased estimation of the tissue volume fractions. We provided a multi-tissue extension of the standard MC model of dMRI in the WM initially proposed by Novikov et al. 2019. It can be constrained in such a way that it models three tissues (intra- and extra-axonal compartments plus CSF), two tissues (WM and CSF) or a single tissue (equivalent to the “old” MC models). Our experiments on both real and synthetic data indicate that the two-tissue formulation is a valid substitute of the old MC formulation, as it can be used directly on single-TE multi-shell dMRI data without further requirements. If additional prior knowledge is available on the T_2 differences between the intra- and extra-axonal compartments, the three tissue model can be employed to further refine the estimation.

Perspectives The proposed model strongly relies on the external estimation of the T_2 or the S_0 of the modelled tissues. Our experiments on real data leveraged the heuristic of Dhollander et al. 2016b to retrieve the S_0 of the WM and the CSF. Understanding how this choice affects the estimation of volume fractions is out of the scope of this work, but the raised question suggests that further efforts should be devoted to researching techniques that estimate tissue-specific S_0 responses using single-TE data. Additionally, analysing the proportion between the S_0 of each tissue in a large cohort of subjects could highlight patterns that could be exploited. If hypothetically the T_2 of extra-axonal was showed to be a constant fraction of the T_2 of the intra-axonal compartment, this could straightforwardly be encoded in the model.

The differences between signal fractions and volume fractions that we highlighted both theoretically and empirically imply that previous studies making use of the concept of volume fraction need to be re-interpreted taking into account those differences. How those differences are expressed in the presence of pathology or group differences remains unexplored and needs to be assessed in future studies.

This work was done in collaboration with Rutger Fick *. This chapter is based in part on the following published works.

- ▶ *Matteo Frigo*, Mauro Zucchelli, Rutger Fick, Samuel Deslauriers-Gauthier, Rachid Deriche. Multi-compartment modelling of diffusion MRI signal shows TE-based volume fraction bias. OHBM 2020.
- ▶ *Matteo Frigo*, Rutger Fick, Mauro Zucchelli, Samuel Deslauriers-Gauthier, Rachid Deriche. Multi Tissue Modelling of Diffusion MRI Signal Reveals Volume Fraction Bias. ISBI 2020.
- ▶ *Matteo Frigo*, Rutger Fick, Mauro Zucchelli, Samuel Deslauriers-Gauthier, Rachid Deriche. Multi Tissue Multi Compartment Modelling of Single-TE Diffusion MRI. bioRxiv 2021.01.29.428843.

* Former PhD student in the team, now at TRIBVN Healthcare, Paris, France

Filtering dMRI-based tractograms

The second contribution of this thesis is motivated by the non-quantitative nature of dMRI-based tractography and the remarkable amount of false positive connections that are produced in the process. The class of the techniques that aim at tackling both problems at the same time in a top-down fashion goes under the name of *tractography filtering techniques* (TFTs). In Chapter 4 we analysed some well known TFTs such as SIFT (R. E. Smith et al. 2013), SIFT2 (R. E. Smith et al. 2015a), COMMIT (Daducci et al. 2014) and COMMIT2 (Schiavi et al. 2020). We performed a systematic comparison of their effects on the graph-theoretical analysis of the resulting connectomes. We showed how these TFTs change the topology of the estimated structural brain networks both in healthy subjects and in patients affected by traumatic brain injury. The analysis was limited to SIFT2 and COMMIT, but, more recently, several other TFTs have been proposed. Further studies are necessary to assess the topological effects of these novelties, including the one presented in this thesis and hereafter mentioned.

Contribution #2

Our third contribution is inspired by the causality relation that links the functioning and the structure of the brain (Deslauriers-Gauthier et al. 2020). In this thesis we proposed a TFT that overtakes the pure structure-based approach followed by the previous TFTs by integrating functional information in the process. In practice, the regularisation term of the optimisation problem associated to the technique has been re-designed to promote the retention of bundles of streamlines that connect regions exhibiting higher functional connectivity. An alternative possibility (that was not explored in this thesis) would be to move the functional prior knowledge from the regularisation term to the data fitting term. This would require to re-design the forward model in such a way that it fits functional data (e.g., fMRI, M/EEG) instead of dMRI or some transformation of it.

Contribution #3

One of the outcomes of our research in the field of TFTs is the TALON Python package, which allows to apply TFTs from the unified TFT framework defined in this work.

Perspectives A significant limitation that we encountered in the process of validating the proposed TFT was the lack of a synthetic dataset explicitly designed for experiments that involve structural and functional connectivity at the same time. Further researches should be addressed towards this direction to provide a fundamental tool that could pave the way to in-vitro studies of the relation between the structure and the functioning of the brain, which is evident but still not fully comprehended. Also, the concept of functional connectivity to be used in the functionally informed tractography filtering technique needs to be adjusted with respect to the one used in the classical resting-state functional connectivity analysis. Among others, the current inability to take into account indirect connections and systems that exhibit relatively low functional connectivity (e.g., the motor system) are two of the challenges that need to be tackled. On this note, the presented FIT framework showed that we have the necessary tools for including any network-based information in the tractography filtering process. This could be exploited to promote

the retention of connections that we know to be biologically plausible. Overall, the performed experiments highlighted how each TFT provides a unique insight into the composition of the brain network. It is possible that an optimal filtering pipeline integrates more than one technique, in such a way to leverage the pros and mitigate the cons of each of them. Further studies are necessary to verify the validity of this hypothesis.

Part of this work was published in the Journal of Neural Engineering and other parts were presented at international conferences.

- ▶ *Matteo Frigo*, Samuel Deslauriers-Gauthier, Drew Parker, Abdol Aziz Ould Ismail, Junghoon John Kim, Ragini Verma, Rachid Deriche. "Diffusion MRI tractography filtering techniques change the topology of structural connectomes." *Journal of Neural Engineering* 17.6 (2020): 065002.
- ▶ *Matteo Frigo*, Guillermo Gallardo, Isa Costantini, Alessandro Dadducci, Demian Wassermann, Rachid Deriche, Samuel Deslauriers-Gauthier. "Reducing false positive connection in tractograms using joint structure-function filtering." 24th Meeting of the Organization for Human Brain Mapping. Singapore, 2018.
- ▶ *Matteo Frigo*, Isa Costantini, Rachid Deriche, Samuel Deslauriers-Gauthier. "Resolving the crossing/kissing fiber ambiguity using Functionally Informed COMMIT." *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, Cham, 2019. In: Elisenda Bonet-Carne et al., (Eds.) *Computational Diffusion MRI: International MICCAI Workshop*, Granada, Spain, September 2018. Springer, 2019.
- ▶ *Matteo Frigo*, Mauro Zucchelli, Rachid Deriche, Samuel Deslauriers-Gauthier. *TALON: Tractograms As Linear Operators in Neuroimaging*. Python package. 2020. gitlab.inria.fr/cobcom/talon

Brain Alignment and Similarity

The fourth and last contribution of this thesis tackles a fundamental question of brain connectivity analysis: given a brain parcellation, does it yield a connectome whose topology is consistent across subjects? To tackle this problem we proposed two novel mathematical objects. First, we defined the graph Jaccard index, a novel similarity measure between networks that exhibits desirable mathematical properties and arises as a natural generalisation of the Jaccard similarity index between sets. The second proposed tool is a new graph alignment technique that relies solely on topological information. Its design is inspired by a generalisation to weighted graphs of the Weisfeiler-Lehman algorithm for the graph isomorphism problem that we provide in the thesis. The algorithm is called WL-align and we showed that it has superior performance with respect to the state-of-the-art technique FAQ. Our results suggest that the choice of the parcellation significantly affects the "alignability" of structural connectomes estimated with dMRI-based tractography. In particular, atlases with higher granularity, i.e., fewer parcels, allow a better alignment between subjects with respect to atlases with a higher number of parcels. Moreover, atlases defined with different criteria (morphology, structural, functional or multi-modal) exhibit substantially different performance, with the Desikan atlas (morphological) performing better than

Contribution #4

the structure-based Gallardo parcellation, whose performance is in line with the multi-modal Glasser atlas and systematically superior to the Schaefer parcellation, which is defined based on the functional organisation of the brain network.

Perspectives We used the assumption that subjects from the same population should have structural connectomes with strongly coherent topology and for each atlas we measured this coherence through the similarity between the aligned networks. Hypothetically, one could turn the argument upside down and design a parcellation that maximises the alignability. Preparing this experiment was out of the scope of this thesis. The study is restricted to structural brain networks, but nothing prevents from investigating the same properties in functional networks, as the proposed methodology applies straightforwardly to static functional connectomes. In case one would want to study the alignability of dynamic functional connectomes, a revision of the whole mathematical framework would be necessary. The presented analysis is based on experiments including only subjects from a single cohort and for a specific selection of the involved parameters. A wider exploration of the parameter space and of the subject cohorts (including pathology) would be necessary to strengthen or confute the conveyed message, but there already is evidence of substantial atlas-based differences in the topology of structural connectomes. From another perspective, the proposed brain alignment framework could be used to evaluate the alignability of connectomes obtained with a certain tractography pipeline. This can be studied by fixing the atlas and analysing the similarity between the networks obtained from tractograms computed with different techniques. In that case, the research question would be: given a tractography pipeline, how robustly does it capture the network topology across different subjects? This would allow to get deeper insights into the effects of tractography filtering techniques on the topology of the retrieved connectomes. Iterating between the two questions (the latter and the one answered in Chapter 5) could in turn help defining the combination of tractography pipeline and cortical atlas that yields connectomes with highest alignability.

A final note on the overall relevance of the proposed methods is due. The graph Jaccard index and the WL-align technique do not make use of any brain-related heuristic. Instead, they provide the mathematical tools that we exploit in the investigation towards answering the initial neuroscientific question. For this reason, the methodological side of what presented in Chapter 5 has the potential to be applied also in other sub-domains of complex systems and network analysis, potentially in combination with domain-specific heuristics.

This work has recently been accepted for publication in *Network Neuroscience*.

- ▶ *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. 2021. In press, *Network Neuroscience*.

- ▶ Emilio Cruciani, *Matteo Frigo*, David Coudert, Rachid Deriche, Emanuele Natale, Samuel Deslauriers-Gauthier. WL-align, Python package. Open Science Framework. 2020. <https://osf.io/depux/>.

Open Science

During the whole process that brought to the preparation of this thesis we did our best to observe the principles of open science that aim at making science more reproducible and accessible. On top of providing unambiguous description of the methodology used in each presented analysis, we made available the code and data necessary to reproduce our most relevant experiments.

- ▶ The employed synthetic datasets were generated using publicly available software:
 - Chapter 3, Microstructure: Dmipy (Fick et al. 2019) for simulations, Dipy (Garyfallidis et al. 2014) for noise generation.
 - Chapter 4, Tractography Filtering: Phantomas (Caruyer et al. 2014) for generation of crossing/kissing phantom and of the ISBI2013 dataset.
- ▶ All chapters include experiments based on real data obtained from the HCP database and publicly available at the Connectome Coordination Facility (WU-Minn Human Connectome Project consortium 2017; Van Essen et al. 2012).
- ▶ Pre-processing of dMRI data was performed using FSL (Jenkinson et al. 2012) and Mrtrix3 (Tournier et al. 2019).
- ▶ Tractography was performed using Mrtrix3 (Tournier et al. 2019).
- ▶ The cortical meshes and the Desikan atlas (Desikan et al. 2006) were computed using Freesurfer (Fischl 2012).
- ▶ Tractography filtering was performed using Mrtrix3 (Tournier et al. 2019), TALON (Frigo et al. 2021c) and COMMIT (Daducci et al. 2014).
- ▶ Brain alignments were computed using an in-house implementation of WL-align available at <https://osf.io/depux/> and the implementation of FAQ provided by Graspologic (J. Chung et al. 2019).
- ▶ Post-processing, analysis and visualisation of the results were curated using Mrtrix3 (Tournier et al. 2019) and Python scripts based on Numpy, Scipy (Harris et al. 2020) and Matplotlib (Hunter 2007).

Appendix **A**

Software contributions

Developing the solutions presented in Chapters 3, 4 and 5 required a significant effort in terms of software design and implementation. For several standard tasks we relied on standard tools such as Mrtrix3 (Tournier et al. 2019), FSL (Jenkinson et al. 2012), Dmipy (Fick et al. 2019) and Dipy (Garyfallidis et al. 2014). Diving into the user experience of these software allowed us to spot some missing features or minor flaws to whose implementation or correction we contributed in first person. Additionally, we designed and implemented a novel software for tractography filtering. In the following paragraphs we are going to present these software contributions.

A.1 TALON: Tractograms As Linear Operators in Neuroimaging

The TALON package provides the necessary tools to design tractography filtering techniques and solve the associated inverse problem. At its core, it provides a way to transform a tractogram into a linear operator, or more precisely a sparse matrix. This matrix can be used in two ways: to generate data and to explain data. In both cases, the type of the data is arbitrary and is specified by the user, not by TALON. In relation to what presented in Chapter 4, TALON provides the framework for defining the forward operator required in the definition of a TFT and solving the associated inverse problem. The TALON package is written in Python 3 (Van Rossum et al. 2009), follows the PEP8 code style standard (Rossum et al. 2001) and includes both unit and functional tests. It is distributed under MIT license and publicly available on PyPI¹. The author of this thesis is the main developer and maintainer of the package together with Samuel Deslauriers-Gauthier. The design of the package is takes inspiration from the COMMIT package (Daducci et al. 2014), of whose development the author of this thesis have contributed in the past, and from the fixel file format of Mrtrix3 (Tournier et al. 2019).

1: <https://pypi.org/project/cobcom-talon/>

In order to filter a tractogram with TALON, one has to design the four elements of the unified tractography filtering framework proposed in Chapter 4:

- ▶ data y to be fitted;
- ▶ forward model mapping streamlines onto data through a forward operator A ;
- ▶ regularisation term promoting solutions having a property that minimises $\Omega : \mathbb{R}^n \rightarrow \mathbb{R}$;
- ▶ constraints encoded through the use of indicator functions ι .

These four elements are then combined in the inverse problem solved by the following optimisation problem:

$$x^* = \underset{x}{\operatorname{argmin}} \|Ax - y\|_2^2 + \lambda\Omega(x) + \iota(x) \quad (\text{A.1})$$

where $\lambda \in \mathbb{R}^+$ acts as regularisation parameter. The optimisation problem is solved with PyUNLocBoX (Defferrard et al. 2017).

Data

The TALON package relies on external tools such as Nibabel (Brett et al. 2020) and the Scipy ecosystem (Harris et al. 2020) for loading the data to be fitted and pre-processing them.

Forward operator

The linear operators in TALON are defined as follows. First, the tractogram is *voxelized* by separating each streamline into voxel elements. If tractography generates streamlines by following the local directions in a 3D image, voxelizing a tractogram is the opposite, i.e., extracting directions from streamlines. Given a template of ν directions $\{\mathbf{u}_k \in \mathbb{R}^3\}$, for each streamline s and voxel ν we estimate the direction $i(\nu, s) = k \in [1, \nu]$ that is closest to the one of s in voxel ν and the length of the corresponding segment $w(\nu, s)$. The two matrices i and w have the same sparsity pattern. The way in which each streamline segment is projected onto the data is obtained by means of the so-called *generators* $g(u) : \mathbb{R}^3 \rightarrow \mathbb{R}^d$ (where d is the number of modelled data points per voxel), which are the response functions associated to each direction \mathbf{u}_k obtained, e.g., via Dmipy (Fick et al. 2019). The signal in voxel ν is then defined as

$$\hat{y}(\nu) = \sum_{s \in \mathcal{S}} x_s \cdot w(\nu, s) g(i(\nu, s)). \quad (\text{A.2})$$

where x_s is the contribution of streamline s to the model. The vertical concatenation of this model for each considered voxel gives the wanted application of the linear operator A to vector x . The (g, i, w) data structure can be used for defining also peak-specific or voxel-specific linear operator as in the MIT paradigm presented in Chapter 4.

The TALON package provides three alternative implementations of this (g, i, w) data structure. A very memory-efficient but slow CPU implementation is included for reference. Also, a slightly less memory-efficient version but much faster CPU implementation and a memory- and time-efficient GPU version are included. The CPU versions exploit the multi-core capabilities of Numpy (Harris et al. 2020), while the GPU version is based on PyOpencl (Klöckner et al. 2012).

Regularisation Term

The regularisation needs to be formulated as a particular case of the following function:

$$\Omega(x) = \sum_{g \in \mathcal{G}} w_g \|x_g\|_2 \quad (\text{A.3})$$

which was already given in Equation (4.8) and includes the ℓ_1 sparsity as a particular case. Setting the regularisation parameter to $\lambda = 0$ yields the standard (possibly non-negative) least squares formulation that characterises COMMIT.

Constraints

In TALON it is possible to include the hard non-negativity constraint $x \geq 0$.

A.2 WL-align

In Chapter 5 we described the WL-align algorithm, which is a graph alignment technique that makes use only of the network's topology. Our Python implementation of such routine was made public as part of the supplementary materials of the journal article where WL-align was first introduced (Frigo et al. 2021b). It is available at the Open Science Framework² (Frigo et al. 2020a). This software was developed in collaboration with Emilio Cruciani.

2: <https://osf.io/depux>

A.3 COMMIT

At the beginning of the doctoral program we designed and implemented the *solvers* module for the COMMIT framework³. The current implementation contains corrections and changes merged by the current maintainers. The author of this thesis discontinued contributing to the package after starting to develop TALON.

3: <https://github.com/daducci/COMMIT/pull/36>

A.4 PyUNLocBoX: Optimization by Proximal Splitting

PyUNLocBoX is a Python package which uses proximal splitting methods to solve non-differentiable convex optimization problems. It is a free software, distributed under the BSD license, and available on PyPI. PyUNLocBoX is a core element of TALON, as it is used for solving the optimisation problem associated to the designed TFT. We contributed to the development of PyUNLocBoX by adding the function Ω of Equation (A.3) to the ones that can be employed directly from the package⁴. The author of this thesis is an occasional contributor to the package.

4: <https://github.com/epfl-lts2/pyunlocbox/pull/26>

A.5 Dmipy: Diffusion Microstructure Imaging in Python

The Dmipy package (Fick et al. 2019) is one of those software without which this thesis would have had a hard time even starting. It is at the core of the works presented in Chapter 3 and of TALON. We contributed by solving several minor compatibility problems and bugs, by adding the multi-tissue formulation of the tortuosity assumption (see Chapter 3) and by initiating the development of the generalised AMICO implementation, which will be presented in future works. The author of this thesis is an active maintainer of the package.

A.6 Mrtrix3

Mrtrix3 is the reference software used for the processing of dMRI data and the extraction of tractograms. Additionally, it has been used for running the experiments involving SIFT2 presented in Chapter 4. We contributed to the development of the 3.0.0 by creating the new program `connectomeedit` for performing basic operations on connectome matrix numerical data, such as transposition, symmetrization and extraction of upper and lower triangular parts⁵. We also reported minor bugs and incompatibilities. The author of this thesis is an occasional contributor to the software.

5: <https://github.com/MRtrix3/mrtrix3/pull/1725>

Appendix **B**

Other works

The study of the state-of-the-art problems and methods of microstructure, tractography and brain network analysis gave us the opportunity to contribute to projects that do not strictly fit in this thesis. These studies are hereafter briefly presented.

B.1 Structure-function mapping

In this work we focused on the ability of mappings based on eigenmodes to predict the complete functional connectivity matrices from the full structural connectivity matrix. Several of these maps based on eigenmodes have been proposed in the past, but the literature lacked a unified view on the problem that allowed to compare these methods. We proposed and analysed a unified framework that yields the pre-existing models as particular cases. The experiments performed on a cohort of healthy subjects from the HCP database highlighted the limitations of such approaches. In particular, we unveiled the presence of a glass ceiling in the prediction performance. We showed how this limitation is related to the high similarity between the structural connectomes of the subjects in the studied cohort, suggesting that eigenmode-based techniques could be exploited in the context of patho-connectomics, where the connectomes of patients and healthy controls are expected to exhibit lower similarity. This work was published in *Medical Image Analysis* Deslauriers-Gauthier et al. 2020.

- ▶ Samuel Deslauriers-Gauthier, Mauro Zucchelli, *Matteo Frigo*, Rachid Deriche. A Unified Framework for Multimodal Structure-function Mapping Based on Eigenmodes. *Medical Image Analysis* 66 (2020): 101799.

B.2 dMRI-PLI resolution gap

A second project concerned the comparison between the homological properties of structural connectomes obtained from dMRI and 3D polarised light imaging (PLI) data. The latter is a technique that allows to reconstruct fiber orientation distribution functions (fODFs) at a resolution which is some order of magnitudes more refined than the one of dMRI Alimi et al. 2020a, therefore closer to the actual resolution of the geometry of the axonal pathways, which makes it a good candidate validation technique for tractography. In this study we explored the persistence of the number of connected components on structural connectomes obtained with data at several resolutions spanning the range from the PLI to the

dMRI resolution. We showed how the weak connections in a structural network obtained at dMRI resolution play a fundamental role in the definition of connectomes with homological properties similar to the ones of connectomes obtained at PLI resolution. This work was presented at the 2020 meeting of the International Society for Magnetic Resonance in Medicine Alimi et al. 2020b.

- ▶ Abib Alimi, *Matteo Frigo*, Samuel Deslauriers-Gauthier and Rachid Deriche. Quantitative assessment of multi-scale tractography: bridging the resolution gap with 3D-PLI. ISMRM 2020.

B.3 Multi-Compartment models of Gray Matter

In an effort to identify tissue microstructure-based biomarkers that are sensitive to the progression of Alzheimer’s disease (AD), the Alzheimer’s Disease Neuroimaging Initiative (ADNI-3) recently introduced multi-shell dMRI protocols (Weiner et al. 2017). We collaborated with Julio Villalon-Reina and Talia Nir under the supervision of Paul Thompson in the systematic evaluation of different dMRI microstructure models (both single-tissue MC models and diffusion-propagator based models) in GM, which is the region where AD can be better observed. We identified which of the derived scalar measures best classify individuals with mild cognitive impairment versus healthy controls. Remarkably, we showed that the simpler DTI-based measures were among the best predictors of cognitive impairment together with the intra-axonal volume fraction estimated via SMT (Kaden et al. 2016). The resulting analysis was presented at ISMRM 2020 (Villalon-reina et al. 2020). Subsequent studies (which we did not co-author) showed how an adaptation on GM of the multi-tissue standard model of dMRI in the WM that we presented in Chapter 3 provides better classification capability between AD patients and healthy controls (Nir et al. 2020; Reina et al. 2020) with respect to the single-tissue correspondent.

- ▶ Julio E. Villalon-Reina, Talia M. Nir, Sophia I. Thomopoulos, Lauren E. Salminen, Neda Jahanshad, Rutger Fick, *Matteo Frigo*, Rachid Deriche, Paul M. Thompson and ADNI, Tracking microstructural biomarkers of Alzheimer’s disease via advanced multi-shell diffusion MRI scalar measures. ISMRM 2020.

B.4 MEMENTO Challenge

During the International Symposium on Biomedical Imaging of 2020 we took part in the diffusion Mri whitE Matter rEcoNstrucTiOn (MEMENTO) challenge*. Most of the software used for solving the problems posed in this challenge is based on Dmipy (Fick et al. 2019), Dipy (Garyfallidis et al. 2014) and Pytorch†. The challenge was organised in three sub-challenges, each dedicated to a specific problem of dMRI signal reconstruction:

1. predict unseen dMRI signal;
2. estimate microstructural measures;

* <https://my.vanderbilt.edu/memento/>

† <https://pytorch.org/>

3. evaluate sensitivity and specificity of potential biomarkers.

The reader will find in [the official website of the challenge](#) a detailed presentation of the submissions.

The first challenge included in-vivo brain data sampled with PGSE sequences (both multi-shell and DSI) in the human volunteer of the MASSIVE dataset (Froeling et al. 2017), and Double Diffusion Encoding (DDE) and Double Oscillating Diffusion Encoding (DODE) mice data acquired ex-vivo. We were provided with a subsampled set of the acquired data and are asked to predict the remaining – unseen – data. For the PGSE data, we extrapolated predictions from several microstructural models such as MAP-MRI (Fick et al. 2016) and the multi-compartment models presented in Chapter 3. Additionally, we trained a simple artificial neural network with a single-layer made of 30 neurons whose input is the table of acquisition parameters and the output is the predicted signal. On the PGSE DSI dataset we obtained the best overall prediction with MAP-MRI. The artificial neural network gave us the first position in the prediction of DDE and DODE signal.

The second challenge aimed at understanding the current ability in modelling WM tissue microstructure with dMRI. The given dataset consisted of a number of simulated WM environments at voxel scale, generated by changing in a controlled fashion a range of microstructural features, including:

- ▶ self-diffusivity (D , um^2/ms);
- ▶ intra-fibre volume fraction (f , unitless);
- ▶ fibre radius index (a , um);
- ▶ permeability (k , um/ms).

The participants were given the simulated PGSE, DDE and DODE signals and asked to estimate the above microstructural features or some biomarker sensitive to changes of these quantities. We submitted estimators of each parameter using multi-compartment models and MAP-MRI-derived indices. We classified first in the estimation of self-diffusivity, while for fiber radius index and permeability we provided the only meaningful (yet strongly biased) submissions.

The third challenge was designed to assess the ability of dMRI signal models to estimate the fiber orientation dispersion and the axonal diameter distribution from simulated PGSE multi-shell data. To estimate the fiber dispersion we modelled the signal as the convolution of a stick compartment with a SH function of order 8, fixed the parallel diffusivity of the stick and fitted the SH coefficients. As far as the axonal diameter distribution is concerned, we modelled the signal as the convolution of a Watson distribution (accounting for fiber dispersion) and a cylinder compartment with Gamma-distributed axonal diameters. We provided the only submissions for both the parts of this sub-challenge.

Bibliography

References in alphabetical order.

- Afzali, Maryam et al. (2020). ‘The sensitivity of diffusion MRI to microstructural properties and experimental factors’. In: *Journal of Neuroscience Methods*, p. 108951. DOI: [10.1016/j.jneumeth.2020.108951](https://doi.org/10.1016/j.jneumeth.2020.108951) (cited on page 30).
- Alexander, Daniel C et al. (2010). ‘Orientationally invariant indices of axon diameter and density from diffusion MRI’. In: *Neuroimage* 52.4, pp. 1374–1389. DOI: [10.1016/j.neuroimage.2010.05.043](https://doi.org/10.1016/j.neuroimage.2010.05.043) (cited on pages 29, 32).
- Alexander, Daniel C et al. (2019). ‘Imaging brain microstructure with diffusion MRI: practicality and applications’. In: *NMR in Biomedicine* 32.4, e3841. DOI: [10.1002/nbm.3841](https://doi.org/10.1002/nbm.3841) (cited on page 30).
- Alimi, Abib et al. (2020a). ‘Analytical and fast Fiber Orientation Distribution reconstruction in 3D-Polarized Light Imaging’. In: *Medical Image Analysis* 65, p. 101760. DOI: [10.1016/j.media.2020.101760](https://doi.org/10.1016/j.media.2020.101760) (cited on page 114).
- Alimi, Abib et al. (2020b). ‘Quantitative assessment of multi-scale tractography: bridging the resolution gap with 3D-PLI’. In: *ISMRM & SMRT Virtual Conference & Exhibition* (cited on page 115).
- Anastasopoulos, C et al. (2014). ‘Local and global fiber tractography in patients with epilepsy’. In: *American Journal of Neuroradiology* 35.2, pp. 291–296. DOI: [10.3174/ajnr.A3752](https://doi.org/10.3174/ajnr.A3752) (cited on page 18).
- Andersson, Jesper LR, Stefan Skare, and John Ashburner (2003). ‘How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging’. In: *Neuroimage* 20.2, pp. 870–888. DOI: [10.1016/S1053-8119\(03\)00336-7](https://doi.org/10.1016/S1053-8119(03)00336-7) (cited on page 63).
- Andersson, Jesper LR and Sotgiu N Sotiropoulos (2016). ‘An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging’. In: *Neuroimage* 125, pp. 1063–1078. DOI: [10.1016/j.neuroimage.2015.10.019](https://doi.org/10.1016/j.neuroimage.2015.10.019) (cited on page 64).
- Andersson, JLR et al. (2012). ‘A comprehensive Gaussian process framework for correcting distortions and movements in diffusion images’. In: *Proceedings of the 20th Annual Meeting of ISMRM*. Vol. 20, p. 2426 (cited on page 63).
- Anwander, Alfred et al. (2007). ‘Connectivity-based parcellation of Broca’s area’. In: *Cerebral cortex* 17.4, pp. 816–825. DOI: [10.1093/cercor/bhk034](https://doi.org/10.1093/cercor/bhk034) (cited on page 22).
- Arsigny, Vincent et al. (2006). ‘Log-Euclidean metrics for fast and simple calculus on diffusion tensors’. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 56.2, pp. 411–421. DOI: [10.1002/mrm.20965](https://doi.org/10.1002/mrm.20965) (cited on page 15).
- Assaf, Yaniv and Peter J Basser (2005). ‘Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain’. In: *Neuroimage* 27.1, pp. 48–58. DOI: [10.1016/j.neuroimage.2005.03.042](https://doi.org/10.1016/j.neuroimage.2005.03.042) (cited on page 29).

- Assaf, Yaniv et al. (2008). ‘AxCaliber: a method for measuring axon diameter distribution from diffusion MRI’. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 59.6, pp. 1347–1354. DOI: [10.1002/mrm.21577](https://doi.org/10.1002/mrm.21577) (cited on page 29).
- Astolfi, Pietro et al. (2020). ‘Tractogram filtering of anatomically non-plausible fibers with geometric deep learning’. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 291–301. DOI: [10.1007/978-3-030-59728-3_29](https://doi.org/10.1007/978-3-030-59728-3_29) (cited on page 52).
- Auzias, Guillaume, Olivier Coulon, and Andrea Brovelli (2016). ‘MarsAtlas: a cortical parcellation atlas for functional mapping’. In: *Human brain mapping* 37.4, pp. 1573–1592. DOI: [10.1002/hbm.23121](https://doi.org/10.1002/hbm.23121) (cited on page 21).
- Ayache, Nicholas and Bernard Faverjon (1987). ‘Efficient registration of stereo images by matching graph descriptions of edge segments’. In: *International Journal of Computer Vision* 1.2, pp. 107–131. DOI: [10.1007/BF00123161](https://doi.org/10.1007/BF00123161) (cited on page 82).
- Bajada, Claude J et al. (2017). ‘A graded tractographic parcellation of the temporal lobe’. In: *Neuroimage* 155, pp. 503–512. DOI: [10.1016/j.neuroimage.2017.04.016](https://doi.org/10.1016/j.neuroimage.2017.04.016) (cited on page 22).
- Barak, Boaz et al. (2019). ‘(Nearly) Efficient Algorithms for the Graph Matching Problem on Correlated Random Graphs’. In: *Advances in Neural Information Processing Systems* 32. Ed. by H. Wallach et al. Curran Associates, Inc., pp. 9190–9198. (Visited on 06/08/2020) (cited on page 82).
- Barakovic, Muhamed et al. (2020). ‘Resolving bundle-specific intra-axonal T2 values within a voxel using diffusion-relaxation tract-based estimation’. In: *NeuroImage*, p. 117617. DOI: [10.1016/j.neuroimage.2020.117617](https://doi.org/10.1016/j.neuroimage.2020.117617) (cited on pages 47, 55).
- Basser, Peter J et al. (2000). ‘In vivo fiber tractography using DT-MRI data’. In: *Magnetic resonance in medicine* 44.4, pp. 625–632. DOI: [10.1002/1522-2594\(200010\)44:4<625::AID-MRM17>3.0.CO;2-0](https://doi.org/10.1002/1522-2594(200010)44:4<625::AID-MRM17>3.0.CO;2-0) (cited on page 18).
- Bassett, Danielle S and Ed Bullmore (2006). ‘Small-world brain networks’. In: *The neuroscientist* 12.6, pp. 512–523. DOI: [10.1177/1073858406293182](https://doi.org/10.1177/1073858406293182) (cited on pages 2, 26).
- (2017). ‘Small-world brain networks revisited’. In: *The Neuroscientist* 23.5, pp. 499–516. DOI: [10.1177/1073858416667720](https://doi.org/10.1177/1073858416667720) (cited on pages 1, 2, 20, 26).
- Batchelor, Phillipp G et al. (2005). ‘A rigorous framework for diffusion tensor calculus’. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 53.1, pp. 221–225. DOI: [10.1002/mrm.20334](https://doi.org/10.1002/mrm.20334) (cited on page 15).
- Bayati, Mohsen et al. (Mar. 2013). ‘Message-Passing Algorithms for Sparse Network Alignment’. In: *ACM Trans. Knowl. Discov. Data* 7.1, 3:1–3:31. DOI: [10.1145/2435209.2435212](https://doi.org/10.1145/2435209.2435212) (cited on page 82).
- Beck, Amir and Marc Teboulle (2009). ‘A fast iterative shrinkage-thresholding algorithm for linear inverse problems’. In: *SIAM journal on imaging sciences* 2.1, pp. 183–202. DOI: [10.1137/080716542](https://doi.org/10.1137/080716542) (cited on pages 54, 58).

- Becker, Cassiano O et al. (2018). 'Spectral mapping of brain functional connectivity from diffusion imaging'. In: *Scientific reports* 8.1, pp. 1–15. DOI: [10.1038/s41598-017-18769-x](https://doi.org/10.1038/s41598-017-18769-x) (cited on page 82).
- Behrens, Timothy EJ et al. (2003). 'Characterization and propagation of uncertainty in diffusion-weighted MR imaging'. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 10.1002/mrm.10609.5, pp. 1077–1088 (cited on pages 29, 32).
- Benjamini, Yoav and Yosef Hochberg (1995). '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, pp. 289–300 (cited on page 66).
- Bloch, Felix (1946). 'Nuclear induction'. In: *Physical review* 70.7-8, p. 460 (cited on pages 9, 10).
- Blondel, Vincent D et al. (2008). 'Fast unfolding of communities in large networks'. In: *Journal of statistical mechanics: theory and experiment* 2008.10, P10008. DOI: [10.1088/1742-5468/2008/10/p10008](https://doi.org/10.1088/1742-5468/2008/10/p10008) (cited on page 65).
- Brett, Matthew et al. (Nov. 2020). 'Nibabel'. In: DOI: [10.5281/zenodo.4295521](https://doi.org/10.5281/zenodo.4295521) (cited on page 111).
- Brodmann, K. (1909). *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Leipzig: Barth (cited on page 22).
- Bullmore, Ed and Olaf Sporns (Mar. 2009). 'Complex brain networks: graph theoretical analysis of structural and functional systems'. eng. In: *Nature Reviews Neuroscience* 10.3, p. 186. DOI: [10.1038/nrn2575](https://doi.org/10.1038/nrn2575) (cited on page 26).
- Caiafa, Cesar F and Franco Pestilli (2017). 'Multidimensional encoding of brain connectomes'. In: *Scientific reports* 7.1, pp. 1–13. DOI: [10.1038/s41598-017-09250-w](https://doi.org/10.1038/s41598-017-09250-w) (cited on page 57).
- Campbell, Alfred Walter (1905). *Histological studies on the localisation of cerebral function*. University Press (cited on page 22).
- Candes, Emmanuel J, Michael B Wakin, and Stephen P Boyd (2008). 'Enhancing sparsity by reweighted ℓ_1 minimization'. In: *Journal of Fourier analysis and applications* 14.5-6, pp. 877–905. DOI: [10.1007/s00041-008-9045-x](https://doi.org/10.1007/s00041-008-9045-x) (cited on page 57).
- Caruyer, Emmanuel et al. (2013). 'Design of multishell sampling schemes with uniform coverage in diffusion MRI'. In: *Magnetic resonance in medicine* 69.6, pp. 1534–1540. DOI: [10.1002/mrm.24736](https://doi.org/10.1002/mrm.24736) (cited on page 72).
- Caruyer, Emmanuel et al. (2014). 'Phantoms: a flexible software library to simulate diffusion MR phantoms'. In: *Ismrm* (cited on pages 72, 73, 80, 109).
- Charikar, Moses S. (2002). 'Similarity Estimation Techniques from Rounding Algorithms'. In: *Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing*. STOC '02. Montreal, Quebec, Canada: Association for Computing Machinery, pp. 380–388. DOI: [10.1145/509907.509965](https://doi.org/10.1145/509907.509965) (cited on page 86).
- Christiaens, Daan et al. (2015). 'Global tractography of multi-shell diffusion-weighted imaging data using a multi-tissue model'. In: *Neuroimage* 123, pp. 89–101. DOI: [10.1016/j.neuroimage.2015.08.008](https://doi.org/10.1016/j.neuroimage.2015.08.008) (cited on pages 18, 51).

- Chuhutin, Andrey, Brian Hansen, and Sune Nørhøj Jespersen (2017). ‘Precision and accuracy of diffusion kurtosis estimation and the influence of b-value selection’. In: *NMR in Biomedicine* 30.11. DOI: [10.1002/nbm.3777](https://doi.org/10.1002/nbm.3777) (cited on page 16).
- Chung, Jaewon et al. (2019). ‘GraSPy: Graph Statistics in Python’. In: *Journal of Machine Learning Research* 20.158, pp. 1–7 (cited on pages 92, 109).
- Chung, Moo K et al. (2017). ‘Topological distances between brain networks’. In: *International Workshop on Connectomics in Neuroimaging*. Springer, pp. 161–170. DOI: [10.1007/978-3-319-67159-8_19](https://doi.org/10.1007/978-3-319-67159-8_19) (cited on pages 82, 84).
- Civier, Oren et al. (2019). ‘Is removal of weak connections necessary for graph-theoretical analysis of dense weighted structural connectomes from diffusion MRI?’ In: *NeuroImage* 194, pp. 68–81. DOI: [10.1016/j.neuroimage.2019.02.039](https://doi.org/10.1016/j.neuroimage.2019.02.039) (cited on pages 62, 69–71).
- Conte, D. et al. (May 2004). ‘Thirty years of graph matching in pattern recognition’. In: *International Journal of Pattern Recognition and Artificial Intelligence* 18.03, pp. 265–298. DOI: [10.1142/S0218001404003228](https://doi.org/10.1142/S0218001404003228) (cited on page 82).
- Cook, Steven J et al. (2019). ‘Whole-animal connectomes of both *Caenorhabditis elegans* sexes’. In: *Nature* 571.7763, pp. 63–71. DOI: [10.1038/s41586-019-1352-7](https://doi.org/10.1038/s41586-019-1352-7) (cited on page 21).
- Côté, Marc-Alexandre et al. (2013). ‘Tractometer: towards validation of tractography pipelines’. In: *Medical image analysis* 17.7, pp. 844–857. DOI: [10.1016/j.media.2013.03.009](https://doi.org/10.1016/j.media.2013.03.009) (cited on pages 52, 72).
- Craddock, R Cameron et al. (2012). ‘A whole brain fMRI atlas generated via spatially constrained spectral clustering’. In: *Human brain mapping* 33.8, pp. 1914–1928. DOI: [10.1002/hbm.21333](https://doi.org/10.1002/hbm.21333) (cited on page 22).
- Daducci, Alessandro and COMMIT contributors (Dec. 2020). *GitHub repository of the COMMIT Python package*. URL: <https://github.com/daducci/commit> (cited on pages 53, 58).
- Daducci, Alessandro et al. (2013). ‘A convex optimization framework for global tractography’. In: *2013 IEEE 10th International Symposium on Biomedical Imaging*. IEEE, pp. 524–527 (cited on page 4).
- Daducci, Alessandro et al. (2014). ‘COMMIT: convex optimization modeling for microstructure informed tractography’. In: *IEEE transactions on medical imaging* 34.1, pp. 246–257. DOI: [10.1109/TMI.2014.2352414](https://doi.org/10.1109/TMI.2014.2352414) (cited on pages 4, 29, 40, 47, 51, 53, 54, 56, 74, 102, 106, 109, 110).
- Daducci, Alessandro et al. (2016). ‘Microstructure informed tractography: pitfalls and open challenges’. In: *Frontiers in neuroscience* 10, p. 247. DOI: [10.3389/fnins.2016.00247](https://doi.org/10.3389/fnins.2016.00247) (cited on page 49).
- Deen, Ben, Naomi B Pitskel, and Kevin A Pelphrey (2011). ‘Three systems of insular functional connectivity identified with cluster analysis’. In: *Cerebral cortex* 21.7, pp. 1498–1506. DOI: [10.1093/cercor/bhq186](https://doi.org/10.1093/cercor/bhq186) (cited on page 22).
- Defferrard, Michaël, Rodrigo Pena, and Nathanaël Perraudin (Dec. 2017). ‘PyUNLocBoX: Optimization by Proximal Splitting’. In: DOI: [10.5281/zenodo.1199082](https://doi.org/10.5281/zenodo.1199082) (cited on page 111).
- Dell’Acqua, Flavio and J-Donald Tournier (2019). ‘Modelling white matter with spherical deconvolution: How and why?’ In: *NMR in Biomedicine* 32.4. DOI: [10.1002/nbm.3945](https://doi.org/10.1002/nbm.3945) (cited on pages 17, 33).

- Deriche, Rachid (2016). 'Computational brain connectivity mapping: A core health and scientific challenge'. In: *Medical Image Analysis* 33, pp. 122–126. DOI: [10.1016/j.media.2016.06.003](https://doi.org/10.1016/j.media.2016.06.003) (cited on page 1).
- Descoteaux, Maxime et al. (2007). 'Regularized, fast, and robust analytical Q-ball imaging'. In: *Magnetic Resonance in Medicine* 58.3, pp. 497–510. DOI: [10.1002/mrm.21277](https://doi.org/10.1002/mrm.21277) (cited on pages 16, 18, 38).
- Descoteaux, Maxime et al. (2008). 'Deterministic and probabilistic tractography based on complex fibre orientation distributions'. In: *IEEE transactions on medical imaging* 28.2, pp. 269–286. DOI: [10.1109/TMI.2008.2004424](https://doi.org/10.1109/TMI.2008.2004424) (cited on pages 16–18).
- Descoteaux, Maxime et al. (2011). 'Multiple q-shell diffusion propagator imaging'. In: *Medical image analysis* 15.4, pp. 603–621. DOI: [10.1016/j.media.2010.07.001](https://doi.org/10.1016/j.media.2010.07.001) (cited on page 16).
- Desikan, Rahul S. et al. (2006). 'An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest'. In: *NeuroImage* 31.3, pp. 968–980. DOI: [10.1016/j.neuroimage.2006.01.021](https://doi.org/10.1016/j.neuroimage.2006.01.021) (cited on pages 21, 65, 73, 91, 109).
- Deslauriers-Gauthier, Samuel et al. (2019). 'White matter information flow mapping from diffusion MRI and EEG'. In: *NeuroImage* 201, p. 116017. DOI: [10.1016/j.neuroimage.2019.116017](https://doi.org/10.1016/j.neuroimage.2019.116017) (cited on pages 25, 26).
- Deslauriers-Gauthier, Samuel et al. (2020). 'A unified framework for multimodal structure–function mapping based on eigenmodes'. In: *Medical Image Analysis* 66, p. 101799. DOI: [10.1016/j.media.2020.101799](https://doi.org/10.1016/j.media.2020.101799) (cited on pages 73, 82, 83, 106, 114).
- Destrieux, Christophe et al. (2010). 'Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature'. In: *NeuroImage* 53.1, pp. 1–15. DOI: [10.1016/j.neuroimage.2010.06.010](https://doi.org/10.1016/j.neuroimage.2010.06.010) (cited on page 21).
- Dhollander, Thijs and Alan Connelly (2016a). 'A novel iterative approach to reap the benefits of multi-tissue CSD from just single-shell (+ b= 0) diffusion MRI data'. In: *Proc ISMRM*. Vol. 24, p. 3010 (cited on pages 3, 19, 44).
- Dhollander, Thijs, David Raffelt, and Alan Connelly (2016b). '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 (cited on pages 35, 44, 46, 47, 65, 105).
- Diez, Ibai et al. (2015). 'A novel brain partition highlights the modular skeleton shared by structure and function'. In: *Scientific reports* 5, p. 10532. DOI: [10.1038/srep10532](https://doi.org/10.1038/srep10532) (cited on page 22).
- Dijkstra, Edsger W (1959). 'A note on two problems in connexion with graphs'. In: *Numerische mathematik* 1.1, pp. 269–271 (cited on page 26).
- Ding, Song-Lin et al. (2016). 'Comprehensive cellular-resolution atlas of the adult human brain'. In: *Journal of Comparative Neurology* 524.16, pp. 3127–3481. DOI: [10.1002/cne.24080](https://doi.org/10.1002/cne.24080) (cited on page 22).
- Dussik, Karl Theo (1942). 'Über die Möglichkeit, hochfrequente mechanische Schwingungen als diagnostisches Hilfsmittel zu verwerten'. In: *Zeitschrift für die gesamte Neurologie und Psychiatrie* 174.1, pp. 153–168 (cited on page 9).
- Economo, Constantin Freiherr von and Georg N Koskinas (1925). *Die cytoarchitektonik der hirnrinde des erwachsenen menschen*. J. Springer (cited on page 22).

- Eichner, Cornelius et al. (2020). 'Increased sensitivity and signal-to-noise ratio in diffusion-weighted MRI using multi-echo acquisitions'. In: *NeuroImage* 221, p. 117172. DOI: [10.1016/j.neuroimage.2020.117172](https://doi.org/10.1016/j.neuroimage.2020.117172) (cited on page 30).
- Eickhoff, Simon B, BT Thomas Yeo, and Sarah Genon (2018). 'Imaging-based parcellations of the human brain'. In: *Nature Reviews Neuroscience* 19.11, pp. 672–686. DOI: [10.1038/s41583-018-0071-7](https://doi.org/10.1038/s41583-018-0071-7) (cited on page 8).
- Eickhoff, Simon B et al. (2008). 'Organizational principles of human visual cortex revealed by receptor mapping'. In: *Cerebral Cortex* 18.11, pp. 2637–2645. DOI: [10.1093/cercor/bhn024](https://doi.org/10.1093/cercor/bhn024) (cited on page 22).
- Einstein, Albert (1905). 'About the movement of suspended particles in liquids at rest as required by the molecular kinetic theory of heat. (Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen)'. In: *Ann Phys* 322, pp. 549–560 (cited on page 13).
- Fan, Lingzhong et al. (2016). 'The human brainnetome atlas: a new brain atlas based on connectional architecture'. In: *Cerebral cortex* 26.8, pp. 3508–3526. DOI: [10.1093/cercor/bhw157](https://doi.org/10.1093/cercor/bhw157) (cited on page 22).
- Färer, Martin (2017). 'On the Combinatorial Power of the Weisfeiler-Lehman Algorithm'. en. In: *Algorithms and Complexity*. Ed. by Dimitris Fotakis, Aris Pagourtzis, and Vangelis Th. Paschos. Lecture Notes in Computer Science. Cham: Springer International Publishing, pp. 260–271. DOI: [10.1007/978-3-319-57586-5_22](https://doi.org/10.1007/978-3-319-57586-5_22) (cited on page 87).
- Farooq, Hamza et al. (2016). 'Microstructure imaging of crossing (MIX) white matter fibers from diffusion MRI'. In: *Scientific reports* 6.1, pp. 1–9. DOI: [10.1038/srep38927](https://doi.org/10.1038/srep38927) (cited on page 29).
- Feizi, Soheil et al. (2019). 'Spectral Alignment of Graphs'. In: *IEEE Transactions on Network Science and Engineering*, pp. 1–1. DOI: [10.1109/TNSE.2019.2913233](https://doi.org/10.1109/TNSE.2019.2913233) (cited on pages 82, 83).
- Fick, Rutger (2017). 'Advanced dMRI signal modeling for tissue microstructure characterization'. PhD thesis. Université Côte d'Azur (cited on page 8).
- Fick, Rutger, Demian Wassermann, and Rachid Deriche (2019). 'The Dmipy Toolbox: Diffusion MRI Multi-Compartment Modeling and Microstructure Recovery Made Easy'. In: *Frontiers in Neuroinformatics* 13.64. DOI: [/10.3389/fninf.2019.00064](https://doi.org/10.3389/fninf.2019.00064) (cited on pages 3, 30, 34, 40, 42, 109–111, 113, 115).
- Fick, Rutger et al. (2016). 'MAPL: Tissue microstructure estimation using Laplacian-regularized MAP-MRI and its application to HCP data'. In: *NeuroImage* 134, pp. 365–385. DOI: [10.1016/j.neuroimage.2016.03.046](https://doi.org/10.1016/j.neuroimage.2016.03.046) (cited on page 116).
- Fischl, Bruce (2012). 'FreeSurfer'. In: *Neuroimage* 62.2, pp. 774–781. DOI: [10.1016/j.neuroimage.2012.01.021](https://doi.org/10.1016/j.neuroimage.2012.01.021) (cited on pages 21, 23, 65, 73, 90, 109).
- Frigo, Matteo et al. (2018a). 'Reducing false positive connection in tractograms using joint structure-function filtering'. In: *OHBM 2018-Organization for Human Brain Mapping* (cited on pages 5, 60).
- Frigo, Matteo et al. (2018b). 'Resolving the crossing/kissing fiber ambiguity using Functionally Informed COMMIT'. In: *Computational Diffusion MRI 2018*. DOI: [10.1007/978-3-030-05831-9](https://doi.org/10.1007/978-3-030-05831-9) (cited on pages 5, 60, 73, 75).

- Frigo, Matteo et al. (2019). ‘Effects of tractography filtering on the topology and interpretability of connectomes’. In: *OHBM 2019* (cited on page 5).
- Frigo, Matteo et al. (2020a). *Code and data for "Network alignment and similarity reveal atlas-based topological differences in structural connectomes"*. DOI: [10.17605/OSF.IO/DEPUX](https://doi.org/10.17605/OSF.IO/DEPUX). URL: <https://osf.io/depux/> (cited on pages 90, 112).
- Frigo, Matteo et al. (2020b). ‘Diffusion MRI tractography filtering techniques change the topology of structural connectomes’. In: *Journal of Neural Engineering* 17.6, p. 065002. DOI: [10.1088/1741-2552/abc29b](https://doi.org/10.1088/1741-2552/abc29b) (cited on pages 4, 5, 24, 47, 51, 58, 63, 68, 102).
- Frigo, Matteo et al. (2020c). ‘Multi Tissue Modelling of Diffusion MRI Signal Reveals Volume Fraction Bias’. In: *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*. IEEE, pp. 991–994. DOI: [10.1109/ISBI45749.2020.9098649](https://doi.org/10.1109/ISBI45749.2020.9098649) (cited on pages 3, 4, 31).
- Frigo, Matteo et al. (2020d). ‘Multi-compartment modelling of diffusion MRI signal shows TE-based volume fraction bias’. In: *OHBM 2020-26th meeting of the Organization of Human Brain Mapping* (cited on pages 4, 31).
- Frigo, Matteo et al. (2021a). ‘Multi-Tissue Multi-Compartment Models of Diffusion MRI’. In: *bioRxiv*. DOI: [10.1101/2021.01.29.428843](https://doi.org/10.1101/2021.01.29.428843) (cited on page 31).
- Frigo, Matteo et al. (2021b). ‘Network alignment and similarity reveal atlas-based topological differences in structural connectomes’. In: *Network Neuroscience*. In Press. DOI: [10.1101/2020.12.16.422501](https://doi.org/10.1101/2020.12.16.422501) (cited on pages 103, 112).
- Frigo, Matteo et al. (2021c). *TALON: Tractograms As Linear Operators in Neuroimaging*. CoBCoM. URL: <https://hal.archives-ouvertes.fr/hal-03116143> (cited on page 109).
- Friston, Karl J (2011). ‘Functional and effective connectivity: a review’. In: *Brain connectivity* 1.1, pp. 13–36. DOI: [10.1089/brain.2011.0008](https://doi.org/10.1089/brain.2011.0008) (cited on pages 26, 101).
- Froeling, Martijn et al. (2017). ‘“MASSIVE” brain dataset: Multiple acquisitions for standardization of structural imaging validation and evaluation’. In: *Magnetic resonance in medicine* 77.5, pp. 1797–1809. DOI: [10.1002/mrm.26259](https://doi.org/10.1002/mrm.26259) (cited on page 116).
- Fukutomi, Hikaru et al. (2019). ‘Diffusion tensor model links to neurite orientation dispersion and density imaging at high b-value in cerebral cortical gray matter’. In: *Scientific reports* 9.1, pp. 1–12. DOI: [10.1038/s41598-019-48671-7](https://doi.org/10.1038/s41598-019-48671-7) (cited on page 34).
- Gallardo, Guillermo et al. (2018a). ‘Groupwise structural parcellation of the whole cortex: A logistic random effects model based approach’. In: *NeuroImage* 170, pp. 307–320. DOI: [10.1016/j.neuroimage.2017.01.070](https://doi.org/10.1016/j.neuroimage.2017.01.070) (cited on pages 22, 82, 91).
- Gallardo, Guillermo et al. (Sept. 2018b). ‘Solving the Cross-Subject Parcel Matching Problem using Optimal Transport’. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention 2018*. Granada, Spain. DOI: [10.1007/978-3-030-00928-1_94](https://doi.org/10.1007/978-3-030-00928-1_94). (Visited on 06/23/2020) (cited on page 82).
- Gallardo Diez, Guillermo Alejandro (2018). ‘Inferring and comparing structural parcellations of the human brain using diffusion MRI’. PhD thesis. Université Côte d’Azur (cited on page 21).

- Ganepola, Tharindu et al. (2018). 'Using diffusion MRI to discriminate areas of cortical grey matter'. In: *NeuroImage* 182, pp. 456–468. DOI: [10.1016/j.neuroimage.2017.12.046](https://doi.org/10.1016/j.neuroimage.2017.12.046) (cited on page 34).
- Garyfallidis, Eleftherios et al. (2012). 'Quickbundles, a method for tractography simplification'. In: *Frontiers in neuroscience* 6, p. 175. DOI: [10.3389/fnins.2012.00175](https://doi.org/10.3389/fnins.2012.00175) (cited on page 79).
- Garyfallidis, Eleftherios et al. (2014). 'Dipy, a library for the analysis of diffusion MRI data'. In: *Frontiers in neuroinformatics* 8, p. 8. DOI: [10.3389/fninf.2014.00008](https://doi.org/10.3389/fninf.2014.00008) (cited on pages 17, 72, 109, 110, 115).
- Girard, Gabriel et al. (2014). 'Towards quantitative connectivity analysis: reducing tractography biases'. In: *Neuroimage* 98, pp. 266–278. DOI: [10.1016/j.neuroimage.2014.04.074](https://doi.org/10.1016/j.neuroimage.2014.04.074) (cited on pages 18, 19, 24).
- Girard, Gabriel et al. (2015). 'AxTract: microstructure-driven tractography based on the ensemble average propagator'. In: *International Conference on Information Processing in Medical Imaging*. Springer, pp. 675–686. DOI: [10.1007/978-3-319-19992-4_53](https://doi.org/10.1007/978-3-319-19992-4_53) (cited on page 51).
- Girard, Gabriel et al. (2017). 'AxTract: Toward microstructure informed tractography'. In: *Human brain mapping* 38.11, pp. 5485–5500. DOI: [10.1002/hbm.23741](https://doi.org/10.1002/hbm.23741) (cited on pages 18, 51, 72).
- Glasser, Matthew F et al. (2013). 'The minimal preprocessing pipelines for the Human Connectome Project'. In: *Neuroimage* 80, pp. 105–124. DOI: [10.1016/j.neuroimage.2013.04.127](https://doi.org/10.1016/j.neuroimage.2013.04.127) (cited on pages 45, 63, 64, 73, 90).
- Glasser, Matthew F. et al. (2016). 'A multi-modal parcellation of human cerebral cortex'. In: *Nature* 536. DOI: [10.1038/nature18933](https://doi.org/10.1038/nature18933) (cited on pages 22, 91).
- Gong, Ting et al. (2020). 'MTE-NODDI: Multi-TE NODDI for disentangling non-T2-weighted signal fractions from compartment-specific T2 relaxation times'. In: *NeuroImage*, p. 116906. DOI: [10.1016/j.neuroimage.2020.116906](https://doi.org/10.1016/j.neuroimage.2020.116906) (cited on pages 3, 30, 35, 36, 104).
- Gray, Henry, Roger Warwick, and Peter Llewellyn Williams (1973). *Gray's Anatomy. Edited by Roger Warwick... Peter L. Williams, Etc.* Longman (cited on pages 7, 8).
- Gutman, Boris et al. (2014). 'Registering cortical surfaces based on whole-brain structural connectivity and continuous connectivity analysis'. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 161–168. DOI: [10.1007/978-3-319-10443-0_21](https://doi.org/10.1007/978-3-319-10443-0_21) (cited on page 22).
- Hagmann, Patric (2005). *From diffusion MRI to brain connectomics*. Tech. rep. EPFL (cited on pages 1, 20, 83).
- Hagmann, Patric et al. (2006). 'Understanding diffusion MR imaging techniques: from scalar diffusion-weighted imaging to diffusion tensor imaging and beyond'. In: *Radiographics* 26.suppl_1, S205–S223. DOI: [10.1148/rg.26si065510](https://doi.org/10.1148/rg.26si065510) (cited on page 9).
- Hagmann, Patric et al. (2008). 'Mapping the structural core of human cerebral cortex'. In: *PLoS Biol* 6.7, e159. DOI: [10.1371/journal.pbio.0060159](https://doi.org/10.1371/journal.pbio.0060159) (cited on page 61).
- Hansen, Per Christian (1992). 'Analysis of discrete ill-posed problems by means of the L-curve'. In: *SIAM review* 34.4, pp. 561–580. DOI: [10.1137/1034115](https://doi.org/10.1137/1034115) (cited on page 74).
- Hara, Shoko et al. (2018). 'Microstructural damage in normal-appearing brain parenchyma and neurocognitive dysfunction in adult moyamoya

- disease'. In: *Stroke* 49.10, pp. 2504–2507. DOI: [10 . 1161 / STROKEAHA . 118 . 022367](https://doi.org/10.1161/STROKEAHA.118.022367) (cited on page 36).
- Harris, Charles R. et al. (2020). 'Array programming with NumPy'. In: *Nature* 585.7825, pp. 357–362. DOI: [10 . 1038 / s41586 - 020 - 2649 - 2](https://doi.org/10.1038/s41586-020-2649-2) (cited on pages 109, 111).
- Hayhoe, Mikhail et al. (May 2019). 'SPECTRE: Seedless Network Alignment via Spectral Centralities'. In: *arXiv:1811.01056 [cs, math]*. arXiv: 1811.01056. (Visited on 03/10/2020) (cited on page 82).
- Heimann, Mark et al. (2018). 'REGAL: Representation Learning-based Graph Alignment'. en. In: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management - CIKM '18*. Torino, Italy: ACM Press, pp. 117–126. DOI: [10 . 1145 / 3269206 . 3271788](https://doi.org/10.1145/3269206.3271788) (cited on page 82).
- Herculano-Houzel, Suzana (2009). 'The human brain in numbers: a linearly scaled-up primate brain'. In: *Frontiers in human neuroscience* 3, p. 31. DOI: [10 . 3389 / neuro . 09 . 031 . 2009](https://doi.org/10.3389/neuro.09.031.2009) (cited on page 21).
- Honey, Christopher J et al. (2007). 'Network structure of cerebral cortex shapes functional connectivity on multiple time scales'. In: *Proceedings of the National Academy of Sciences* 104.24, pp. 10240–10245. DOI: [10 . 1073 / pnas . 0701519104](https://doi.org/10.1073/pnas.0701519104) (cited on pages 20, 73).
- Horwitz, Barry (2003). 'The elusive concept of brain connectivity'. In: *Neuroimage* 19.2, pp. 466–470. DOI: [10 . 1016 / s1053 - 8119 \(03\)00112 - 5](https://doi.org/10.1016/s1053-8119(03)00112-5) (cited on page 25).
- Hounsfield, Godfrey N (1973). 'Computerized transverse axial scanning (tomography): Part 1. Description of system'. In: *The British journal of radiology* 46.552, pp. 1016–1022 (cited on page 9).
- Hunter, J. D. (2007). 'Matplotlib: A 2D graphics environment'. In: *Computing in Science & Engineering* 9.3, pp. 90–95. DOI: [10 . 1109 / MCSE . 2007 . 55](https://doi.org/10.1109/MCSE.2007.55) (cited on page 109).
- Huth, Alexander G et al. (2016). 'Natural speech reveals the semantic maps that tile human cerebral cortex'. In: *Nature* 532.7600, pp. 453–458. DOI: [10 . 1038 / nature17637](https://doi.org/10.1038/nature17637) (cited on page 22).
- Jacobi, Carl Gustav Jacob (1890). 'The reduction to normal form of a non-normal system of differential equations. De aequationum differentialium systemate non normali ad formam normalem revocando'. In: *III. C. G. J. Jacobi manuscriptis posthumis in medium protulit A. Clebsch* (cited on page 89).
- Jbabdi, Saad and Heidi Johansen-Berg (2011). 'Tractography: where do we go from here?' In: *Brain connectivity* 1.3, pp. 169–183. DOI: [10 . 1089 / brain . 2011 . 0033](https://doi.org/10.1089/brain.2011.0033) (cited on pages 4, 19, 20, 23, 49, 50, 102).
- Jelescu, Ileana O and Matthew D Budde (2017). 'Design and validation of diffusion MRI models of white matter'. In: *Frontiers in physics* 5, p. 61 (cited on page 30).
- Jelescu, Ileana O et al. (2016). 'Degeneracy in model parameter estimation for multi-compartmental diffusion in neuronal tissue'. In: *NMR in Biomedicine* 29.1, pp. 33–47 (cited on page 33).
- Jenatton, Rodolphe et al. (2010). 'Proximal Methods for Sparse Hierarchical Dictionary Learning.' In: *ICML*. 2010. Citeseer, pp. 487–494 (cited on page 56).
- Jenatton, Rodolphe et al. (2011). 'Proximal methods for hierarchical sparse coding'. In: *The Journal of Machine Learning Research* 12, pp. 2297–2334 (cited on page 56).

- Jenkinson, Mark et al. (2012). 'Fsl'. In: *Neuroimage* 62.2, pp. 782–790 (cited on pages 63, 109, 110).
- Jeurissen, Ben et al. (2014). 'Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data'. In: *NeuroImage* 103, pp. 411–426 (cited on pages 18, 19, 30, 37, 40, 46, 47, 65, 73, 91, 105).
- Jeurissen, Ben et al. (2019). 'Diffusion MRI fiber tractography of the brain'. In: *NMR in Biomedicine* 32.4, e3785 (cited on pages 2, 18–20, 24, 50).
- Johansen-Berg, Heidi and Timothy EJ Behrens (2006). 'Just pretty pictures? What diffusion tractography can add in clinical neuroscience'. In: *Current opinion in neurology* 19.4, p. 379 (cited on page 19).
- Jones, Derek K (2008). 'Tractography gone wild: probabilistic fibre tracking using the wild bootstrap with diffusion tensor MRI'. In: *IEEE transactions on medical imaging* 27.9, pp. 1268–1274 (cited on page 18).
- (2010). *Diffusion mri*. Oxford University Press (cited on page 9).
- Jones, Derek K, Thomas R Knösche, and Robert Turner (2013). 'White matter integrity, fiber count, and other fallacies: the do's and don'ts of diffusion MRI'. In: *Neuroimage* 73, pp. 239–254 (cited on pages 4, 19, 20, 24, 49–51).
- Just, M and M Thelen (1988). 'Tissue characterization with T1, T2, and proton density values: results in 160 patients with brain tumors'. In: *Radiology* 169.3, pp. 779–785 (cited on page 34).
- Kaden, Enrico et al. (2016). 'Multi-compartment microscopic diffusion imaging'. In: *NeuroImage* 139, pp. 346–359 (cited on pages 29, 55, 115).
- Kang, James M (2008). *Voronoi Diagram*. (Cited on page 23).
- Kay, Kevin R et al. (2013). 'Studying synapses in human brain with array tomography and electron microscopy'. In: *Nature protocols* 8.7, pp. 1366–1380 (cited on page 8).
- Klößner, Andreas et al. (2012). 'PyCUDA and PyOpenCL: A Scripting-Based Approach to GPU Run-Time Code Generation'. In: *Parallel Computing* 38.3, pp. 157–174. DOI: [10.1016/j.parco.2011.09.001](https://doi.org/10.1016/j.parco.2011.09.001) (cited on page 111).
- Konopleva, Lidia et al. (2018). 'Model-free global tractography'. In: *NeuroImage* 174, pp. 576–586 (cited on page 18).
- Korula, Nitish and Silvio Lattanzi (2014). 'An efficient reconciliation algorithm for social networks'. In: *Proceedings of the VLDB Endowment* 7.5, pp. 377–388. DOI: [10.14778/2732269.2732274](https://doi.org/10.14778/2732269.2732274) (cited on page 82).
- Kreher, BW, I Mader, and VG Kiselev (2008). 'Gibbs tracking: a novel approach for the reconstruction of neuronal pathways'. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 60.4, pp. 953–963 (cited on pages 18, 51).
- Kuhn, Harold W (1955). 'The Hungarian method for the assignment problem'. In: *Naval research logistics quarterly* 2.1-2, pp. 83–97. DOI: [10.1002/nav.3800020109](https://doi.org/10.1002/nav.3800020109) (cited on page 89).
- Lampinen, Björn et al. (2017). 'Neurite density imaging versus imaging of microscopic anisotropy in diffusion MRI: A model comparison using spherical tensor encoding'. In: *Neuroimage* 147, pp. 517–531 (cited on pages 29, 33, 104).
- Lampinen, Björn et al. (2019). 'Searching for the neurite density with diffusion MRI: challenges for biophysical modeling'. In: *Human brain mapping* 40.8, pp. 2529–2545 (cited on pages 3, 30, 35, 36, 104).

- Lampinen, Björn et al. (2020). ‘Towards unconstrained compartment modeling in white matter using diffusion-relaxation MRI with tensor-valued diffusion encoding’. In: *Magnetic Resonance in Medicine* 84.3, pp. 1605–1623 (cited on pages 3, 30, 35, 36).
- Latora, Vito and Massimo Marchiori (2001). ‘Efficient behavior of small-world networks’. In: *Physical review letters* 87.19, p. 198701 (cited on page 27).
- Lauterbur, Paul C (1973). ‘Image formation by induced local interactions: examples employing nuclear magnetic resonance’. In: (cited on page 9).
- Le Bihan, Denis and E Breton (1985). ‘Imagerie de diffusion in vivo par résonance magnétique nucléaire’. In: *Comptes rendus de l’Académie des sciences. Série 2, Mécanique, Physique, Chimie, Sciences de l’univers, Sciences de la Terre* 301.15, pp. 1109–1112 (cited on page 13).
- Le Bihan, Denis et al. (1988). ‘Separation of diffusion and perfusion in intravoxel incoherent motion MR imaging’. In: *Radiology* 168.2, pp. 497–505 (cited on page 29).
- Legarreta, Jon Haitz et al. (2020). ‘Tractography filtering using autoencoders’. In: *arXiv preprint arXiv:2010.04007* (cited on page 52).
- Li, Chaozhuo et al. (Oct. 2018). ‘Distribution Distance Minimization for Unsupervised User Identity Linkage’. In: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management. CIKM ’18*. Torino, Italy: Association for Computing Machinery, pp. 447–456. DOI: [10.1145/3269206.3271675](https://doi.org/10.1145/3269206.3271675) (cited on page 82).
- Liewald, Daniel et al. (2014). ‘Distribution of axon diameters in cortical white matter: an electron-microscopic study on three human brains and a macaque’. In: *Biological cybernetics* 108.5, pp. 541–557 (cited on page 8).
- Liu, Li et al. (2016). ‘Aligning Users across Social Networks Using Network Embedding’. In: *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence. IJCAI’16*. New York, New York, USA: AAAI Press, pp. 1774–1780 (cited on pages 82, 83).
- Maier-Hein, Klaus H et al. (2017). ‘The challenge of mapping the human connectome based on diffusion tractography’. In: *Nature communications* 8.1, pp. 1–13 (cited on pages 4, 18–20, 24, 49–51, 72).
- Manjón, José V et al. (2013). ‘Diffusion weighted image denoising using overcomplete local PCA’. In: *PloS one* 8.9, e73021 (cited on page 64).
- Mann, Henry B and Donald R Whitney (1947). ‘On a test of whether one of two random variables is stochastically larger than the other’. In: *The annals of mathematical statistics*, pp. 50–60 (cited on page 66).
- Mardia, Kanti V and Peter E Jupp (1990). *Directional statistics*. Vol. 494. John Wiley & Sons (cited on page 32).
- Marnier, Lisbeth et al. (2003). ‘Marked loss of myelinated nerve fibers in the human brain with age’. In: *Journal of comparative neurology* 462.2, pp. 144–152 (cited on page 1).
- Mars, R.B. et al. (2016). ‘Comparing brains by matching connectivity profiles’. In: *Neurosci. Biobehav.* 60, pp. 90–97 (cited on page 82).
- McKay, Brendan D. and Adolfo Piperno (Jan. 2014). ‘Practical graph isomorphism, II’. In: *Journal of Symbolic Computation* 60, pp. 94–112. DOI: [10.1016/j.jsc.2013.09.003](https://doi.org/10.1016/j.jsc.2013.09.003) (cited on page 88).
- Milgram, Stanley (1967). ‘The small world problem’. In: *Psychology today* 2.1, pp. 60–67 (cited on page 26).
- WU-Minn Human Connectome Project consortium (2017). *1200 Subjects Data Release Reference Manual*. URL: <https://www.humanconnectome.org/>

- [org/storage/app/media/documentation/s1200/HCP_S1200_Release_Reference_Manual.pdf](#) (cited on pages 63, 90, 95, 109).
- Mohammadian, Mehrbod et al. (2017). 'High angular resolution diffusion-weighted imaging in mild traumatic brain injury'. In: *NeuroImage: Clinical* 13, pp. 174–180 (cited on page 63).
- Moreno-Dominguez, David, Alfred Anwander, and Thomas R Knösche (2014). 'A hierarchical method for whole-brain connectivity-based parcellation'. In: *Human brain mapping* 35.10, pp. 5000–5025 (cited on page 22).
- Mori, Susumu et al. (1999). 'Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging'. In: *Annals of Neurology: Official Journal of the American Neurological Association and the Child Neurology Society* 45.2, pp. 265–269 (cited on page 18).
- Moyer, Daniel et al. (2016). 'A continuous model of cortical connectivity'. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 157–165 (cited on pages 22, 23).
- Mueller, Susanne G et al. (2005). 'The Alzheimer's disease neuroimaging initiative'. In: *Neuroimaging Clinics* 15.4, pp. 869–877 (cited on page 30).
- Nassar, Huda et al. (2018). 'Low Rank Spectral Network Alignment'. en. In: *Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW '18*. Lyon, France: ACM Press, pp. 619–628. DOI: [10.1145/3178876.3186128](#). (Visited on 03/05/2020) (cited on page 82).
- Newman, Mark EJ (2004). 'Analysis of weighted networks'. In: *Physical review E* 70.5, p. 056131 (cited on page 27).
- Nir, Talia M et al. (2020). 'Improved Neurite Density Estimation in Alzheimer's Disease with Multi-Tissue Multicompartment Diffusion MRI Modeling'. In: *EMBC* (cited on page 115).
- Novikov, Dmitry S, Valerij G Kiselev, and Sune N Jespersen (2018). 'On modeling'. In: *Magnetic resonance in medicine* 79.6, pp. 3172–3193 (cited on page 14).
- Novikov, Dmitry S et al. (2019). 'Quantifying brain microstructure with diffusion MRI: Theory and parameter estimation'. In: *NMR in Biomedicine* 32.4, e3998 (cited on pages 32, 36, 41, 105).
- O'Donnell, Lauren J and Carl-Fredrik Westin (2011). 'An introduction to diffusion tensor image analysis'. In: *Neurosurgery Clinics* 22.2, pp. 185–196 (cited on pages 15, 16).
- Ocampo-Pineda, Mario et al. (2021). 'Hierarchical Microstructure Informed Tractography'. In: *Brain Connectivity*. DOI: [10.1089/brain.2020.0907](#) (cited on page 56).
- Opsahl, Tore and Pietro Panzarasa (2009). 'Clustering in weighted networks'. In: *Social networks* 31.2, pp. 155–163 (cited on page 27).
- Osmanlioğlu, Yusuf et al. (2019). 'System-level matching of structural and functional connectomes in the human brain'. In: *NeuroImage* 199, pp. 93–104 (cited on page 82).
- Panagiotaki, Eleftheria et al. (2012). 'Compartment models of the diffusion MR signal in brain white matter: a taxonomy and comparison'. In: *Neuroimage* 59.3, pp. 2241–2254 (cited on page 54).
- Panagiotaki, Eleftheria et al. (2014). 'Noninvasive quantification of solid tumor microstructure using VERDICT MRI'. In: *Cancer research* 74.7, pp. 1902–1912 (cited on pages 29, 31, 32).
- Pandit, Anand S et al. (2013). 'Traumatic brain injury impairs small-world topology'. In: *Neurology* 80.20, pp. 1826–1833 (cited on page 63).

- Papo, David et al. (2016). 'Beware of the small-world neuroscientist!' In: *Frontiers in human neuroscience* 10, p. 96 (cited on page 26).
- Parisot, S. et al. (2015). 'Tractography-Driven Groupwise Multi-scale Parcellation of the Cortex'. In: *Inf. Process. Med. Imaging* 24, pp. 600–612 (cited on page 82).
- Parisot, Sarah et al. (2017). 'A flexible graphical model for multi-modal parcellation of the cortex'. In: *Neuroimage* 162, pp. 226–248 (cited on page 22).
- Parker, Drew et al. (2020). 'Freewater estimator using iNtErpolated initialization (FERNET): Characterizing peritumoral edema using clinically feasible diffusion MRI data'. In: *Plos one* 15.5, e0233645. DOI: [10.1371/journal.pone.0233645](https://doi.org/10.1371/journal.pone.0233645) (cited on page 29).
- Pestilli, Franco et al. (2014). 'Evaluation and statistical inference for human connectomes'. In: *Nature methods* 11.10, pp. 1058–1063 (cited on pages 4, 51, 58, 102).
- Plewes, Donald B (1994). 'Contrast mechanisms in spin-echo MR imaging'. In: *Radiographics* 14.6, pp. 1389–1404 (cited on page 34).
- Pollak, B (1953). 'Experiences with Planography'. In: *Diseases of the chest* 24.6, pp. 663–669 (cited on page 9).
- Power, Jonathan D et al. (2012). 'Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion'. In: *Neuroimage* 59.3, pp. 2142–2154 (cited on page 74).
- Preti, Maria Giulia, Thomas AW Bolton, and Dimitri Van De Ville (2017). 'The dynamic functional connectome: state-of-the-art and perspectives'. In: *Neuroimage* 160, pp. 41–54 (cited on pages 25, 62, 101).
- Rabinowitz, Amanda R et al. (2018). '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, pp. 237–246 (cited on page 64).
- Reina, Julio E Villalon et al. (2020). 'Evaluating NODDI-based biomarkers of Alzheimer's disease'. In: *2020 Alzheimer's Association International Conference*. ALZ (cited on page 115).
- Reisert, Marco et al. (2011). 'Global fiber reconstruction becomes practical'. In: *Neuroimage* 54.2, pp. 955–962 (cited on page 18).
- Reus, Marcel A de and Martijn P Van den Heuvel (2013). 'The parcellation-based connectome: limitations and extensions'. In: *Neuroimage* 80, pp. 397–404 (cited on page 22).
- Rheault, Francois et al. (2020a). 'Common misconceptions, hidden biases and modern challenges of dMRI tractography'. In: *Journal of Neural Engineering* (cited on pages 19, 20, 49–51).
- Rheault, Francois et al. (2020b). 'Tractostorm: The what, why, and how of tractography dissection reproducibility'. In: *Human Brain Mapping* 41.7, pp. 1859–1874 (cited on page 19).
- Röntgen, Wilhelm Conrad (1896). 'On a new kind of rays'. In: *Science* 3.59, pp. 227–231 (cited on page 9).
- Rosenke, Mona et al. (2018). 'A cross-validated cytoarchitectonic atlas of the human ventral visual stream'. In: *Neuroimage* 170, pp. 257–270 (cited on page 22).
- Rossum, Guido van, Barry Warsaw, and Nick Coghlan (2001). *Style Guide for Python Code*. PEP 8 (cited on page 110).

- Rubinov, Mikail and Olaf Sporns (2010). ‘Complex network measures of brain connectivity: uses and interpretations’. In: *Neuroimage* 52.3, pp. 1059–1069 (cited on pages 26, 65).
- Sanz Leon, Paula et al. (2013). ‘The Virtual Brain: a simulator of primate brain network dynamics’. In: *Frontiers in neuroinformatics* 7, p. 10 (cited on page 73).
- Sarubbo, Silvio and Laurent Petit (2019). ‘Organization of the white matter anatomy in the human brain’. In: *Frontiers in Neuroanatomy* 13 (cited on page 19).
- Sarwar, Tabinda, Kotagiri Ramamohanarao, and Andrew Zalesky (2019). ‘Mapping connectomes with diffusion MRI: deterministic or probabilistic tractography?’ In: *Magnetic resonance in medicine* 81.2, pp. 1368–1384 (cited on pages 19, 62).
- Schaefer, Alexander et al. (July 2017). ‘Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI’. In: *Cerebral Cortex* 28.9, pp. 3095–3114. DOI: [10.1093/cercor/bhx179](https://doi.org/10.1093/cercor/bhx179) (cited on pages 22, 91).
- Scherrer, Benoit and Simon K Warfield (2010). ‘Why multiple b-values are required for multi-tensor models. Evaluation with a constrained log-Euclidean model’. In: *2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. IEEE, pp. 1389–1392 (cited on pages 30, 34).
- Scherrer, Benoit et al. (2016). ‘Characterizing brain tissue by assessment of the distribution of anisotropic microstructural environments in diffusion-compartment imaging (DIAMOND)’. In: *Magnetic resonance in medicine* 76.3, pp. 963–977 (cited on page 29).
- Schiavi, Simona et al. (2020). ‘A new method for accurate in vivo mapping of human brain connections using microstructural and anatomical information’. In: *Science advances* 6.31 (cited on pages 4, 51, 53, 56, 57, 79, 106).
- Schilling, Kurt G et al. (2019a). ‘Challenges in diffusion MRI tractography—Lessons learned from international benchmark competitions’. In: *Magnetic resonance imaging* 57, pp. 194–209 (cited on pages 18, 50).
- Schilling, Kurt G et al. (2019b). ‘Limits to anatomical accuracy of diffusion tractography using modern approaches’. In: *NeuroImage* 185, pp. 1–11 (cited on pages 19, 20, 49, 50).
- Schilling, Kurt G et al. (2020a). ‘Brain connections derived from diffusion MRI tractography can be highly anatomically accurate—if we know where white matter pathways start, where they end, and where they do not go’. In: *Brain Structure and Function* 225.8, pp. 2387–2402 (cited on page 19).
- Schilling, Kurt G et al. (2020b). ‘Tractography dissection variability: what happens when 42 groups dissect 14 white matter bundles on the same dataset?’ In: *bioRxiv* (cited on page 19).
- Schotten, Michel Thiebaut de et al. (2011). ‘Atlasing location, asymmetry and inter-subject variability of white matter tracts in the human brain with MR diffusion tractography’. In: *Neuroimage* 54.1, pp. 49–59 (cited on page 19).
- Schreiber, Jan et al. (2014). ‘Plausibility tracking: a method to evaluate anatomical connectivity and microstructural properties along fiber pathways’. In: *Neuroimage* 90, pp. 163–178. DOI: [10.1016/j.neuroimage.2014.01.002](https://doi.org/10.1016/j.neuroimage.2014.01.002) (cited on pages 18, 51).

- Schubotz, Ricarda I et al. (2010). ‘Anatomical and functional parcellation of the human lateral premotor cortex’. In: *Neuroimage* 50.2, pp. 396–408. DOI: [10.1016/j.neuroimage.2009.12.069](https://doi.org/10.1016/j.neuroimage.2009.12.069) (cited on page 22).
- Seth, Anil K, Adam B Barrett, and Lionel Barnett (2015). ‘Granger causality analysis in neuroscience and neuroimaging’. In: *Journal of Neuroscience* 35.8, pp. 3293–3297 (cited on page 25).
- Shen, Rui Sherry et al. (2020). ‘Graph Matching Based Connectomic Biomarker with Learning for Brain Disorders’. In: *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Graphs in Biomedical Image Analysis*. Springer, pp. 131–141 (cited on page 82).
- Singh, Rohit, Jinbo Xu, and Bonnie Berger (Sept. 2008). ‘Global alignment of multiple protein interaction networks with application to functional orthology detection’. en. In: *Proceedings of the National Academy of Sciences* 105.35, pp. 12763–12768. DOI: [10.1073/pnas.0806627105](https://doi.org/10.1073/pnas.0806627105) (cited on page 82).
- Sizemore, Ann E and Danielle S Bassett (2018). ‘Dynamic graph metrics: Tutorial, toolbox, and tale’. In: *Neuroimage* 180, pp. 417–427. DOI: [10.1016/j.neuroimage.2017.06.081](https://doi.org/10.1016/j.neuroimage.2017.06.081) (cited on page 101).
- Smith, Robert E et al. (2012). ‘Anatomically-constrained tractography: improved diffusion MRI streamlines tractography through effective use of anatomical information’. In: *Neuroimage* 62.3, pp. 1924–1938 (cited on pages 19, 65, 73, 74, 90, 91).
- (2013). ‘SIFT: Spherical-deconvolution informed filtering of tractograms’. In: *Neuroimage* 67, pp. 298–312 (cited on pages 4, 51, 54, 106).
 - (2015a). ‘SIFT2: Enabling dense quantitative assessment of brain white matter connectivity using streamlines tractography’. In: *Neuroimage* 119, pp. 338–351 (cited on pages 4, 51, 53, 102, 106).
 - (2015b). ‘The effects of SIFT on the reproducibility and biological accuracy of the structural connectome’. In: *Neuroimage* 104, pp. 253–265 (cited on page 51).
- Smith, Stephen M (2002). ‘Fast robust automated brain extraction’. In: *Human brain mapping* 17.3, pp. 143–155 (cited on page 64).
- Smith, Stephen M et al. (2013). ‘Functional connectomics from resting-state fMRI’. In: *Trends in cognitive sciences* 17.12, pp. 666–682 (cited on page 25).
- Sotiropoulos, Stamatios N and Andrew Zalesky (2019). ‘Building connectomes using diffusion MRI: why, how and but’. In: *NMR in Biomedicine* 32.4, e3752 (cited on page 19).
- Sotiropoulos, Stamatios N et al. (2013). ‘Advances in diffusion MRI acquisition and processing in the Human Connectome Project’. In: *Neuroimage* 80, pp. 125–143 (cited on pages 44, 63).
- Sporns, O., G. Tononi, and R. Kötter (2005). ‘The human connectome: A structural description of the human brain’. In: *PLoS Computational Biology* 1.4 (4), pp. 245–251 (cited on pages 1, 20, 83).
- Sporns, Olaf et al. (2004). ‘Organization, development and function of complex brain networks’. In: *Trends in cognitive sciences* 8.9, pp. 418–425 (cited on page 26).
- Sra, Suvrit, Sebastian Nowozin, and Stephen J Wright (2011). *Optimization for machine learning*. Mit Press (cited on page 59).
- Stanisz, Greg J et al. (1997). ‘An analytical model of restricted diffusion in bovine optic nerve’. In: *Magnetic Resonance in Medicine* 37.1, pp. 103–111 (cited on page 29).

- Stejskal, Edward O and John E Tanner (1965). 'Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient'. In: *The journal of chemical physics* 42.1, pp. 288–292 (cited on page 13).
- Stephan, Klaas Enno and Karl J Friston (2010). 'Analyzing effective connectivity with functional magnetic resonance imaging'. In: *Wiley Interdisciplinary Reviews: Cognitive Science* 1.3, pp. 446–459 (cited on page 25).
- Stratos, Karl L. (2020). 'A graph from the viewpoint of algebraic topology'. In: (cited on page 91).
- Sudlow, Cathie et al. (2015). 'UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age'. In: *Plos med* 12.3, e1001779 (cited on page 30).
- Suzuki, Hideaki et al. (2017). 'Abnormal brain white matter microstructure is associated with both pre-hypertension and hypertension'. In: *PLoS One* 12.11, e0187600 (cited on page 36).
- Szafer, A et al. (1995). 'Diffusion-weighted imaging in tissues: theoretical models'. In: *NMR in Biomedicine* 8.7, pp. 289–296. DOI: [10.1002/nbm.1940080704](https://doi.org/10.1002/nbm.1940080704) (cited on page 33).
- Szafer, Aaron, Jianhui Zhong, and John C Gore (1995). 'Theoretical model for water diffusion in tissues'. In: *Magnetic resonance in medicine* 33.5, pp. 697–712 (cited on page 33).
- Tae, Woo-Suk et al. (2018). 'Current clinical applications of diffusion-tensor imaging in neurological disorders'. In: *Journal of Clinical Neurology* 14.2, pp. 129–140 (cited on page 16).
- Tariq, Maira et al. (2016). 'Bingham–NODDI: mapping anisotropic orientation dispersion of neurites using diffusion MRI'. In: *NeuroImage* 133, pp. 207–223 (cited on page 29).
- Thiebaut de Schotten, Michel et al. (2017). 'Rostro-caudal architecture of the frontal lobes in humans'. In: *Cerebral Cortex* 27.8, pp. 4033–4047 (cited on page 22).
- Thomas, Cibu et al. (2014). 'Anatomical accuracy of brain connections derived from diffusion MRI tractography is inherently limited'. In: *Proceedings of the National Academy of Sciences* 111.46, pp. 16574–16579 (cited on page 72).
- Tononi, Giulio, Olaf Sporns, and Gerald M Edelman (1994). 'A measure for brain complexity: relating functional segregation and integration in the nervous system'. In: *Proceedings of the National Academy of Sciences* 91.11, pp. 5033–5037 (cited on page 26).
- Torrey, Henry C (1956). 'Bloch equations with diffusion terms'. In: *Physical review* 104.3, p. 563 (cited on page 9).
- Tournier, J-Donald, Fernando Calamante, and Alan Connelly (2007). 'Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution'. In: *Neuroimage* 35.4, pp. 1459–1472 (cited on pages 16, 18, 19, 38, 65).
- (2012). 'MRtrix: diffusion tractography in crossing fiber regions'. In: *International journal of imaging systems and technology* 22.1, pp. 53–66 (cited on page 18).
- Tournier, J-Donald et al. (2004). 'Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution'. In: *NeuroImage* 23.3, pp. 1176–1185 (cited on pages 65, 73).

- Tournier, J-Donald et al. (2019). 'MRtrix3: A fast, flexible and open software framework for medical image processing and visualisation'. In: *NeuroImage* 202, p. 116137 (cited on pages 18, 19, 35, 40, 44, 45, 52, 58, 60, 65, 73, 90, 109, 110).
- Tzourio-Mazoyer, Nathalie et al. (2002). 'Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain'. In: *Neuroimage* 15.1, pp. 273–289 (cited on page 21).
- Van Den Heuvel, Martijn P and Hilleke E Hulshoff Pol (2010). 'Exploring the brain network: a review on resting-state fMRI functional connectivity'. In: *European neuropsychopharmacology* 20.8, pp. 519–534 (cited on pages 25, 61, 62, 101).
- Van Essen, David C et al. (2012). 'The Human Connectome Project: a data acquisition perspective'. In: *Neuroimage* 62.4, pp. 2222–2231 (cited on pages 6, 30, 44, 63, 90, 109).
- Van Rossum, Guido and Fred L. Drake (2009). *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace (cited on page 110).
- Venkatesh, Manasij, Joseph Jaja, and Luiz Pessoa (2020). 'Comparing functional connectivity matrices: A geometry-aware approach applied to participant identification'. In: *NeuroImage* 207, p. 116398 (cited on page 84).
- Veraart, Jelle, Dmitry S. Novikov, and Els Fieremans (2018). 'TE dependent Diffusion Imaging (TEdDI) distinguishes between compartmental T2 relaxation times'. In: *NeuroImage* 182, pp. 360–369. doi: [10.1016/j.neuroimage.2017.09.030](https://doi.org/10.1016/j.neuroimage.2017.09.030) (cited on pages 3, 30, 34–36, 104).
- Verma, Ragini, Yusuf Osmanlioglu, and Abdol Aziz Ould Ismail (2018). 'Multimodal Patho-Connectomics of Brain Injury'. In: *International MIC-CAI Brainlesion Workshop*. Springer, pp. 3–14 (cited on page 63).
- Verma, Ragini et al. (2019). 'Neuroimaging Findings in US Government Personnel With Possible Exposure to Directional Phenomena in Havana, Cuba'. In: *Jama* 322.4, pp. 336–347 (cited on pages 20, 63).
- Vestergaard-Poulsen, Peter et al. (2007). 'Microstructural changes in ischemic cortical gray matter predicted by a model of diffusion-weighted MRI'. In: *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine* 26.3, pp. 529–540 (cited on page 36).
- Villalon-reina, Julio E et al. (2020). 'Tracking microstructural biomarkers of Alzheimer's disease via advanced multi-shell diffusion MRI scalar measures'. In: *ISMRM* (cited on pages 34, 115).
- Villareal-Haro, Juan Luis, Alonso Ramirez-Manzanares, and Juan Antonio Pichardo-Corpus (2020). 'A community-based topological distance for brain-connectome classification'. In: *Journal of Complex Networks* 8.4, cnaa034 (cited on page 82).
- Vogelstein, Joshua T. et al. (Apr. 2015). 'Fast Approximate Quadratic Programming for Graph Matching'. In: *PLoS ONE* 10.4. doi: [10.1371/journal.pone.0121002](https://doi.org/10.1371/journal.pone.0121002) (cited on pages 83, 84, 92).
- Washington University School of Medicine (Sept. 2020). *The Brain Analysis Library of Spatial Maps and Atlases (BALSA)*. URL: <https://balsa.wustl.edu/WN56> (cited on page 91).
- Watts, Duncan J and Steven H Strogatz (1998). 'Collective dynamics of 'small-world' networks'. In: *nature* 393.6684, pp. 440–442 (cited on page 26).
- Wedeen, Van J et al. (2008). 'Diffusion spectrum magnetic resonance imaging (DSI) tractography of crossing fibers'. In: *Neuroimage* 41.4, pp. 1267–1277 (cited on page 16).

- Weiner, Michael W et al. (2017). ‘The Alzheimer’s Disease Neuroimaging Initiative 3: Continued innovation for clinical trial improvement’. In: *Alzheimer’s & Dementia* 13.5, pp. 561–571 (cited on page 115).
- Weisfeiler, B Yu and A A Leman (1968). ‘The reduction of a graph to canonical form and the algebra which appears therein’. en. English translation of the original paper published in Russian. (cited on pages 83, 87).
- Westin, C.-F. et al. (Apr. 1997). ‘Geometrical Diffusion Measures for MRI from Tensor Basis Analysis’. In: *ISMRM ’97*. Vancouver Canada (cited on page 16).
- Wilcoxon, Frank (1945). ‘Individual Comparisons by Ranking Methods’. In: *Biometrics Bulletin* 1.6, pp. 80–83 (cited on pages 43–45, 94, 98).
- Wilson, Richard C. and Ping Zhu (Sept. 2008). ‘A study of graph spectra for comparing graphs and trees’. In: *Pattern Recognition* 41.9, pp. 2833–2841. DOI: [10.1016/j.patcog.2008.03.011](https://doi.org/10.1016/j.patcog.2008.03.011) (cited on page 82).
- Xie, Sangma et al. (2015). ‘How does B-value affect HARDI reconstruction using clinical diffusion MRI data?’ In: *PloS one* 10.3, e0120773 (cited on page 16).
- Yeh, Chun-Hung et al. (2016). ‘Correction for diffusion MRI fibre tracking biases: The consequences for structural connectomic metrics’. In: *Neuroimage* 142, pp. 150–162 (cited on page 63).
- Yeh, Chun-Hung et al. (2019). ‘Connectomes from streamlines tractography: Assigning streamlines to brain parcellations is not trivial but highly consequential’. In: *Neuroimage* 199, pp. 160–171. DOI: [/10.1016/j.neuroimage.2019.05.005](https://doi.org/10.1016/j.neuroimage.2019.05.005) (cited on pages 19, 24).
- Yeo, BT Thomas et al. (2011). ‘The organization of the human cerebral cortex estimated by intrinsic functional connectivity’. In: *Journal of neurophysiology* (cited on pages 22, 91).
- Yeo, Thomas (Sept. 2020). *Computational Brain Imaging Group (CBIG) repository*. URL: https://github.com/ThomasYeoLab/CBIG/tree/master/stable_projects/brain_parcellation/Schaefer2018_LocalGlobal/Parcellations/HCP (cited on page 91).
- Yo, Ting-Shuo et al. (2009). ‘Quantifying brain connectivity: a comparative tractography study’. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 886–893 (cited on page 16).
- Youden, William J (1950). ‘Index for rating diagnostic tests’. In: *Cancer* 3.1, pp. 32–35 (cited on page 77).
- Yu, Shan et al. (2008). ‘A small world of neuronal synchrony’. In: *Cerebral cortex* 18.12, pp. 2891–2901 (cited on page 26).
- Zalesky, Andrew et al. (2016). ‘Connectome sensitivity or specificity: which is more important?’ In: *Neuroimage* 142, pp. 407–420 (cited on pages 4, 72).
- Zarghami, Tahereh S and Karl J Friston (2020). ‘Dynamic effective connectivity’. In: *Neuroimage* 207, p. 116453 (cited on pages 25, 26).
- Zeng, Hongkui and Joshua R Sanes (2017). ‘Neuronal cell-type classification: challenges, opportunities and the path forward’. In: *Nature Reviews Neuroscience* 18.9, pp. 530–546 (cited on page 7).
- Zhang, Bin and Steve Horvath (2005). ‘A general framework for weighted gene co-expression network analysis’. In: *Statistical applications in genetics and molecular biology* 4.1 (cited on page 27).

- Zhang, Hui et al. (2012). 'NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain'. In: *Neuroimage* 61.4, pp. 1000–1016 (cited on pages 29, 32, 33, 55).
- Zou, Cunlu, Katherine J Denby, and Jianfeng Feng (2009). 'Granger causality vs. dynamic Bayesian network inference: a comparative study'. In: *BMC bioinformatics* 10.1, pp. 1–17 (cited on page 26).