



**University of
Nottingham**

UK | CHINA | MALAYSIA

Mapping Connections in the Neonatal Brain with Magnetic Resonance Imaging

Elinor Thompson

Thesis submitted to the University of Nottingham
for the degree of Doctor of Philosophy

December 2020

Abstract

The neonatal brain undergoes rapid development after birth, including the growth and maturation of the white matter fibre bundles that connect brain regions. Diffusion MRI (dMRI) is a unique tool for mapping these bundles in vivo, providing insight into factors that impact the development of white matter and how its maturation influences other developmental processes. However, most studies of neonatal white matter do not use specialised analysis tools, instead using tools that have been developed for the adult brain.

However, the neonatal brain is not simply a small adult brain, as differences in geometry and tissue decomposition cause considerable differences in dMRI contrast. In this thesis, methods are developed to map white matter connections during this early stage of neurodevelopment. First, two contrasting approaches are explored: ROI-constrained protocols for mapping individual tracts, and the generation of whole-brain connectomes that capture the developing brain's full connectivity profile. The impact of the gyral bias, a methodological confound of tractography, is quantified and compared with the equivalent measurements for adult data.

These connectomes form the basis for a novel, data-driven framework, in which they are decomposed into white matter bundles and their corresponding grey matter terminations. Independent component analysis and non-negative matrix factorisation are compared for the decomposition, and are evaluated against in-silico simulations. Data-driven components of dMRI tractography data are compared with manual tractography, and networks obtained from resting-state functional MRI. The framework is further developed to provide corresponding components between groups and individuals. The data-driven components are used to generate cortical parcellations, which are stable across subjects.

Finally, some future applications are outlined that extend the use of these methods beyond the context of neonatal imaging, in order to bridge the gap between functional and structural analysis paradigms, and to chart the development of white matter throughout the lifespan and across species.

Acknowledgements

This PhD was part of the Oxford-Nottingham Biomedical Imaging CDT, funded by the EPSRC. I would like to thank Professor Peter Jezzard and Professor Penny Gowland for putting together such a great course, and for your personal advice and support, especially during the uncertainty of the pandemic. I would also like to thank everyone in the ONBI 2016 cohort, for the fun and camaraderie we shared during the training year and beyond.

I would like to thank my supervisors Dr. Stam Sotiropoulos and Dr. Matteo Bastiani for all your guidance and help over the course of this PhD. You have always been generous with your time and advice, and encouraged me to produce my best work. Our meetings have always given me fresh ideas and enthusiasm, and I have learned a huge amount from working with you both. I would also like to thank Professor Matt Brookes, for sparking my interest in brain imaging back when I first came to SPMIC on work experience, and for encouraging me to apply to the CDT.

I would like to express my gratitude to all the co-authors who collaborated with me on conference abstracts and offered feedback and advice on paper drafts, and to Rebecca Woodrow for your help with the baby-XTRACT protocols. I would also like to thank my examiners, Professor Daniel Alexander and Professor Richard Bowtell, for the interesting discussions and helpful comments during the viva.

To everybody in the Radiological Sciences Group, thank you all for making it such a warm and friendly place to work, all the coffee breaks and snack trips have been much missed over the past months. Cheryl, thank you for keeping us all organised and for everything you've done to keep lab morale up during the pandemic. To Shaun especially, I owe a huge thank you - for all the tractography chats, for all the times you put things in perspective when I was stressed, and more recently cheered me up with GIFs and jokes when thesis writing got a bit much. I hope I can return the favour in the New Year.

Thank you to all my friends who have been so supportive during this PhD. To all

the TAFKAPers: Eleanor, George, Stefan, Tiago and Poppy, thanks for being such lovely housemates, especially during the lockdown. Those weekly potlucks really kept me going during that weird time. George, thanks for all the laughs and the deep conversations at the kitchen table, and for being such a great friend. To Sarina, thanks for all the great tram journeys, and all the fun times we've had together over the past few years. It's been so inspiring to watch you grow your business this year and I can't wait until we can be back together hitting the Beeston pub quiz scene once again! To Issie, thanks for your (20 years of !) support and friendship, it really means a lot to me.

To Louis, thank you for sticking with me through long-distance weekends and lockdowns. It's been a crazy journey but I wouldn't have done it with anyone else.

The biggest thank you of all goes to my family: Nain and Taid, Mum and Izzie, thank you for being there for me every step of the way. I'm so grateful that I got to spend more time with you all while I was studying in Nottingham. You've been endlessly supportive, and always there to give me a boost when I needed it. This thesis is dedicated to you.

List of Publications

Journals

- E. Thompson, A.R. Mohammadi-Nejad, E.C. Robinson, J.L.R. Andersson, S. Jbabdi, M.F. Glasser, M. Bastiani, S.N. Sotiropoulos. Non-negative data-driven mapping of structural connections with application to the neonatal brain, *Neuroimage*, 222:117273, 2020.
- E. Thompson, J. Dubois, S.N. Sotiropoulos, M. Bastiani. baby-XTract: Standardised protocols for automated tractography in the neonatal brain, *In preparation*

Conferences

- E. Thompson, S. Jbabdi, M.F. Glasser, M. Bastiani, S.N. Sotiropoulos. Non-Negative decomposition of structural connectivity in the developing brain. *OHBM Annual Meeting*, Online, June 2020.
- E. Thompson, S. Jbabdi, M. Bastiani, S.N. Sotiropoulos. Data-driven mapping of structural connectivity patterns in the neonatal brain. *OHBM Annual Meeting*, Rome, Italy, June 2019.
- E. Thompson, E.C. Robinson, J. Božek, S. Jbabdi, M. Bastiani, S.N. Sotiropoulos. Exploring the gyral bias on white matter tractography in neonates. *OHBM Annual Meeting*, Rome, Italy, June 2019.
- E. Thompson, S. Jbabdi, M. Bastiani, S.N. Sotiropoulos. Data-driven exploration of structural connectivity patterns in the neonatal brain. *British Chapter ISMRM Postgraduate Symposium*, London, UK, April 2019.
- E. Thompson, M. Bastiani, M. Brookes, S. Jbabdi, S.N. Sotiropoulos. Asynchronous macrostructural changes in white matter tracts of the developing brain. *Proceedings of the ISMRM Annual Meeting*, Paris, France, June 2018.

Contents

Abstract	i
Acknowledgements	ii
List of Publications	iv
I Introduction	1
1 Introduction	2
1.1 Organisation of the Thesis	4
1.2 Software	6
II Background	7
2 Measuring Connections in the Brain Using MRI	8
2.1 What is Structural Connectivity?	9
2.2 What is Functional Connectivity?	10
2.3 Diffusion MRI	12
2.3.1 Basic concepts of Magnetic Resonance Imaging	13
2.3.2 Diffusion of Water Molecules	15
2.3.3 Using MRI to Measure Diffusion	16
2.3.4 The Diffusion Tensor	18
2.3.5 Modelling Crossing Fibres - Beyond the Diffusion Tensor	23
2.3.6 Tractography	30
2.3.7 Summary	35

3	Imaging the Early Development of Connectivity	37
3.1	The Early Development of Structural Connectivity	38
3.2	Using dMRI to Map Early Brain Development In Vivo	39
3.2.1	Delineating Tracts	40
3.2.2	Estimating Tract Microstructure	42
3.2.3	Whole Brain Connectivity Analysis	44
3.2.4	Impact of External Factors on Structural Connectivity	46
3.3	The Challenges of Imaging Early Brain Development In Vivo	47
3.4	Big Data Projects	48
3.4.1	The Developing Human Connectome Project	49
3.4.2	The Baby Connectome Project	50
3.4.3	Aims of this Work	51
III	Original Research	53
4	Tractography Approaches to Map Connectivity During Early Development	54
4.1	Introduction	56
4.2	Methods: Automated Tractography Protocols	59
4.2.1	Protocol Definitions	60
4.2.2	Data	68
4.2.3	Model Fitting	68
4.2.4	Tractography and Atlases	69
4.2.5	Comparison with Existing Protocols	70
4.2.6	Microstructural Analysis	72
4.3	Methods: Gyral Bias	72
4.3.1	Data	72
4.3.2	Surface Alignment	73
4.3.3	Predicted Bias	74

4.3.4	Tractography Bias	74
4.4	Results	76
4.4.1	Automated Tractography Protocols	76
4.4.2	The Impact of the Gyral Bias on the Neonatal Connectome	80
4.5	Discussions	85
4.5.1	Standardised Protocols for Neonatal Tractography	86
4.5.2	The Impact of the Gyral Bias on the Neonatal Connectome	87
4.6	Concluding Remarks	88
4.7	Appendix	89
5	Using Data-Driven Methods To Map Structural Connectivity in the Neonatal Brain	92
5.1	Introduction	94
5.2	Theory	96
5.2.1	Matrix Decomposition with ICA and NMF	96
5.3	Methods	99
5.3.1	Generating Connectivity Matrices	100
5.3.2	Dimensionality Reduction and Back-Projection	102
5.3.3	Simulations	102
5.3.4	In Vivo Data Decompositions	106
5.3.5	Comparison to Tractography-Derived White Matter Tracts	107
5.3.6	Comparison with fMRI	107
5.4	Results	109
5.4.1	Simulations	109
5.4.2	In Vivo Data Results	113
5.5	Discussion	127
5.5.1	Validation using Simulations	127
5.5.2	In Vivo Results	128
5.5.3	Decomposition Domain	130

5.6	Concluding Remarks	131
6	Reliability and Applications of Data-Driven Connectivity Mapping	132
6.1	Introduction	134
6.2	Theory	135
6.2.1	From Group to Individual Subject Decompositions - Non-Negative Dual Regression	135
6.3	Methods	137
6.3.1	Split-half Reliability Analysis	137
6.3.2	Non-Negative Dual Regression	138
6.3.3	Cortical Parcellations	139
6.4	Results	140
6.4.1	Assessing the Reproducibility of the Decompositions	140
6.4.2	Parcellations	144
6.5	Discussion	148
6.5.1	Reproducibility	148
6.5.2	Parcellations	149
6.5.3	Limitations and Future Work	152
6.6	Concluding Remarks	153
7	Conclusions and Future Directions	154
7.1	Summary and Conclusions	155
7.2	Future Directions	157
7.2.1	Cross-Modal Connectivity Analysis	157
7.2.2	Applications of Non-Negative Dual Regression	158
7.2.3	Exploring the Ontogeny and Phylogeny of Brain Connections	160
7.2.4	Connectivity-Based Alignment and Parcellation	165
7.3	Final Remarks	166
	Bibliography	168

Acronyms

189

List of Tables

4.1	Forty-two tracts included in baby-XTRACT	61
4.2	Streamlines per seed used for each of the forty-two tracts in the baby-XTRACT protocols	71

List of Figures

1.1	Comparison between representative neonatal and adult brains	4
2.1	Diagram of a neuron	10
2.2	Spin echo pulse sequence	14
2.3	Illustration of diffusion profiles in different tissues	16
2.4	Pulsed gradient spin echo sequence used in diffusion MRI	17
2.5	Illustration of the effective diffusion ellipsoid	20
2.6	Example of diffusion tensors modelled from real MRI data.	21
2.7	Different DTI parameters in an adult subject	22
2.8	Schematic illustration of the spherical deconvolution process.	25
2.9	Comparison of different extensions of the ball and stick model	30
2.10	Schematic diagram of the streamline tractography process	32
2.11	Comparison of deterministic and probabilistic tractography.	33
3.1	Timeline of the different maturational processes in the brain.	39
3.2	Average microstructure maps from different age groups.	42
3.3	T2-weighted brain atlases from neonates and adults.	48
4.1	Diagram illustrating the anatomical basis of the gyral bias	58
4.2	Schematic diagram of the baby-XTRACT masks and their functions .	59
4.3	Comparison of baby-XTRACT results with varying numbers of stream- lines per seed point.	70

4.4	Schematic illustration of the two seeding strategies used to investigate the gyral bias	75
4.5	Percentage average maps of the tracts from the baby-XTRACT protocols, for different age groups within the cohort.	77
4.6	Results from baby-XTRACT in individual subjects	78
4.7	Comparison of baby-XTRACT results with those from adult humans and macaques	79
4.8	Boxplots showing correlations between spatial maps from XTRACT and AutoPtx.	80
4.9	Scatter plots of average tract FA against subjects' age at scan	81
4.10	Scatter plots of average tract MD against subjects' age at scan	82
4.11	Barchart of beta values from GLM of tract microstructure	83
4.12	Average predicted bias maps, based on the cortical volume per unit area of the WGB.	83
4.13	Average streamline density maps using different seeding approaches, for neonates and adults.	84
4.14	Violin plots to compare the predicted and measured streamline density in the gyri and sulci, for different seeding approaches	85
4.15	Violin plots of the correlation scores between the predicted and measured streamline density, for neonates and adults	85
4.16	Comparison of baby-XTRACT and AutoPtx protocols for the acoustic radiation.	89
4.17	Comparison of baby-XTRACT and AutoPtx protocols for the cingulum.	90
4.18	Comparison of baby-XTRACT and AutoPtx protocols for the fornix.	90
4.19	Comparison of baby-XTRACT and AutoPtx protocols for the optic radiation.	91
5.1	Data-driven decomposition of connectivity matrices	97

5.2	Schematic of back projection step to obtain full mixing matrix after PCA	103
5.3	Simulation results with varying L1-norm regularisation	110
5.4	Simulation results with varying numbers of sources	111
5.5	Simulation results with varying noise levels	112
5.6	Comparison between data-driven connectivity maps and virtual dissection	119
5.7	Detailed comparison of ICA and NMF components	121
5.8	Hierarchical clustering of NMF components	122
5.9	Data-driven callosal segments	123
5.10	Unassigned components	124
5.11	A comparison of grey matter structural components with resting-state networks from fMRI	125
5.12	The reconstruction error of the in vivo data-driven components at different model orders	126
5.13	Sparsity of the in vivo data-driven components	126
5.14	Comparing results from decompositions applied to transpose connectivity matrix	131
6.1	Illustration of standard dual regression technique.	136
6.2	Illustration of novel non-negative dual regression technique.	137
6.3	Split-half analysis	141
6.4	Boxplots illustrating the consistency of the NMF results across different group sizes.	141
6.5	Data-driven components from different numbers of subjects	142
6.6	Comparison of results from standard and non-negative dual regression approaches.	143

6.7	Quantitative comparison of dual regression results with subject-level decompositions across the cohort	144
6.8	Qualitative comparison of dual regression results with subject-level decompositions from two representative subjects	145
6.9	Example components from a group-level dual regression	146
6.10	Data-driven parcellations	147
6.11	Analysis of subject-level parcellations	148
6.12	Split-half analysis of data-driven parcellation scheme	149
6.13	Silhouette score of data-driven parcellations	150
6.14	The effect of the gyral bias on data-driven parcellations	151
7.1	Statistical analysis of dual-regressed structural components	159
7.2	Generating connectivity fingerprints	161
7.3	Cross-species comparison of connectivity patterns from XTRACT	162
7.4	Comparing connectivity profiles from different groups using KL Divergence	163
7.5	Comparison of data-driven NMF components from adult and neonatal subjects.	164
7.6	Multi-modal contributions to parcellations and surface-based registration	166

Part I

Introduction

Chapter 1

Introduction

Contents

1.1	Organisation of the Thesis	4
1.2	Software	6

Introduction

During the months after birth, the human brain undergoes a period of rapid development and growth, increasing in size by around 1% each day during the first three months of life (Holland et al., 2014). Many complex processes drive this growth, including synaptogenesis, growth of axons and dendrites, and myelination of the subcortical white matter (Kostović et al., 2019). This occurs concurrently with fibre pruning, which then continues into adolescence (Huttenlocher and Dabholkar, 1997). These processes are impacted by prematurity (Batalle et al., 2017; Anjari et al., 2007) and the early life environment (Lautarescu et al., 2020; Rifkin-Graboi et al., 2015), and in turn have downstream effects on later outcomes (Counsell et al., 2008; Ball et al., 2015; Girault et al., 2019).

Post-mortem dissection provides valuable ground-truth data about the early growth and maturation of brain connections (Kinney et al., 1988). However, in order to study the longitudinal impacts of early life experiences on brain development, and the links between brain structure and function, we need tools that can be used to image the brain non-invasively and in vivo. Magnetic Resonance Imaging (MRI) has been used for biomedical imaging since the 1970s (Mansfield and Maudsley, 1977; Lauterbur, 1974). MRI is an ideal modality for imaging neonates because it does not use harmful ionising radiation, and provides multiple contrasts that allow us to study different aspects of development. One of these is diffusion MRI (dMRI), which measures the diffusive motion of water molecules to provide information about their environment. For example, a molecule in the ventricles will diffuse freely in all directions, whereas a water molecule in the fibrous white matter will be more likely to diffuse parallel to the orientation of the fibre bundle than perpendicular to it. Using this information to trace the likely route of white matter fibre bundles through the brain is called tractography. This has been crucial for mapping early white matter development, for example in elucidating the asynchronous maturation

trajectories of different tracts (Dubois et al., 2008b).

However, the analysis paradigms used to study neonatal white matter have been mainly based on the analysis pipelines developed for adult brains. As we can see from figure 1.1, there are substantial differences when imaging adult and neonatal brains, including tissue composition, relative resolution, image contrast and morphological differences. This means that standard approaches are not necessarily suited to the specific challenges associated with imaging the neonatal brain. In this thesis, we devise new methods that are tailored for mapping the brain and its connectional architecture during this unique period of development.

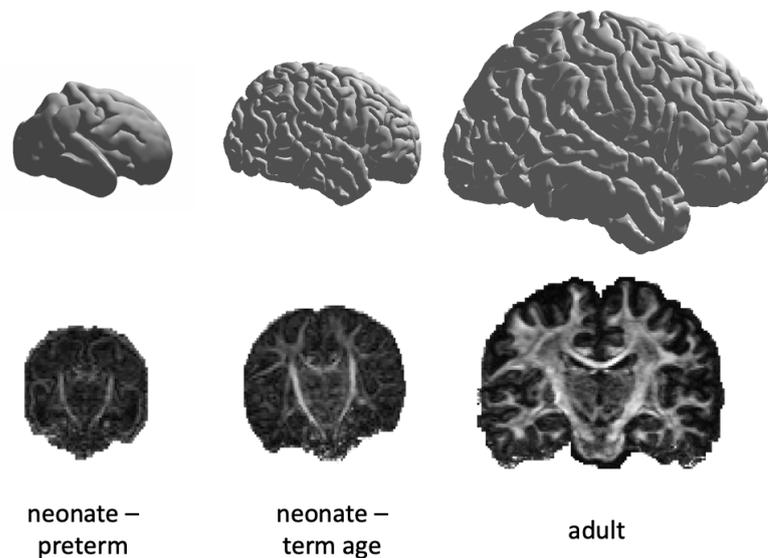


Figure 1.1: *A comparison between representative neonatal and adult brains. Top: pial surface reconstruction, bottom: fractional anisotropy maps from diffusion MRI.*

1.1 Organisation of the Thesis

This thesis is organised into seven chapters. The next two chapters of the thesis provide a background on the current methods used to map early white matter development and the main findings to date. The following three chapters describe the new contributions of the thesis: first exploring tractography approaches in the neonatal brain, devising new protocols for tract delineation, but also for whole-brain

connectomes in a large neonatal cohort; then the development of a new, data-driven framework for mapping white matter connections and grey matter networks from dMRI data. The final chapter provides an overview and some ideas for future work.

Chapter 2 introduces and explains the concepts of structural and functional connectivity in the brain, and the techniques available to measure them, with particular focus on diffusion MRI. An overview is given of the physical basis for dMRI, different methods for modelling fibre orientations, as well as the different tractography techniques that are used to map white matter bundles.

Chapter 3 outlines how brain imaging has been used to chart the early development of structural connectivity in neonates. It details the key findings, and the unique challenges of studying this population. It also discusses the advantages of population-level, or “Big Data” projects, such as the developing Human Connectome Project.

Chapter 4 begins the section of original work, exploring different tractography approaches to map neonatal white matter, from standardised tractography protocols to whole brain connectomes. It describes a new set of tractography protocols that have been developed, which facilitate the mapping of 42 white matter tracts. It also explores different ways of generating whole brain connectomes, and the impact of the gyral bias on these, which is an important confound for mapping connections in adult brains.

Chapter 5 introduces data-driven methods for extracting white matter tracts and their corresponding grey matter networks from whole-brain connectomes. Two approaches are compared: independent component analysis and non-negative matrix factorisation, using both simulated data and in vivo neonatal MRI data. A new method for non-negative dual regression is introduced, to generate subject-level versions of the group-level structural connectivity patterns. Data-driven white matter bundles are validated against results from standard tractography protocols, and their corresponding grey matter networks are compared to resting-state networks

from functional MRI.

Chapter 6 provides further analysis and applications of the non-negative approach outlined in the previous chapter. Subject-level components are compared against the group-level maps, and the stability of the results are assessed across different subject groups. The grey matter networks are used to generate a novel cortical parcellation scheme, based on structural connectivity. The reliability and validity of the parcellations are assessed.

Chapter 7 summarises the findings of the thesis, and outlines some future applications for the methods that have been developed.

1.2 Software

Routines for data-driven connectivity mapping with Independent Component Analysis (ICA) or Non-Negative Matrix Factorisation (NMF) can be found on GitHub: (<https://github.com/ethompson93/Data-driven-tractography>), including code for regressing group components to individuals. Standardised neonatal tractography protocols are also available on GitHub: (<https://github.com/ethompson93/baby-XTRACT>).

Statistical analysis and new data-driven tractography methods were implemented using the Python programming language (version 3.7.2) (Python Software Foundation, <https://www.python.org>). Connectome Workbench (version 1.3.2), developed by the Van Essen Laboratory at Washington University, was used for processing of surface data, and for visualisation of data on the cortical surface. The FMRI software library, FSL (version 6.0.3) was used for all other processing of imaging data, including probabilistic tractography (Jenkinson et al., 2012). FSL's image viewer, FSLeyes, was used for the visualisation of volumetric data.

Part II

Background

Chapter 2

Measuring Connections in the Brain Using MRI

Contents

2.1	What is Structural Connectivity?	9
2.2	What is Functional Connectivity?	10
2.3	Diffusion MRI	12
2.3.1	Basic concepts of Magnetic Resonance Imaging	13
2.3.2	Diffusion of Water Molecules	15
2.3.3	Using MRI to Measure Diffusion	16
2.3.4	The Diffusion Tensor	18
2.3.5	Modelling Crossing Fibres - Beyond the Diffusion Tensor	23
2.3.6	Tractography	30
2.3.7	Summary	35

2.1 What is Structural Connectivity?

The human brain contains around 86 billion neurons (Herculano-Houzel, 2009), which are connected via synapses to form a complex network that facilitates the rich and varied functioning of the brain. We can broadly divide the brain into two tissue types: grey matter and white matter. Grey matter is mostly composed of neuronal cell bodies and is located in the cerebral cortex, the folded layer of grey matter surrounding the cerebral hemispheres, and the subcortical grey matter structures. The grey matter is where functional processing takes place, as neurons respond to electro-chemical stimuli from other neurons, organs in the body, or sensory input.

The white matter mainly consists of myelinated axons. An axon connects the cell body to the axon terminals, where the neuron can form a synaptic connection with another cell, as illustrated in figure 2.1. It is these axons that form the structural connections in the brain, transmitting electrical impulses, known as axon potentials, from one neuron to the next. Although individual axons are only up to a few micrometers in diameter, they group together in thousands to form large macroscopic bundles, or white matter tracts, which connect functional brain regions (Walhovd et al., 2014). Connections between different brain regions are crucial for integration and information transfer, analogous to the wiring in a computer that links the different components together. Figure 2.1 also shows the myelin sheath that encloses the axon. Myelin is a fatty substance that encloses the axons and insulates them from each other, increasing the speed at which information can be transmitted.

A major goal of modern neuroscience is the mapping of the connectome, a comprehensive description of the connections between different brain regions (Sporns et al., 2005). Mapping the connectome requires the definition of nodes and edges. At the microscopic level, nodes would be individual neurons and edges individual axons. At the macroscopic systems level, nodes can be thought of as discrete brain

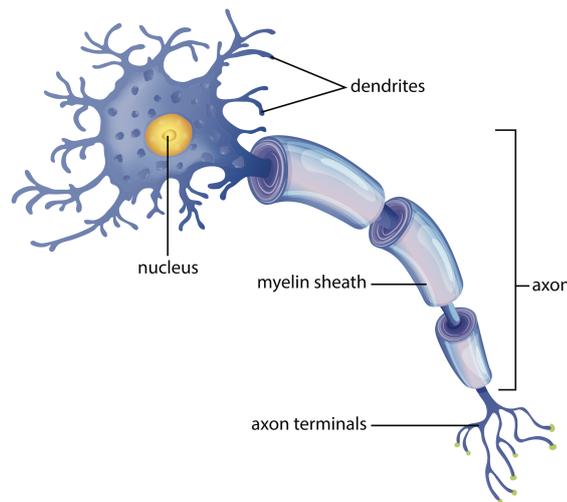


Figure 2.1: A schematic diagram of a neuron. (Reproduced from vecteezy.com)

regions and edges the white matter bundles that link them. Different tools can be used to study the brain’s structural connectome at these different length scales, such as post-mortem histology for connections between individual neurons (at the microscale) or neuronal tracers for neural assemblies (at the mesoscale). However, non-invasive measurement of connections between brain structures is only possible at the macroscale with diffusion MRI (dMRI) (Jbabdi et al., 2015).

2.2 What is Functional Connectivity?

Neurons in the brain are constantly communicating with each other via electrochemical signals called action potentials. Groups of neurons activate or deactivate together to perform different functions, and these networks can be measured “at rest” when the subject is not attending to a task (Biswal et al., 1995). Two regions are assumed to be functionally connected if their measured spontaneous (resting) activity profiles exhibit some statistical dependence (Aertsen et al., 1989; Friston et al., 1993). Reversing this definition, if the temporal profile of the activity of different regions exhibit similarities, they are thought to be part of the same network and therefore structurally connected. This indirect, yet powerful, way of describing

connections is called functional connectivity.

The electromagnetic fields induced by these synchronised neuronal currents can be measured directly on the scalp using electroencephalography (EEG) or magnetoencephalography (MEG), or indirectly through the alterations in blood flow induced by neuronal activity in the case of functional MRI (fMRI). Specifically, following activation in a brain region there is a local drop in the vascular concentration of deoxyhaemoglobin. This arises because the increase in cerebral blood flow outweighs oxygen metabolism following brain activity. As deoxyhaemoglobin is paramagnetic, this leads to a difference in susceptibility between the blood vessel and the surrounding tissue, which can be measured with MRI (Buxton, 2013). This is referred to as the blood oxygenation level-dependent (BOLD) signal (Ogawa et al., 1990). fMRI enables researchers to measure activity across the whole brain completely non-invasively, without the need for ionising radiation.

Functional and structural connectivity can be seen as complementary approaches to estimate the connectome, as regions that are connected by white matter pathways are likely to have statistical dependence in their functional activity. Therefore, structural connections can be inferred between pairs of regions with strong correlations between their functional activity profiles (Jbabdi et al., 2015). However, statistical dependence in functional activity does not guarantee a direct structural connection between regions: intermediate connections or a common input can cause two regions to have correlated activity when in fact they are not directly connected. Physiological noise can also affect the BOLD signal and introduce spurious correlations between brain regions (Jbabdi et al., 2015).

This thesis focuses on structural connectivity, and how we can use diffusion MRI to map its developmental trajectory in the neonatal brain. The next section overviews the theory and methods typically used to analyse diffusion MRI data, which can be used to explore the development of the brain connectome in vivo.

2.3 Diffusion MRI

Prior to the invention of MRI, researchers were limited to post-mortem methods for studying brain anatomy. Gross dissection was first used to identify macroscopic brain structures and to differentiate between grey and white matter. The invention of the microscope by Hans and Zacharius Janssen in 1590 paved the way for cell theory and allowed scientists to see individual nerve fibres for the first time (Catani et al., 2013b), and later, staining techniques developed by Golgi and Cajal elucidated the cell structure of the neuron for the first time (Morecraft et al., 2013; Glickstein, 2006). More recently, histo-chemical tracer studies in non-human primates have been used to visualise connections at the resolution of single cells (Morecraft et al., 2013). At the microscale, the introduction of electron microscopy in the 1950s has enabled the quantification of white matter structure, such as myelin thickness and axon density (Walhovd et al., 2014).

The invasive methods described above have brought about huge advances in our understanding of brain anatomy. However, dMRI is the only tool that can measure the white matter structure of the brain non-invasively and *in vivo*. Although dMRI is an indirect method, the ability to image structural connectivity in the living brain opens a much wider field of research than would be possible with invasive methods alone. For example, dMRI facilitates longitudinal studies; studies of the effect of environmental factors on brain structure; and multi-modal investigations of the relationship between brain structure and function. In addition, dMRI allows us to study the very early development of structural connections in the neonatal brain, which is the subject of this thesis. The following sections provide an introduction to the physics underlying the phenomenon of nuclear magnetic resonance (NMR), and the basic principles of MRI. This is followed by an overview of diffusion and how we can use MRI to measure diffusive motion of water molecules, and then use this information to reconstruct white matter pathways in the brain.

2.3.1 Basic concepts of Magnetic Resonance Imaging

Protons have an intrinsic quantum mechanical property called “spin”, which means that they have a magnetic dipole moment. This causes them to behave like tiny bar-magnets that will tend to align to an external magnetic field. When there is no magnetic field present, an ensemble of spins will be randomly aligned, and so their overall net magnetisation will be zero. However, when a magnetic field is applied, which we call B_0 , the spins align with the field and precess around it. The frequency at which the spins precess is called the Larmor frequency ω , and is given by the following equation:

$$\omega = \gamma B \tag{2.1}$$

Where γ is the gyromagnetic ratio (42.58 MHz/T for protons), and B is the magnitude of the applied field. In other words, the frequency of the precession is proportional to the strength of the applied field. The spins precess at the same speed but with different phases, so the transverse components of the magnetisation cancel out, yielding a net magnetisation vector \mathbf{M} in the sample aligned to B_0 .

To induce magnetic resonance, a second magnetic field B_1 is applied, perpendicular to B_0 , this time as a radio frequency pulse with frequency ω . This transfers energy to the protons and introduces phase coherence in the precession of the spins, introducing a component of the magnetic field in the transverse plane to B_0 , rotating with a frequency ω (Bloch, 1946). When a coil is placed around the sample, the oscillating magnetic field induces a proportional current in the coil as a consequence of Faraday’s law, which forms the magnetic resonance signal.

In time, \mathbf{M} returns to its original orientation. This is caused by two processes: T1 relaxation and T2 relaxation, sometimes referred to as spin-lattice and transverse relaxation, respectively (Bloch, 1946; Purcell et al., 1946). Energy is released through interactions between the spins and their environment, which causes \mathbf{M} to

return to its original orientation with a time constant T_1 . T_2 relaxation results from the magnetic dipole of each spin interacting with the B_0 field, which causes local inhomogeneities in the magnetic field. This has the net effect that the spins become dephased, reducing the transverse component of \mathbf{M} . Additional inhomogeneities in the magnetic field, due to scanner itself or due to susceptibility differences in the sample, contribute to a quicker dephasing of the signal than would occur from T_2 processes alone (Bloembergen et al., 1948), this observed transverse relaxation is called T_2^* . T_1 and T_2 relaxation coefficients vary across tissue types, which is the basis for tissue contrast in structural MRI.

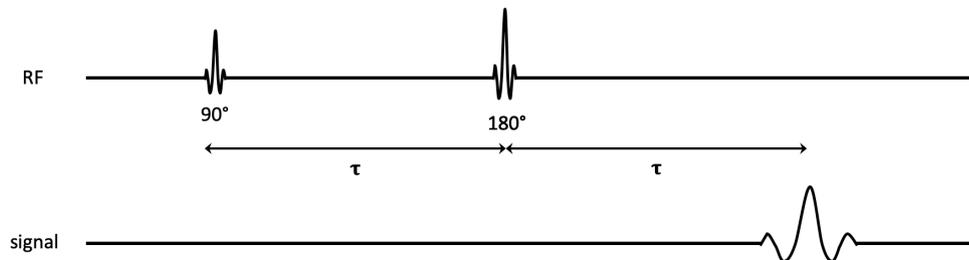


Figure 2.2: A pulse sequence diagram of a spin echo sequence.

Relaxation processes can be probed using the spin echo sequence, one of the basic and most widely used pulse sequences for MRI (Hahn, 1950). The pulse sequence is shown in figure 2.2. It starts with a 90° radio frequency pulse B_1 that flips \mathbf{M} into the transverse plane (the angular values describe the angle through which \mathbf{M} is rotated relative to B_0). The spins begin to dephase during the time τ , after which a 180° pulse is applied. This rephasing pulse reverses the T_2^* processes, so that an echo signal can be read out at the echo time, $TE = 2 * \tau$ after the original 90° pulse.

The measured signals provide information about the T_1 and T_2 relaxation properties of the sample, but to reconstruct an image we need some way of differentiating the signals from different locations. This is achieved by applying spatially-varying magnetic field gradients across the sample. A magnetic field gradient causes the

spins at different locations to precess at slightly different frequencies, thereby encoding spatial information into the measured signal (Mansfield and Maudsley, 1977).

2.3.2 Diffusion of Water Molecules

Diffusion MRI uses the phenomenon of diffusion to introduce extra contrast into a spin-echo experiment. But what is diffusion? Diffusion is the random, thermally-driven motion of particles in a fluid. Freely diffusing water molecules undergo random Brownian motion, with a mean-squared diffusion distance $\langle r^2 \rangle$ (Einstein, 1905):

$$\langle r^2 \rangle = 2D\Delta \quad (2.2)$$

Where D is the diffusion coefficient of the medium, and Δ is the time over which the diffusion occurs. D depends on the size of the diffusing particles, their temperature, and the microstructural properties of the environment. It is this dependence that enables us to probe the microstructure of the brain using diffusion MRI. In a medium in which particles can diffuse freely, equation 2.2 gives rise to a Gaussian distribution of displacements, with a width proportional to D (D. C. Alexander, 2006). However, in biological tissues, diffusion is restricted by interactions with tissue compartments and macromolecules, so that the measured diffusion coefficient is much lower than it would be in free water (Beaulieu, 2002). We call this the apparent diffusion coefficient (ADC) (Le Bihan et al., 1986).

In the brain, water molecules diffuse in a way that is determined by their local tissue environment. This is illustrated schematically in figure 2.3. For example, water molecules in the cerebro-spinal fluid in the ventricles can diffuse freely and isotropically. Water molecules in the white matter, on the other hand, have much more hindered motion, particularly along the axes perpendicular to the axons.

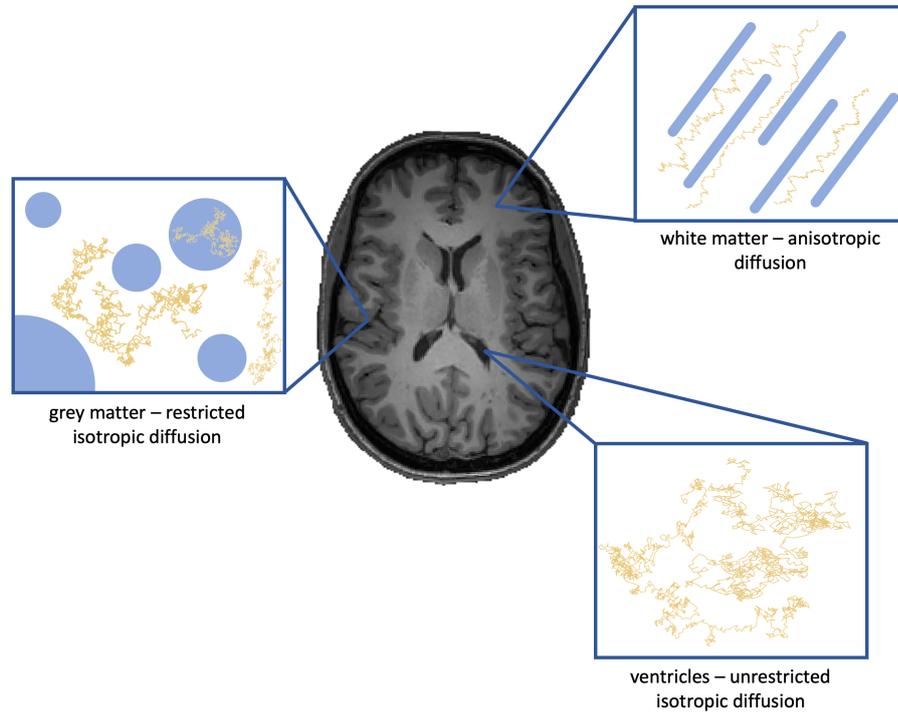


Figure 2.3: Illustration of diffusion profiles in different brain tissues. The boxes show illustrative trajectories of water molecules (yellow) in the presence of different cell types (blue). Axons are illustrated by blue cylinders, and cell bodies by blue circles.

2.3.3 Using MRI to Measure Diffusion

As described above, the diffusion of water molecules in the brain is dependent on their local tissue environment. Using MRI pulse sequences that are sensitive to diffusion, we can infer the structural properties of the underlying tissue.

In 1965, Stejskal and Tanner introduced the pulsed gradient spin echo sequence, a modification of the spin echo sequence in figure 2.2, in which two diffusion-encoding gradient pulses $\mathbf{G}(t)$ are applied after the initial excitation pulse, either side of the 180° refocussing pulse. Each gradient pulse lasts for time δ , separated by a time interval of Δ , as shown in figure 2.4. The gradients \mathbf{G} can be applied in any direction $[G_x(t), G_y(t), G_z(t)]$, which determines the direction of the diffusive motion that will be measured.

From equation 2.1, the phase accumulated during the first pulse by a spin located at position \mathbf{x}_1 along the direction of $\mathbf{G}(t)$, is given by:

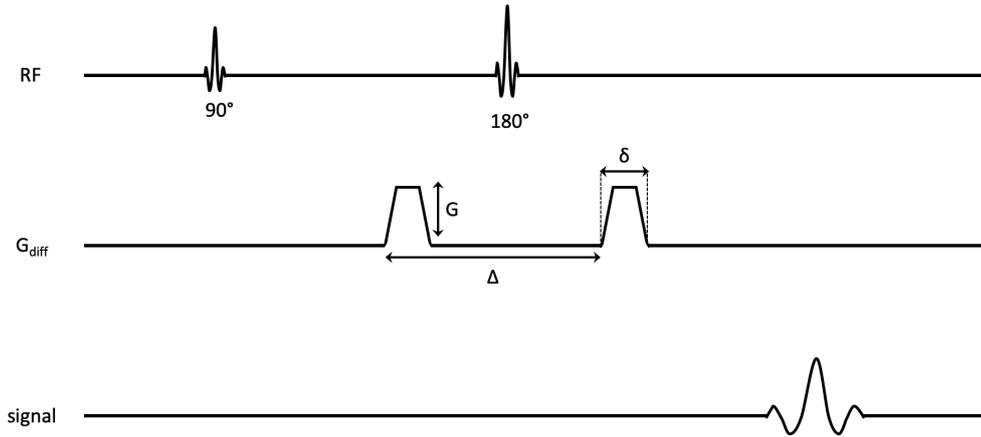


Figure 2.4: A diagram of the pulsed gradient spin echo sequence used to obtain diffusion-weighted MRI images. \mathbf{G}_{diff} describes a diffusion-encoding gradient along an arbitrary direction.

$$\begin{aligned}\phi_1 &= - \int_0^\infty \omega(t) dt \\ &= - \gamma \mathbf{x}_1 \int_0^\infty \mathbf{G}(t) dt\end{aligned}\tag{2.3}$$

We can introduce the term \mathbf{q} , which summarises the experimental parameters:

$$\mathbf{q} = \gamma \int_0^\infty \mathbf{G}(t) dt\tag{2.4}$$

So that the expression for the change in phase simplifies to $\phi_1 = -\mathbf{q} \cdot \mathbf{x}_1$. When the 180° pulse is applied, the phase change induced by the first pulse is reversed. Therefore, a spin that has remained stationary over the time interval Δ will experience no net change in phase. However, a spin that moves during this time to position \mathbf{x}_2 , will experience a net phase change given by:

$$\begin{aligned}\phi &= \phi_2 - \phi_1 \\ &= -\mathbf{q} \cdot (\mathbf{x}_2 - \mathbf{x}_1)\end{aligned}\tag{2.5}$$

In other words, the spin experiences a change in phase proportional to its displacement along the direction of the applied gradient. When this effect is averaged over the random displacement of the spins present in an imaging voxel, it leads to phase dispersion and attenuation of the observed signal, compared to the signal that would have been measured without the application of diffusion-sensitising gradients, i.e. $\mathbf{G} = 0$. This attenuation is proportional to the displacement of the spins along the gradient direction, during the time period between the two pulses.

The signal attenuation attributed to diffusion can be expressed as a function of the diffusion weighting \mathbf{q} :

$$E(\mathbf{q}) = \frac{S(\mathbf{q})}{S_0} \quad (2.6)$$

Here, S_0 is the signal measured in the absence of any diffusion encoding gradients. This baseline allows us to disregard the signal dependence on other factors. The signal attenuation is linked to the diffusion propagator $P(\mathbf{r}, \tau)$ via a Fourier transform (Stejskal and Tanner, 1965):

$$E(\mathbf{q}, \tau) = \int P(\mathbf{r}, \tau) \cdot e^{-i\mathbf{q} \cdot \mathbf{r}} d\mathbf{r} \quad (2.7)$$

The diffusion propagator $P(\mathbf{r}, \tau)$, is a Green's function that describes the probability for a given particle to diffuse by a displacement \mathbf{r} in a time τ (Stejskal and Tanner, 1965). This indicates a reciprocal relationship between the strength of the magnetic field gradients and the length scales of the measured diffusion.

2.3.4 The Diffusion Tensor

For simple isotropic Gaussian diffusion, the signal attenuation will also have a Gaussian profile. Placing a Gaussian diffusion propagator P with a variance proportional to D (equation 2.2) into equation 2.7, results in a signal attenuation:

$$E = e^{-bD} \quad (2.8)$$

We have introduced another term b , sometimes called the “b-value”, which reflects the strength and timing of the diffusion gradients from the pulsed spin echo sequence: $b = f(\mathbf{q}, \Delta, \delta)$, where the exact relationship depends on the shape of the gradients used (Stejskal and Tanner, 1965). For anisotropic diffusion, the scalar diffusion coefficient D is replaced by a diffusion tensor \mathbf{D} (i.e. the Gaussian propagator with a scalar variance is replaced by a covariance matrix proportional to a tensor \mathbf{D}), and the measured attenuation depends on the direction of the diffusion encoding gradients \mathbf{g} :

$$E(\mathbf{q}) = e^{-b\mathbf{g}^T \cdot \mathbf{D} \cdot \mathbf{g}} \quad (2.9)$$

Given sufficient measurements with different gradient directions \mathbf{g} , we can solve this equation to calculate the diffusion tensor, \mathbf{D} (Basser et al., 1994). \mathbf{D} can be estimated from the data using multiple linear least squares (Basser et al., 1994), or non-linear methods (Koay et al., 2006).

$$\mathbf{D} = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{bmatrix} \quad (2.10)$$

The diagonal terms describe the diffusion coefficients along each of the coordinate axes. The off-diagonal terms are the covariance terms, and are symmetric about the diagonal ($D_{ij} = D_{ji}$). Thus, there are six unique terms in the tensor. The coordinate axes (x, y, z) are the axes defined by the scanner gradients, but these may not provide the most useful reference frame to describe the diffusion profile. Instead, we can diagonalise the tensor to obtain its eigenvalues λ_i , for $i = 1, 2, 3$.

The eigenvalues describe the diffusion along the tensor eigenvectors ϵ_i , which are given by the eigenvector relationship: $\mathbf{D}\epsilon_i = \lambda_i\epsilon_i$. The eigenvectors provide a new coordinate system that defines the principal axes of diffusion, with the corresponding eigenvalues giving the apparent diffusivities along each axis.

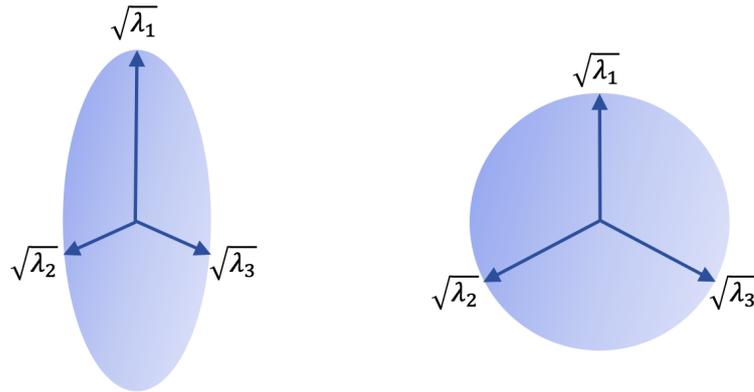


Figure 2.5: Illustration of two diffusion tensor ellipsoids. The left represents anisotropic diffusion, in which $\lambda_1 \gg \lambda_2, \lambda_3$. The right represents isotropic diffusion, in which $\lambda_1 \approx \lambda_2 \approx \lambda_3$.

A convenient way to visualise the information contained in the diffusion tensor is by considering the effective diffusion ellipsoid (Basser, 1995). This is an ellipsoid with axes of length $\sqrt{\lambda_i}$, oriented along the directions given by ϵ_i . In a physical sense, it is the surface encompassing the locations to which a spin is equally likely to diffuse within a given time period τ . Two such ellipsoids are illustrated in figure 2.5. The prolate effective diffusion ellipsoid describes anisotropic diffusion, whereas the spherical effective diffusion ellipsoid describes isotropic diffusion.

Figure 2.6 shows an example of diffusion tensors in real MRI data. The enlarged segment covers three different tissue types: the white matter of the corpus callosum, cerebrospinal fluid in the ventricles, and grey matter in the thalamus. The ellipsoids in the white matter are highly anisotropic, and all oriented with their principal diffusion direction aligned to the orientation of the white matter fibres. The ellipsoids in the cerebrospinal fluid (CSF) are large and isotropic, reflecting the free motion of the water molecules there. The ellipsoids in the grey matter are still

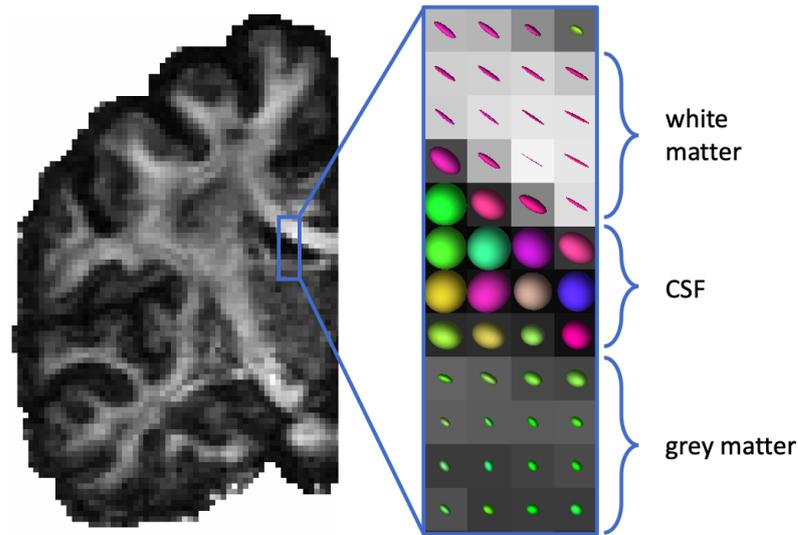


Figure 2.6: Examples of diffusion tensors calculated from real MRI data. The effective diffusion ellipsoids are overlaid on a fractional anisotropy map in the enlarged segment. The ellipsoids are colour coded according to their principal diffusion direction (red: medial-lateral, green: anterior-posterior, blue: superior-inferior).

fairly isotropic, but much smaller in size, as the cell bodies restrict the magnitude of the diffusion in this region.

Scalar Invariants of the Diffusion Tensor

From the eigenvalues of the diffusion tensor several parameters can be derived that describe various aspects of the diffusion profile. These are often used to characterise microstructural properties of the underlying tissue. One of the most commonly used is fractional anisotropy (FA), which is the normalised variance of the tensor eigenvalues (Pierpaoli and Basser, 1996) (where $\hat{\lambda}$ is the mean of the three eigenvalues):

$$FA = \sqrt{\frac{3}{2}} \frac{\sqrt{(\lambda_1 - \hat{\lambda})^2 + (\lambda_2 - \hat{\lambda})^2 + (\lambda_3 - \hat{\lambda})^2}}{\sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}} \quad (2.11)$$

FA reflects the level of anisotropy in the diffusion profile of a given voxel. If the diffusion in a given voxel is strongly anisotropic, the FA value will be close to one, whereas for a voxel in which the diffusion is isotropic, such as in the CSF, the FA value will be close to zero. The FA map in figure 2.7 shows the higher values in the

white matter, and low values in the grey matter and the ventricles. FA is a useful metric, as it reflects axonal ordering (Pierpaoli and Basser, 1996), axonal density (Takahashi et al., 2002) and the degree of myelination in a voxel (S.-K. Song et al., 2002). It has also been shown to be sensitive to microstructural changes during development (Cohen et al., 2016; Dubois et al., 2008b; Geng et al., 2012).

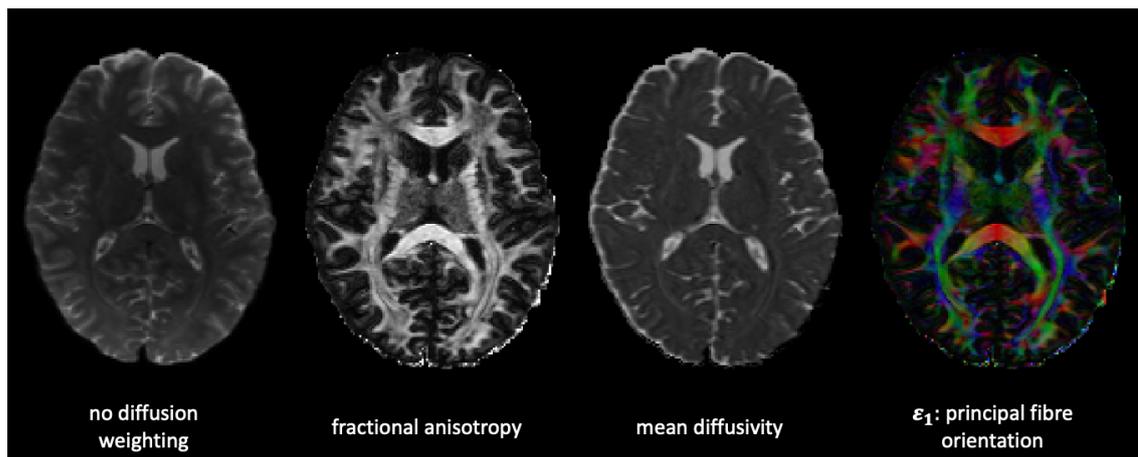


Figure 2.7: Images derived from DTI in an axial slice from a single adult subject. From left to right: no diffusion weighting, fractional anisotropy map, mean diffusivity map, FA colour coded according to the principal direction of the diffusion tensor (red: medial-lateral, green: anterior-posterior, blue: superior-inferior).

The mean diffusivity (MD) (also called the apparent diffusion coefficient) is the average of the eigenvalues of the diffusion tensor, a scalar measure of the bulk diffusivity in a voxel (Basser and Pierpaoli, 1996):

$$MD = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3} \quad (2.12)$$

MD reflects the amount of hindrance to diffusion within a voxel. As shown in figure 2.7, MD is highest in the CSF, where diffusion is less restricted, but does not show as much contrast between grey and white matter as FA.

Scalar invariants of the diffusion tensor are convenient to use because they are rotationally invariant and give insight into the intrinsic properties of the tissue (Pierpaoli and Basser, 1996). However, they are not specific measures of microstructure so care must be taken not to over-interpret their relationship to myelination or “fibre

integrity” (Jones et al., 2013).

2.3.5 Modelling Crossing Fibres - Beyond the Diffusion Tensor

Although the tensor model is a powerful and simple way to characterise tissue structure, it fails in regions with crossing fibres, since it only describes one major orientation for each voxel. More than 90% of voxels in white matter are estimated to contain crossing fibre populations (Jeurissen et al., 2013), so this drawback seriously limits the accuracy of the diffusion tensor for mapping fibre pathways, leading to false negatives in tractography and inaccurate fibre reconstructions. More complex models are required to characterise the complex fibre architecture in the brain. These are more computationally intensive than DTI and require data to be acquired over many gradient directions and often at different b-values, in order to obtain the necessary angular resolution. These protocols can be classified as High Angular Resolution Diffusion Imaging (HARDI) protocols (Tuch et al., 2002).

There are two broad approaches for modelling fibre orientations: non-parametric methods and parametric methods (D. C. Alexander, 2006). Non-parametric approaches, such as the diffusion propagator or spherical deconvolution, estimate the fibre orientations directly from the data, exploiting the Fourier relationship between the diffusion propagator and the measured signal. Parametric approaches use an explicit biophysical model for the fibre orientation distribution function (fODF) and fit the model parameters from the data. These are discussed in more detail below.

Diffusion Propagator Techniques and Spherical Deconvolution

As shown in equation 2.7, there is a Fourier relationship between the diffusion propagator and the diffusion-weighted MR signal. Therefore, the propagator $P(\mathbf{r}, \tau)$ can be estimated directly via an inverse Fourier transform of the diffusion weighted signal

$S(\mathbf{q}, \tau)$. Once the diffusion propagator has been estimated, the diffusion orientation distribution function (dODF) can be calculated as its radial projection $\Psi(\mathbf{u}, \tau)$:

$$\Psi(\mathbf{u}, \tau) = \int_0^\infty P(r\mathbf{u}, \tau)r^2 dr \quad (2.13)$$

Where the displacement vector \mathbf{r} is separated into a unit direction vector \mathbf{u} and its length r . The orientations of the fibre populations are assumed to coincide with the peaks of the dODF.

Diffusion Spectrum Imaging (DSI) was the first method presented to estimate $P(r\mathbf{u}, \tau)$ in this fashion, by sampling q-space over a Cartesian grid (Wedeen et al., 2005). However, this comprehensive sampling leads to long acquisition times, which has limited the use of DSI (Dell’Acqua and Tournier, 2019). Similarly to DSI, Q-ball imaging aims to estimate the dODF directly from the dMRI data (Tuch, 2004). Q-ball imaging uses a more time-efficient HARDI protocol, in which q-space is sampled over a sphere (i.e. with a constant b-value). This means that the Funk-Radon transform can be used to recover the dODF, rather than the full Fourier Transform. This is much simpler to implement but neglects the r^2 term in the dODF, which reduces its angular resolution relative to DSI (Aganj et al., 2010; Dell’Acqua and Tournier, 2019). Despite this, Q-ball imaging has been widely adopted and many extensions have been proposed, including using a spherical harmonic basis to improve reconstruction times (Descoteaux et al., 2007; Hess et al., 2006). Other non-parametric methods include persistent angular structure (PAS) (Jansons and D. C. Alexander, 2003) and diffusion orientation transform (Özarslan et al., 2006).

Diffusion ODF methods assume that directions of maximum diffusivity coincide with fibre orientations. If one is interested in orientations, approaches exist that estimate the fibre ODF directly. These are based on a spherical deconvolution approach, in which the measured diffusion signal S is the convolution of the fibre orientation distribution function F and the fibre response function R (Anderson,

2005; Tournier et al., 2004; Tournier et al., 2007):

$$S(\theta, \phi) = F(\theta, \phi) \otimes R(\theta) \quad (2.14)$$

The diffusion signal is the summation of the signal contributions of different fibre populations, as illustrated in figure 2.8. The fibre response function can be based on an axially symmetric tensor model (Dell’Acqua et al., 2007), a perfectly symmetric tensor (Behrens et al., 2007), or drawn from a canonical fibre distribution from the data itself (Tournier et al., 2004; Tax et al., 2014).

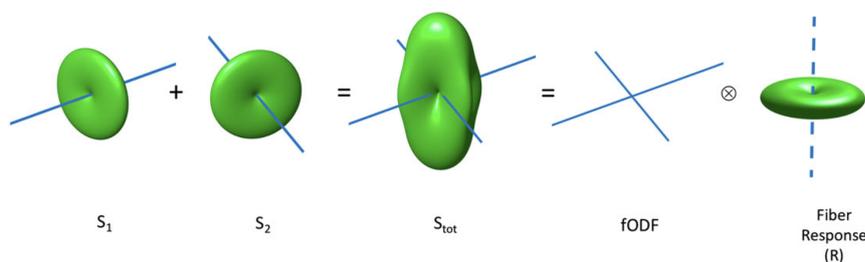


Figure 2.8: Signals from two fibre populations (S_1 and S_2) can be summed to produce the measured diffusion signal S_{tot} , which is also the convolution of the fibre response function with the fODF. Reproduced from (Dell’Acqua and Tournier, 2019).

The fODF is then obtained by deconvolving the measured signal with the fibre response function. The spherical harmonic basis set is usually used to reconstruct the signal, which is Fourier series over the surface of the sphere (Anderson, 2005; Tournier et al., 2004). The deconvolution can be solved using a constrained spherical deconvolution (CSD) approach (Tournier et al., 2007) or the Richardson-Lucy algorithm (Dell’Acqua et al., 2010). Multi-tissue multi-shell CSD is an extension of the original CSD approach for multi-shell data (Jeurissen et al., 2014). It exploits the varying dependencies of the different tissue types on the b-value, which allows separate compartments to be modelled for each tissue type, each with a different response function. This improves the precision of ODF estimates and reduces spurious peaks in the fODF for voxels containing grey matter or CSF (Jeurissen et al., 2014).

More sophisticated approaches use information from neighbouring voxels to improve the accuracy of the estimates, by informing asymmetric fODFs (Bastiani et al., 2017; Reisert et al., 2012) and fibre continuity (Reisert and Kiselev, 2011). Although these methods are not yet widely used, they show promising avenues for increasing the accuracy of fODFs, by helping to differentiate between ambiguous fibre geometries (Dell’Acqua and Tournier, 2019).

Non-parametric approaches have the advantage of being model free. By calculating the diffusion ODF or the fibre ODF, they are aiming to capture the signal shape, rather than making any assumptions about how the underlying microstructure influences or generates the diffusion-weighted signal.

Parametric Methods

This set of methods are sometimes referred to as mixture or multi-compartment models, because the signal is modelled as a mixture of separate compartments (D. C. Alexander et al., 2019). These compartments can correspond to distinct fibre populations (sometimes referred to as fixels (R. E. Smith et al., 2012)), or to isotropic diffusion in the CSF and grey matter. The model parameters are then chosen that best describe the measured data, which provide estimates of the underlying tissue structure. Typically, these parameters include the orientation of the fibre populations (described by spherical angles θ and ϕ), and their volume fraction f .

A natural extension of the diffusion tensor model is the multi-tensor model, first demonstrated by Tuch et al. (2002). This model is based on three assumptions: 1) the inhomogeneity in the diffusion signal is made up of a finite number of discrete compartments; 2) the exchange between the compartments is negligible; and 3) the diffusion within each compartment is Gaussian (Tuch et al., 2002). Therefore, the diffusion signal can be expressed as a finite sum of the signals from each of the i Gaussian compartments, each comprising a volume fraction f_i of the voxel:

$$S(\mathbf{g}) = S(0) \left[\sum_{i=1}^N f_i \exp(-b\mathbf{g}^T \mathbf{D}_i \mathbf{g}) \right] \quad (2.15)$$

Where $S(\mathbf{g})$ is the signal corresponding to the diffusion-weighted gradient with direction \mathbf{g} and b-value b , and $S(0)$ is the signal in the absence of any diffusion-encoding gradients. This model is not identifiable in general, at least from data obtained in standard acquisitions (Scherrer et al., 2016), therefore constrained versions of it have been proposed. The ball and stick model is a multi-tensor model, except that the tensors modelling the fibres have their second and third eigenvalues reduced to zero (i.e. perfectly anisotropic “sticks”), while partial volume is represented as a spherical tensor (i.e. a perfectly anisotropic “ball”) (Behrens et al., 2003). The advantage of reducing the model parameters in this way is that more fibre populations can robustly be modelled from the data. Given N anisotropic compartments, or “sticks”, each with volume fraction f_i and orientation given by \mathbf{v}_i , the signal equation for the ball and stick model is:

$$S(\mathbf{g}) = S(0) \left[\left(1 - \sum_{i=1}^N f_i \right) \exp(-bd) + \sum_{i=1}^N f_i \exp(-bd(\mathbf{g} \cdot \mathbf{v}_i)^2) \right] \quad (2.16)$$

The multi-tensor and ball and stick models belong to the family of methods that parameterise the orientation information as a discrete set of fibre populations (Assaf and Basser, 2005; Hosey et al., 2008; Scherrer et al., 2016). To provide more flexibility in the fODF, some methods model a discrete set of fibre populations with dispersion associated with each one, which leads to a continuous fODF (Jian et al., 2007; Kaden et al., 2007; Sotiropoulos et al., 2012; Zhang et al., 2012). Other methods, more similar to the non-parametric approaches described above, model the fODF as a continuous distribution (Jian and Vemuri, 2007; Dell’Acqua et al., 2007; Tournier et al., 2007), parameterised by a fixed number of coefficients.

A number of extensions of the ball and stick model are used in this thesis,

they are therefore reviewed here briefly. The method by Jbabdi et al. (2012) was developed specifically for multi-shell data. The use of multi-shell data allows us to combine information from the high SNR at low b-values with the improved angular contrast of high b-values. However, the original ball and stick model overfits to multi-shell data, in that extra fibre orientations are fit to account for the non-monoexponential signal decay at higher b-values. The solution proposed by Jbabdi et al. (2012) uses a Gamma distribution of diffusivities to model a continuous distribution of diffusion coefficients. The Gamma distribution has a shape parameter α and a scale parameter $\frac{1}{\beta}$. This model gives a signal for a diffusion encoding gradient with direction \mathbf{g} and a b-value b_k :

$$S(b_k, \mathbf{g}) = S(0) \left[\left(1 - \sum_{i=1}^N f_i \right) \left(\frac{\beta}{\beta + b_k} \right)^\alpha + \sum_{i=1}^N f_i \left(\frac{\beta}{\beta + b_k (\mathbf{g} \cdot \mathbf{v}_i)^2} \right)^\alpha \right] \quad (2.17)$$

Using this model on multi-shell data was shown to reduce the amount of spurious fibre estimates, particularly at the interface between white and grey matter, while the only adding one additional scalar parameter to the regular ball and stick model (Jbabdi et al., 2012).

Another extension of the ball and stick model especially suited to dMRI data from the neonatal brain is presented in (Sotiropoulos et al., 2016). This uses the Gamma distribution approach from (Jbabdi et al., 2012), to characterise the isotropic diffusion. Instead of a stick compartment, the fibre response function is the signal attenuation from an axially symmetric tensor, or “zeppelin”, which better characterises the lower anisotropy in the developing brain (Bastiani et al., 2019; Panagiotaki et al., 2012). The tensor has eigenvalues $\lambda_1 \gg \lambda_2, \lambda_2 = \lambda_3$, and a principal diffusion direction \mathbf{v} . However, due to the variety of maturational states in the different tracts in this age group (covered further in chapter 3), a single fibre response function is not appropriate. Instead, the tensor is allowed to vary in

shape between voxels. Specifically, λ_R , the ratio of the major and minor axes of the tensor, is a model parameter that is estimated on a voxel-by-voxel basis. The use of voxel-specific fibre response kernels facilitates the estimation of fibre crossings in white matter regions with low anisotropy. Combining the Gamma distribution of isotropic diffusivities with the new convolution kernel gives the following expression for the signal, (where λ_m is the mean of the tensor eigenvalues):

$$S(b_k, \mathbf{g}) = S(0) \left[\left(1 - \sum_{i=1}^N f_i \right) \left(\frac{\beta}{\beta + b_k} \right)^\alpha + \sum_{i=1}^N f_i \exp \left(-b_k \frac{3\lambda_m}{2\lambda_R + 1} \left((1 - \lambda_R) (\mathbf{g} \cdot \mathbf{v}_i)^2 + \lambda_R \right) \right) \right] \quad (2.18)$$

We fit these two models to a neonatal dataset from the developing Human Connectome Project (figure 2.9). Focussing on the centrum semiovale, we can see that the stick response kernel is unable to resolve the crossing fibre configurations, whereas the zeppelin response kernel supports the multiple fibre populations required to describe this complex fibre architecture. Furthermore, the average anisotropic volume fraction estimated by the stick model is much lower than for the zeppelin model, which indicates that the stick model is failing to capture the diffusion contrast in these neonatal data.

As with all models, both deterministic and stochastic inference approaches can be used to estimate the model parameters. Bayesian inference (and Markov-Chain Monte-Carlo sampling) is used to estimate the parameters of this model (Behrens et al., 2003). This asks what parameters best fit the model, given the evidence from the data, and some prior knowledge about the parameters. For example, in the above model, λ_R has a Gaussian prior, whose mean and standard deviation are estimated from the data. Bayesian modelling provides a probability density function, called the posterior distribution, associated with each of the estimated parameters, which

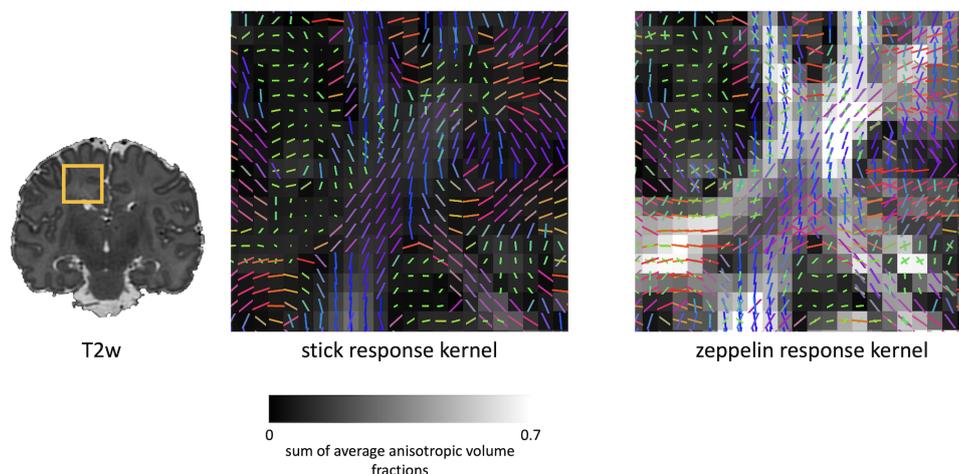


Figure 2.9: Comparison of fibre orientations estimated using the stick and zeppelin (axially symmetric tensor) response kernels, in the same neonatal dataset. The zoomed-in portion is indicated in the T2 weighted image on the left. Fibre orientations are coloured according to their orientations (red: medial-lateral, green: anterior-posterior, blue: superior-inferior).

helps to quantify their uncertainty. It also allows online “model selection” using Automatic Relevance Determination (ARD) priors (Behrens et al., 2007). ARD are shrinkage priors applied to volume fractions f_i (where $i \geq 2$) that a-priori penalise the existence of more than one fibre compartment. Multiple fibres are included in the model (i.e. fibre crossings), only when these are strongly supported by the data.

Overall, parametric methods have the advantage over non-parametric methods that they can provide a more concise characterisation of the fODF via parameters that provide a phenomenological representation of the signal. Furthermore, they can directly model the fODF peaks, which are relevant for tractography applications.

2.3.6 Tractography

Tractography describes the process of using information about local fibre orientations (i.e. fODFs) derived from dMRI to map white matter connections in the brain (Mori et al., 1999). The estimated fibre trajectories are called “streamlines”, which are tracked through the vector field formed by the orientation estimates. Mathematically, a streamline $\mathbf{r}(s)$ is a curve of arc length s , tangent to the orientation field

at each point \mathbf{v} , as illustrated in figure 2.10. This can be expressed mathematically as follows:

$$\frac{d\mathbf{r}(s)}{ds} = \mathbf{v}[\mathbf{r}(s)] \quad (2.19)$$

This is an ordinary differential equation that can be solved by integration, with streamlines starting at a seed point, $\mathbf{r}(0) = \mathbf{r}_0$:

$$\mathbf{r}(s) = \int_0^s \mathbf{v}[\mathbf{r}(s)] ds \quad (2.20)$$

The integration is carried out numerically. The simplest method is Euler integration, in which the following procedure is iterated for each step \mathbf{r}_i along the path, separated by step size δ (Conturo et al., 1999):

$$\mathbf{r}_{i+1} = \mathbf{r}_i + \delta \mathbf{v}(\mathbf{r}_i) \quad (2.21)$$

In practise, higher order methods such as Runge-Kutta are more commonly used as they are less susceptible to integration errors (Basser et al., 2000). The step size δ in these algorithms is typically smaller than the voxel size, so interpolation is needed to ensure smoothly varying estimates of fibre orientation. Most algorithms use trilinear interpolation, which reduces interpolation errors compared to nearest-neighbour interpolation (Conturo et al., 1999). Streamlines are propagated through the vector field of orientation estimates until a stopping criterion is reached, such as low amplitude in the fODF or high streamline curvature (Jeurissen et al., 2019). These criteria are based on the anatomical plausibility of the fibres and aim to reduce false positives.

There are many different approaches to tractography (Jeurissen et al., 2019). We briefly overview here the difference between the two most commonly used types, deterministic and probabilistic.

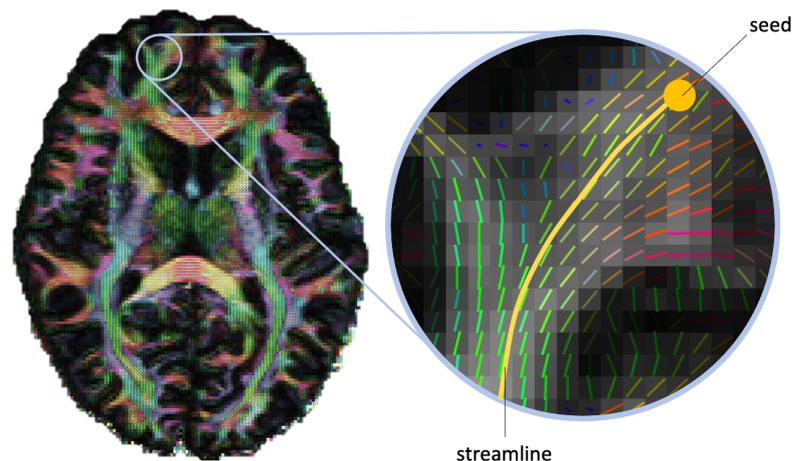


Figure 2.10: *Diagram showing how streamlines are propagated through the vector field of fibre orientation estimates (colour coded here according to their direction). The streamline begins at the seed point and continues tangentially to the fibre orientations.*

Deterministic vs. Probabilistic

Tractography estimates the spatial distribution of paths least-hindered to diffusion from the seed point (Behrens et al., 2014). The type of estimation (deterministic or stochastic) is the difference between the two main types of tractography. Deterministic tracking returns a point estimate of this spatial distribution. Probabilistic returns a spatial histogram of the path distribution (i.e. most likely paths and associated uncertainty around them). An example of the difference between the deterministic and probabilistic approach is illustrated in figure 2.11, where both have been applied in the same dataset with the same seed. Both outputs cover spatially similar regions of the brain, however the deterministic output is a binary map, whereas the probabilistic output is a probability map. If a deterministic approach is taken, the resultant streamline map is a binary descriptor of the presence or absence of a diffusion path, given the data.

Probabilistic tractography, on the other hand, provides a probability distribution on the most probable location of the underlying fibre bundle (Behrens et al., 2007). This provides a quantitative estimate of the uncertainty in the estimated trajectory. Many streamlines are propagated from each seed point to build a distri-

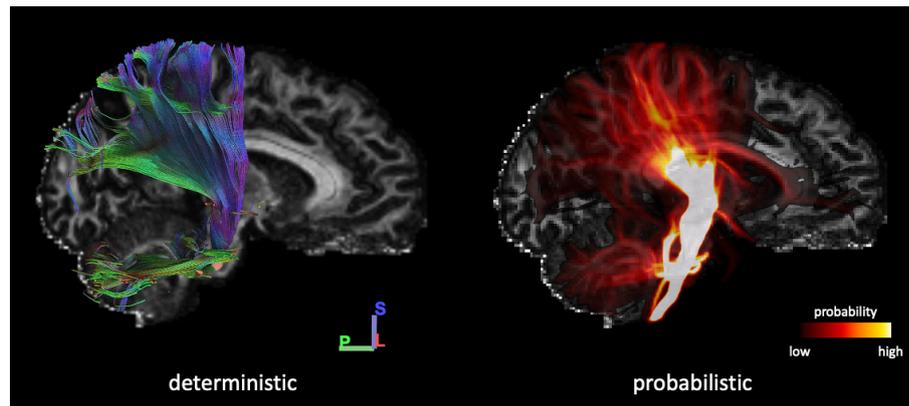


Figure 2.11: Results from deterministic (Yeh et al., 2013) and probabilistic (Behrens et al., 2007) tractography algorithms. Streamlines were seeded from the same seed ROI in the brainstem in both cases, in a single adult subject. Probabilistic results are displayed as a maximum intensity projection.

bution on the trajectory of the tract. There are two ways to generate and interpret these distributions. Firstly, the local trajectory of the individual streamlines can be drawn from a new ODF called the uncertainty orientation distribution function (uODF). This characterises the uncertainty in the fibre orientation due to noise and model errors and it is estimated by stochastic inference measures, such as Bayesian inference of the ball and stick model described above. Propagating streamlines according to the uODF yields a final distribution that indicates the measurement uncertainty in the fibre trajectory (Behrens et al., 2007). Alternatively, the pathways can be drawn from the fODF, (Tournier et al., 2012), in which case the results are assumed to reflect the underlying dispersion of the white matter fibres themselves, under the strong assumption that the width of the fODF truly encodes fibre dispersion. We should point out that the width of the fODF is typically much larger than the respective uODF, so results from the latter approach yield considerably higher dispersion (Jeurissen et al., 2011).

Using Tractography to Map Connections in the Brain

There are different ways to use tractography algorithms to probe structural connectivity, depending on the application. To map specific fibre bundles, it is common

to use an “in-vivo dissection” approach, in which regions of interest (ROIs) are delineated to specify the streamlines that are included (Catani et al., 2002). This can either be an ROI that encompasses the whole tract, or seed and target ROIs that specify the route that the streamlines must take to be included in the tract, combined with rules to minimise false positives and constrain the tractography based on anatomical plausibility (Warrington et al., 2020). Examples of the applications of this approach include surgery planning (Romano et al., 2009) or to examine variation in tract microstructure or geometry across different subjects (Dubois et al., 2008b; Thiebaut de Schotten et al., 2011b). The advantage of this approach is that bundles can be accurately reproduced by incorporating prior anatomical knowledge into the protocols. The protocols can be determined on a subject-by-subject basis, to ensure good alignment to individual data (Conturo et al., 1999). This however is prohibitively time-consuming, especially for large data-sets such as the Human Connectome Projects (Hughes et al., 2017; Van Essen et al., 2013b; Howell et al., 2019). Alternatively, protocols can be defined in a template space and then registered to the individual data to generate subject-level results (Catani and Schotten, 2008; Yendiki, 2011; Warrington et al., 2020). This has been shown to generate results that are robust across subjects, while still respecting the differences in individual anatomy.

An alternative approach is to seed streamlines across the brain to sample the whole connectome, the comprehensive matrix of brain connections. The resultant connectivity matrix can be used to examine global properties, through for example graph-based analyses (Bullmore and Sporns, 2009; Bassett and Bullmore, 2006), or as a basis for data-driven methods, which aim to identify white matter bundles from whole brain connectivity data in an unsupervised fashion. One way to do this is by using clustering techniques (Garyfallidis et al., 2012), such as hierarchical clustering (Siless et al., 2018; Ding et al., 2003) and spectral clustering (O’Donnell and Westin, 2007), to identify fibre bundles from the whole-brain connectivity data.

Matrix decomposition methods can also be applied to the connectivity matrix to decompose it into a set of data-driven components. This is a commonly used technique for the analysis of fMRI data, where Independent Component Analysis (ICA) is typically applied to the matrix of functional timecourses to yield a set of functional networks in grey matter (Mckeown et al., 1998). In the context of structural connectivity, ICA has been applied to tractography data seeded from the thalamus (O’Muircheartaigh et al., 2011) and throughout the whole brain (L. Wu et al., 2015) to map white matter bundles and their corresponding grey matter networks (O’Muircheartaigh and Jbabdi, 2017; Mars et al., 2019). These networks have shown to be sensitive to changes in structural connectivity brought about by disease (L. Wu et al., 2015), and robust across subject groups (O’Muircheartaigh and Jbabdi, 2017).

2.3.7 Summary

This chapter provides an overview of the concepts of structural and functional connectivity in the brain and the tools used to measure them, with particular focus on diffusion MRI, the only method available to measure white matter connections non-invasively and in vivo. The basic physics underlying dMRI has been described, followed by the methods used to estimate tissue microstructure and fibre orientations from the data. In particular, a parametric method for fibre orientation modelling has been described that is especially suited to this age-group (see section 2.3.5). This will be used in the work presented in Part III (Original Research), because it is able to estimate fibre crossings even in areas of low anisotropy, which are prevalent in the neonatal brain. We have also given an overview of tractography, and some of the different ways that tractography algorithms can be used map white matter connections in the brain. The focus of the rest of the thesis will be the application of these methods to the neonatal brain. The next chapter provides an overview of tractography studies in neonates, and subsequent chapters describe work done to

address some of the existing problems in this context.

Chapter 3

Imaging the Early Development of Connectivity

Contents

3.1	The Early Development of Structural Connectivity . . .	38
3.2	Using dMRI to Map Early Brain Development In Vivo	39
3.2.1	Delineating Tracts	40
3.2.2	Estimating Tract Microstructure	42
3.2.3	Whole Brain Connectivity Analysis	44
3.2.4	Impact of External Factors on Structural Connectivity . .	46
3.3	The Challenges of Imaging Early Brain Development In Vivo	47
3.4	Big Data Projects	48
3.4.1	The Developing Human Connectome Project	49
3.4.2	The Baby Connectome Project	50
3.4.3	Aims of this Work	51

In the previous chapter, the basic principles of MRI and MRI-based mapping of brain connectivity were presented. We described how diffusion MRI provides unique opportunities to study brain organisation non-invasively and *in vivo*. In this chapter, existing literature is presented on early white matter development and how dMRI can be used to probe the growth and maturation of structural connections in the brain during this important period.

3.1 The Early Development of Structural Connectivity - Insights from Ex Vivo Studies

The maturation of the human brain comprises many processes that occur asynchronously from shortly after conception until well into early adulthood (figure 3.1). Post-mortem studies have shown that certain white matter bundles have formed as early as 12 weeks post-menstrual age (PMA) (Radoš et al., 2006). The limbic fibres are the first to be established, followed by the projection and callosal fibres during the early preterm period (26 - 30 weeks PMA) (Kostović et al., 2019). During the late preterm period (31 - 36 weeks PMA), the longer commissural and association fibres develop, so that by full term all the long range fibre bundles are in place (Kostović et al., 2019). In the months after birth there is an overproduction of neurons and synapses, followed by an elimination of redundant connections by cellular apoptosis and axonal pruning in response to early life experiences (Huttenlocher and Dabholkar, 1997; Kostović and Jovanov-Milošević, 2006). Myelination occurs from the second trimester through to adolescence, improving the conduction of the nerve impulse for more efficient information transfer. At 20 weeks PMA microstructural myelin is present in the central limbic regions, and by 40 weeks there is mature myelin in the cerebellum and external capsule (Kinney et al., 1988; Brody et al., 1987). The myelination process is asynchronous across different regions, following a central-to-peripheral and posterior-to-anterior trajectory. It also reflects the

neurodevelopmental hierarchy. For example, myelination occurs earlier and more rapidly in the sensory regions than in the motor regions (Dubois et al., 2014). Figure 3.1 illustrates the timeline of these complex maturational processes that occur during brain development.

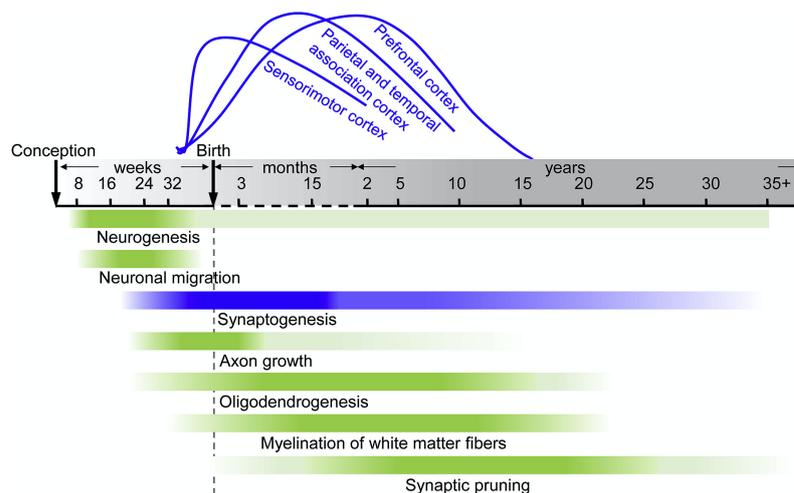


Figure 3.1: *Timeline of the different maturational processes in the human brain. Time axis is in post-conceptual weeks (before birth), postnatal months (until 24 months), and postnatal years (after 2 years). The color intensity in each bar corresponds to the rate of developmental changes. The spatial progression across brain regions is illustrated using synaptogenesis (blue bar) as an example, illustrated by the blue curves above the time axis. Reproduced with permission from (Ouyang et al., 2019).*

3.2 Using dMRI to Map Early Brain Development In Vivo

Although post-mortem studies provide an accurate and valuable insight into the early development of structural connectivity, in vivo imaging is required to answer questions about longitudinal development, the impacts of early life factors on brain development, and the diagnosis of neurodevelopmental disorders. Diffusion MRI has been used in previous studies to map brain connections and organisation in neonates, from delineating white matter bundles and their microstructure to mapping whole brain connectomes. A representative set of related studies is overviewed here.

3.2.1 Delineating Tracts

A simple approach for tract delineation is to define a white matter atlas, from which template tract ROIs can be registered to individual subjects' data. A multi-contrast MRI atlas for neonates was developed by Oishi et al. (2011), based on the JHU-MNI adult atlas (Oishi et al., 2008). The atlas includes T1-weighted, T2-weighted and DTI contrasts, and was constructed using data from 25 subjects scanned between 38 and 41 weeks PMA. The segmentation includes 38 cerebral white matter structures that were manually delineated. A downside of this approach is that it only assigns one label to each voxel, when in reality there can be several different fibre bundles present in a given voxel (Jeurissen et al., 2013). To deal with this issue, a multi-label white matter atlas was developed by Ratnarajah and Qiu (2014). Generated from a training set of 20 subjects, each voxel is classified using a multi-label k-nearest neighbour algorithm in Riemannian diffusion tensor space. However, these atlas-based approaches do not provide in principle the same accuracy as tractography, as they do not model the trajectory of white matter pathways based on the subject-specific directional information from dMRI.

A number of tractography studies of the neonatal brain exist. Due to the challenges associated with scanning neonates (see also next section), most of these tractography studies have used DTI, due to their shorter acquisition requirements compared to HARDI protocols (Hüppi and Dubois, 2006; Ouyang et al., 2019). However, as described in the previous chapter, the diffusion tensor model can lead to inaccurate tractography because it cannot identify multiple fibre populations within a voxel. Recent developments, such as multiband imaging, have enabled HARDI data to be acquired in shorter time frames, so more recent studies have been able to take advantage of the better angular resolution afforded by HARDI data for fibre orientation modelling (Anblagan et al., 2015; Bastiani et al., 2019; Zöllei et al., 2019).

Many neonatal tractography studies have relied on manual segmentation, either

of the whole bundle (Partridge et al., 2004; Huang et al., 2006) or manual placement of ROIs to constrain tractography on a subject-by-subject basis (Dubois et al., 2008b; Liu et al., 2010; Kaur et al., 2014; Cohen et al., 2016; Akazawa et al., 2016). This ensures an accurate characterisation of individual anatomy, but is labour intensive, especially for large numbers of subjects or many white matter tracts. Some alternative approaches to manual segmentation are described below, which provide automated tractography in neonatal subjects.

The University of North Carolina-Utah National Alliance for Medical Imaging Computing (UNC-Utah NA-MIC) framework provides an end-to-end toolbox for tractography analysis of neonates and infants (Verde et al., 2014). The framework generates a study-specific atlas, in which users interactively define protocols for deterministic tractography. Interactive fibre cleaning steps allow the user to refine the results. This has been used in many studies (Geng et al., 2012; Gupta et al., 2015; S. J. Lee et al., 2017; Rasmussen et al., 2017; Swanson et al., 2018), but does not provide a canonical set of ROIs for the tractography step so these may be defined differently across studies.

A variety of methods have been used to map white matter bundles using probabilistic tractography. Anblagan et al. (2015) used probabilistic neighbourhood tractography, a method that uses reference tracts to automatically delineate seed points in individual subjects' data, and then uses maximum likelihood to retain streamlines that are closely matched to the reference. A standardised set of neonatal tractography protocols were defined by Bastiani et al. (2019), using the tractography recipes from the AutoPtx toolbox for probabilistic tractography (De Groot et al., 2013). Another automated probabilistic tractography tool (TRACULInA) was developed by Zöllei et al. (2019). This method uses global tractography, and tracts are segmented based on a Bayesian approach, using their anatomical neighbourhood as a prior.

3.2.2 Estimating Tract Microstructure

At a global level, the neonatal brain has a higher water content and lower myelination than the adult brain, leading to lower anisotropy in the diffusion signal and higher diffusivity overall (Dubois et al., 2014). These properties can be seen in figure 3.2, which also illustrates the rapid increases in FA and reduction in MD during this period. In particular, FA is shown to increase in the internal capsule and the corpus callosum, while MD decreases are especially prominent in the frontal white matter.

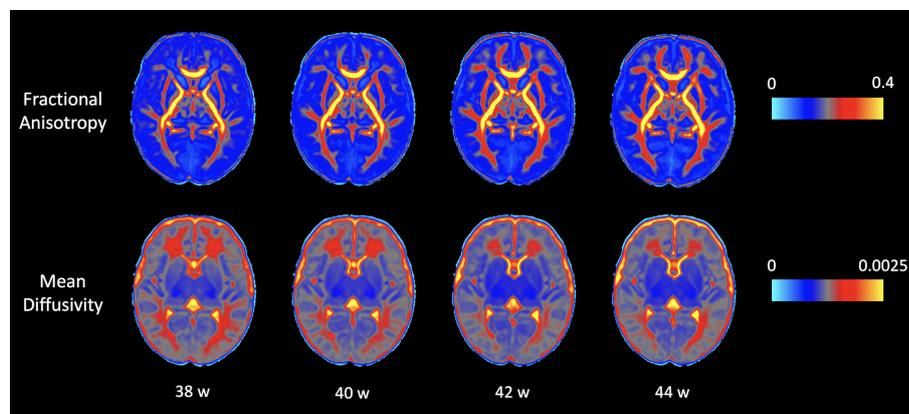


Figure 3.2: Average FA and MD maps illustrating the microstructural changes across the neonatal period. Ages are given in postmenstrual weeks at scan. Each map is the average of 20 subjects. Data were obtained from the public data release of the developing Human Connectome Project.

Many studies have used tracts as ROIs and computed average DTI metrics within each bundle as measures of tract maturity. These studies have supported the findings from histology, describing maturation occurring over a central-to-peripheral and a posterior-to-anterior trajectory (Dubois et al., 2014). Various models have been developed to characterise the maturational sequence of the different fibre bundles. For example, Dubois et al. (2008b) defined four relative stages of maturation, by comparing both the rates of change and the absolute values of DTI metrics in each tract against the average across all tracts, and between infants and adults. In this model, the anterior limb of the internal capsule and the cingulum were found to be the least mature, while the cortico-spinal tract was the most mature bundle. Sadeghi et al. (2017) used nonlinear mixed effects modelling to look at the

maturation of white matter bundles during the first two years of life, and to assess the maturational differences between twins and singletons. Changes in DTI metrics were described by a Gompertz function, which parameterises the asymptote, delay and speed of the maturation. Only small differences were found between twins and singletons, in the delay parameter of axial diffusivity in the anterior limb of the internal capsule and anterior corona radiata.

Tract-based spatial statistics (TBSS) is another method for analysing microstructural measures across subject groups, in which data is projected onto a mean FA tract skeleton (S. M. Smith et al., 2006). This method has been used to study microstructural changes in neonates using specific neonatal FA templates (Anjari et al., 2007; Ball et al., 2010; Counsell et al., 2008). However, a drawback of TBSS is a lack of anatomical specificity. Tract-specific analysis is a similar approach that instead uses individual tract skeletons, which was used by Pecheva et al. (2017) to study microstructural changes in tracts with age at scan.

More sophisticated microstructure models can be used to assess white matter maturation, such as the Neurite Orientation Dispersion and Density Imaging (NODDI) model, which models the density of neurites (dendrites and axons) and their orientation dispersion, disentangling these two factors that contribute to FA (Zhang et al., 2012; Dean et al., 2017). For example, recent work by Batalle et al. (2019) uses the NODDI model to investigate changes in cortical microstructure such as dendritic arborisation and neurite growth. Other work has used this method to measure the impact of clinical encephalopathy on neonatal white matter (Kansagra et al., 2016).

Lateralisation of tract volume and microstructural measures has also been explored. The arcuate fasciculus, which is associated with language processing, is left lateralised in neonates (Dubois et al., 2009; Liu et al., 2010), as it is in adults, although this finding was not replicated in a separate study (J. W. Song et al., 2015). Some studies have also indicated left lateralisation in the maturation and volume

of the cortico-spinal tract (Dubois et al., 2009; Liu et al., 2010), the inferior longitudinal fasciculus (J. W. Song et al., 2015; Cohen et al., 2016), and the cingulum (J. W. Song et al., 2015; Cohen et al., 2016).

3.2.3 Whole Brain Connectivity Analysis

More recently, improvements in tractography methods and data acquisition have enabled the use of dMRI to investigate whole-brain connectivity in neonates. Graph based analysis can be applied to structural connectivity data to understand the network properties of brain organisation (Bullmore and Sporns, 2009). This requires the definition of nodes and edges, with an edge representing the connection between a pair of nodes. Nodes are often defined according to parcellations and edges represent connectivity measures between them. The number of streamlines connecting two nodes is a common metric for the strength of a structural connection, and can be normalised with respect to the total number of streamlines or the node size (Hagmann et al., 2008; Sotiropoulos and Zalesky, 2019). Microstructural measures can also be used as a proxy for connection efficiency (Heuvel et al., 2015; Ball et al., 2015).

Small world topology is a commonly used measure, as it facilitates both integrated and segregated processing. It is characterised by high levels of local clustering, with short paths that globally link all nodes (Bassett and Bullmore, 2006). Small world topology is present at birth, even in preterm subjects, and the small world property increases with age (Yap et al., 2011; Tymofiyeva et al., 2013; Ratnarajah et al., 2013; Brown et al., 2014; Heuvel and Hulshoff Pol, 2010). Local efficiency of the network structure increases with age, but global efficiency remains similar from birth to two years (Yap et al., 2011). The left hemisphere has higher structural efficiency than the right (Ratnarajah et al., 2013), which aligns with findings from functional data that indicate left lateralisation of language networks emerges during the first year (Emerson et al., 2016).

Rich club organisation is also observed in neonatal brain networks. This consists of a backbone of well-connected hubs, in which connectivity is resilient to the removal of any individual hub. It is thought to be the basis for efficient global information transfer (Ball et al., 2014). Rich club architecture was also observed in preterm infants scanned at 31 weeks PMA (Ball et al., 2014). This was characterised by a network of densely connected hub regions in the deep grey matter and frontal and parietal cortices. There is an increase in connectivity between the core hubs and the rest of the brain between 31 and 41 weeks PMA, mainly in the fronto-parietal regions (Ball et al., 2014; Heuvel et al., 2015). This is in line with the microstructural maturation over the same period (Kinney et al., 1988). Increases in network integration were also observed in another study that studied changes from birth through to adulthood (Tymofiyeva et al., 2013).

Network analysis has provided useful insight into the early development of structural connectivity. However, the definition of nodes in these populations is not straightforward, due to high inter-subject variability and the lack of well-validated templates. The approach taken for some of the earlier work in this field was to non-linearly register adult atlases to the data (Yap et al., 2011; Fan et al., 2011; Lewis et al., 2014). This made it easier to interpret the connectivity measures assigned to different areas and facilitated cross-subject comparisons. However, due to the significant differences in anatomy between adults and babies, this approach is likely to lead to registration errors and inaccurate characterisation of different regions. The multi-contrast neonatal atlas detailed above (Oishi et al., 2011) is better suited to this population, and has been used to define nodes in several studies (Ratnarajah et al., 2013; Brown et al., 2014; Heuvel et al., 2015).

Overall, graph theoretical analysis promises new insights into brain organisation and the potential for new biomarkers that cannot be characterised by other methods (Ball et al., 2014). However, the results can be very sensitive to the initial choice of nodes and it is difficult to link some of the measures to the underlying neurobiology,

so care must be taken with their interpretation (Hallquist and Hillary, 2018).

3.2.4 Impact of External Factors on Structural Connectivity

In vivo imaging studies further allow us to investigate the impact of early environmental factors on white matter development. Factors such as breastfeeding (Deoni et al., 2013), increased fatty acid levels (Tam et al., 2016), and even music (Sa de Almeida et al., 2020) have been shown to have a positive effect on white matter maturation, whereas infants with chronic lung disease have significantly reduced FA in several tracts (Ball et al., 2010). Antenatal maternal anxiety is also linked to reduced FA in areas related to emotional functioning (Rifkin-Graboi et al., 2015).

Premature birth is a risk factor for many neurodevelopmental disorders, and so it is important to understand how this affects brain development (Johnson and Marlow, 2014). Many studies have probed the effect of premature birth on white matter development (Pannek et al., 2014). For example, it was found that premature birth has a significant impact on the strength of thalamocortical connections (Ball et al., 2013) and that this reduced connectivity had a downstream effect on cognitive scores at two years (Ball et al., 2015).

Furthermore, measures of white matter development can be used to predict later outcomes. A deep-learning approach has been used to predict cognitive scores at two years from whole-brain tractography connectomes (Girault et al., 2019). Connections within the frontal lobe, and between the frontal lobe and the rest of the brain were found to be most important for prediction. This illustrates the potential for biomarkers based on early structural connectivity measures.

3.3 The Challenges of Imaging Early Brain Development In Vivo

Despite significant efforts and progress overviewed in the previous sections, imaging early brain development brings its own unique set of challenges. At a practical level, it is difficult to ensure that a baby stays asleep throughout the scanning process, and even if this is achieved, motion remains a significant problem for neonatal imaging (Barkovich et al., 2019). This is a particular issue for dMRI and can cause issues such as signal drop-out and spin history effects (Bastiani et al., 2019). Reductions in scan duration, dedicated ear protection, and specially designed protocols with graduated noise have all been used as approaches to deal with this issue (Cordero-Grande et al., 2018; Hutter et al., 2018a). Another practical consideration is that standard head coils designed for adults are too big to achieve good signal-to-noise ratio (SNR) on a baby (the neonatal brain is typically one third of the volume of the adult brain (Holland et al., 2014)). Therefore, the use of specially designed head coils provides a large improvement in SNR for neonatal imaging (Hughes et al., 2017; Lopez Rios et al., 2018). The small neonatal head size also means that a high image resolution is required to avoid partial voluming effects.

The microstructural characteristics of the neonatal brain, such as low myelination and high water content, lead to a very different image contrast compared to the adult brain. The immature fibre architecture also reduces the diffusion contrast, which hinders tractography. Structurally, the cortical folding of the neonatal brain is not as developed as in the adult brain (Dubois et al., 2008a). Moreover, these factors develop rapidly during the weeks and months after birth (Dubois et al., 2014), which makes it difficult to define a single representative template. Thus, multiple templates have been proposed that characterise the neonatal brain on a week-by-week basis (Kuklisova-Murgasova et al., 2012; Serag et al., 2012; Schuh et al., 2018), including the example shown in figure 3.3.

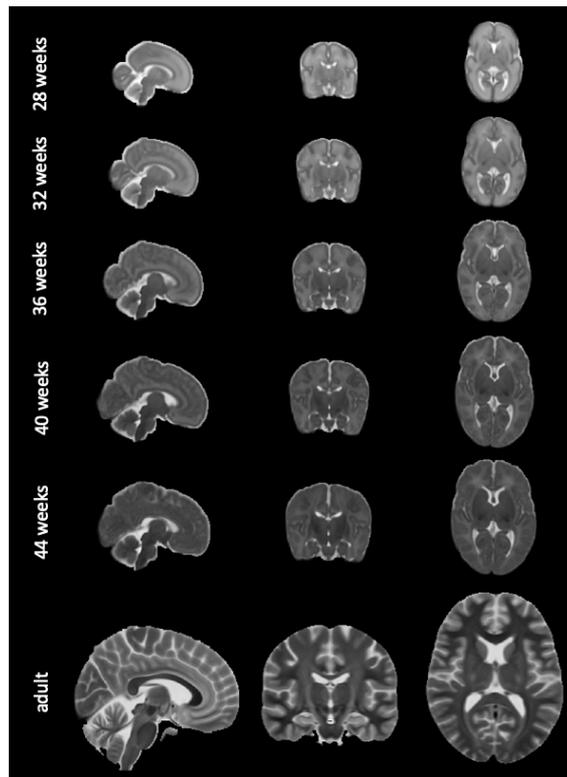


Figure 3.3: *T2-weighted atlases from different timepoints during the neonatal period (Serag et al., 2012), displayed alongside an adult atlas. Changes in image contrast, gyri-fication and brain size can be seen over the four month period shown.*

3.4 Big Data Projects

The last decade has seen the advent of “big data” neuroimaging projects, such as the UK Biobank (Miller et al., 2016) and the Human Connectome Project (HCP) (Van Essen et al., 2013b), that aim to acquire high quality data from large cohorts (1000s of participants). The data acquired from these projects is made available to the wider research community as a resource to help further our understanding of the brain. More recently, several such projects have been started with a focus on the neonatal period (Hughes et al., 2017; Howell et al., 2019). By optimising scanning hardware, protocols, and analysis pipelines for neonates, they aim to address some of the challenges of imaging this population that are described in the previous section.

3.4.1 The Developing Human Connectome Project

The developing Human Connectome Project (dHCP) aims to create a dynamic map of human brain connectivity during the first weeks of life (Hughes et al., 2017). To date, over 800 neonatal and over 250 fetal MRI scans have been acquired. Neonates were scanned between 24 and 45 weeks PMA. Structural, functional and diffusion MRI have all been collected, with scanning protocols tailored for the unique challenges of this population.

A number of MRI acquisition developments have been made for the dHCP to mitigate some of the problems associated with imaging neonates, including a specialised neonatal head coil and a dedicated MR compatible trolley to minimise disturbance when placing the babies in the scanner (Hughes et al., 2017).

The effects of motion can also be reduced by minimising scan times and avoiding abrupt changes in noise levels. Functional and diffusion data were acquired for the dHCP using a multiband EPI sequence designed to have reduced noise and scan length compared to standard sequences (Hutter et al., 2018a), while maintaining image quality and SNR. In addition, a novel multi-shell HARDI protocol has been developed (Hutter et al., 2018b), which incorporates multiband capacity to reduce scanning times, and has flexible sampling of both diffusion and phase encoding to maximise the usability of data even after loss due to motion or scan disruption. Reconstruction methods have also been developed to cope with within-plane and through-plane motion (Cordero-Grande et al., 2018).

Preprocessing pipelines have been developed for the dHCP that are specifically tailored to data from this challenging population (Bastiani et al., 2019; Makropoulos et al., 2018; Fitzgibbon et al., 2019). The cortical surface reconstruction pipeline builds upon the segmentation and cortical surface extraction techniques from the adult HCP (Makropoulos et al., 2018), using tools from FreeSurfer (Fischl, 2012). However, the pipeline has been tailored to specific contrasts and resolution of the developing brain, yielding more accurate segmentation than other neonatal seg-

mentation models. The functional preprocessing pipeline incorporates motion and susceptibility distortion corrections, boundary based registration of functional to structural data, and ICA denoising (Fitzgibbon et al., 2019). Diffusion preprocessing and quality control is detailed in (Bastiani et al., 2019). The dMRI data were simultaneously corrected for eddy currents, susceptibility induced distortions, and motion, both within volumes and across volumes (Andersson and Sotiropoulos, 2016; Andersson et al., 2017). This provides a state-of-the-art distortion correction, which helps to mitigate some of the issues with data quality described above. Automated quality control is built into both the functional and diffusion pipelines.

3.4.2 The Baby Connectome Project

The Baby Connectome Project (BCP) is one of the US NIH Lifespan Connectome Projects, building on the approach of the Human Connectome Project (HCP) (Van Essen et al., 2013b). Its primary aim is to characterise brain and behavioural developments in 500 typically developing infants during the first five years of life (Howell et al., 2019). As far as possible, the scans and behavioural tests replicate those used for the young adult HCP, although necessary changes have been made to suit the needs of the younger subjects in the BCP.

One of the main features of the BCP is the hybrid longitudinal and cross-sectional study design. 285 subjects will be scanned between four and six times, from shortly after birth, until 44 months. The remainder (215 subjects) will be scanned once, between the ages of 3 and 60 months. The dense longitudinal sampling improves the ability to detect non-linear trajectories in functional and structural development (Howell et al., 2019).

The structural and functional imaging parameters are mainly similar to those of the other HCP projects. However, the dMRI acquisition parameters were modified for this age range. A six-shell sampling scheme was chosen, as this was found to facilitate the reconstruction of local fibre orientations with increased accuracy,

compared to the three-shell scheme used for the dHCP (Howell et al., 2019).

3.4.3 Aims of this Work

The projects described above will provide a wealth of high-quality data to the neonatal imaging community, but improved analysis methods tailored to this early phase of development are needed to make the most of these data. In the next chapters, new frameworks are presented for mapping brain connectivity from diffusion MRI data in the neonatal brain.

First, in chapter 4, a new and extensive set of tractography protocols are defined for the neonatal brain. These describe a much larger set of tracts than existing protocols, and can be used for reproducible analysis of tract maturation. The protocols are generalisable across the lifespan, ensuring seamless links between the analysis of neonatal and adult brain connectivity. Different methods are then compared for generating whole-brain connectomes for the neonatal brain. We focus on the impact of the gyral bias on these, as this is an important confound for adult tractography that has not yet been explored in the developing brain.

In chapter 5, a novel data-driven approach is devised to simultaneously extract white matter bundles and grey matter networks from these whole-brain connectomes. This method is protocol-free, so does not rely on a particular template or any prior assumptions about the rapidly changing neonatal anatomy. In chapter 6, the stability of these results are explored, and we demonstrate how the group-level results can be mapped onto individual subjects or subgroups. This approach can also be used to generate parcellations of the neonatal cortex that are based on structural connectivity.

The final chapter summarises these methods and provides an outline of their potential future applications. The frameworks developed provide a springboard for a combined analysis of neonatal brain connectivity with adults and other non-human primates, while still being sensitive to the neonatal anatomy, which will be useful

for exploring the ontogeny and phylogeny of white matter connections. They also provide routes to joint analysis of structural and functional connectivity in the brain.

Part III

Original Research

Chapter 4

Tractography Approaches to Map Connectivity During Early Development

Contents

4.1	Introduction	56
4.2	Methods: Automated Tractography Protocols	59
4.2.1	Protocol Definitions	60
4.2.2	Data	68
4.2.3	Model Fitting	68
4.2.4	Tractography and Atlases	69
4.2.5	Comparison with Existing Protocols	70
4.2.6	Microstructural Analysis	72
4.3	Methods: Gyral Bias	72
4.3.1	Data	72
4.3.2	Surface Alignment	73
4.3.3	Predicted Bias	74
4.3.4	Tractography Bias	74
4.4	Results	76
4.4.1	Automated Tractography Protocols	76
4.4.2	The Impact of the Gyral Bias on the Neonatal Connectome	80
4.5	Discussions	85
4.5.1	Standardised Protocols for Neonatal Tractography	86
4.5.2	The Impact of the Gyral Bias on the Neonatal Connectome	87
4.6	Concluding Remarks	88
4.7	Appendix	89

Foreword

This chapter presents two complementary approaches for mapping white matter connections in the neonatal brain: ROI-based tracking of individual tracts, and the generation of whole-brain connectivity matrices. First, a new set of tractography protocols is presented, that allow us to reliably map 42 white matter tracts in the neonatal brain. The protocols are defined to be functionally consistent with an existing set of protocols for the adult human and the macaque (Warrington et al., 2020), which aids comparison between these groups. The results are used to analyse changes in microstructure across different tracts over the neonatal period.

We then explore the impact of the gyral bias on neonatal connectomes. The gyral bias is a methodological confound of tractography, which causes streamlines to be over-represented at the gyral crowns, compared to the sulcal walls and fundi. Results are compared between two different seeding strategies, and between neonates and adults, to determine the optimal seeding strategy to use in future work.

4.1 Introduction

Tractography has been a valuable tool for elucidating many details of early white matter development and the complex factors that impact the development of white matter pathways, as demonstrated in the previous chapter. In this chapter, traditional tractography approaches are extended and explored in two ways, relevant for connectivity mapping in the neonatal brain. Firstly, automated tractography protocols are developed for major white matter bundle reconstruction, tailored to neonatal brain anatomy. Secondly, the fidelity of whole-brain connectome mapping approaches are further explored in newborns with respect to the gyral bias, a major confound in connectome mapping of the adult brain.

Many of the studies referenced in the previous chapter have used ROI-based tractography protocols to map individual white matter tracts. However, in most of these studies ROIs are delineated on a subject-by-subject basis (Dubois et al., 2008b; Liu et al., 2010; Kaur et al., 2014; Cohen et al., 2016; Akazawa et al., 2016), which is time consuming for large cohorts. Even with neuroanatomists who have extensive training in the anatomy of this difficult population, there is some inter-rater variability in manually delineated tracts (Kaur et al., 2014). Standardised tractography protocols facilitate the consistent and reproducible extraction of white matter bundles for large numbers of subjects, using prior anatomical knowledge to reduce false positives in fibre tracking (Catani et al., 2002; Wakana et al., 2004). However there is no widely used set of protocols that can be used to map a comprehensive set of white matter tracts in neonates.

We have developed and present a set of automated tractography protocols for 42 white matter tracts in the neonatal brain, building upon a library of protocols that has recently been devised and used in large cohorts of subjects (Warrington et al., 2020). These protocols (XTRACT for X-species TRACTography) have been designed to be generalisable across adult humans and macaques. We adjust, aug-

ment and test them for the neonatal brain, ensuring correspondence with their adult counterparts. We then use them on data from the developing Human Connectome Project, to examine early changes in tract microstructure. The consistency between these neonatal tractography protocols and those for the adult human and macaque means that they will provide a valuable resource for the study of changes in structural connectivity across the lifespan and evolutionary development.

The second part of this chapter is concerned with mapping the whole connectome in the neonatal brain. There are many different choices to be made during the data processing that can affect the results and introduce bias. We specifically examine the effects of different streamline seeding strategies on the gyral bias. The gyral bias is a major confound for adult brain connectomes and describes the tendency of tractography streamlines to terminate preferentially at the gyral crowns rather than the gyral walls or sulcal fundi (Van Essen et al., 2013b; Schilling et al., 2018). When this was first observed, it was thought to reflect the underlying anatomy (Nie et al., 2012; Chen et al., 2013), however comparisons with tracer results from the macaque (Van Essen et al., 2013a) and histology (Schilling et al., 2018) have shown that this is a tendency intrinsic to tractography itself (Van Essen et al., 2013a).

Analysis of the gyral bias is complicated by the fact that there is some innate tendency towards the gyri to be expected in axonal terminations, based on the geometry of the cortex (Van Essen et al., 2013b). This is partly because the cortex tends to be thicker along the gyral crowns than the sulcal fundi. In addition, if we look through a cross section of a gyrus, as in figure 4.1, and imagine wedge shaped sections of equal volume in the cortical grey matter at the sulci and gyri, we can see that the surface area bounding the white matter is greater relative to the volume of the cortical grey matter at the sulcal fundi than it is at the gyral crown. If a constant density of axons per unit of cortical volume is assumed, as has been observed in histology (Van Essen et al., 2013b), then there must be a higher density of axons terminating at the gyri than at the sulci, due to the smaller surface

area available at the gyral crown. The fluctuations in streamline density due to this effect can be predicted by calculating the cortical volume associated with each unit of surface area of the white matter/grey matter boundary (WGB), and using this method it has been demonstrated that the gyral bias in tractography exceeds the predictions made based on anatomy alone (Van Essen et al., 2013b; Donahue et al., 2016).

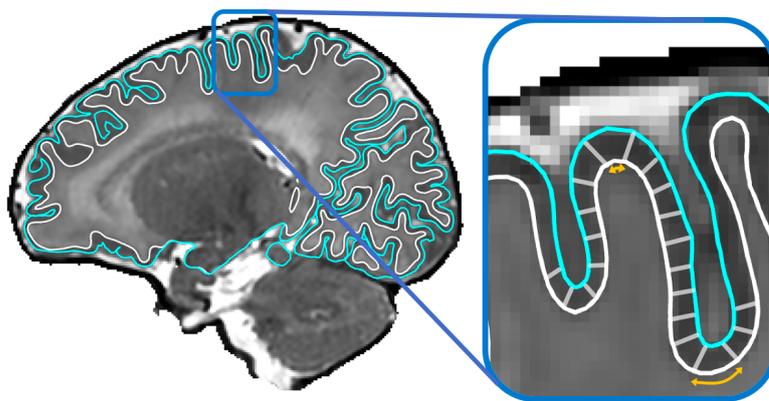


Figure 4.1: *A diagram illustrating how cortical folding patterns introduce expected variations in the density of axonal terminations at the cortex. The surface area of the WGB bounding a wedge of cortex with equal volume is greater at the sulcal fundi than it is at the gyral crown (yellow arrows). This accounts for some of the measured gyral bias, but is not sufficient to explain the magnitude of the variation in streamline density across the cortex that is observed in tractography.*

Given that the cortical folding patterns in neonates are less mature than in adults, we explore the effects of the gyral bias on tractography results in the neonatal brain with respect to the theoretically expected bias. Two seeding strategies are compared: seeding from the white matter/grey matter boundary (WGB), and seeding streamlines throughout the brain volume. Previous work has indicated that whole-brain seeding leads to a larger effect of the gyral bias in adults compared to WGB seeding (Van Essen et al., 2013b; Schilling et al., 2018), however the impact of the gyral bias has not been measured in neonatal tractography before. Exploring the effects of the gyral bias in the neonatal brain, where gyrification is less developed, and comparing against adults, allows us to determine optimal seeding strategies for

whole-brain connectomes and the extent to which gyral bias is an issue to consider when interpreting neonatal connectome studies.

4.2 Methods: Automated Tractography Protocols

Automated tractography protocols provide a set of instructions for tractography, to reconstruct white matter tracts in a reproducible way across subjects. These instructions take the form of a set of masks in a standard space that provide logical operations for the propagation and retention of streamlines, based on anatomical knowledge. These masks are then warped to the subject's native space to perform tractography. Figure 4.2 illustrates the different masks used in our protocols. Streamlines all start at the seed mask, and only those which reach the target mask are included in the final result (so streamline *a* is rejected because it does not propagate through the target). An exclusion mask can be added to remove streamlines that enter it, so streamline *b* is also rejected, despite reaching the target mask. Streamlines *c* and *d* are included in the final result because they both propagate from the seed mask to the target, without being impeded by any exclusion masks. However, streamline *c* is cut short upon reaching the stop mask, which terminates any streamlines which enter it.

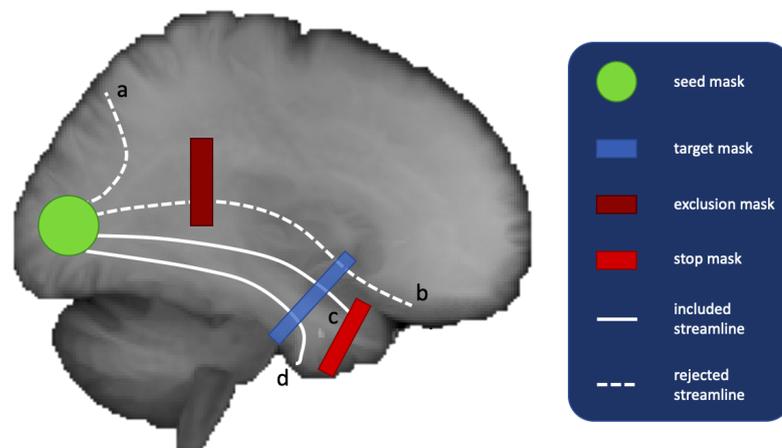


Figure 4.2: Schematic diagram to illustrate the different masks used in the baby-XTRACT protocols.

In this work we use the 40-week template from a spatio-temporal neonatal atlas developed by Schuh et al. (2018), as our standard space. The atlas was created using a group-wise method applied to data from 275 healthy neonates scanned between 35 and 44 weeks PMA from the dHCP. The resultant atlases were shown to preserve cortical details better than other atlases that are available for this age-group (Schuh et al., 2018).

Protocols were defined in this standard space. Masks are transformed from standard space to the subject’s native space using a non-linear registration warp field (Avants et al., 2011). Tractography is performed in native space and the results are directly resampled back to the standard space, to allow for between-subject comparison and averaging. This process also minimises the resampling required.

4.2.1 Protocol Definitions

Tractography protocols were defined following the adult XTRACT protocols developed by Warrington et al. (2020), hence we refer to them collectively as baby-XTRACT protocols. The underlying idea is that all these protocols (human adult, macaque and human neonate) use equivalent functional definitions, while preserving the shape and size differences across these diverse brains. To define the protocols, a non-linear warp field was used to roughly align the adult XTRACT protocol masks to a 40 week PMA neonatal template (Schuh et al., 2018). Using the adult protocols as a starting point, the registered masks were then manually redrawn to ensure good alignment and correspondence to the neonatal anatomy. Seed and target masks were enforced to have equal volumes in each hemisphere, for bilateral tracts, to avoid artificial lateralisation. Some additional changes were made to the protocols, based on preliminary results, in order to optimise the results for the neonatal anatomy, as described below.

A list of the tracts included in the protocols is shown below in table 4.1. Some of the protocols use a reverse-seeding approach, in which the protocol is run twice,

with the roles of the seed and target masks exchanged. The resultant streamline distributions are then added together.

Category	Tract Name	Bilateral	Reverse Seeding	Adjustments to Adult Protocols
Association	Superior longitudinal fasciculus 1, 2 and 3 (SLF1, SLF2, SLF3)	X		Additional exclusions, larger seed masks
	Arcuate fasciculus (AF)	X	X	Change to target position, additional exclusion
	Middle longitudinal fasciculus (MdLF)	X	X	
	Inferior longitudinal fasciculus (ILF)	X	X	
	Inferior fronto-occipital fasciculus (IFO)	X	X	
	Uncinate fasciculus (UF)	X		
	Frontal aslant tract (FA)	X		
	Vertical occipital fasciculus (VOF)	X	X	
Commissural	Middle cerebellar peduncle (MCP)		X	Additional exclusion
	Splenium of the corpus callosum (FMA)		X	
	Genu of the corpus callosum (FMI)		X	
	Anterior commissure (AC)		X	
Limbic	Cingulum subsections: temporal, dorsal and peri-genua (CBT, CBD, CBP)	X		Changes to exclusions
	Fornix (FX)	X		
Projection	Corticospinal tract (CST)	X		Additional exclusion
	Anterior thalamic radiation (ATR)	X		
	Superior thalamic radiation (STR)	X		
	Acoustic radiation (AR)	X	X	
	Optic radiation (OR)	X	X	

Table 4.1: Forty-two tracts included in baby-XTRACT

The sections below describe the protocol for each tract in detail. With the exception of the commissural tracts, all protocols include the midline sagittal plane as an exclusion mask to restrict fibres to the ipsilateral hemisphere.

Association Fibres

Superior Longitudinal Fasciculus (SLF) 1/2/3: The superior longitudinal fasciculus is a longitudinal parieto-frontal tract associated with visuospatial attention

(Thiebaut de Schotten et al., 2011a). It can be separated into three bundles: a dorsal superior longitudinal fasciculus SLF1, middle SLF2 and ventral SLF3. The three branches of the SLF are reconstructed using an extension of the approach taken by Thiebaut de Schotten et al. (2011a). In each case a coronal plane in the region of the central sulcus within the frontal/parietal cortex is used as a seed along with two target masks. Frontally, target masks for the first, second, and third branches of the SLF were coronal sections through the superior, middle, and inferior frontal gyri, respectively, placed at the level of the posterior end of the genu of the corpus callosum. Posteriorly, a large coronal target mask in the superior parietal lobule, immediately posterior to the margin of the cingulate gyrus is used for SLF1. For SLF2 and SLF3, the second target masks are placed in the angular gyrus and supramarginal gyrus respectively. In each case, seed placement reflects the placement of the second target whilst being moved anteriorly into the region of the central sulcus. For each protocol, an axial exclusion mask was placed underneath the parietal cortex and one blocking subcortical areas prevented leaking into ventrally oriented fibres. A coronal exclusion mask through subcortical areas posterior to the caudal end of the genu of the corpus callosum prevented leaking into ventral longitudinal tracts. The SLFs are still maturing during the neonatal period (Dubois et al., 2014), which makes them challenging to delineate in this age-group. Therefore, extra exclusion masks were added compared to the adult protocols, to help constrain the tractography. For each branch, the target masks for each of the other SLF branches were also included as exclusion masks. An additional exclusion mask was also placed in the cingulate gyrus for the SLF1, to prevent leakage into this region. The seed and target masks were also increased in size compared to the adult masks, which made the results more robust.

Arcuate Fasciculus (AF): The arcuate fasciculus connects Broca's area and Wernicke's area, and is associated with language functioning (Eichert et al., 2019). The AF was reconstructed with a seed in the supramarginal gyrus (SMG), a tempo-

ral target mask was in the white matter encompassing the superior temporal gyrus (STG) and middle temporal gyrus (MTG), and an anterior target at the level of the ventral premotor cortex, posterior to the inferior frontal gyrus (IFG) and anterior to the precentral sulcus. An axial target mask was placed in the parietal-temporal white matter posterior to the caudal end of the Sylvian fissure. An additional axial plane was placed in the IFG. The anterior target was moved posteriorly to be closer to the seed than in the adult protocol, as we found that this improved our ability to segment the frontal part of this bundle. An axial exclusion mask was also added for the neonatal protocols to prevent fibres running through the internal capsule.

Middle/Inferior Longitudinal Fasciculus (MdLF, ILF): The middle and inferior longitudinal fasciculi are tracts within the lateral posterior cortex of the temporal lobe. The MdLF was seeded in the anterior part of the superior frontal gyrus (SFG) (Makris et al., 2008), and ILF in the middle and inferior temporal gyri. For the MdLF, large axial and coronal planes covering the white matter in the temporo-parietal-occipital junction were used as targets, based on anatomical descriptions from Makris et al. (2013). For ILF, a coronal plane in middle and inferior temporal gyrus is used as a target. For both protocols, exclusion masks were placed axially through the brainstem, coronally through the fornix, axially through the cingulum bundle posterior to the corpus callosum and through the entire frontal cortex. In addition, the seed mask of MdLF served as an exclusion mask for ILF and vice versa, and the ILF target mask was used as an exclusion mask in the MdLF. Additionally, for the ILF, a coronal exclusion mask was placed in the in the centrum semiovale and an axial exclusion mask covering the white matter of the SMG was used.

Inferior Fronto-Occipital Fasciculus (IFO): In contrast to MdLF and ILF, the inferior fronto-occipital fasciculus, also termed the extreme capsule fibre complex (Mars et al., 2016), runs more medially and courses into the frontal cortex through the extreme capsule. Extending the recipe of (Wakana et al., 2007), the seed was a coronal plane through the anterior part of the occipital cortex, the target a coronal

plane through the frontal cortex anterior to the genu of the corpus callosum. An exclusion mask just behind the anterior commissure excluded all fibres except those running through the extreme capsule.

Uncinate Fasciculus (UF): The uncinata fasciculus is a hook-shaped bundle that connects the frontal lobe with the anterior temporal lobe. The tract was reconstructed using a seed in the STG at the first location where the temporal and frontal cortex are separated, a target through the ventral part of the extreme capsule, and an exclusion mask layer between the seed and the target to force the curve. An additional coronal exclusion mask prevented leaking into the fibres running longitudinally through the temporal lobe.

Frontal Aslant (FA): The frontal aslant is a short tract running in the frontal lobe between the posterior part of the inferior and superior frontal gyri (Catani et al., 2012). The seed was placed sagittally in the white matter of the IFG, the target axially in that of the SFG. A posterior coronal exclusion mask prevented leakage into longitudinal fibres.

Vertical Occipital Fasciculus (VOF): The vertical occipital fasciculus runs in a predominantly dorsal-ventral orientation in the occipital lobe. The original protocol was adapted from Takemura et al. (2017). An axial seed mask was placed in the lateral part of the ventral occipital white matter posterior to the anterior occipital sulcus (Petrides et al., 2012). A larger axial target mask was placed dorsally at the level of the lateral occipital sulcus. A coronal plane just posterior to the corpus callosum served as an exclusion mask to prevent leakage into anterior-posterior tracts.

Commissural Fibres

Middle Cerebellar Peduncle (MCP): The middle cerebellar peduncle connects the cerebellum to the pons. This tract was seeded in the cerebellar white matter with a

target in the opposite hemisphere (and their inverses). Exclusion masks were placed sagittally along the cerebellar midline and axially through the thalamus.

Corpus Callosum Splenium (FMA) and Genu (FMI): Callosal connections to the occipital lobe were constructed via the splenium of corpus callosum (forceps major, FMA) and to the frontal lobe via the genu of corpus callosum (forceps minor, FMI) using recipes based on those defined by Wakana et al. (2007). Seed and target masks (and their inverse) for the FMA were defined as coronal sections through the occipital lobe at the posterior end of the parietal occipital sulcus. The sagittal exclusion mask was confined to the occipital cortex and the subcortex. Additional exclusion masks through the inferior fronto-occipital white matter and a coronal plane through the pons prevented leakages to longitudinal fibres. Seed and target masks (and their inverse) for the FMI were defined as coronal sections through the frontal lobe at the anterior end of the pregenual cingulate sulcus. The midsagittal exclusion mask was interrupted at the level of the anterior third of the corpus callosum and an additional coronal exclusion mask at the same level prevents posterior projections.

Anterior Commissure (AC): The anterior commissure connects the temporal lobes of the two hemispheres across the midline. It was seeded in the left-right oriented fibres on the midline, with a target mask covering the white matter lateral to the globus pallidus. Stop masks were placed directly underneath and lateral to the two amygdalae. A large axial exclusion mask was placed dorsal to the seed through the entire subcortex. Compared to the adult protocols, an axial exclusion mask was added covering the optic chiasm to prevent leakage into the optic nerve.

Limbic fibres

Cingulum subsections (CBT, CBP, CBD): The cingulum facilitates communication between different parts of the limbic system. It projects from the cingulate gyrus to the entorhinal cortex (Heilbronner and Haber, 2014). Protocols were defined for three distinct sections of the cingulum, based on a recent segmentation by Heilbron-

ner and Haber (2014). The temporal part (CBT) was seeded in the posterior part of the temporal lobe at a section where the fibres of the cingulum are mostly oriented in the anterior-posterior direction. The target was placed posteriorly to the amygdala and stop masks were placed posteriorly and anteriorly to the seed and target masks, respectively. An exclusion mask prevented leaking into the fornix. The dorsal segment (CBD) was seeded just above the posterior part of the corpus callosum and had a target at the start of the genu of the corpus callosum. A sagittal exclusion mask in the anterior limb of the internal capsule prevented leakage into the temporal lobe. Finally, the peri-genu part of the cingulum bundle (CBP) was seeded anteriorly above the corpus callosum and a target placed below the sub-genu callosum with a stop mask placed inferior and anterior to the target. A callosal plane at the level of the rostral end of the Sylvian fissure prevented leakage into the CBD.

Fornix (FX): The fornix connects the hippocampus with the mammillary bodies, the anterior thalamic nuclei, and the hypothalamus (Catani et al., 2013a). The tract was reconstructed using a seed in the body of the fornix at the level of the middle of the corpus callosum and a target in the hippocampus. A callosal plane at the anterior end of the occipital cortex prevented leakage into posterior tracts and bilateral sagittal planes around the midline, at the level of the anterior tip of the thalamus prevented lateral propagation to the anterior limb of the internal capsule. These exclusion masks were reduced in size relative to the adult protocol as they were found to hinder the fibre tracking. To prevent leakage into the cingulum, an axial exclusion mask posterior to the splenium of the corpus callosum and a small axial exclusion covering the parahippocampal gyrus region of the cingulum are also used.

Projection Fibres

Corticospinal Tract (CST): The corticospinal, or pyramidal, tract extends from the spinal cord through the midbrain and distributes to motor cortex, premotor cortex

and somatosensory cortex. The tract is seeded from the pons with a large target covering the motor, premotor and somatosensory cortices. An axial exclusion mask is used to restrict tracking to the cerebral peduncle of the midbrain. In addition, the exclusion mask includes two coronal planes, anterior and posterior to the target, to exclude tracking to the prefrontal cortex and occipital cortex respectively and a plane preventing leakage into the cerebellar peduncles.

Anterior and Superior Thalamic Radiations (ATR, STR): The anterior and superior thalamic radiations connect the thalamus to the frontal lobe and pre-/post-central gyrus respectively. The anterior thalamic radiation is seeded using a coronal mask through the anterior part of the thalamus (Wakana et al., 2007) with coronal target mask at the anterior thalamic peduncle. In addition, the exclusion mask contains an axial plane covering the base of the midbrain, a coronal plane preventing leakage via the posterior thalamic peduncle and a coronal plane preventing leakage via the cingulum. A coronal stop mask covers the posterior part of the thalamus, extending from the base of the midbrain to the callosal sulcus. An additional axial exclusion mask was added, compared to the adult protocol, at the level of the cingulate gyrus, to prevent contamination from fibres in the STR. The superior thalamic radiation is seeded using a mask covering the whole thalamus and a target axial plane covering the superior thalamic peduncle. An axial plane is used as a stop mask ventrally to the thalamus. The exclusion mask includes two coronal planes, anterior and posterior to the target, to exclude tracking to the prefrontal cortex and occipital cortex respectively.

Acoustic Radiation (AR): The acoustic radiation connects the medial geniculate nucleus (MGN) of the thalamus to the auditory cortex. It was seeded from the transverse temporal gyrus with a target covering the MGN of the thalamus. The exclusion mask consists of two coronal planes, anterior and posterior to the thalamus, and an axial plane superior to the thalamus. In addition, the exclusion mask contains the brainstem and a horizontal region covering the optic tract.

Optic Radiation (OR): The optic radiation consists of fibres from the lateral geniculate nucleus (LGN) of the thalamus to the primary visual cortex. It was seeded in the LGN and the target mask consisted of a coronal plane through the anterior part of the calcarine fissure. Exclusion masks consisted of an axial block in the brainstem, a coronal block of fibres directly posterior to the LGN to select fibres that curl around dorsally, and a coronal plane anterior to the seed to prevent leakage into longitudinal fibres.

4.2.2 Data

The above protocols were applied to a large number of subjects to obtain neonatal tractography atlases. Diffusion MRI data were used from 489 subjects born at median (range) 39.0 (24.6 – 42.3) and scanned at 40.4 (29.3 – 45.1) weeks post-menstrual age, made available by the second data release of the developing Human Connectome Project (Hughes et al., 2017). Briefly, data were acquired during natural sleep on a 3T Philips Achieva with a dedicated neonatal imaging system, including a neonatal 32 channel head coil (Hughes et al., 2017; Hutter et al., 2018a). Diffusion MRI data were acquired over a spherically optimised set of directions on three shells ($b = 400, 1000$ and 2600 s/mm²). A total of 300 of volumes were acquired per subject, including 20 with $b = 0$ s/mm². For each volume, 64 interleaved overlapping slices were acquired (in-plane resolution = 1.5 mm, thickness = 3 mm, overlap = 1.5 mm). The data were then super-resolved along the slice direction to achieve isotropic resolution of 1.5 mm³ (Kuklisova-Murgasova et al., 2012).

4.2.3 Model Fitting

Preprocessing was carried out according to the dHCP’s preprocessing pipeline (Bastiani et al., 2019), as described in chapter 3. Fibre orientations were modelled (up to 3 per voxel) with FSL’s BEDPOSTX, using a model-based deconvolution against

a zeppelin response kernel, to accommodate for the low anisotropy inherent in data from this age group (Bastiani et al., 2019; Sotiropoulos et al., 2016; Hernandez-Fernandez et al., 2013) (see also chapter 3, figure 2.9).

4.2.4 Tractography and Atlases

Probabilistic tractography was performed using FSL's PROBTRACKX (Behrens et al., 2007; Hernandez-Fernandez et al., 2019), with streamlines seeded from and constrained by the protocol masks, as described above. A curvature threshold of $\pm 80^\circ$ was used, the maximum number of streamline steps was 2000, and subsidiary fibres were considered above a volume threshold of 1%. The step size was 0.5 mm. The resultant path distributions were normalised by the total number of viable streamlines.

When choosing the number of streamlines to propagate from each seed point, there is a trade-off between using sufficient samples to map the full extent of the tract, but not an unnecessarily high number to hinder the speed of the reconstruction. We began by using the same number of streamlines as in the adult protocols (Warrington et al., 2020), but checked if they needed to be increased to account for the difficulty of tracking some bundles in the neonatal brain. For this we used data from ten subjects, all born and scanned at 40 weeks PMA, and the protocols were run with the original number of streamlines per seed, and then repeated with a higher number of streamlines for comparison (50,000). The correlations between the resultant spatial maps from each method are shown below in figure 4.3.

These results indicate that the results for the SLFs and the fornix are altered by using a greater number of streamlines. Visible inspection of the results confirmed that using more streamlines improved the reconstruction of these tracts. Using additional streamlines did not affect the results from the other tracts so the number of streamlines were kept the same as in the adult protocols. Further analysis of the SLFs showed that the SLF 2 would benefit from a further increase in streamlines

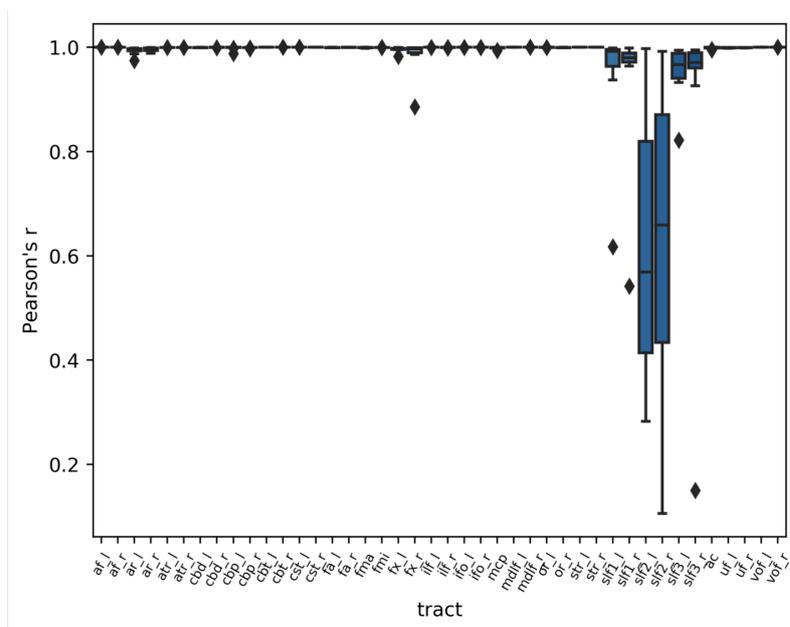


Figure 4.3: Comparison of baby-XTRACT results with a large number of streamlines per seed point (50,000), and the standard number from the adult protocols. Results are shown from ten neonatal subjects, scanned at 40 weeks PMA.

per seed, so 100,000 were used for this tract. The final numbers used are shown in table 4.2.

Tractography results from large groups of subjects were used to obtain tract atlases, in the form of population percentage overlap. The normalised path distributions for each tract were binarised at a threshold value of 0.1%. These binary masks were then averaged across subjects. The resultant spatial maps describe the percentage of subjects for which a given tract is present at a given voxel. Tract atlases were generated for three different age-groups, composed of 70 subjects each, with ages of 36-39 weeks, 39-42 weeks and 42-45 weeks PMA, as well as for the full cohort of 489 subjects.

4.2.5 Comparison with Existing Protocols

The results from baby-XTRACT were compared against results from another set of independently-defined tractography protocols for neonates (AutoPtx) (Bastiani et al., 2019), which however does not ensure correspondence with adult tracts, is

Tract	Streamlines per Seed
Anterior commissure	5000
Arcuate fasciculus	10,000
Acoustic radiation	5000
Anterior thalamic radiation	1000
Cingulum (dorsal segment)	10,000
Cingulum (peri-genua)	10,000
Cingulum (temporal)	10,000
Cortico-spinal tract	3000
Frontal aslant tract	3000
Forceps major	500
Forceps minor	500
Fornix	50,000
Inferior longitudinal fasciculus	3000
Inferior fronto-occipital fasciculus	3000
Middle cerebellar peduncle	3000
Middle longitudinal fasciculus	5000
Optic radiation	10,000
Superior thalamic radiation	1000
Superior longitudinal fasciculus 1	50,000
Superior longitudinal fasciculus 2	100,000
Superior longitudinal fasciculus 3	50,000
Uncinate fasciculus	1000
Vertical occipital fasciculus	10,000

Table 4.2: Streamlines per seed used for each of the forty-two tracts in the baby-XTRACT protocols

defined in an older neonatal template space, and contains only a subset of the protocols developed here: the acoustic radiation, anterior thalamic radiation, cingulum (dorsal and temporal segments), cortico-spinal tract, forceps minor, forceps major, fornix, inferior fronto-occipital fasciculus, middle cerebellar peduncle, optic radiation, superior thalamic radiation and the uncinate fasciculus. To perform the comparison, the baby-XTRACT results were registered to the same template space as the AutoPtx maps (Serag et al., 2012), using ANTs registration tools (Avants et al., 2011) with B-spline interpolation. The Pearson’s correlation coefficient was calculated between the corresponding tracts from each method, for each subject, to assess the similarity of the results.

4.2.6 Microstructural Analysis

FSL's DTIFIT tool was used to fit the diffusion tensor model on the $b = 1000$ s/mm² shell. Fractional anisotropy (FA) and mean diffusivity (MD) were calculated from the diffusion tensor in each voxel. The percentage average maps for each tract from the whole group were thresholded at 30% and used as ROIs for the microstructural analysis. Median FA and MD were calculated within binarised masks, registered to each subjects' native diffusion space. A general linear model was used to assess the relationship between these metrics and the subjects' gestational age at scan in weeks, with the additional regressors included: birth weight (g), head circumference at scan (cm), tract volume in native space (mm³), and data quality control (QC) score (Bastiani et al., 2019). The QC score is the z-score average of the signal-to-noise ratio and contrast-to-noise ratio of the subject's dMRI data.

4.3 Methods: Gyral Bias

To explore the impact of the gyral bias, connectivity matrices were generated for both neonatal and adult subjects using white matter/grey matter boundary (WGB) seeding and whole brain seeding, which are compared in figure 4.4. The theoretically expected bias (under a constant density per unit cortical volume, as described in the Introduction and figure 4.1) was calculated and was compared with the measured streamline density from tractography, as described in more detail below.

4.3.1 Data

Diffusion MRI data were used from 36 neonatal subjects and 36 young adult subjects, obtained from the dHCP and the HCP, respectively. Neonatal subjects were scanned at a median (range) age of 39 (37 - 44) weeks PMA. Neonatal dMRI data were acquired as described in the previous section. The adult HCP data acquisition is outlined in (Sotiropoulos et al., 2013; Van Essen et al., 2013b).

Preprocessing and model fitting of the neonatal data were carried out as above. The crossing fibre model from FSL's BEDPOSTX was also fit to the adult data, with fibre populations modelled as sticks of variable diffusivities (Jbabdi et al., 2012), rather than the zeppelin model used for the neonates.

White matter and pial surfaces were obtained for adults and neonates using the surface extraction pipelines from the developing and adult human connectome projects (Makropoulos et al., 2018; Glasser et al., 2013). These both use information from T1 and T2 weighted MRI to segment the different tissue types and fit a surface mesh to the tissue boundaries.

4.3.2 Surface Alignment

Anatomical surfaces were aligned to ensure consistent seed points for tractography across subjects, using multi-modal surface matching (MSM) (Robinson et al., 2014). MSM is a surface-based registration tool that uses a discrete optimisation framework to align surfaces with a flexible selection of similarity metrics (Robinson et al., 2014). A specialised pipeline was used that has been developed to align neonatal data to the dHCP surface templates (Bozek et al., 2018), which can be found here: https://github.com/ecr05/dHCP_template_alignment. This uses a newer version of MSM, which uses a regularisation term based on biomechanical models of tissue deformation, leading to more anatomically plausible deformations than the first-order regularisation penalties used in the original version (Robinson et al., 2018). A non-linear transform between each neonatal surface's native T2 space and the 40-week template was estimated with MSM, based on cortical folding. The native surfaces were then resampled to the surface topology of the template (the FS_LR32k space). A rigid registration was then applied to align the surfaces to the subjects' dMRI data.

4.3.3 Predicted Bias

Under the assumption of constant axonal density per unit cortical volume, some of the regional variation in streamline density at the WGB surface can be attributed to cortical folding and the variation in cortical thickness across the cortex (Van Essen et al., 2013a; Donahue et al., 2016) (figure 4.1 and Introduction). This expected bias was predicted by calculating the cortical volume associated with each unit of surface area of the WGB, using the HCP’s Connectome Workbench Tools (Marcus et al., 2011). Specifically, the volume per-vertex between the WGB and the pial surfaces was calculated, and the surface area associated with each vertex on the WGB surface was calculated as one third of the area of each triangle it is a part of. The ratio of these two quantities was calculated to give the predicted bias.

4.3.4 Tractography Bias

Fibre orientation estimates were used to carry out whole-brain probabilistic tractography in subject-specific dMRI space for each subject. 64k x 64k connectomes were obtained with two different seeding strategies, which are illustrated below in figure 4.4.

WGB Seeding

10,000 streamlines were seeded from each vertex on the WGB. A visitation count was recorded when a streamline terminated at another point on the WGB, connecting two grey matter vertices.

Whole Brain Seeding

3000 streamlines were seeded from each voxel in a 2 mm downsampled whole brain mask, with the WGB surface used as a target mask. In this case, if a streamline propagated from a seed point in both directions hits the WGB at two locations,

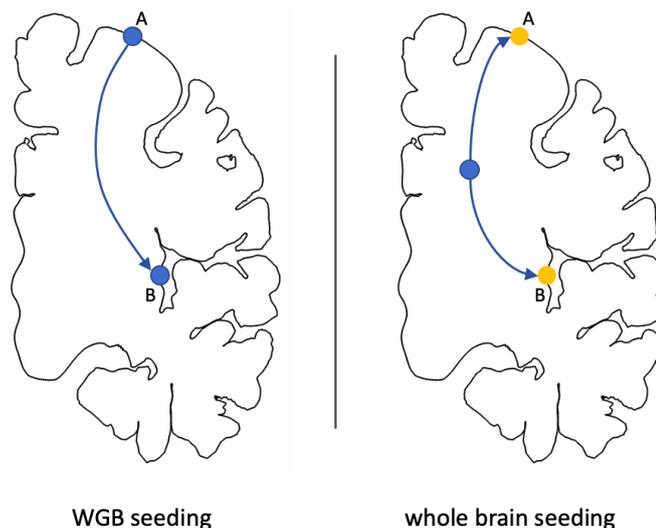


Figure 4.4: Schematic diagram illustrating the seeding strategies used for investigating the impact of the gyral bias. Selected seed points shown in blue. Both diagrams would result in a structural connection recorded between points A and B on the WGB. For WGB seeding, streamlines are seeded on the WGB, and a connection is recorded between seed points A and B, if a streamline from A hits the WGB at point B (or vice versa). In the case of whole-brain seeding, streamlines are seeded throughout the brain volume. A connection is recorded between points A and B if separate streamlines from a seed point hit both A and B.

a visitation count is recorded in the row and column of the connectivity matrix corresponding to these grey matter locations.

In both cases, the pial surface was used as a stop mask and streamlines were allowed to cross the WGB no more than twice. The resultant connectivity matrices were summed along their rows to give the density of streamlines terminating at each point on the WGB.

Normalisation

All quantities (predicted bias and measured streamline densities) were normalised with respect to their average values in zero curvature regions. This describes the density in a given region compared to the average density in zero-curvature regions for that method and enables easy comparison between the different methods. Zero curvature regions were defined as those with curvature = 0 ± 0.025 .

4.4 Results

4.4.1 Automated Tractography Protocols

The tractography protocols were applied to 489 subjects from the dHCP. Percentage averages were generated by binarising the thresholded maps for each subject, and averaging across subjects to obtain atlases for three age groups, containing subjects scanned at 36-39, 39-42 and 42-45 weeks PMA. The percentage average maps from these groups are displayed in figure 4.5. All tracts are well characterised by the protocols. The results are very similar between the three age groups, which indicates that the protocols provide robust results over this age range.

In figure 4.6, results are displayed from two subjects in the 39 - 42 week age-group. To illustrate the effect of data quality on the results, we show results from one subject with a low QC score (9th percentile of the cohort), and another with a high QC score (93rd percentile). The QC score is the z-score average of the signal-to-noise ratio and contrast-to-noise ratio of the subject's dMRI data. These subjects were selected randomly to be representative of the variation in data quality across the cohort. All the tracts are reproduced in both subjects, with only small differences in the subject with low QC score, such as a reduction in the frontal projections of the ATR. Overall, these results show that the protocols are robust to differences in data quality.

Figure 4.7 shows qualitatively a key feature of these protocols; correspondence with the adult brain, and even with a different non-human primate brain. The same tractography definitions seem to be generalisable across these three very diverse groups.

Comparison with Independently-Defined Protocols

To gain further evidence on the fidelity of these protocols, the results from the baby-XTRACT protocols were compared with those from another recently-defined set of

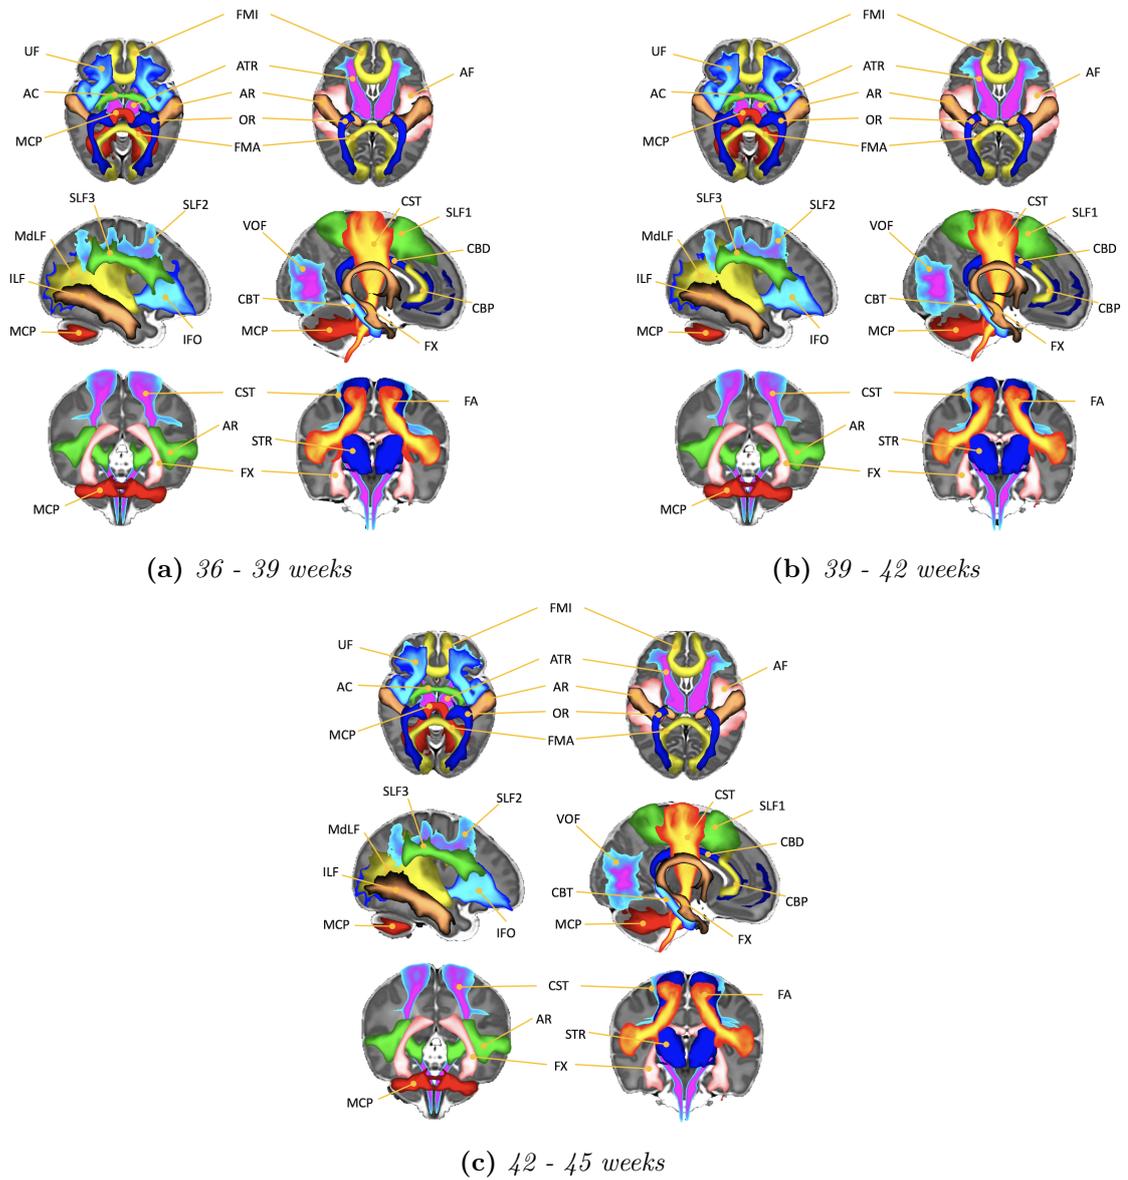


Figure 4.5: Percentage average maps of the tracts from the baby-XTRACT protocols, for different age groups within the cohort. The tracts are displayed as maximum intensity projections, for ease of visualisation.

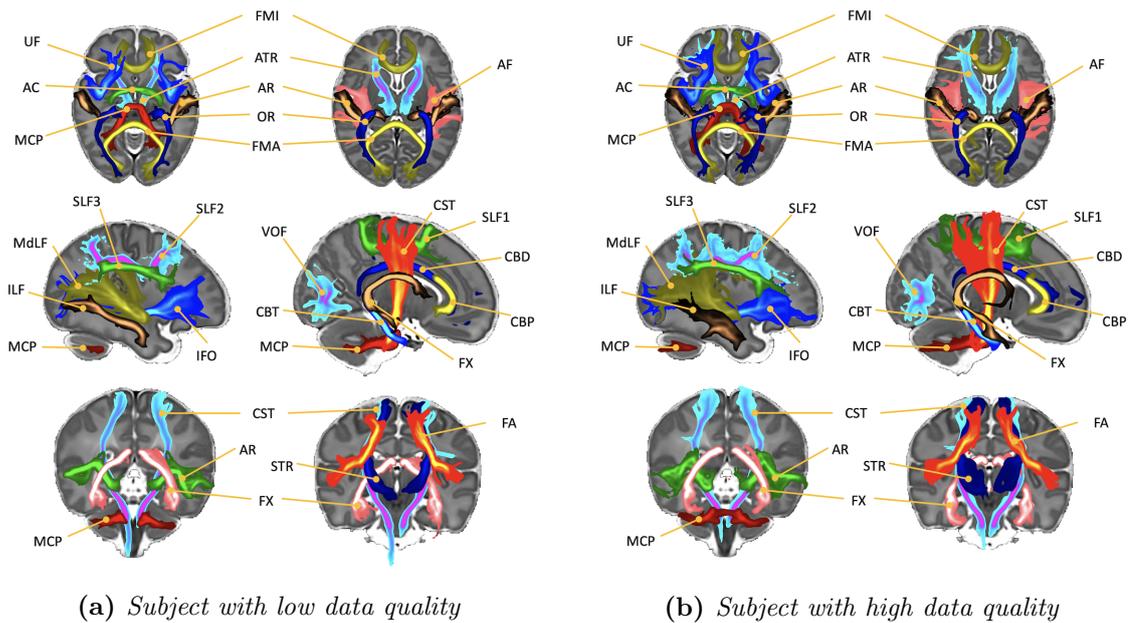


Figure 4.6: *Baby-XTRACT* results from two subjects, with *QC* scores at the lower and higher end of the cohort, respectively. Tracts are displayed as maximum intensity projections.

protocols (Bastiani et al., 2019) (referred to as AutoPtx), on the same subjects. Correlation scores between the spatial maps from the two approaches are shown below in figure 4.8.

The two protocols give fairly similar results. Both protocols are defined relative to different atlases, so some differences will be because of the inconsistencies between the templates, or inaccuracies in the registration between them. In addition, seeds and targets of bilateral tracts were ensured to have equal volumes in baby-XTRACT, but not in AutoPtx, which may inform some of the inter-hemispheric differences. A number of tracts had less agreement than others and differences in these protocols are presented in detail in the Appendix. Summarising, we obtain reasonable similarity with an independently-defined set of tracts, while baby-XTRACT protocols have the advantage of inherent adult and cross-species correspondence.

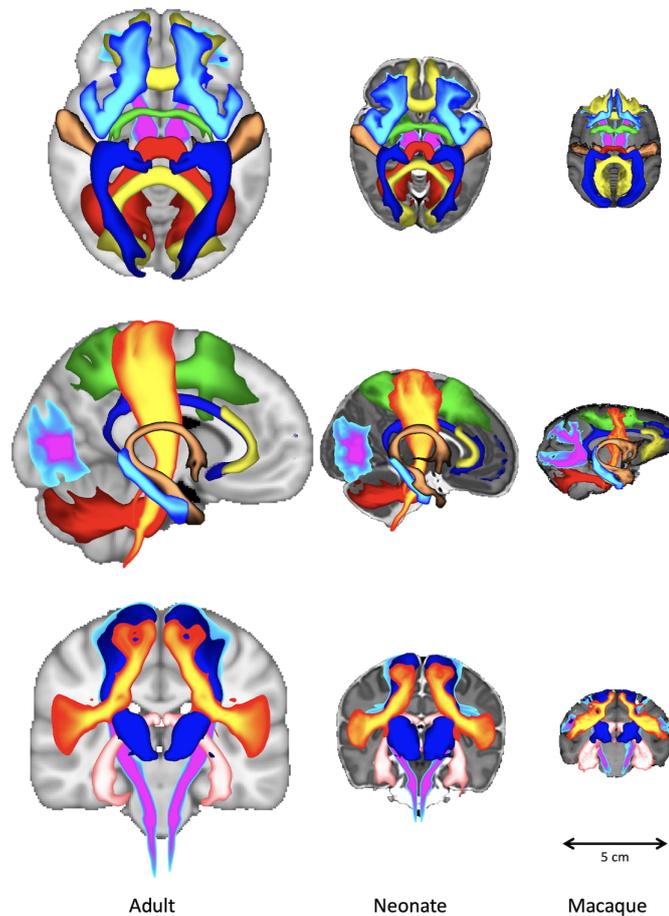


Figure 4.7: Comparison of group-average baby-XTRACT results with those from adult humans and macaques. The adult and macaque tract atlases were obtained from (Warrington et al., 2020). Results are displayed as maximum intensity projections.

Microstructural Analysis

The subject-level results were used as ROIs to investigate the changes in FA and MD with age at scan in these tracts. All tracts showed a significant increase in FA and reduction in MD over the age-range studied, even when corrected for birth weight, head circumference, tract volume, and QC score. Scatter plots of FA and MD against age at scan are shown below in figures 4.9 and 4.10. For bilateral tracts, the tract in each hemisphere is given its own data point. The high values of MD in the fornix for a few subjects of older ages are probably caused by partial voluming effects, due to the tract's proximity to the ventricles.

A linear regression analysis was performed, correcting for birth weight, head cir-

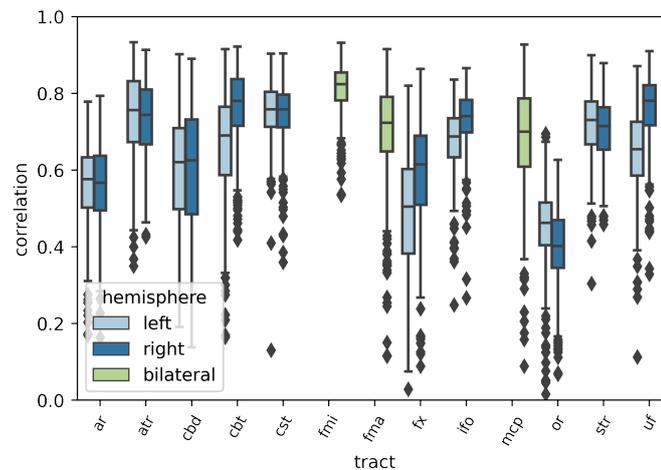


Figure 4.8: Correlations between matching tracts from XTRACT and AutoPtX across 489 subjects, aged between 29.3 and 45.1 weeks PMA at scan.

cumference, tract volume, and QC score. The regression coefficients corresponding to age at scan are plotted in figure 4.11. This reflects the asynchronous maturation of different tracts (Dubois et al., 2008b).

In general, projection fibres are maturing more quickly over this period, followed by commissural and association fibres, with limbic fibres showing the slowest rates of maturation. This agrees with the trends shown in other studies (Dubois et al., 2008b; Bastiani et al., 2019), and affirms that the results from baby-XTRACT can capture accurate anatomical information.

4.4.2 Assessing the Impact of the Gyral Bias on the Neonatal Connectome

Whole-brain connectivity matrices were used to explore the impact of the gyral bias in neonates and adults, and to compare two different seeding approaches: WGB seeding and whole-brain seeding.

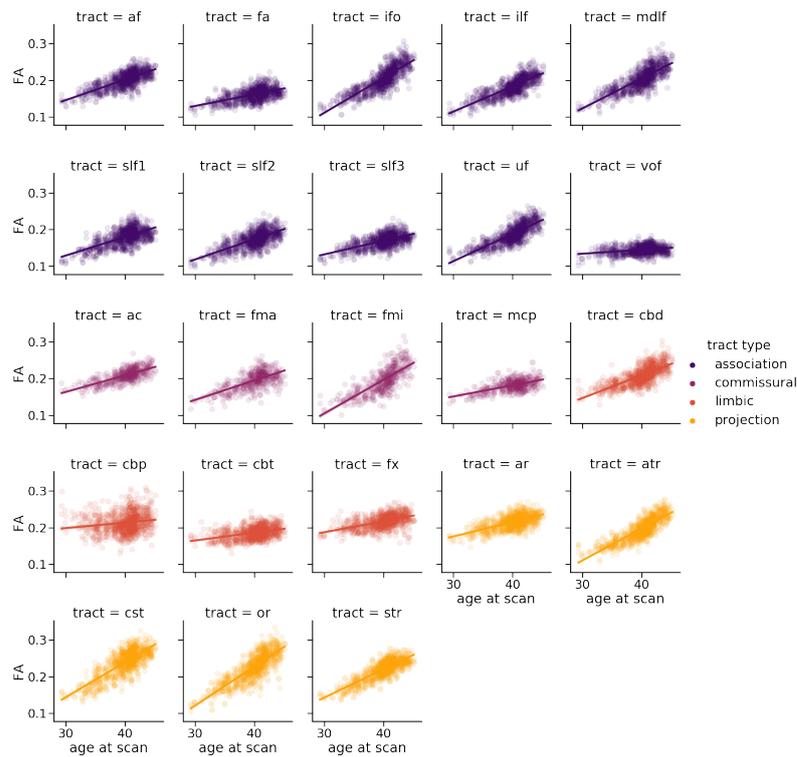


Figure 4.9: Scatter plots of median tract FA against subjects' age at scan (in weeks PMA), with linear regression line overlaid. Plots are colour coded according to tract type.

Predicted Bias

First, the level of anatomically-justifiable bias as predicted by cortical geometry alone was calculated as the cortical volume associated with each unit of surface area of the WGB. Figure 4.12 shows the average predicted bias, for the adult and neonatal subjects. The maps are normalised with respect to areas of zero curvature, and then log-scaled.

Axonal density is predicted to be higher in the gyri than the sulci, with the relative difference between them greater in adults than in neonates.

Measured Bias from Tractography

Connectivity matrices were generated for 36 adult and 36 neonatal subjects, seeding streamlines both from the WGB and across the whole brain. Each of these matrices were summed across their rows to give the streamline density at each vertex on the

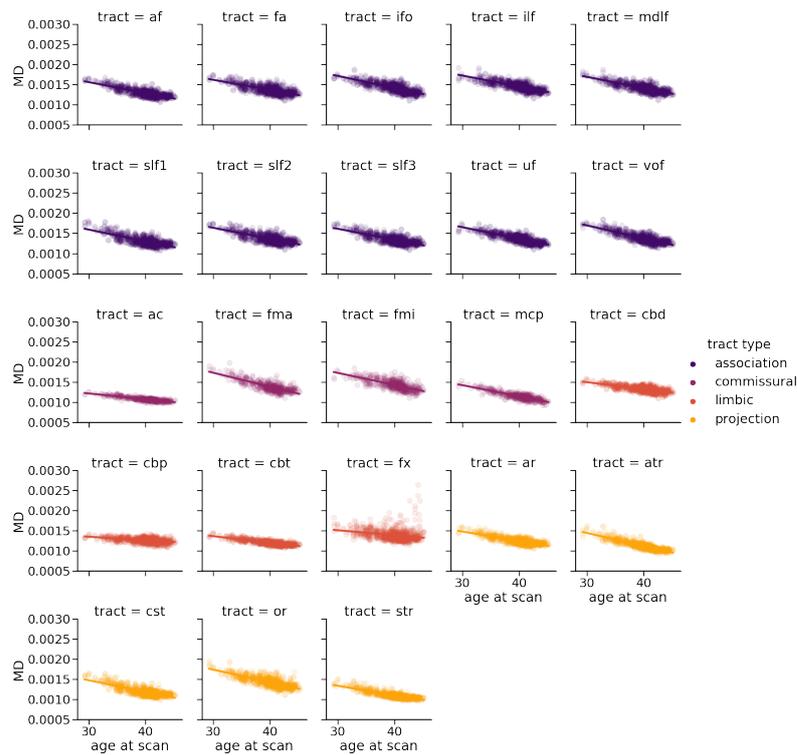


Figure 4.10: Scatter plots of median tract MD against subjects' age at scan (in weeks PMA), with linear regression line overlaid. Plots are colour coded according to tract type.

WGB, then normalised and log-scaled. The average maps are shown in figure 4.13. They indicate a much larger bias towards gyral terminations in streamlines from whole-brain seeding (note the different colour-scale on these maps).

To provide a more quantitative estimate of the gyral bias, the median predicted bias and streamline density were calculated within the gyri and sulci. Gyri were defined as regions in which the local curvature on the WGB exceeded the 66th percentile, and sulci as regions in which the local curvature was less than the 33rd percentile.

In general, as shown in figure 4.14, WGB seeding causes an overestimation of the axonal density in the sulci, whereas whole brain seeding causes an overestimation of the relative axonal density in the gyri. These effects are more pronounced in the adult brain, where gyrification is higher.

Spatial correlation was calculated between the predicted and measured density maps, to see how the spatial patterns compared between the empirical predictions

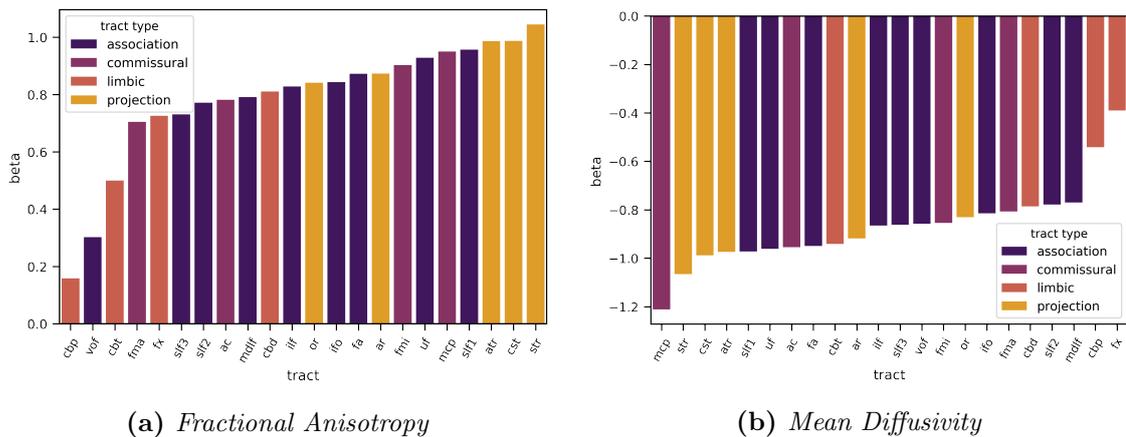


Figure 4.11: Bar-charts of beta values from the GLM analysis corresponding to age at scan. Bars are colour-coded according to tract type.

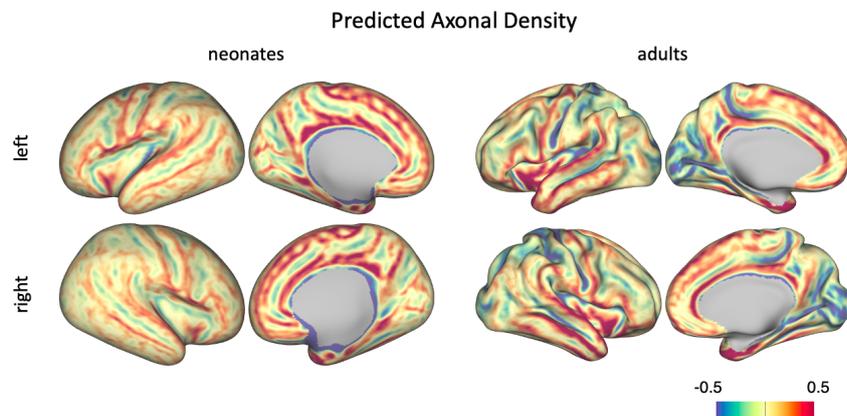


Figure 4.12: Log-scaled average predicted bias maps, based on the cortical volume per unit area of the WGB.

and the measured streamline density. The results of this analysis are shown in figure 4.15.

The results indicate that the spatial patterns from WGB seeding are much more closely aligned with prediction than the maps from whole brain seeding, in terms of their spatial pattern. Correlations are similar between predicted and measured density for WGB seeding in adults and neonates, but the matching between prediction and measurement is much less for whole brain seeding in neonates than in adults. These results indicate that the type of seeding has significant downstream effects for whole-brain connectome mapping in the neonatal brain, as in adults, and needs to be carefully considered. WGB seeding results to less gyral bias than whole-brain

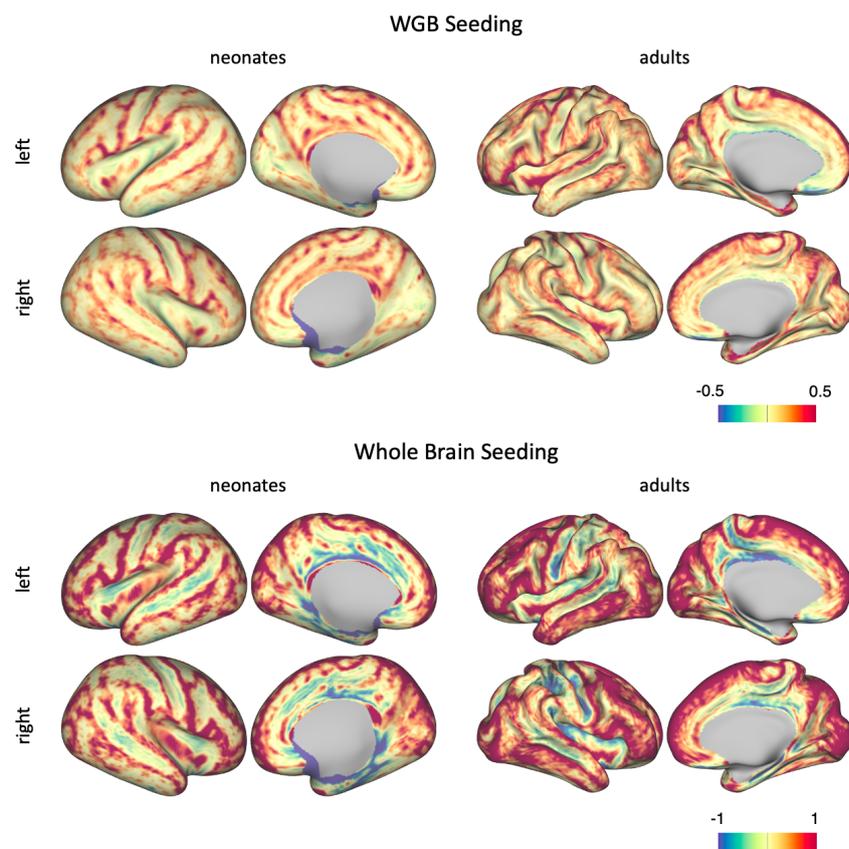


Figure 4.13: Average streamline density maps using different seeding approaches, for neonates and adults. Density values have been normalised relative to regions where curvature = 0 ± 0.025 , and log-scaled.

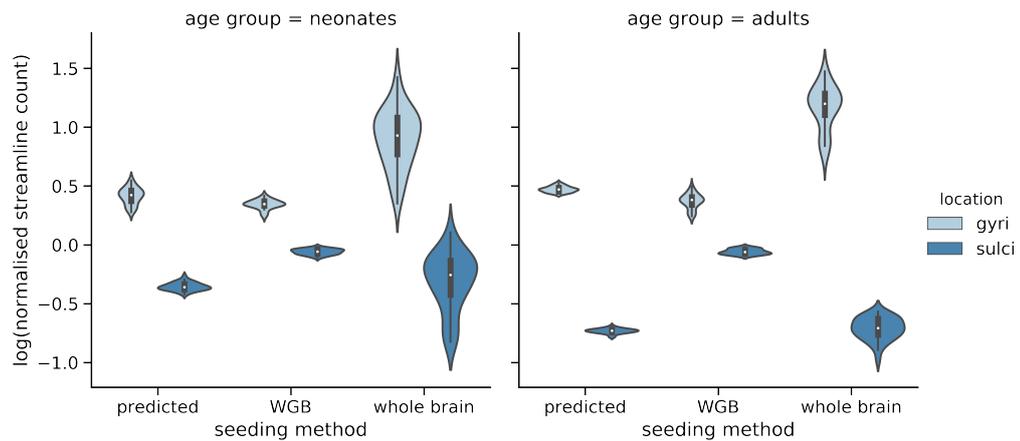


Figure 4.14: Violin plots to compare the predicted and measured streamline density in the gyri and sulci, for different seeding approaches

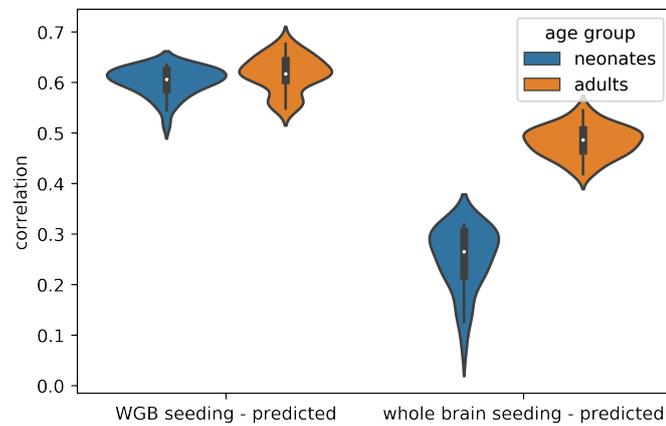


Figure 4.15: Violin plots of the correlation scores between the predicted and measured streamline density, for neonates and adults

seeding, and this bias is smaller in magnitude with WGB seeding in neonates than it is in adults.

4.5 Discussions

In this chapter, two different approaches have been explored for mapping white matter connections in the neonatal brain: virtual dissection and whole brain connectomes. These methods offer different and complementary insights into structural connectivity.

4.5.1 Standardised Protocols for Neonatal Tractography

First, a set of tractography protocols were developed that facilitate the standardised mapping of 42 tracts in the neonatal brain. The results were consistent across subjects of different ages and provided similar results, for a subset of tracts, to a set of independently-defined protocols for neonates (Bastiani et al., 2019). A key benefit of our protocols is that they are matched to the existing XTRACT protocols for the adult human and the macaque (Warrington et al., 2020). This provides a valuable resource for the comparison of homologous white matter tracts across species and across the lifespan.

These protocols also provide a much more extensive set of tracts than has previously been available for this age-group. Neonatal tractography studies have tended to focus on the cortico-spinal tract, the corpus callosum, and the thalamic radiations, while other pathways have been studied less frequently (Zöllei et al., 2019). To the best of our knowledge, the three branches of the SLF have not been identified in neonates before, so these protocols provide a unique opportunity to study the early development of these pathways, which are important for visuospatial processing (Thiebaut de Schotten et al., 2011a).

Microstructural metrics from DTI were used to compare the maturational trajectories of the different tracts and also to provide indirect evidence on the fidelity of the reconstructed tracts. Different fibre bundles are expected to have asynchronous patterns of maturation, with the limbic tracts showing the slowest rates of change, and the projection fibres the quickest (Dubois et al., 2008b). Using the baby-XTRACT protocols we were able to replicate this expected behaviour. Furthermore, extracting microstructure information using the corresponding adult tracts would provide additional insight into the developmental trajectories of these tracts.

The protocols have been tested here on data from the dHCP, which is of very high quality. In future, it would be useful to test how well these protocols are able to map white matter tracts in more typical neonatal dMRI data, with lower resolution

and SNR.

4.5.2 Investigating the Impact of the Gyral Bias on the Neonatal Connectome

Tractography streamlines have been shown to preferentially terminate at the gyral crowns, rather than the sulcal walls or fundi, a limitation that is particularly problematic when using tractography to map cortical connections at a high resolution (Van Essen et al., 2013a; Schilling et al., 2018). This gyral bias has been characterised in adult brains (Van Essen et al., 2013a; Schilling et al., 2018), but has not been explored for newborn human brains. Characterising the gyral bias in neonates allows us to measure its impact on connectivity estimates in this population where gyrification is much less developed than in adults (Dubois et al., 2008a).

Connectivity matrices were generated using a surface registration method optimised for this age group (Robinson et al., 2018). The streamline density on the cortical surface was compared for different seeding strategies, and between adults and neonates. The bias in streamline termination towards the gyri was found to be greater in adults, which is expected from their increased gyrification. Adults also have more developed superficial white matter systems, which have been shown to impede accurate detection of cortico-cortical connections with tractography (Reveley et al., 2015). Seeding streamlines from the WGB was found to reduce the impact of the gyral bias, compared to whole brain seeding, which is in agreement with previous work (Van Essen et al., 2013a; Schilling et al., 2018). In addition, a quantitative analysis of streamline density in different regions revealed that the gyral bias from whole brain seeding is the result of an overestimation of connections to the gyri, whereas WGB seeding overestimates the connections in the sulci.

One of the main reasons proposed for the gyral bias is that streamlines entering a gyral blade have to make sharp angular turns if they are to terminate at the sulcal

wall, which is penalised by the curvature thresholds in most tractography algorithms (Van Essen et al., 2013a). Improving tractography techniques to mitigate the effects of the gyral bias is an active area of research. The use of asymmetric fODFs has been shown to generate streamlines that can curve more sharply to terminate at the sulcal wall, compared to conventional fODFs (Y. Wu et al., 2020). Surface enhanced tractography is another method that has been shown to reduce the gyral bias in tractography (St-Onge et al., 2018). It uses a surface flow from the WGB mesh to model the superficial white matter streamlines, which leads a more realistic trajectory near the cortex. Nevertheless, our work suggests that the gyral bias is an issue of smaller magnitude for neonatal tractography compared to the adult counterpart (Figure 4.14), particularly when WGB seeding is used.

4.6 Concluding Remarks

Two contrasting approaches were explored for using tractography to map white matter connections in neonates. First, a set of standardised tractography protocols were developed that provide a reliable mapping of 42 tracts in the neonatal brain. These protocols were used to investigate microstructural maturation of these bundles, and the results were in agreement with previous literature. An advantage of these protocols is that they are consistently defined in line with protocols for the adult human and macaque from the XTRACT framework (Warrington et al., 2020). This resource can be used in future to investigate the development of structural connectivity, both in terms of ontogeny (development across the lifespan) and phylogeny (evolutionary development).

Whole-brain connectivity matrices were then used to explore the impact of the gyral bias in neonates. The gyral bias is a methodological confound that leads to tractography streamlines terminating preferentially at the gyral crowns, rather than the sulcal walls or fundi. It was found that the gyral bias is less prevalent in neonates

than in adults, and that its impact can be reduced by seeding streamlines from the white matter/grey matter boundary, rather than from the brain volume. These connectivity matrices form the basis of the work in the next chapter, in which a new framework for data-driven tractography is developed.

4.7 Appendix

Specific Protocol Differences Between Baby-XTRACT and AutoPtx

Below, specific differences in the protocols for the acoustic radiation, cingulum (dorsal segment), fornix, and optic radiation are detailed, which give rise to the discrepancies observed in these tracts.

Acoustic Radiation

Both protocols seed the AR from the transverse temporal gyrus. However, the target in the baby-XTRACT protocol is the medial geniculate nucleus of the thalamus, whereas the target in the AutoPtx protocol covers the entire thalamus.

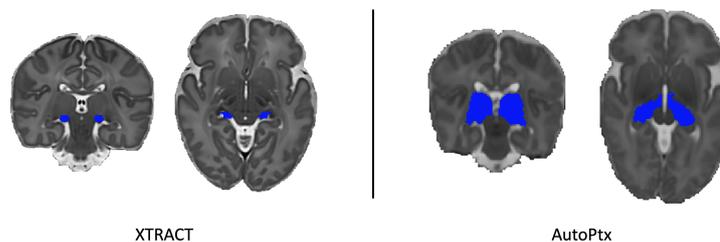


Figure 4.16: Comparison of baby-XTRACT and AutoPtx target masks for the acoustic radiation, shown in blue.

Cingulum Bundle - Dorsal Segment

Seeds, targets and exclusion masks are the same in both protocols. However, AutoPtx uses additional stop masks to constrain the tractography, consisting of a ring

surrounding the anterior section of the tract, a coronal plane posterior to the thalamus and an axial plane inferior to the anterior part of the tract.

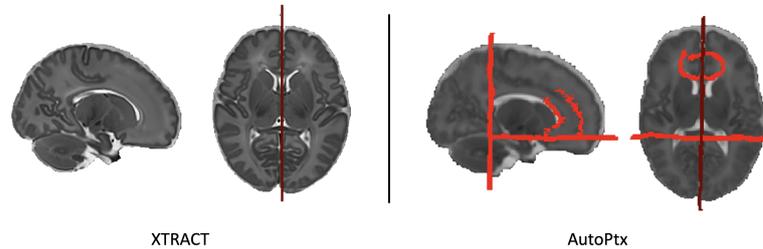


Figure 4.17: Comparison of the exclusion and stop masks for the cingulum protocol from baby-XTRACT and AutoPtx. Exclusion masks are shown in dark red, stop masks in red.

Fornix

The baby-XTRACT seed is in the body of the fornix at the level of the middle of the corpus callosum, whereas the AutoPtx seed is larger and more anterior, lying between the ventricles. The baby-XTRACT target is in the hippocampus and the AutoPtx target is located in the crus of the fornix.

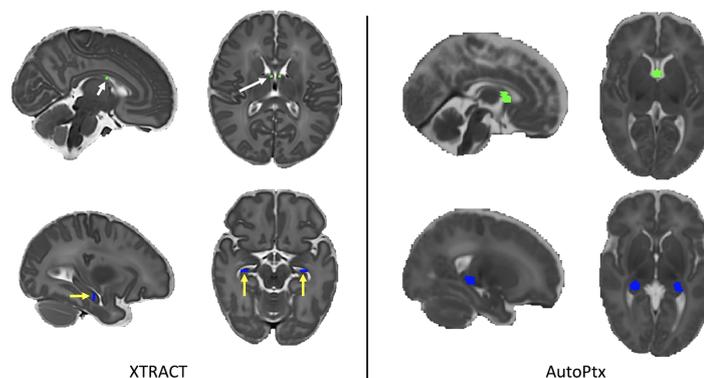


Figure 4.18: Comparison of baby-XTRACT and AutoPtx seed and target masks for the fornix. The seed masks are shown on the top row, in green (white arrows on baby-XTRACT), while the target masks are shown on the bottom row, in blue (yellow arrows on baby-XTRACT).

Optic Radiation

Similarly to the acoustic radiation, the baby-XTRACT protocol has a more specific thalamic mask than AutoPtx: the seed is in the lateral geniculate nucleus of the

thalamus, whereas the AutoPtx seed is a coronal plane through the centre of the thalamus. The target ROIs are both coronal planes through the anterior part of the calcarine fissure, however in the AutoPtx protocol this plane just extends through the occipital lobe, whereas in the baby-XTRACT protocol it covers the whole brain.

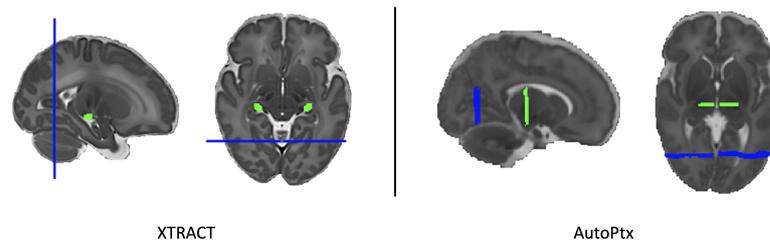


Figure 4.19: Comparison of baby-XTRACT and AutoPtx seed and target masks for the optic radiation. The seed masks are shown in green, the target masks are shown in blue.

Chapter 5

Using Data-Driven Methods To Map Structural Connectivity in the Neonatal Brain

Contents

5.1	Introduction	94
5.2	Theory	96
5.2.1	Matrix Decomposition with ICA and NMF	96
5.3	Methods	99
5.3.1	Generating Connectivity Matrices	100
5.3.2	Dimensionality Reduction and Back-Projection	102
5.3.3	Simulations	102
5.3.4	In Vivo Data Decompositions	106
5.3.5	Comparison to Tractography-Derived White Matter Tracts	107
5.3.6	Comparison with fMRI	107
5.4	Results	109
5.4.1	Simulations	109
5.4.2	In Vivo Data Results	113
5.5	Discussion	127
5.5.1	Validation using Simulations	127
5.5.2	In Vivo Results	128
5.5.3	Decomposition Domain	130
5.6	Concluding Remarks	131

Foreword

In this chapter we present a novel, data-driven technique for mapping white matter connections in the brain, with no prior assumptions or constraints placed on their trajectories. Whole-brain connectivity matrices are decomposed into a set of components and a mixing matrix, which describe linked cortical regions and the white matter pathways that connect them. Two decomposition methods are compared, both in simulated data and neonatal dMRI data: Independent Component Analysis (ICA) and Non-negative Matrix Factorisation (NMF). We find that the non-negativity constraints in NMF provide a more accurate characterisation of the non-negative structural connectivity patterns found in the brain. When applied to neonatal dMRI data, the results from the framework are shown to accurately reproduce white matter pathways, and the grey matter components show similarity to resting-state networks from fMRI.

5.1 Introduction

The previous chapters have illustrated the benefits of the “virtual dissection” approach for mapping individual white matter bundles from neonatal diffusion MRI data, using standardised regions of interest as seeds and targets to constrain tractography. For example, the baby-XTRACT protocols that were developed in chapter 4 provide a robust set of 42 tracts with direct correspondence to similar protocols for the adult human and macaque (Warrington et al., 2020).

Here, we present an alternative approach for mapping whole-brain connections that does not require an a priori defined set of tracts and protocols, but extracts networks and their underlying connecting patterns in a data-driven manner. This paradigm can be particularly powerful when applied to the neonatal brain, where the traditional virtual dissection approach can be challenging. Neonatal brains are not simply small adult brains (Batalle et al., 2018). The rapid growth and changes in brain morphology during the neonatal period, as well as fast alterations in tissue composition that alter imaging contrast over time (Bastiani et al., 2019), render it challenging to ensure correspondence between template-driven ROIs and tractography protocols at different stages of development (Serag et al., 2012). Such protocols also focus on long range fibres, neglecting the complex network of short fibres that connect neighbouring gyri (Catani et al., 2012). They also require prior anatomical knowledge for a specific set of tracts, which given the dynamic and rapid changes that occur during this period are not straightforward to obtain. The protocols defined in the previous chapter assume that they convey generalisable information across the different stages of development. It is challenging to confirm how reasonable such an assumption is.

An alternative is to use data-driven methods to identify connections from whole-brain connectivity matrices, which are model-free and therefore expected to be more immune to the challenges described above. Different data-driven methods include

clustering (Garyfallidis et al., 2012; Siless et al., 2018; Ding et al., 2003; O’Donnell and Westin, 2007), graph-based methods (Bassett and Bullmore, 2006; Bullmore and Sporns, 2009), or matrix decompositions, which will be the focus of this chapter. Independent Component Analysis (ICA) has been a commonly used matrix decomposition technique for identifying brain networks from resting-state functional MRI (fMRI) data (Mckeown et al., 1998), and recent work has shown that it can also be applied to data from tractography. Its first use in this context was to cluster tractography data seeded from the thalamus, yielding a set of thalamic connections and a thalamic parcellation (O’Muircheartaigh et al., 2011). L. Wu et al. (2015) applied ICA to whole brain voxel-to-voxel connectivity matrices, resulting in a dual parcellation of spatially independent sources and their corresponding connectivity profiles. Although this method showed sensitivity to connectivity alterations in schizophrenia, the paired volumetric components are difficult to interpret. A different approach was used by O’Muircheartaigh and Jbabdi (2017); using grey matter-to-whole brain connectivity matrices, yielding pairs of components that correspond to white matter bundles and their cortical terminations. The resultant connectivity patterns showed strong spatial similarity to tracts from standard tractography protocols.

One limitation of applying ICA to tractography data is that the estimated independent components and the respective mixing matrix can contain both positive and negative values, whereas structural connectivity data are inherently non-negative. This leads to challenges in the interpretation of negative weights. To address this problem, an alternative data-driven method is presented in this chapter that can be used to identify non-negative connectivity components. This approach is based on Non-Negative Matrix Factorisation (NMF) (D. D. Lee and Seung, 1999). Like ICA, NMF is an unsupervised technique that estimates a pre-defined number of components from the data. However, the elements and their weights are constrained to take non-negative values. Sparsity constraints in the decomposition allow identifiability and further provide an indirect means of requiring independence between the

estimated components. This results in a set of components whose weighted summation represents the whole system. Due to these advantageous properties, NMF has been recently used to identify networks of structural covariance from MRI data (Ball et al., 2019; Sotiras et al., 2015; Sotiras et al., 2017).

In this chapter, an NMF-based framework is devised for extracting connectivity components from diffusion MRI data for the first time. This approach is used to map patterns of structural connections in new-born babies, aged 37 to 44 weeks post-menstrual age (PMA) at scan, using publicly-released data provided by the developing Human Connectome Project (dHCP) (Hughes et al., 2017; Hutter et al., 2018a). First, we describe the theory for decomposing whole-brain tractography-induced connectivity matrices into grey matter networks and their corresponding white matter bundles. Then, simulations are used to quantitatively evaluate the behaviour of the method and assess its performance against Independent Component Analysis (ICA). The validity and interpretability of the automatically detected white matter patterns are measured against results from standard tractography protocols (Bastiani et al., 2019), and the grey matter patterns against components obtained from data-driven mapping of resting-state fMRI in the same subjects.

5.2 Theory

5.2.1 Matrix Decomposition with ICA and NMF

Let \mathbf{X} be an $M \times N$ dense¹ connectivity matrix, with $X_{ij}\{i = 1 : M, j = 1 : N\}$ carrying information on the likelihood of structural connections existing between locations i and j in the brain (figure 5.1). Without loss of generality, let us assume that locations i represent the whole brain and comprise of all imaging voxels, and that locations j represent grey matter and reside on the cortical white/grey matter

¹By “dense” we refer to voxel-wise / vertex-wise representations rather than areal-wise nodes, i.e. N and M are in the order of thousands.

boundary (WGB) and in subcortical grey matter. Diffusion MRI tractography can provide such a matrix if we seed streamlines from N seeds on the WGB and subcortical nuclei, and record visitation counts to M voxels across the brain, such that each column of \mathbf{X} describes the connectivity profile of a grey matter location j . A data-driven decomposition of \mathbf{X} can identify K components based on similarity of connectivity profiles, such that a (linear) combination of these components return back \mathbf{X} . Different numbers of components can be obtained depending on the desired properties of the estimated components, and different methods can be used to estimate the mixing of the components to reconstruct the input matrix.

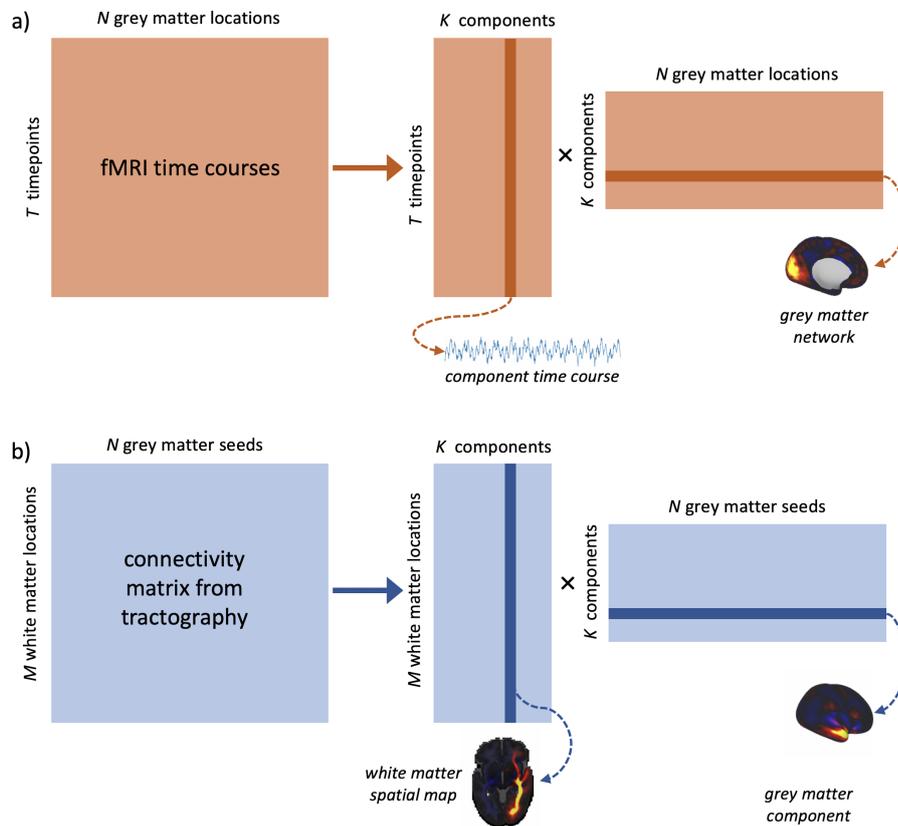


Figure 5.1: Data-driven matrix decomposition methods applied to resting-state functional MRI and structural connectivity data. a) N functional time-courses of length T are recorded from points in the grey matter. We can apply a matrix decomposition technique, such as ICA, to this matrix, yielding an $T \times K$ mixing matrix of time courses and a $K \times N$ matrix of spatial components. b) an $M \times N$ connectivity matrix describes the likelihood of structural connections existing between each of N grey matter seeds and M locations in the brain. The equivalent decomposition applied to this matrix gives us an $M \times K$ mixing matrix of spatial maps, and a $K \times N$ matrix of components in the grey matter.

Independent Component Analysis (ICA) imposes statistical independence between the components to perform a linear decomposition. An observed matrix \mathbf{X} is represented as $\mathbf{X} = \mathbf{A}\mathbf{S}$, where \mathbf{S} is the independent sources matrix (each row k corresponds to a component) and \mathbf{A} the weights or mixing matrix (each column k corresponds to the weights of source k). As this is an ill-posed problem in general, ICA uses source independence to estimate an un-mixing matrix \mathbf{W} , that best approximates \mathbf{A}^{-1} , to recover the original sources from the observed data: $\mathbf{W}\mathbf{X} \approx \mathbf{S}$. This process is entirely driven by the statistical properties of the mixture, with no prior knowledge of the mixing matrix or the signals. The first step of all ICA algorithms is to centre and whiten the data, for normalisation. This can be achieved with a Principal Component Analysis (PCA) (Wold et al., 1987) or singular value decomposition (SVD) (De Lathauwer et al., 2000). Then we seek an orthogonal rotation \mathbf{V} to apply to the whitened data to optimise the statistical independence of the estimated components. This cannot be done analytically but there are a number of different methods of solving the problem iteratively. The FastICA algorithm (Hyvärinen and Oja, 2000), which uses non-Gaussianity as a proxy for independence, is one of the typically used algorithms. ICA has been used to identify networks from resting-state fMRI data (McKeown et al., 1998), where $M = T$, the number of timepoints, and the decomposition results in K spatial maps (covering all N brain voxels), each with a weight vector of length T (figure 5.1a). Each weight w_{ik} represents how much component k contributes to activity recorded at time point i . ICA has also been used recently in the case of dMRI tractography, where N is the number of seed locations (O’Muircheartaigh and Jbabdi, 2017). In that case, the decomposition provides K spatial maps (covering all N points on the grey matter), each representing a component with shared connectivity profile through white matter, associated with a weight vector of length M . Each weight w_{ik} represents in this case how much component k contributes to the connection patterns of voxel i .

Non-Negative Matrix Factorisation is an alternative decomposition technique,

where a matrix \mathbf{X} is factorised into the product of two matrices \mathbf{A} and \mathbf{H} , under the constraint that all three contain only positive values (D. D. Lee and Seung, 1999). This is more naturally suited for use with structural connectivity data, which is inherently non-negative. In general, NMF is an ill-posed problem (as is ICA) and there exist multiple solutions in most cases. The linear superposition of components, combined with the non-negativity constraint, lead to an implicit sparsity constraint in the algorithm (requesting a signal to be explained as a linear combination of non-negative regressors will lead to many weights close to zero). Additional explicit sparsity constraints can be applied to further constrain the solution space and improve the identifiability of the decomposition (Hoyer, 2004). Specifically, the cost function C to minimise is of the form:

$$C = \frac{1}{2} \|\mathbf{X} - \mathbf{AH}\|_F + \alpha_1 \|\mathbf{A}\|_{L1} + \alpha_2 \|\mathbf{H}\|_{L1} \quad (5.1)$$

where $\|X\|_F$ is the Frobenius norm, $\|x\|_{L1}$ is the L1-norm, used to explicitly promote sparsity, and α_1 and α_2 are tuning parameters that allow us to control the degree of regularisation on the mixing matrix and component matrix, respectively. Higher α values lead to more sparsity in the resultant decomposition. The NMF can be initialised with a non-negative SVD, which has been shown to improve the accuracy of the decompositions (Boutsidis and Gallopoulos, 2008). Most NMF algorithms use a two-block coordinate descent approach to optimise C over \mathbf{A} and \mathbf{H} alternatively, while keeping the other fixed. Each block is a convex problem that can be solved using non-negative least squares (Cichocki and Phan, 2009).

5.3 Methods

The above frameworks were developed to map structural connectivity in the neonatal brain. First, an overview is given of the data employed. A set of simulations is then described that allows principled evaluation of the decomposition frameworks.

Finally, the methods used to illustrate the benefits of this approach are described.

Data and Preprocessing

Structural and dMRI data were used from the dHCP. Diffusion MRI data were acquired and preprocessed as outlined in the previous chapter. Cortical surface reconstruction was carried out from T2w structural images with an isotropic resolution of 0.5 mm^3 , using a pipeline specifically adapted for neonatal structural MRI data (Makropoulos et al., 2018). Data were considered from a group of 161 subjects born at term age (94 male, 67 female). Median (range) birth age was 40.1 (37.1, 42.1) postmenstrual weeks and age at scan 40.9 (37.4, 44.4) weeks.

Preprocessed data were further analysed to obtain structural connectivity matrices. To ensure alignment between subjects, the anatomical surfaces were registered to a representative template space before performing tractography. First, a surface registration pipeline was used, based on the multi-modal surface matching (MSM) algorithm (Robinson et al., 2014; Robinson et al., 2018). Cortical folding was used to drive the alignment of neonatal WGB, cortical mid-thickness, and pial surfaces to the dHCP 40-week PMA surface templates (Bozek et al., 2018). This aligned vertices on the WGB surface to ensure consistent seed points for tractography across subjects. A previously computed non-linear volumetric registration (Avants et al., 2011) was then applied to all MSM-derived surfaces to register them to 40-week PMA volumetric template space (Serag et al., 2012). This step was necessary to ensure that the tractography seeds were aligned to the target space, because the volumetric and surface-based neonatal templates are not aligned (Bozek et al., 2018; Serag et al., 2012).

5.3.1 Generating Connectivity Matrices

Once the surfaces were aligned, connectivity matrices \mathbf{X} were obtained for each subject, by performing whole-brain probabilistic tractography using FSL (Behrens et

al., 2007; Hernandez-Fernandez et al., 2019). Fibre orientations (up to 3 per voxel) were estimated using a model-based deconvolution against a zeppelin response kernel, to accommodate for the low anisotropy inherent in data from this age group (Bastiani et al., 2019; Hernandez-Fernandez et al., 2013; Sotiropoulos et al., 2016) (also, see figure 2.9). Subsequently, 10,000 streamlines were seeded from each of 58,551 vertices on the WGB of both hemispheres (average vertex spacing 1.2 mm, excluding the medial wall) and from each of 2548 subcortical 2 mm³ voxels (bilateral amygdala, caudate, hippocampus, putamen and thalamus), giving a total of $N = 61,099$ seeds. WGB seeding was used as this had been shown to reduce the impact of the gyral bias on connectome estimation in this agegroup (see chapter 4). Visitation counts were recorded between each seed point and each of $M = 50,272$ voxels in a whole-brain mask with the ventricles removed, down-sampled to 2 mm³. The pial surface was used as a termination mask to prevent streamlines from crossing between gyri, and streamlines were not allowed to cross the WGB more than twice (once at the seed point and again at termination), to reduce false positives (Hernandez-Fernandez et al., 2019; R. E. Smith et al., 2012). All masks and surfaces (seeds, targets, exclusions) were defined in a 40 week PMA volumetric template space (Serag et al., 2012), however tractography was carried out in native space with results resampled directly to template space. Visitation counts were multiplied by the length of the pathway to correct for compound uncertainty in the estimated trajectories (O’Muircheartaigh and Jbabdi, 2017). The resulting dense matrices describe the likelihood of a white matter connection between each grey matter seed and the rest of the brain. The connectivity matrices were normalised by the total number of viable streamlines in each subject before being averaged across the group. Connectivity matrices were saved and averaged in a sparse format to reduce computation time and memory requirements.

5.3.2 Dimensionality Reduction and Back-Projection

Data-driven connection mapping was evaluated using both ICA and NMF. Large M (i.e. a large number of imaging voxels) can pose computational and numerical convergence challenges for ICA. Therefore, PCA was used to reduce the $M \times N$ matrix \mathbf{X} , into a $P \times N$ matrix \mathbf{X}_r of principal components. Applying ICA to this reduced matrix results in a $K \times N$ set of components \mathbf{S} , and a $P \times K$ mixing matrix in PCA space \mathbf{A}_r . In order to obtain the mixing matrix in the original space of M imaging voxels, the pseudoinverse of the component matrix \mathbf{S} was obtained and projected back onto the original data to obtain the tract space mixing matrix, i.e. $\mathbf{A} = \mathbf{X}\mathbf{S}^\dagger$, where \mathbf{S}^\dagger denotes the pseudoinverse of \mathbf{S} . This approach was followed for the ICA analysis in both simulations and on real data, to cope with excessive memory requirements of decomposing the full connectivity matrix. However, the dimensionality reduction step was not necessary for NMF.

5.3.3 Simulations

The decomposition frameworks were evaluated (using NMF and ICA) in numerically simulated data, before applying them to real data. We simulated datasets with a known number of underlying sources, to observe how the behaviour of the decompositions over different model orders reflects the true dimensionality of the data. The purpose of the simulations was to compare ICA and NMF in a situation where the underlying sources in the data are known, and can be compared directly to the derived components. Furthermore, the simulations allow us to analyse the impact of different noise levels in the data, and adjusting the regularisation term in NMF, where there is a known ground truth to validate the results.

To find a realistic generative distribution to use for our sources, the spatial maps from standard tractography protocols in the neonatal brain (Bastiani et al., 2019) were used to generate connectivity blueprints (Mars et al., 2018) as proxies

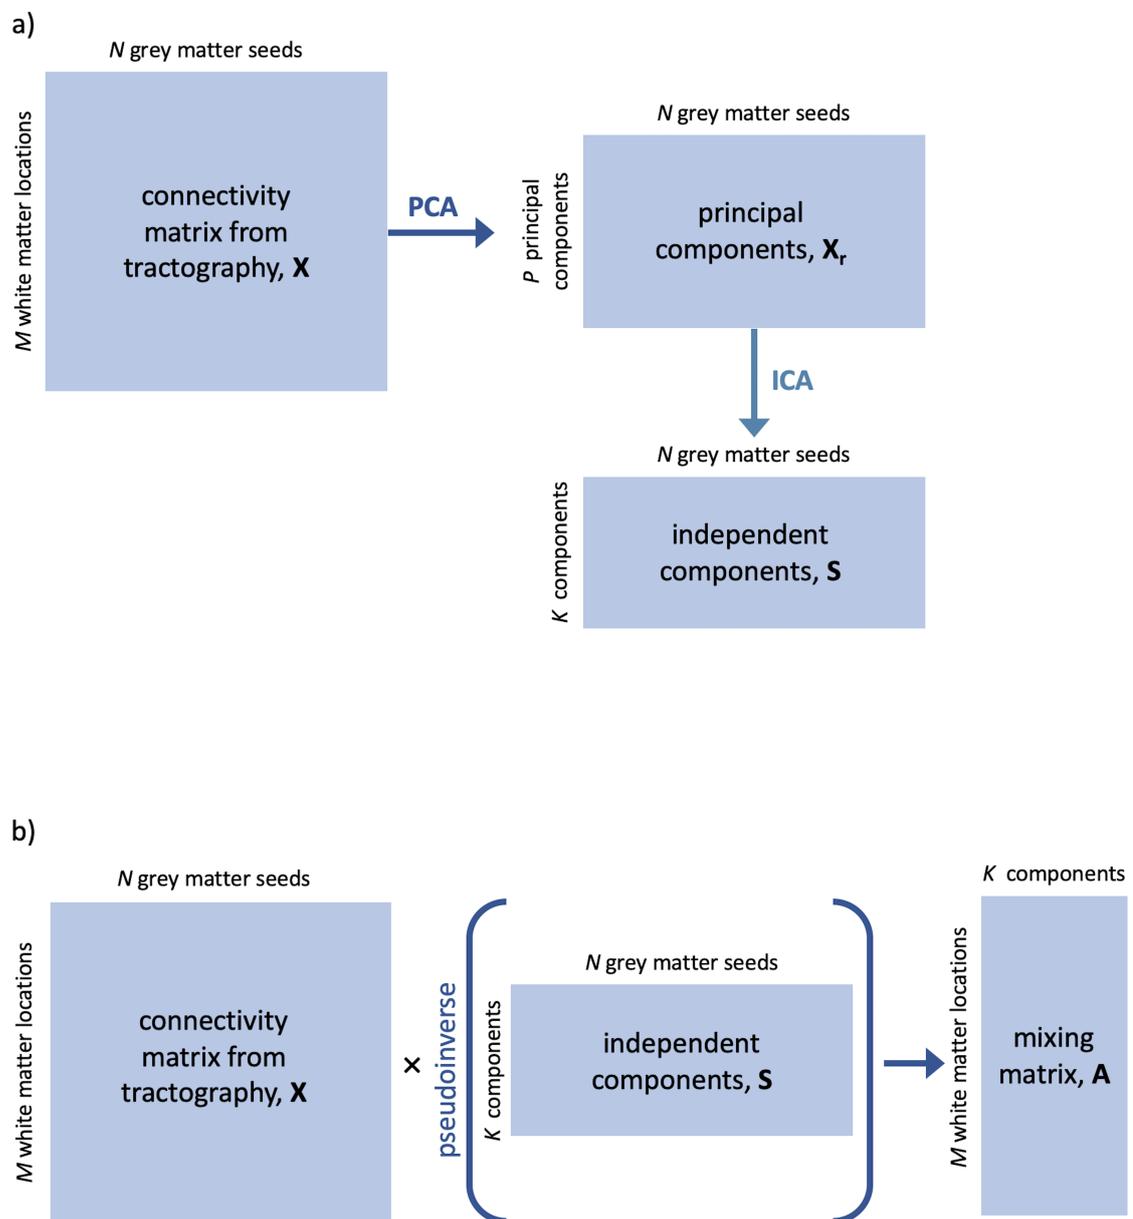


Figure 5.2: Illustration of a back projection step to obtain a white matter mixing matrix after data reduction by PCA. a) Data are first reduced to P principal components, and ICA applied to the data in the reduced subspace. b) Independent components are regressed onto the original data to obtain a mixing matrix in white matter space.

for the source spatial maps in grey matter space. Several distributions were fit to the intensities of these maps (unwrapped to 1D). It was found that log-beta distributions best described the data. The sources were therefore drawn from log-beta distributions, whose parameters in turn were drawn from Gaussian distributions according to the fits to the measured data. These sources are random and sparse, features that indirectly ensure a high degree of independence. Sources were scaled to lie in the range 0 - 1. The mixing matrix was randomly generated, normalised so the columns sum squared to 1. The simulated data, \mathbf{X} were calculated as the product of the mixing matrix \mathbf{A} with the source matrix \mathbf{S} . Zero-mean, additive Gaussian noise was applied to that product via a logit transform, to maintain non-negativity.

Varying L1-norm Regularisation in NMF

The NMF decomposition can be regularised with L1-norm terms to promote sparsity in the components (see equation 5.1) (Févotte and Idier, 2011). NMF was first tested on the simulated data with varying levels of regularisation to assess the effect of regularisation on the accuracy and robustness of the decomposition. Data were simulated with $K = 50$ sources, and overall dimensions of $N = 1200$ and $M = 1000$, with noise added with $\sigma^2 = 0.05$ to best match the real data. The same regularisation parameter was used for the mixing matrix and the components, i.e. $\alpha_1 = \alpha_2 = \alpha$, following the implementation in scikit-learn (Pedregosa et al., 2011). NMF was applied with model orders from 1 to 100 and with regularisation parameters, $\alpha = 0, 0.1, 0.25, 0.5$. This process was repeated with 100 noisy realisations of the data in each case.

Varying Number of Sources

The simulations were performed with varying number of sources in the data to check how this affects the results. The data were generated with $\sigma^2 = 0.05$ and with $K = 25, 50$ and 75 sources. ICA and NMF were applied with model orders from

1 to 100. For NMF, a regularisation parameter of $\alpha = 0.1$ was used (see Results for justification). This was also repeated 100 times. ICA was first initialised with a PCA with $P = 100$ components, as described in the Dimensionality Reduction section.

Varying Noise Levels

Finally, the decompositions were tested on data with varying noise levels. Data were simulated as above. Gaussian noise was added to the data with varying $\sigma^2 = 0.0005, 0.005, 0.05,$ and 0.5 . 100 noisy realisations were generated in each case. The data were decomposed with ICA and NMF, with model orders K from 1 to 100. ICA was first initialised with a PCA with $P = 100$ components, as above.

Assessing Performance

Three different metrics were used to assess the success of the decompositions on the simulated data: i) Reconstruction error: the sum of squared errors between the reconstructed data after decomposition and the original data: i.e. $\Sigma(\mathbf{X} - \mathbf{AS})^2$. This provides a measure of the information lost through the decomposition. ii) Source-component correlation: the correlation between each original source and the estimated components. The best-matched component to each source was identified and the mean of the maximum correlation values for each component was considered. This describes how well the decompositions have characterised the underlying signals in the data, and is sensitive to overfitting, as redundant components that are not well matched to sources will bring the value down. iii) Sparsity: Following the approach in (Hoyer, 2004; Sotiras et al., 2015), we used a sparsity measure for the derived components based on the relationship between the L1-norm and the L2-norm:

$$\text{sparsity}(x) = \frac{\sqrt{N} - \sum |x_i| / \sqrt{\Sigma x_i^2}}{\sqrt{N} - 1} \quad (5.2)$$

This returns values between 0 and 1, with 1 signifying a maximally sparse component with only one non-zero element. This was calculated for each component vector in \mathbf{S} , and the mean value across all components is reported. Sparse components are desirable because they provide an easily interpretable representation of the data with minimal redundant information. In the case of NMF, sparsity constraints also make results more reproducible, by constraining the solution space. For ICA, sparsity can be thought of as a proxy for independence.

5.3.4 In Vivo Data Decompositions

For real data, the group-average tractography matrices were decomposed using ICA and NMF, with a range of model orders K . ICA was initialised with regular PCA, in which the first 500 components were retained (explaining 97% of the total variance). ICA was applied to the reduced dataset using the FastICA algorithm (Hyvärinen and Oja, 2000), with independence imposed in the seed domain. The pseudo-inverse of this matrix was projected back onto the group-level connectivity matrix to yield the corresponding components in target space. To deal with the sign ambiguity of ICA, components that were negative in the long tail of their distribution were sign-flipped, for consistency with the other methods (i.e. so that the main mass of the distribution was in the positive valued domain). NMF was performed with a coordinate descent algorithm (Cichocki and Phan, 2009), a Frobenius norm cost function (see equation 5.1), and an L1-norm regularisation parameter $\alpha = 0.1$. In NMF, the matrix is decomposed directly into the $M \times K$ mixing matrix and the $K \times N$ component matrix so there is no need for the back-projection step that was carried out for ICA after the PCA. All decompositions were implemented using scikit learn (Pedregosa et al., 2011). An NMF decomposition on a group average matrix takes around 2 hours and 80 GB of RAM on a single CPU.

5.3.5 Comparison to Tractography-Derived White Matter Tracts

To assess the validity and interpretability of the extracted components, the automatically extracted white matter connectivity patterns were compared with results obtained from standard, template-driven tractography protocols, developed for neonatal subjects, as described in (Bastiani et al., 2019). 28 tracts (13 bilateral) were mapped in each subject. The tracts included in this analysis were: acoustic radiation (AR), anterior thalamic radiation (ATR), cingulate gyrus part of cingulum (CGC), parahippocampal part of cingulum (CGH), cortico-spinal tract (CST), forceps minor (FMI), forceps major (FMA), fornix (FOR), inferior fronto-occipital fasciculus (IFO), inferior longitudinal fasciculus (ILF), medial lemniscus (ML), posterior thalamic radiation (PTR), superior longitudinal fasciculus (SLF), superior thalamic radiation (STR), and uncinate fasciculus (UNC). These were registered to a 40-week PMA template (Serag et al., 2012) and down sampled to 2 mm for comparison with the tract space representations of the data-driven components. Each tract was averaged across all subjects, and the Pearson’s correlation coefficient was calculated between each of the average tracts with each of the data-driven components from the $K = 100$ decompositions. A one-to-one matching was performed between the standard tractography results and the component maps, based on the correlation scores.

5.3.6 Comparison with fMRI

The cortical patterns of structural connectivity from the NMF components were compared with resting state networks from fMRI. For this analysis, a group of 55 subjects were selected, all born and scanned between 40 weeks and 41 weeks PMA (i.e. all subjects within this age range who had both structural and functional data available). The dHCP’s resting-state functional MRI data were acquired for

15 minutes (TE/TR = 38/392 ms, 2300 volumes) with an acquired resolution of 2.15 mm isotropic. fMRI preprocessing was carried out as detailed in (Fitzgibbon et al., 2019), with an automated pipeline including fieldmap preprocessing to estimate susceptibility distortion; registration steps; susceptibility and motion correction; and denoising with ICA-FIX.

The functional data were first mapped onto the cortical surface, broadly following the fMRISurface pipeline outlined in (Glasser et al., 2013). The native WGB, midthickness and pial surfaces were affine registered to the same space as the functional data. The fMRI timeseries were then mapped onto the cortical surface using a partial volume weighted, ribbon-constrained volume to surface mapping algorithm, as implemented in HCP’s connectome workbench (Marcus et al., 2011). These data were then downsampled from the native mesh and registered to the 32k resolution template (using the same MSM transform as for the WGB surface used to seed tractography). Spatial smoothing was applied over the cortical surface with a Gaussian kernel, with FWHM = 2 mm.

Temporally-concatenated group-ICA was performed on the functional data using FSL’s Melodic (Beckmann and S. M. Smith, 2004), with Melodic’s Incremental Group PCA (MIGP) for the PCA step (S. M. Smith et al., 2014). MIGP uses an incremental approach to closely approximate PCA of very large datasets but with a reduction in the amount of memory required. 50 independent components were specified. NMF was performed on the group-averaged structural connectivity matrices of the same group of subjects, with $K = 50$, for comparison. The similarity between the resultant grey matter spatial maps was assessed using Pearson’s correlation coefficient.

5.4 Results

5.4.1 Simulations

Simulations were performed to evaluate the performance of ICA and NMF decompositions on a synthetic dataset in which the underlying sources were known. First, the degree of L1-norm regularisation in the NMF cost function was varied to find an optimal value to use for later experiments. The effect of the number of sources and the noise level in the data were then examined.

Varying L1-norm regularisation

Increasing the regularisation parameter α , increases sparsity, but also increases the reconstruction error, as shown in figure 5.3. The NMF decomposition breaks down for high regularisation ($\alpha = 0.5$), with high error and very low source-component correlation. Smaller amounts of regularisation improve the agreement between the components and sources and reduce the reconstruction error at the cost of reducing sparsity. A good middle-ground solution is shown ($\alpha = 0.1$), balancing reconstruction accuracy and sparsity. A regularisation level of $\alpha = 0.1$ was therefore used for subsequent experiments.

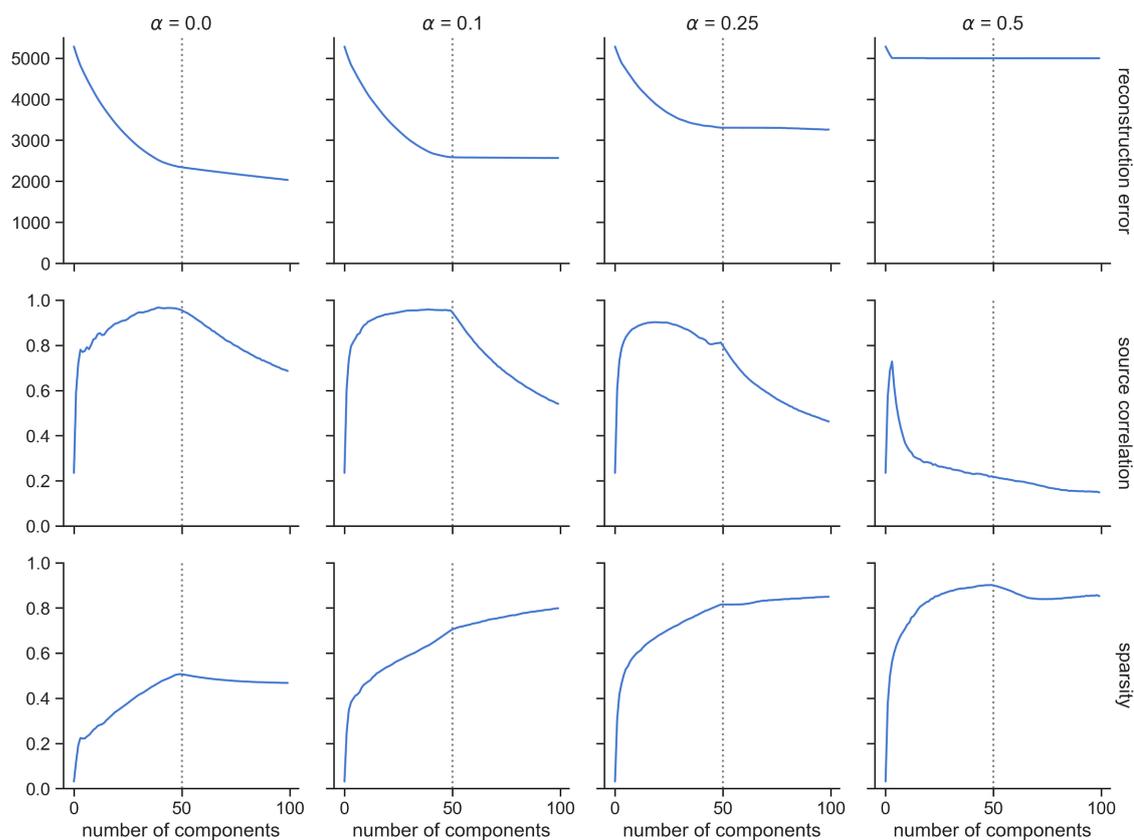


Figure 5.3: Simulation experiment to assess the effect of $L1$ -norm regularisation on NMF. The degree of regularisation increases from left to right across the plots ($\alpha = 0.0, 0.1, 0.25, 0.5$). Different metrics are shown from top to bottom: reconstruction error, correlation between derived components and underlying sources, sparsity of components. The true number of underlying sources ($K = 50$) is denoted by a vertical dashed line. Noise variance was $\sigma^2 = 0.05$. Results are shown averaged over 100 noisy realisations of the data.

Varying numbers of sources

The decompositions were carried out on data with varying numbers of underlying sources. Figure 5.4 shows that reconstruction error increases with the number of sources, so more information is lost between the decomposition and the original data as the data become more complex. For the source-component correlation, there are two different regimes. When the number of components, N , is lower than the true number of sources in the data, K , the average correlation between the components and the true sources rises quickly for very low N , then plateaus until $N = K$. When $N > K$, the extra components overfit to the noise and bring down the average

correlation with the sources. NMF achieves overall very high correlations between the reconstructed components and the true non-negative sources. NMF component sparsity increases rapidly for low N , then increases more slowly once the number of components exceeds the number of sources. In the case of ICA, sparsity reaches a peak when the number of components is equal to the number of underlying sources, then decreases.

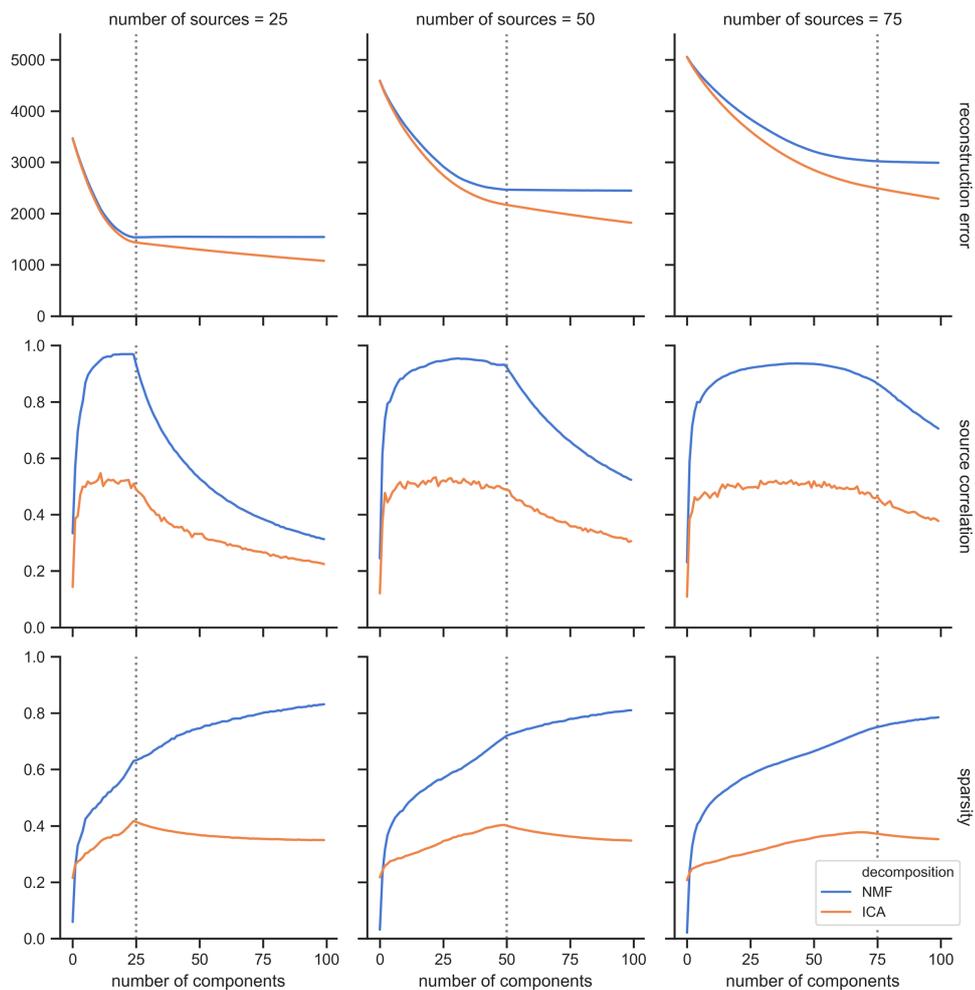


Figure 5.4: Simulation results to show how decompositions vary with differing numbers of underlying sources. The dotted vertical line shows the number of underlying sources in each case (from left to right: $K = 25, 50, 75$). Results are shown from ICA and NMF decompositions, in orange and blue, respectively. The noise variance $\sigma^2 = 0.05$ and the regularisation parameter for NMF $\alpha = 0.1$.

Varying SNR

Results from data with varying SNR are shown in figure 5.5. Overall, reconstruction error increases with noise level. In general, reconstruction error decreases as the model order approaches K , the true number of underlying sources, and then plateaus for higher model orders. The mean correlation between the components and the underlying non-negative sources increases as the number of components approaches K , and then decreases as the models overfit to noise. The sparsity of the components exhibits a relatively stable pattern for low and mid-levels of noise, but it becomes considerably reduced in the high noise scenario ($\sigma^2 = 0.5$).

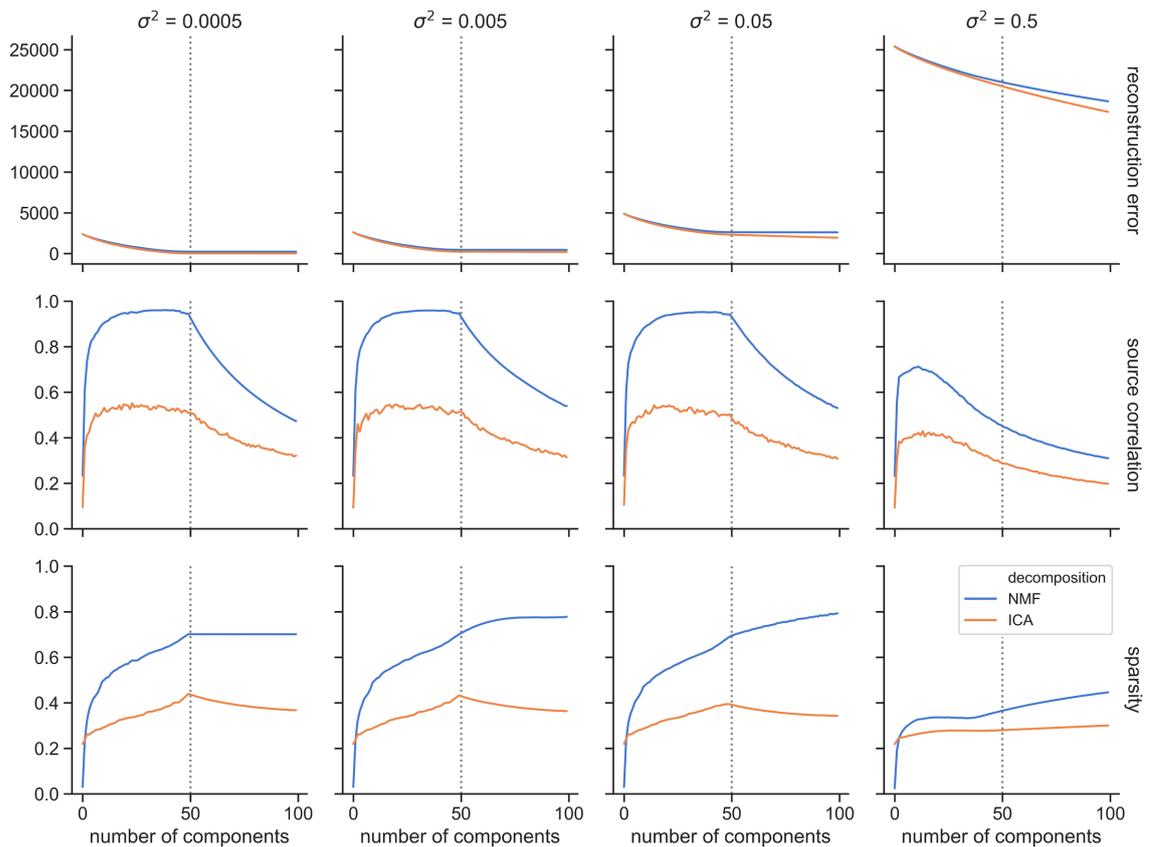


Figure 5.5: Simulation results to assess the effect of varying noise levels on the ICA (orange) and NMF (blue) decompositions. The noise level increases from left to right across the plots ($\sigma^2 = 0.0005, 0.005, 0.05, 0.5$). Different metrics are shown from top to bottom: reconstruction error, correlation between derived components and underlying sources, sparsity of components. The true number of underlying sources ($K = 50$) is denoted by a vertical dashed line.

Figures 5.4 and 5.5 also enable us to compare the performances of ICA and NMF on simulated, non-negative data. ICA has a consistently lower reconstruction error than NMF, particularly when model order exceeds the number of true sources. This could, however, signify that ICA is overfitting to noise more than NMF, particularly since ICA also exhibits a lower correlation between its components and the underlying sources than NMF, at all model orders. This reflects the better suitability of NMF for identifying inherently non-negative patterns within the data, in contrast to ICA, which generates components that contain both positive and negative values. NMF also generates components with consistently higher sparsity than those from ICA.

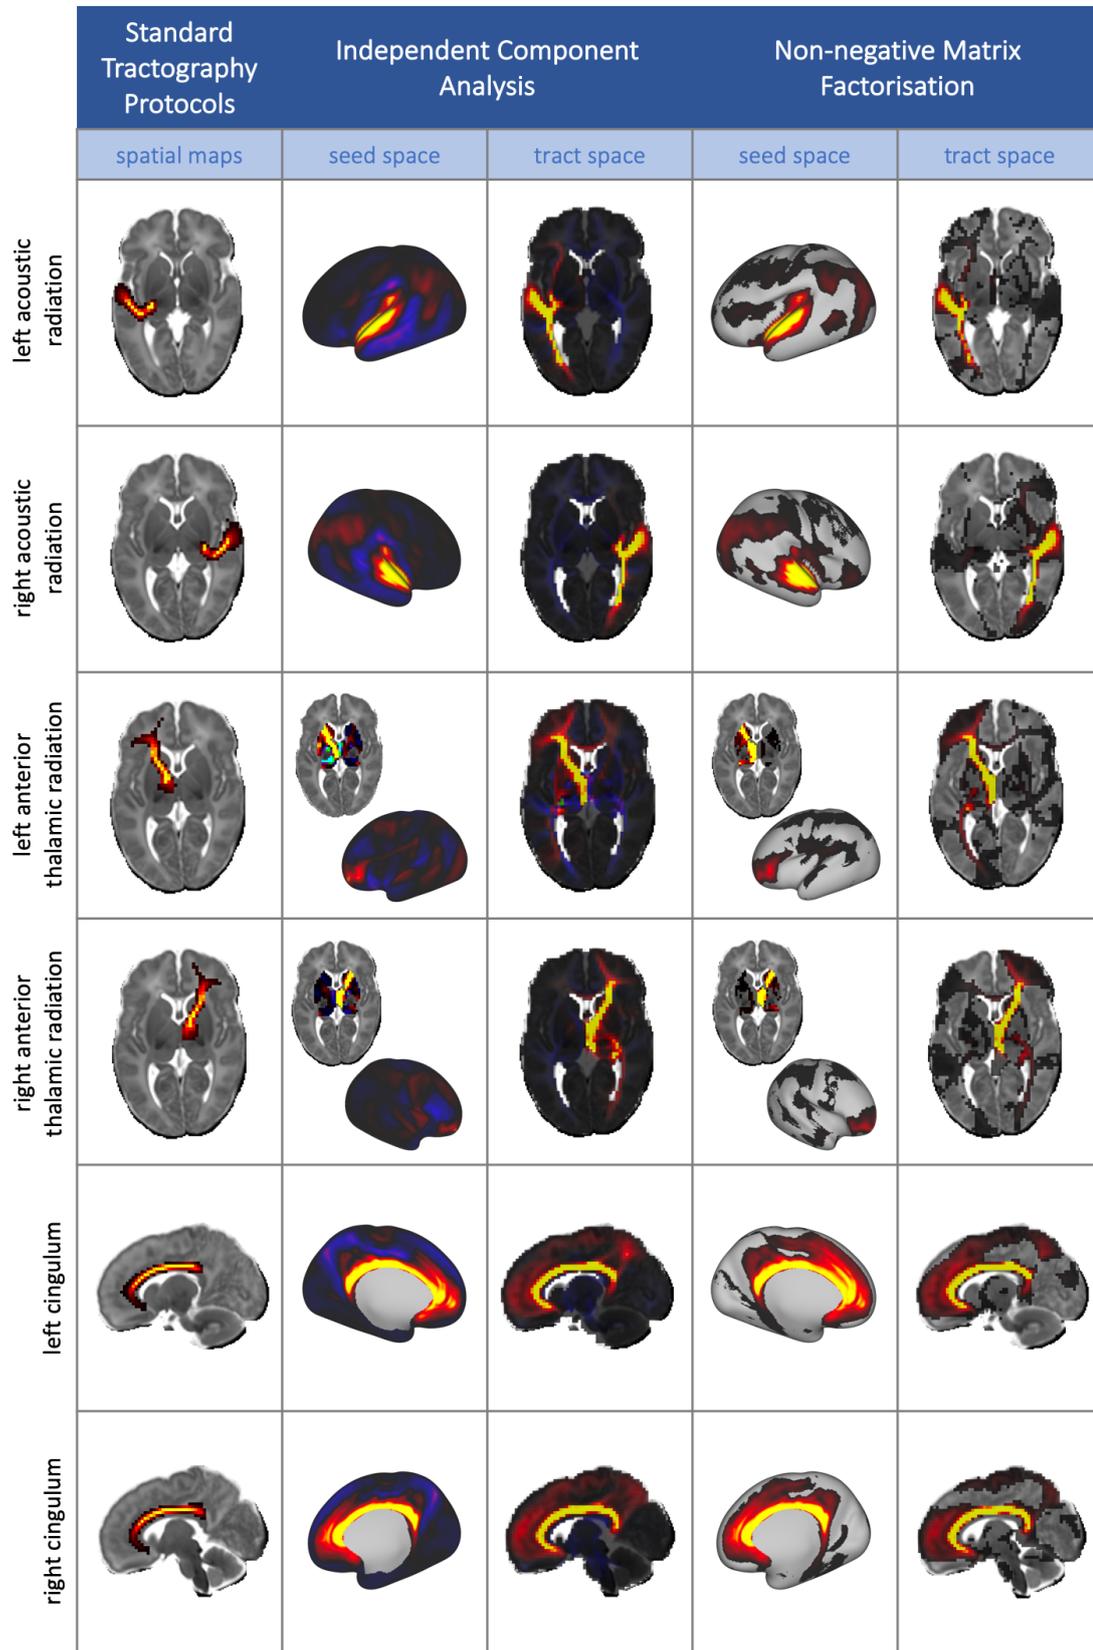
To summarise, the performance of ICA and NMF were evaluated on a simulated dataset with non-negative sources. Based on the results of these simulations, an L1-norm regularisation parameter of $\alpha = 0.1$ was chosen for NMF to use on the real data, as this promotes sparsity in the components, without compromising too much accuracy in the reconstruction. NMF has a number of advantages over ICA for non-negative data: it generates components that are more closely matched to the real sources, with higher sparsity and potentially less overfitting to noise.

5.4.2 In Vivo Data Results

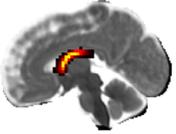
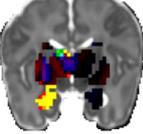
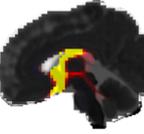
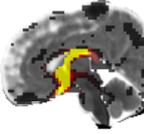
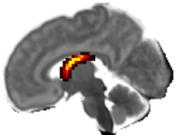
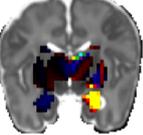
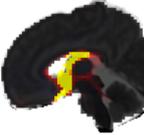
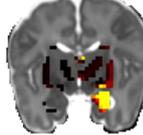
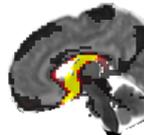
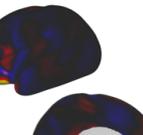
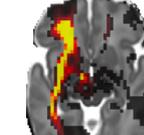
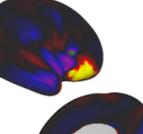
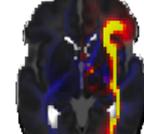
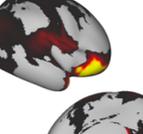
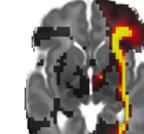
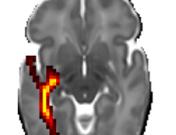
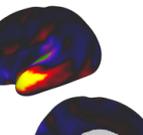
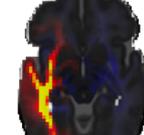
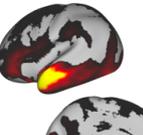
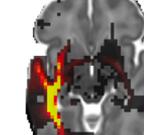
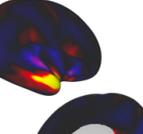
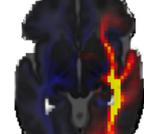
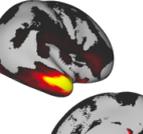
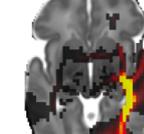
Comparison with Standard Tractography

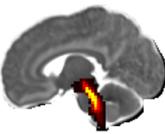
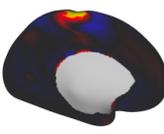
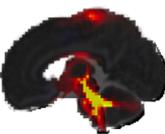
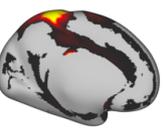
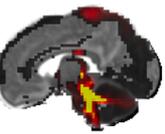
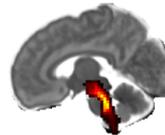
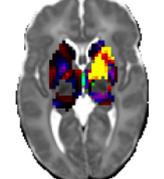
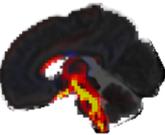
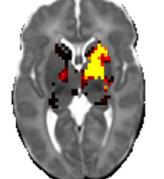
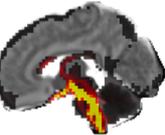
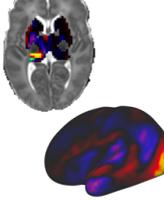
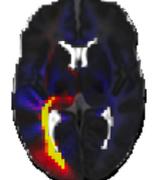
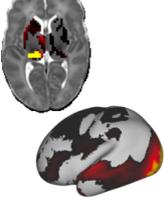
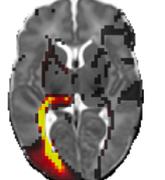
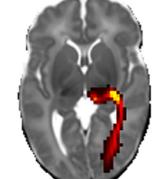
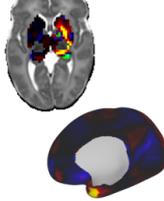
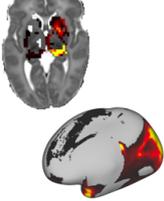
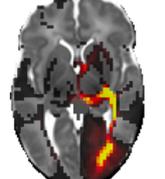
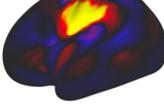
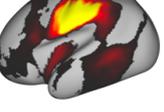
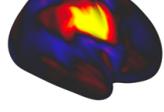
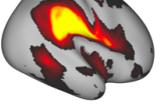
To investigate the interpretability and validity of the extracted components, the white matter components from both ICA and NMF were compared with the group-averaged results from standard tractography protocols. A number of our data-driven components exhibit strong spatial similarity to known white matter pathways (figure 5.6). In fact, all the considered 28 tracts have well-matching components. Both ICA and NMF are able to identify spatially separate regions of grey matter (i.e. networks), along with their underlying white matter connections, for example

in the forceps minor, the inferior longitudinal fasciculus and the various thalamic projections.



	Standard Tractography Protocols		Independent Component Analysis		Non-negative Matrix Factorisation	
	spatial maps	seed space	tract space	seed space	tract space	
left cingulum (hippocampal part)						
right cingulum (hippocampal part)						
left cortico-spinal tract						
right cortico-spinal tract						
forceps major						
forceps minor						

		Standard Tractography Protocols		Independent Component Analysis		Non-negative Matrix Factorisation	
		spatial maps	seed space	tract space	seed space	tract space	
left fornix							
							
left inferior fronto-occipital fasciculus							
							
left inferior longitudinal fasciculus							
							

	Standard Tractography Protocols		Independent Component Analysis		Non-negative Matrix Factorisation	
	spatial maps	seed space	tract space	seed space	tract space	
left medial lemniscus						
right medial lemniscus						
left posterior thalamic radiation						
right posterior thalamic radiation						
left superior longitudinal fasciculus						
right superior longitudinal fasciculus						

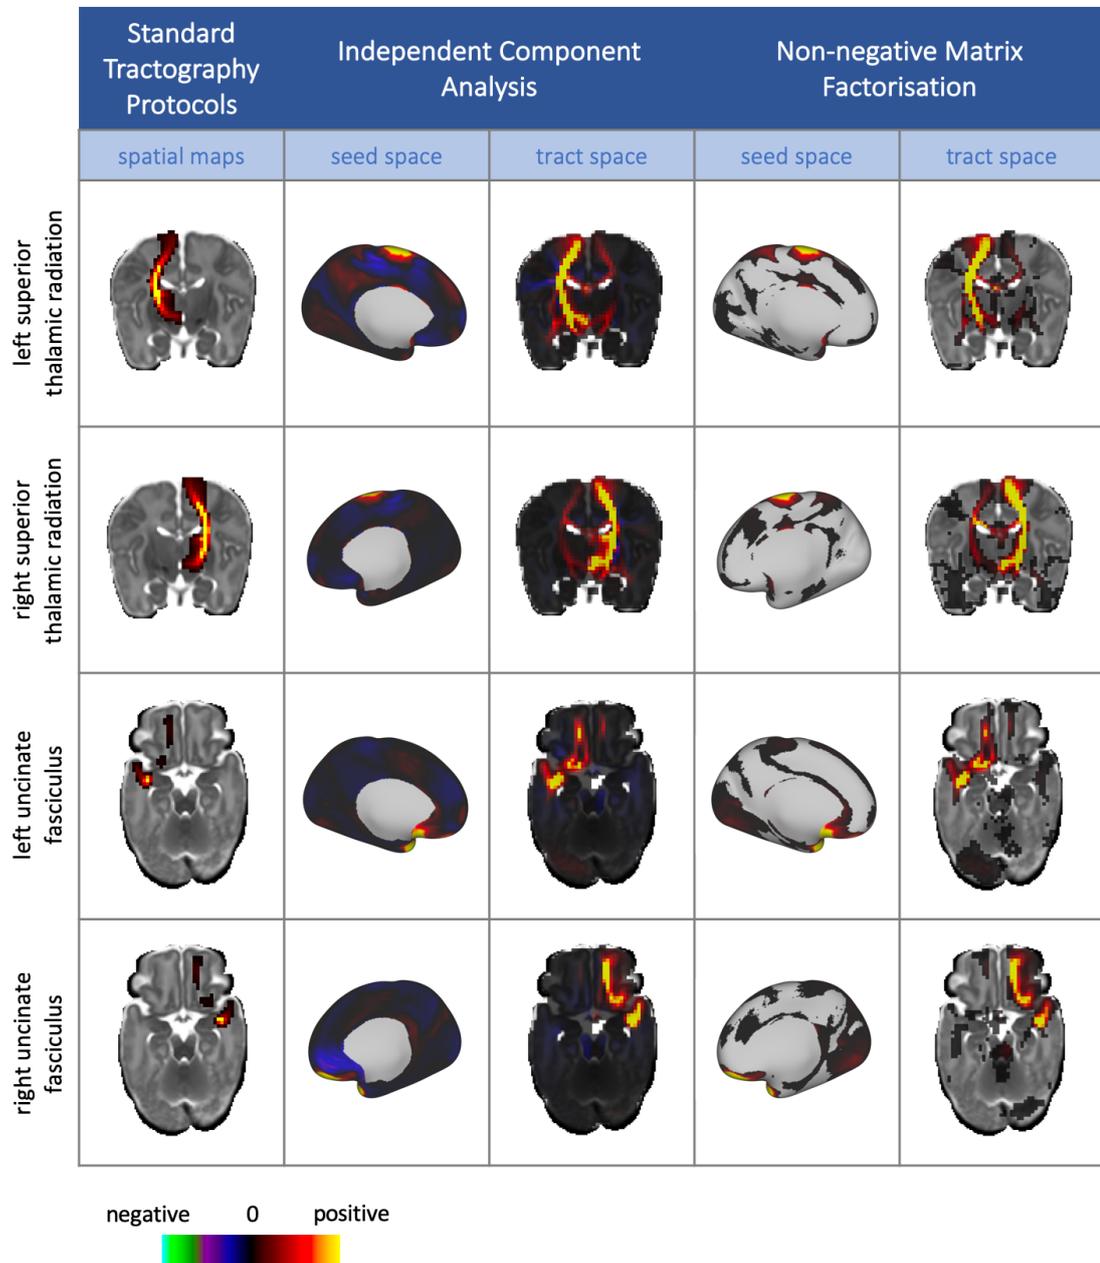


Figure 5.6: The full set of 28 tracts from the standardised protocols (Bastiani et al., 2019), alongside their corresponding components from ICA and NMF. Data-driven components are unthresholded to enable the comparison between the negative values in the ICA components and the sparse, non-negative representations from NMF, whereas the maps from standard protocols are lower thresholded at 0.001 for clearer visualisation of the tract.

The examples in figure 5.6 demonstrate the advantages of using NMF over ICA. NMF components are inherently more sparse (ICA-derived spatial maps typically cover the whole brain) and by construction non-negative. The main body of the anatomically relevant information conveyed by ICA components is present with NMF decompositions but in an inherently non-negative manner. This suggests that the NMF sparsity constraints effectively enforce independence in the composition, similarly to ICA. In addition, there are qualitative improvements of NMF over ICA for a number of tracts. For instance, the NMF component corresponding to the right IFO has a stronger peak in the occipital lobe than the equivalent ICA component, and NMF has fewer false positive frontal projections in the left ILF. Further detailed comparison between NMF and ICA components with differences between matched pairs is shown in figure 5.7. These examples further illustrate that the NMF results convey different information than the ICA results, even when the latter are thresholded to only retain positive values (3rd column) or are thresholded to retain the areas with the strongest weights (4th column). These examples also suggest that different thresholds would be needed across the ICA components to match the relevant information from NMF, as a high threshold that removes the false projection in row a), for example, would remove the full extent of the tract in row c).

The results in figure 5.8 explore how increasing the model order in the decomposition affects the splitting of components. Equivalent components were identified across model orders by calculating the correlations between their spatial maps. The more coarse-grained connectivity patterns from the low dimensionality decompositions are broken down into more sparse, fine-grained spatial maps as the number of components are increased. For example, the left panel of figure 5.8 shows an NMF component and the associated white matter spatial map from the $K = 5$ decomposition that delineates the left pyramidal tract. As the the number of components is increased from $K = 5$ to $K = 50$, this bundle splits into sub-components that

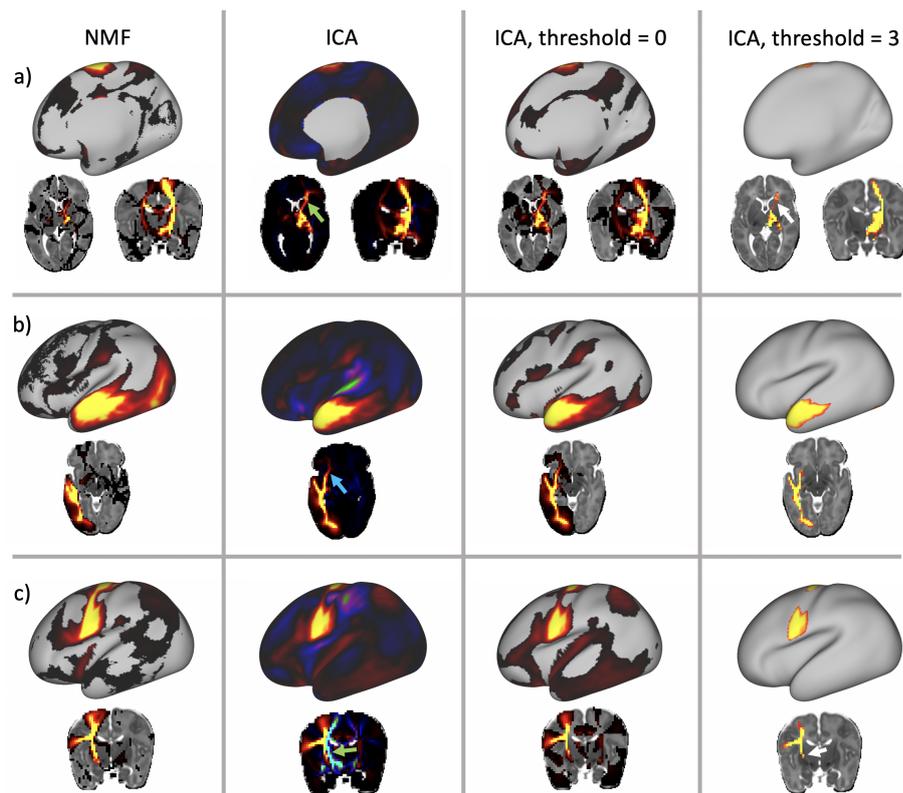


Figure 5.7: *Examples of matched pairs of NMF and ICA components from the $K = 100$ decomposition that illustrate differences between the two methods. Components have all been variance normalised for consistent scaling. a) In the right superior thalamic radiation, there is a false positive anterior projection in the ICA component (green arrow) that is not present in the NMF component. This is still present at a high threshold for ICA (white arrow). b) In this component, corresponding to the left inferior longitudinal fasciculus, there is an anterior projection in the ICA component that is not seen in the NMF result (blue arrow). c) The ICA component has strong negative and positive aspects that obscure the main part of the tract (green arrow). High thresholding, such as in the right column, is too conservative and the tract itself is thresholded out (white arrow).*

characterise different parts of corona radiata projections. An increase in sparsity is also seen from the low to the higher order components, which agrees with the quantitative results in figure 5.13.

Interpretability can be also illustrated for components that do not match any tracts from the set reconstructed using standard tractography protocols. An example is demonstrated in figure 5.9, where 10 components from the $K = 100$ NMF decomposition have been identified as corresponding to different segments of the corpus callosum. For each component, the grey matter (seed space) map is shown, along with the WM spatial map (tract space).

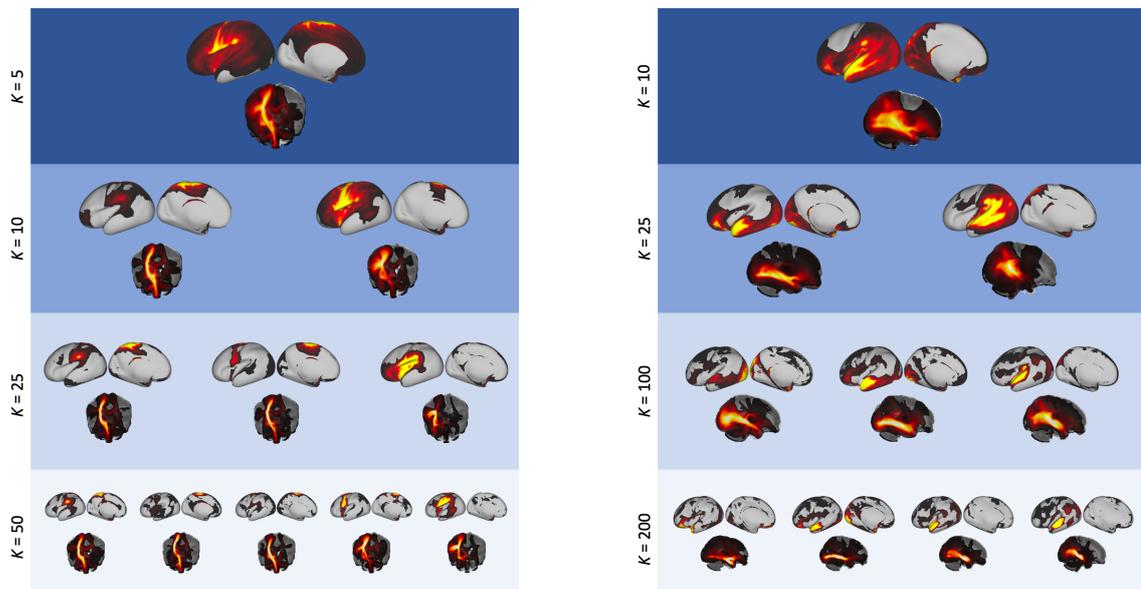


Figure 5.8: *Similar components across different model orders, demonstrating the hierarchical nature of the decomposition. Starting with a single component from a lower dimensionality decomposition, we show components from higher model orders that have high spatial correlation with the original component, in tract space ($r > 0.5$). Tract space results are displayed as maximum intensity projections. Left: A component from the $K = 5$ decomposition showing the left cortico-spinal tract, which is split into more localised sub-components for higher K . Right: A component from the $K = 10$ decomposition that includes several different association fibres in the left hemisphere. At $K = 200$, this has been split into the uncinate fasciculus, inferior longitudinal fasciculus and middle longitudinal fasciculus.*

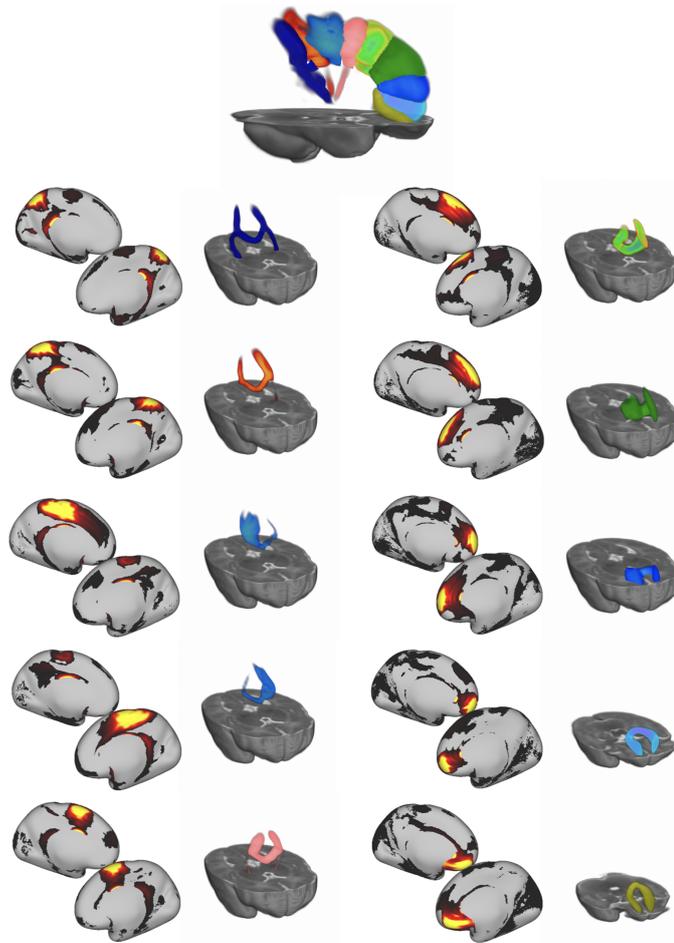


Figure 5.9: Ten components from the $K = 100$ NMF decomposition that correspond to segments of the corpus callosum. For each component, the grey matter (seed space) map is shown, along with the white matter spatial map (tract space) rendered in 3D to aid visualisation. All rendered white matter segments are shown together at the top.

In addition to these callosal bundles, additional components are seen at higher model orders that do not have matching predefined tracts from the standard protocols. Examples of these “unassigned” NMF components are shown in figure 5.10. Many of these components are bilateral, and show short range connections in the frontal lobe, such as the fronto-marginal tract (bottom row, first two columns) (Catani et al., 2012). Others may reflect false positive connections, such as the thalamic loops in the fifth row.

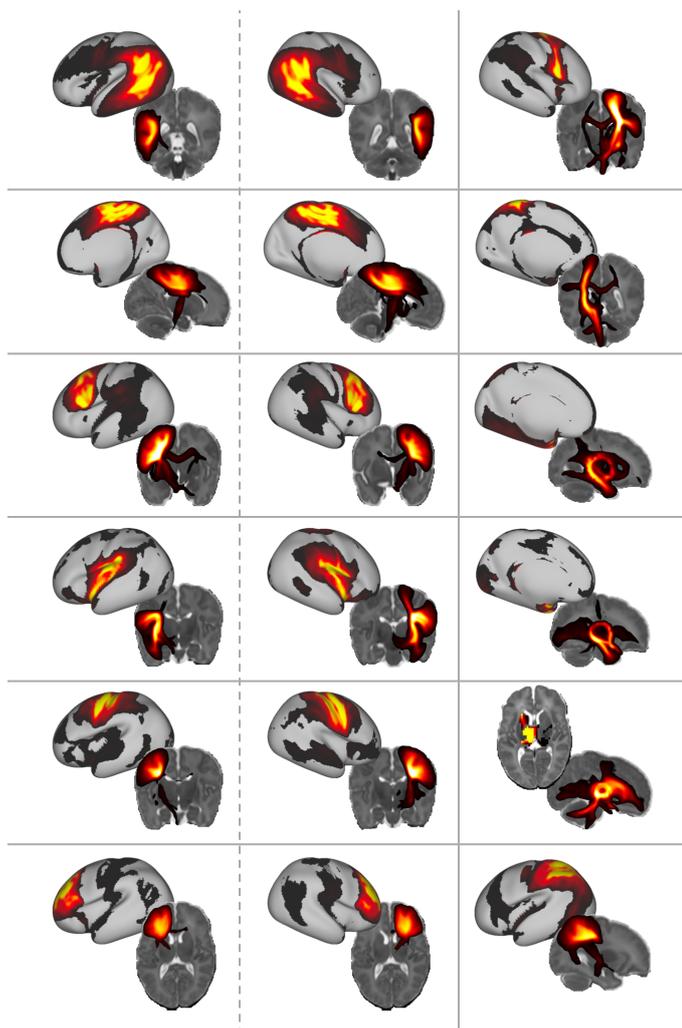


Figure 5.10: *Unassigned components from the $K = 50$ NMF decomposition. One-to-one matching was performed with the 28 standard tracts, leaving 22 unassigned components. Four of these were callosal fibres, similar to those shown in figure 5.9, leaving 18 additional components that we show here. The components in the left and centre columns are bilateral, unpaired components are in the right column.*

Comparison with fMRI

As an additional indirect validation, the grey matter maps from the NMF decompositions of the tractography data were compared with resting-state networks (RSNs) obtained from fMRI data. Group-level ICA ($K = 50$) was performed on fMRI data from 55 subjects and the resultant resting-state networks were compared to those from a $K = 50$ NMF decomposition of the structural connectivity data from the same subjects. Through visual inspection, 24 of the functional components were

found to contain noise or artefacts, so were discarded. The similarity of the remaining 26 RSNs to our structural grey matter components was measured using Pearson's correlation coefficient, r , to identify the best matching pairs.

Most functional components were well matched to at least one structural component, with the lowest correlation value between an RSN and a tractography component being $r = 0.2$. Over half (14 out of the 26 networks identified) had a correlation value $r > 0.5$ with their best-matched structural component. The correlation matrix in figure 5.11 is sparse, which indicates that there is specificity in the matching. Where RSNs were strongly associated with multiple structural components, this was either a bilateral network split into the two hemispheres (e.g. figure 5.11b and c) or structural networks that overlapped with different regions of the RSN (5.11a and d).

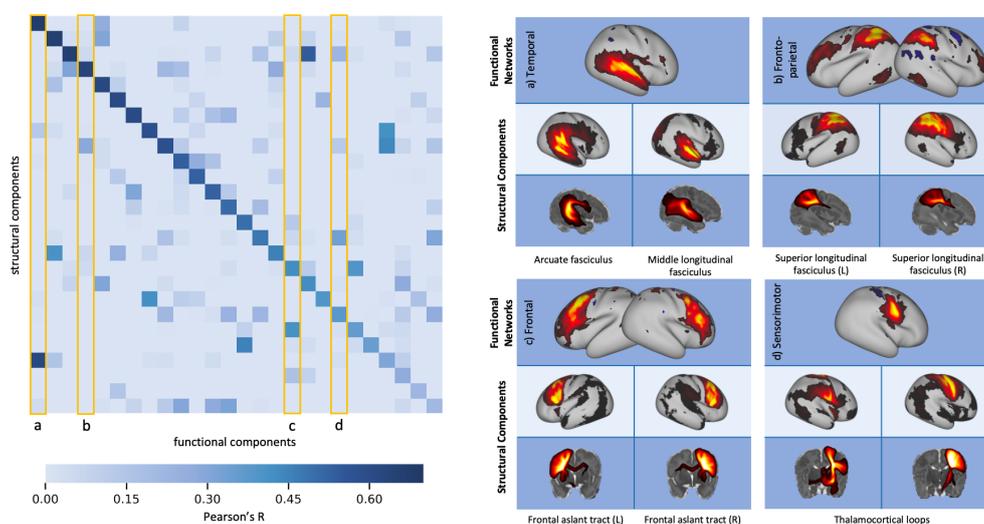


Figure 5.11: *Left: correlation matrix between the fMRI RSNs and their 26 best-matched tractography NMF components. Right: examples of the functional networks and their most spatially similar grey matter components from structural NMF. These correspond to the columns outlined in yellow on the correlation matrix. The corresponding white matter patterns are shown as maximum intensity projections.*

Reconstruction Error and Sparsity

The reconstruction error and component sparsity were also computed for the in vivo results. In line with the results from the simulations, reconstruction error decreases with increasing numbers of components, with ICA having slightly higher reconstruction accuracy than NMF (figure 5.12). Sparsity is much higher for NMF than for ICA (figure 5.13). Sparsity increases rapidly from 5 to 50 components and then plateaus after 100 components.

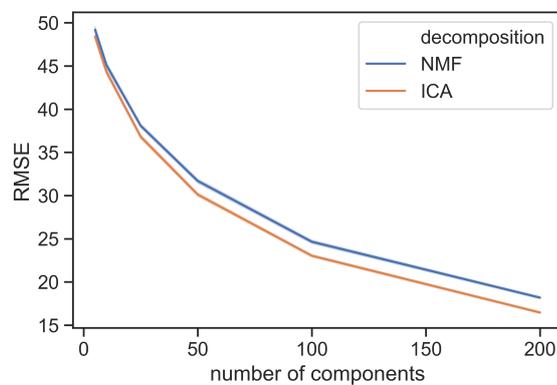


Figure 5.12: *The reconstruction error in the ICA and NMF components.*

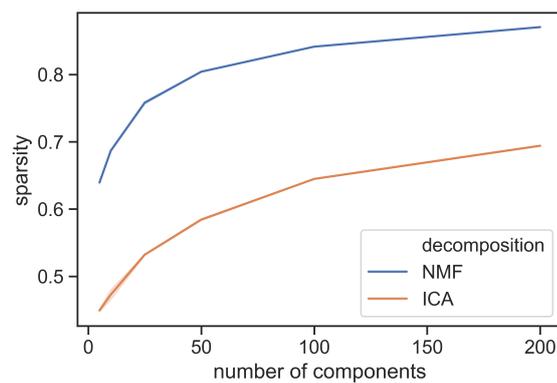


Figure 5.13: *The sparsity of the derived components, calculated according to equation 5.2.*

5.5 Discussion

A non-negative framework was developed for simultaneously mapping white matter connections and corresponding grey matter networks from diffusion MRI data in a data-driven manner. Non-Negative Matrix Factorisation (NMF) is a powerful alternative to traditional tract delineation that has no parametric assumptions, no dependence on predefined ROIs and masks in a template space, and is inherently suited to the non-negative nature of tractography data.

This work falls within the family of other data-driven approaches for mapping structural connections from whole-brain tractograms, such as (Garyfallidis et al., 2012; O’Donnell and Westin, 2007; Siless et al., 2018). The approach extends these efforts by allowing simultaneous reconstructions of white matter bundles, but also the corresponding grey-matter networks that these bundles connect. Furthermore, none of the previous data-driven approaches have been applied for mapping connections from diffusion MRI data of the neonatal brain, as shown here.

5.5.1 Validation using Simulations

Simulations were used to investigate the behaviour of the decompositions in controlled scenarios, in which the ground truth was known. In order to generate realistic simulations for such a decomposition framework, the sources were based on the properties of the distributions learned from in vivo data, and mixed randomly to generate synthetic data with a known number of components. First, the effect of adding an L1-norm regularisation term to the objective function for NMF was investigated (see equation 5.1). Increasing the regularisation reduces the accuracy of the data reconstruction, but a small amount ($\alpha = 0.1$) improves the correlations between the sources and the components at lower model orders and promotes component sparsity. Therefore, it was decided that an alpha value of 0.1 would be used for subsequent work, as this was deemed to be a good compromise between

higher component sparsity and source reproduction, with only a minimal impact on reconstruction accuracy. Increasing the sparsity of components has been shown to generate features that are inherently more independent, while constraining the NMF solution space to make the decomposition more reliable (Hoyer, 2004).

The effect of different noise levels in the data was also investigated. As expected, the reconstruction error of the decompositions increased with increasing noise, but the correlation between components and true sources was fairly stable, particularly at low model orders. Comparing the results from ICA and NMF, both were able to reconstruct the original data (using the dot product of the mixing matrix and component matrix) with good accuracy, but the components from ICA were less well matched to the true non-negative sources themselves than those from NMF. This is because the components from ICA contain negative values that are not found in the real sources, although mutual cancellation of positive and negative values in the components and mixing matrix allows the data matrix to be reconstructed accurately.

5.5.2 In Vivo Results

White matter spatial maps of the NMF components show strong spatial similarity to known white matter pathways (figures 5.6, 5.8, 5.9). Each of the 28 tracts that were considered had a corresponding component from the $K = 100$ decomposition. The tractography-matched patterns from ICA and NMF have similarities, as seen in figure 5.6. This hints towards NMF being able to separate spatially independent components, in an analogous manner to ICA, despite not having independence constraints enforced explicitly. This is because the sparsity constraint on the NMF decomposition implicitly promotes non-Gaussianity in the resultant components, which is used as a proxy for independence in the FastICA algorithm (Hyvärinen and Oja, 2000). Indeed, sparsity and independence criteria have previously been shown to generate very similar basis sets across several different data types (Saito et

al., 2000). Despite the overall similarity between the results from the two methods, there are some noticeable differences between the spatial maps from ICA and NMF, shown in figures 5.6 and 5.7. For example, the component corresponding to the forceps major extends more strongly into the right hemisphere in the NMF component than in the ICA component. In addition, the components from ICA corresponding to the inferior longitudinal fasciculus extends into the frontal lobe, mixing with the inferior fronto-occipital fasciculus, which is not seen in the NMF result. Figure 5.7 shows further examples and illustrates the effects different levels of thresholding on the ICA results. These results show that a) the NMF results convey different information than the ICA results, even when the latter are thresholded to only retain positive values, and b) different ICA components would require different levels of thresholding to match the results from NMF.

A range of model orders were explored from 5 to 200. The lower model orders generate more distributed components that contain multiple white matter bundles, whereas the higher model orders give more specificity, as shown in figure 5.8. The components from lower model orders (eg. $K = 5$) are split into smaller constituent parts for higher model orders, providing a component hierarchy as K increases.

There are also some tracts which are not so well-characterised by either method, such as the acoustic radiation, which contains a mixture of the middle longitudinal fasciculus, and the superior longitudinal fasciculus, which does not have separate lobes in the grey matter components. However, it is worth noting that the data-driven methods presented here are not meant to replace tractography for major bundle delineation, particularly in cases where we have well-defined tractography protocols. Instead, they can provide complementary ways to concurrently extract grey matter and white matter connectivity patterns from all the data simultaneously, particularly for cases where this delineation of bundles is challenging or incomplete. This can potentially be a powerful novel way of summarising the information content of tractography data for applications other than bundle delineation, such as

connectivity-driven functional localisation (for example (Mars et al., 2018)).

The grey matter maps of the NMF components were also shown to align well to resting-state networks from fMRI. This provides further evidence that these data-driven results are anatomically meaningful. It also highlights how this new framework provides a similar analysis paradigm for structural connectivity to that commonly used for the analysis of fMRI data. This could provide a route to analyse structural and functional connectivity data within a unified framework.

The sparsity and reconstruction error of the components were calculated, shown in figures 5.13 and 5.12. Reconstruction error decreases and the sparsity of the components increases, with increasing numbers of components. This reflects the higher degree of freedom afforded by more components that permits a more detailed reconstruction of the original data, and components that are more tightly localised around fine-grained regions of similar connectivity. NMF components are more sparse than those from ICA, which indicates that the former is able to localise connectivity patterns more effectively, disregarding redundant information and keeping non-negativity in the reconstruction.

5.5.3 Decomposition Domain

In the results presented here, decompositions have been applied in the seed domain, allowing white matter tract overlap. Decompositions were also applied to the transpose of the connectivity matrix, \mathbf{X}^\top , which meant decomposing (and in the case of ICA enforcing independence) in the tract domain. ICA and NMF were performed on the transpose of the group-average connectivity matrix, with $K = 50$. Figure 5.14 illustrates the similarity between the results from each method, which shows that the ICA components are most affected by this change. Most NMF components are nearly identical to the original results. This agrees with expectations, as in NMF the sparsity and non-negativity constraints are enforced in both the mixing matrix and the components (see equation 5.1).

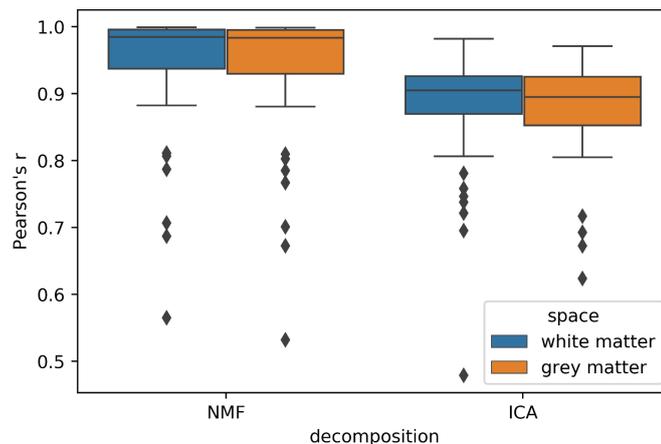


Figure 5.14: Correlations between decompositions applied to connectivity matrices in both the ($seed \times white\ matter$) and ($white\ matter \times seed$) configuration. Highest Pearson's r is plotted for each component or column of the mixing matrix with the equivalent matrix from the transposed decomposition.

5.6 Concluding Remarks

This chapter has demonstrated a data-driven framework for mapping white matter connectivity patterns and their corresponding grey matter networks from neonatal MRI data. This follows up previous work that has been done in adults using ICA (O'Muircheartaigh and Jbabdi, 2017), and provides an alternative approach using NMF that has improved interpretability. In simulations, NMF was shown to provide a more accurate recovery of non-negative sources than ICA. Furthermore, the results from neonatal MRI data show that the connectivity patterns are well matched to tracts from standard tractography protocols, and provide additional information about the terminations of these tracts in the grey matter. Nevertheless, we only showed group results in this chapter. The following chapter presents a framework for projecting group results to individual data. It further analyses the stability of these decompositions across different subject groups, and shows how the grey matter components can be used to parcellate the cortex into regions of similar structural connectivity to the rest of the brain.

Chapter 6

Reliability and Applications of Data-Driven Connectivity Mapping

Contents

6.1	Introduction	134
6.2	Theory	135
6.2.1	From Group to Individual Subject Decompositions - Non-Negative Dual Regression	135
6.3	Methods	137
6.3.1	Split-half Reliability Analysis	137
6.3.2	Non-Negative Dual Regression	138
6.3.3	Cortical Parcellations	139
6.4	Results	140
6.4.1	Assessing the Reproducibility of the Decompositions	140
6.4.2	Parcellations	144
6.5	Discussion	148
6.5.1	Reproducibility	148
6.5.2	Parcellations	149
6.5.3	Limitations and Future Work	152
6.6	Concluding Remarks	153

Foreword

In this chapter we evaluate and extend the work presented in the previous chapter. We investigate whether the data-driven patterns are reproducible across different subject groups, and derive a new non-negative method for dual regression. This provides a means of obtaining subject-level representations of the original group-level components, while maintaining their advantageous non-negative properties.

The grey matter components are used to generate parcellations of the neonatal cortex, in which vertices within a parcel share common connectivity patterns to the rest of the brain. The reproducibility of these parcellations is also assessed, and we show how the dual-regression framework can be used to generate parcellations tailored to the connectivity profiles of individual subjects.

6.1 Introduction

In the previous chapter a new framework was presented, in which whole-brain connectivity matrices are decomposed into a set of white matter spatial maps and their corresponding grey matter termination points, in a data-driven fashion. Two different methods were compared for the decomposition: ICA and NMF. By comparing these methods on a simulated dataset, it was found that NMF provides a more accurate characterisation of non-negative signals, of the kind found in structural connectivity data. NMF was applied to connectivity matrices from neonatal dMRI data and was shown to generate white matter spatial maps and grey matter networks that reflect anatomically relevant information.

In this chapter, the reproducibility of the NMF results across independent subject groups is investigated. Similarity is assessed in the NMF components between two halves of the cohort and across groups with different numbers of subjects. That way, the stability of our approach in providing connectivity estimates at the group-level is assessed.

An approach is then devised that allows correspondence between group NMF results and individual subjects' data. The ability to derive subject-level components provides a gateway to more individualised analysis of white matter development. Our new method obtains subject-level projections of the group-level components, in a way that maintains the non-negativity and sparsity of the original components.

The chapter concludes by presenting an application of the data-driven components, which are used to generate cortical parcellations. The reproducibility and clustering accuracy of the parcellation is assessed, relative to a benchmark of random parcellations. Most existing cortical parcellation schemes for neonates are based on cortical folding landmarks (Adamson et al., 2020; Gousias et al., 2012), so this new connectivity-driven approach can provide additional insight into the organisation of the neonatal cortex. It further provides an additional piece of evidence to indirectly

validate the NMF approach.

6.2 Theory

6.2.1 From Group to Individual Subject Decompositions - Non-Negative Dual Regression

When considering whole-brain tractography matrices \mathbf{X} from multiple subjects (e.g., by averaging across subjects in the simplest case), the components and mixing matrices will represent the group. A relevant question is whether these group components can be projected to the data of individual subjects or other subgroups. This has the important benefit of providing correspondence of components across a set of subjects (such that the n^{th} component represents the same network and bundle across all considered subjects), compared to performing individual subject analysis separately, which would not ensure such correspondence.

To solve this problem of group-subject correspondence in fMRI ICA, dual regression frameworks have been proposed and used to generate subject-level representations of the group components and mixing matrices (Beckmann et al., 2009; Nickerson et al., 2017). Dual regression comprises of two steps:

- i) Identify the subject-specific mixing matrix $\tilde{\mathbf{A}}$ from the group-level components \mathbf{S} , using the subject-level data matrix $\tilde{\mathbf{X}}$:

$$\tilde{\mathbf{X}} = \tilde{\mathbf{A}}\mathbf{S} \rightarrow \tilde{\mathbf{A}} = \tilde{\mathbf{X}}\mathbf{S}^\dagger$$

(where \mathbf{S}^\dagger denotes the pseudoinverse of \mathbf{S})

- ii) Find the subject-level components $\tilde{\mathbf{S}}$, using the subject-specific mixing matrix $\tilde{\mathbf{A}}$:

$$\tilde{\mathbf{X}} = \tilde{\mathbf{A}}\tilde{\mathbf{S}} \rightarrow \tilde{\mathbf{S}} = \tilde{\mathbf{A}}^\dagger\tilde{\mathbf{X}}$$

In previous work, this multivariate regression has been achieved by taking the pseudoinverse of the group-level components and the subject-level mixing matrix (O’Muircheartaigh and Jbabdi, 2017), as illustrated in figure 6.1.

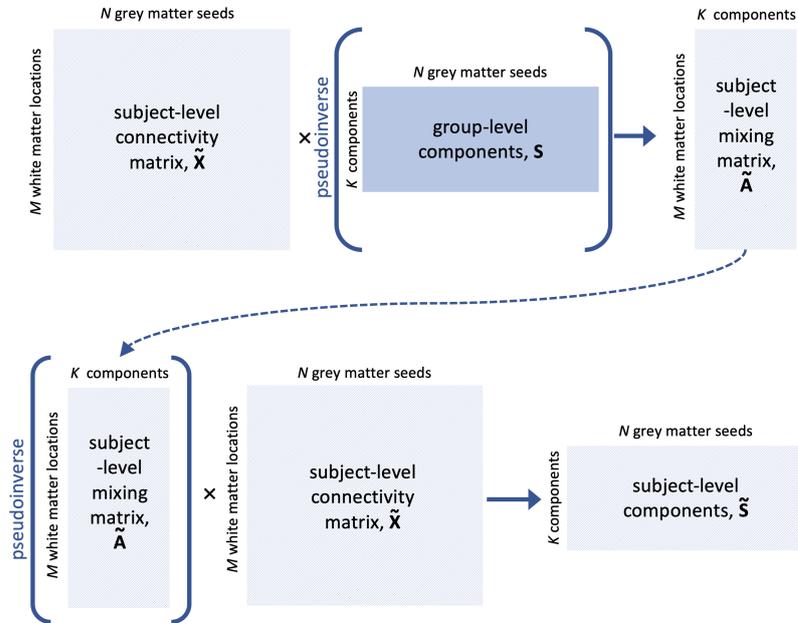


Figure 6.1: Schematic illustration of the dual regression step used for ICA, that generates subject-level representations of the group-level components. The group-level grey matter components are first regressed onto the subject’s connectivity matrix to obtain the subject-level representations of the components in white matter. The pseudoinverse of this mixing matrix is then used to obtain the subject-level grey matter components.

Even though the above approach can be applied directly to NMF components and mixing matrices, taking the pseudoinverse introduces negative values into the components and their weights, which leads to mixed-sign subject-level representations of the original non-negative group-level components. Instead, a “non-negative dual regression” technique has been developed and is presented for back projecting NMF results, using non-negative least squares (NNLS) (Ling et al., 1977) for the regression steps. NNLS solves an equation of the form $\operatorname{argmin}_{\mathbf{x}} \|\mathbf{B}\mathbf{x} - \mathbf{y}\|_2$ subject to $\mathbf{x} \geq 0$, in which \mathbf{x} and \mathbf{y} are vectors, and \mathbf{B} is a matrix. Thus, the optimisation has to be performed separately for each target voxel in step (i) and each grey matter seed in step (ii) (see figure 6.2), but this process can be parallelised to reduce computation time. This provides an entirely non-negative framework for dual regression

that retains the sparse characteristics of the group-level NMF components.

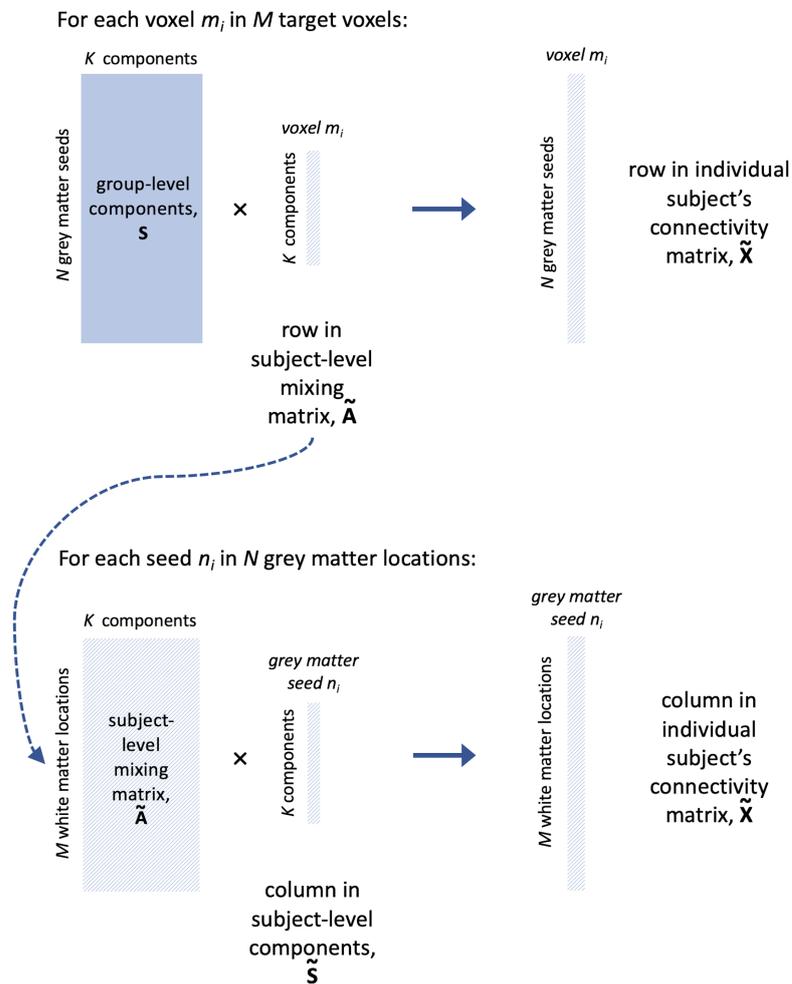


Figure 6.2: Schematic illustration of the non-negative dual regression step used for NMF, which generates subject-level representations of the group-level components. The first equation is solved by non-negative least squares, using the group-level grey matter components and the subject's connectivity matrix to solve for each row in the subject-level mixing matrix. This mixing matrix is then used to find the subject-level grey matter components by the same method.

6.3 Methods

6.3.1 Split-half Reliability Analysis

A split-half analysis was performed on a cohort of 323 term-age subjects born at term age (175 male, 148 female). Median (range) birth age was 40.1 (37.0, 42.3)

postmenstrual weeks and age at scan was 40.9 (37.4, 44.4) weeks. ICA and NMF were applied to the group-averaged connectivity matrices from each half of the cohort with a number of model orders: $K = 5, 10, 25, 50, 100, 200$. For each value of K , a one-to-one matching was performed of components across the split-half, based on the Pearson's correlation coefficients of their spatial maps, recording the correlation coefficients of the matched pairs as a measure of their similarity. This was repeated for the grey matter and white matter maps.

The consistency of the decompositions was also tested for subject groups of different sizes. NMF was applied to a single subject's connectivity matrix and average matrices for groups of 5, 10, 50 and 100 subjects with $K = 50$. The resultant components were compared with the results from split 1 of the main cohort.

6.3.2 Non-Negative Dual Regression

Group to Individual

Non-negative dual regression was applied to all subjects in split 1 of the cohort, projecting the $K = 50$ group level components back on the subject-level data as described in section 6.2. NMF was also applied directly to the individual subjects' connectivity matrices for comparison. Dual regressed group components for a subject are expected to match in general components obtained from individual NMF analysis of that subject. However, the ordering of the components across subjects is preserved with non-negative dual regression.

Group to Subgroup

Dual regression can also be used to project components from one group to another. The $K = 50$ group level components from split 1 of the cohort were projected onto averaged connectivity matrices from four different groups, each containing 15 subjects of different age ranges (37-38, 39-40, 41-42 and 43-44 weeks PMA at scan).

6.3.3 Cortical Parcellations

The grey matter components were used to generate cortical parcellations, in which vertices were grouped according to their structural connectivity profiles. This was done using a “winner-takes-all” approach (Behrens et al., 2003), where each vertex is labelled according to its highest weighted component. This provides a hard cortical parcellation, where vertices in a given parcel can be linked to a specific white matter bundle.

The stability of the parcellations across subjects was assessed with a split-half analysis, using the Dice coefficient between parcellations generated on each of the split-halves as a similarity metric. Dice coefficients measure the overlap between two clusters X and Y , normalised by the number of elements in each cluster:

$$Dice = \frac{2|X \cap Y|}{|X| + |Y|} \quad (6.1)$$

In this case, $|x|$ describes the number of vertices within a parcel, and $|X \cap Y|$ the number of shared vertices between the two parcels considered. The Dice score was recorded between each parcel and its highest matching parcel in the parcellation from the other split.

The Silhouette score was used as a measure of how well the parcellations clustered similarly connected vertices (Rousseeuw, 1987). Silhouette score was calculated using the mean within-cluster distance a and the mean nearest-cluster distance b , with $(1 - \text{Pearson's } r)$ used as a distance metric:

$$Silhouette = \frac{b - a}{\max(a, b)} \quad (6.2)$$

This assesses a successful parcellation as one which groups vertices with similar connectivity profiles, which are distinct from the connections in other parcels.

The group-level parcellations were benchmarked against random Voronoi parcellations of the same model order (Aurenhammer, 1991). Voronoi tessellations with

N parcels were first generated over the surface of a sphere, by randomly distributing N seed points across the sphere. Each vertex on the sphere was then labelled according to its closest seed point on the spherical surface. The spherical parcellations were then projected onto the cortical surface. The resultant parcellations have, by definition, N spatially contiguous, randomly distributed regions, which provide a good benchmark with which to compare our data-driven parcellations.

Subject-level parcellations were generated from the non-negative dual regressed grey matter components from each subject in split 1 of the cohort. The variability of these subject-level parcellations was assessed by calculating the Dice coefficient between equivalent parcels in the subject-level and group-level parcellations.

6.4 Results

6.4.1 Assessing the Reproducibility of the Decompositions

To assess the reproducibility of the derived components, a split-half reliability analysis was performed for the ICA and NMF decompositions. Figure 6.3 presents histograms of correlations between the best-matching components across the split-halves, for both ICA and NMF. In all cases, the median correlation value lies above 0.8, which shows that both methods are robust across different subject groups. Even if patterns are more variable for lower model orders ($K < 25$), both methods perform similarly for higher K . Similar behaviour is observed for grey matter components and white matter mixing matrices.

Having ascertained the reliability of the data-driven framework for a large group of subjects, the stability of the components was explored for smaller groups. A $K = 50$ decomposition was performed on a single subject's data, and then for groups of 5, 10, 50 and 100 subjects. A quantitative analysis of the similarity between the small group-size results and the components from split 1 of the cohort is shown in figure 6.4, from which we can see that components from groups of more than

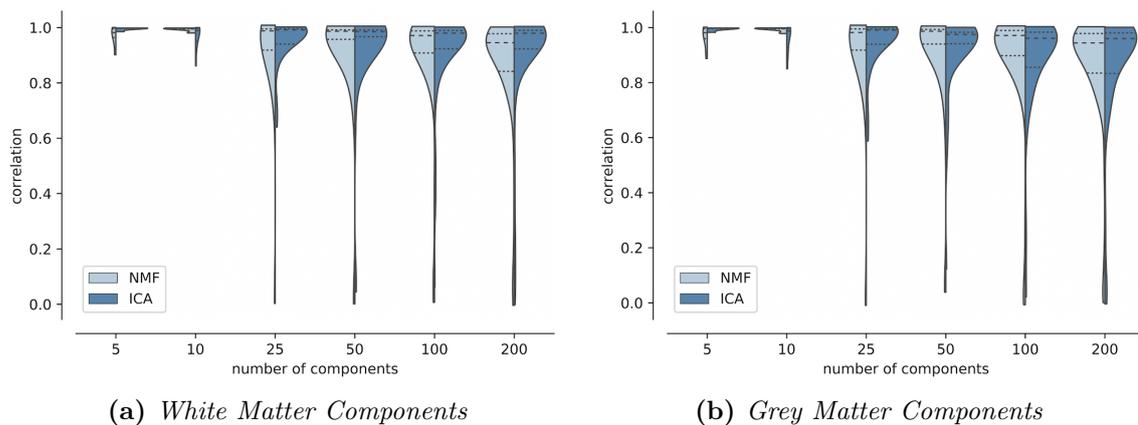


Figure 6.3: Split-half reliability analysis for ICA and NMF. Pearson's correlation scores were calculated between the best-matched components in each split for the white matter spatial maps (a) and the grey matter maps (b). The dotted lines on the violin plots indicate the 25th and 75th percentiles and the median is represented by a dashed line.

ten subjects have very strong correspondence with the full cohort, while even the single-subject results are reasonably similar.

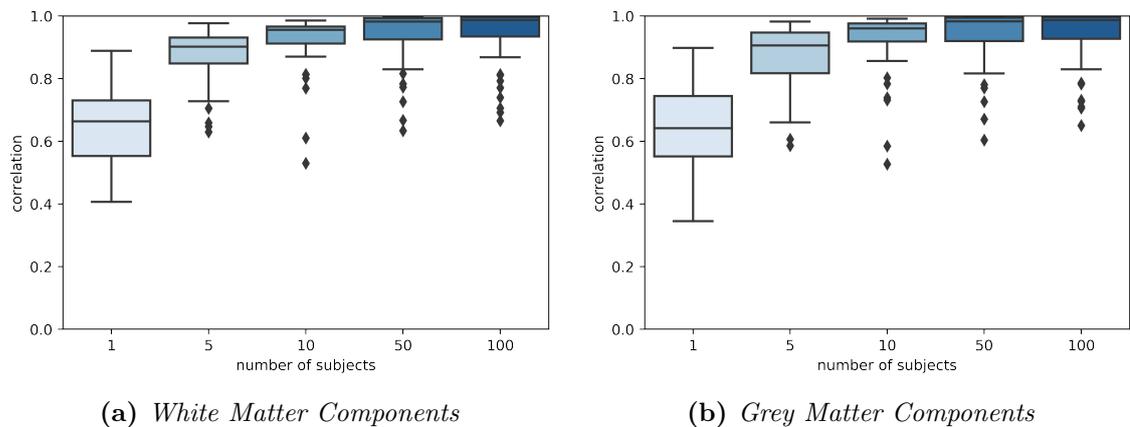


Figure 6.4: Boxplots illustrating the correlations between the white matter (a) and grey matter (b) spatial maps from the $K = 50$ decompositions of groups with varying numbers of subjects, compared to the results from 161 subjects in split 1.

Figure 6.5 shows some example components from these group-level decompositions, to illustrate visual similarities between them. This shows that the patterns are robust even at the single-subject level, although the patterns are noisier with fewer subjects.

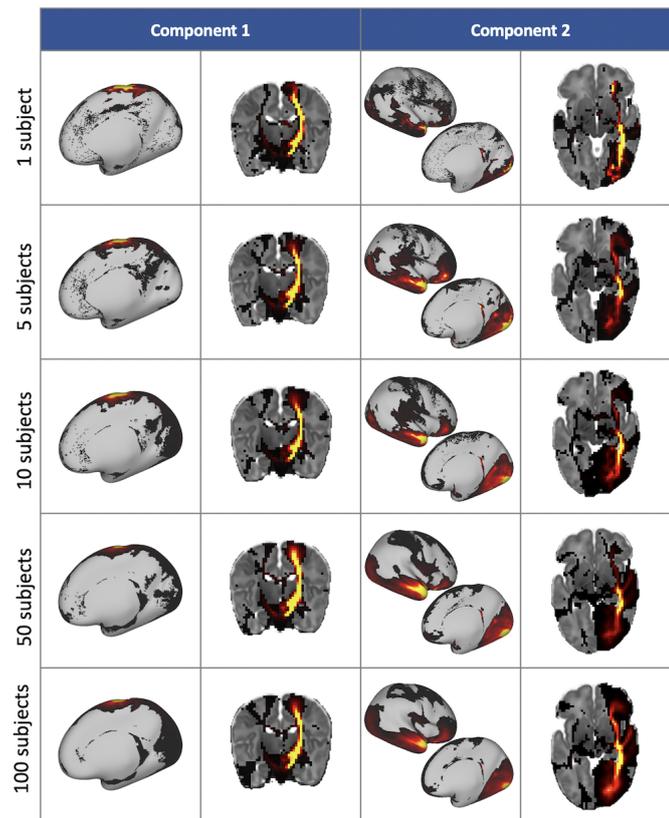


Figure 6.5: *Two components and their corresponding white matter pathways from $K = 50$ group-level decompositions with varying numbers of subjects. Component 1 correlates well with the tractography-delineated cortico-spinal tract, and component 2 with the inferior longitudinal fasciculus.*

Dual Regression Examples

Figure 6.6 compares the results from standard dual regression and non-negative dual regression for an example component in two subjects. It demonstrates how standard dual regression generates subject-level components with negative values in both the grey matter components and the white matter mixing matrices, which loses the interpretability of the non-negative group-level components. On the other hand, the novel non-negative dual regression technique generates components that have the advantageous sparsity and non-negativity of the group-level spatial maps, whilst retaining the individual features of the different subjects.

The results from the non-negative dual regression to individual subjects were compared with the single-subject NMF decompositions, as shown in figures 6.7 and

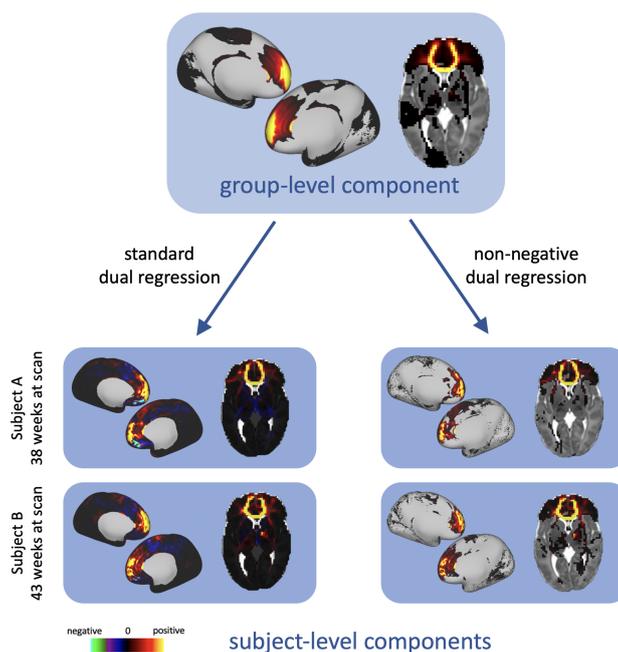


Figure 6.6: Example dual regression results for a component from a $K = 50$ NMF decomposition. On the left, the component has been dual regressed onto two subjects' data with the standard approach using the pseudoinverse. On the right, the component has been dual regressed with our non-negative method that uses non-negative least squares.

6.8. There is a strong agreement between the component maps obtained from these different approaches, which is reassuring and highlights the benefit of using non-negative dual-regression against a group decomposition in ensuring consistency in the components between subjects, but also preserving individual subject features. A small number of cases (lower end of the depicted distribution in figure 6.7) exhibit relatively lower agreement between the two sets of results. It is possible that imperfections in registration and/or alignment of the surfaces to the volumetric template are reflected in these disagreements; but even in these cases (Subject B as representative example), the spatial maps of the components do not look too dissimilar, demonstrating the robustness of the approach.

Non-negative dual regression was also used to project the components from the main group onto smaller subgroups of the cohort of different ages. Two example components from these results are shown in figure 6.9. Even though the spatial maps of the dual-regressed components are broadly similar across the different groups

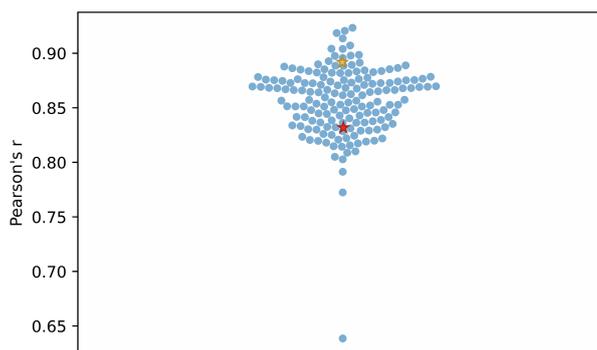


Figure 6.7: *Swarm-plot of the median correlations between single-subject decompositions and results from non-negative dual regression. Correlations were calculated for each of the $K = 50$ components in the subjects from split 1 of the cohort.*

for the presented examples, the results show slight variations in the branching of both tracts between the 37-38 week and 43-44 week age groups (more confined for the former). The results also demonstrate the feasibility of back-projecting group components to subgroups, and the potential for subsequently exploring effects in corresponding white and grey matter regions.

6.4.2 Parcellations

The grey matter components from NMF were used to generate hard parcellations of the cortex, using a winner-takes-all approach. This process was carried out on each of the split-half groups to assess how robust the parcellations are to different groups of subjects. Figure 6.10 illustrates the parcellation results for different values of K . There is high reproducibility of the parcels between the two split-halves, and parcellation schemes are robust across different model orders. Figure 6.10 also shows a subject-specific parcellation generated from the results of a non-negative dual regression that demonstrates qualitatively how the group results correspond to single subjects. In order to quantify the variability of these group parcellations across subjects, the Dice coefficient was calculated between the equivalent parcels in the group-level and subject-level parcellations. The average coefficient for each parcel of the $K = 100$ parcellation is shown in figure 6.11, alongside two example

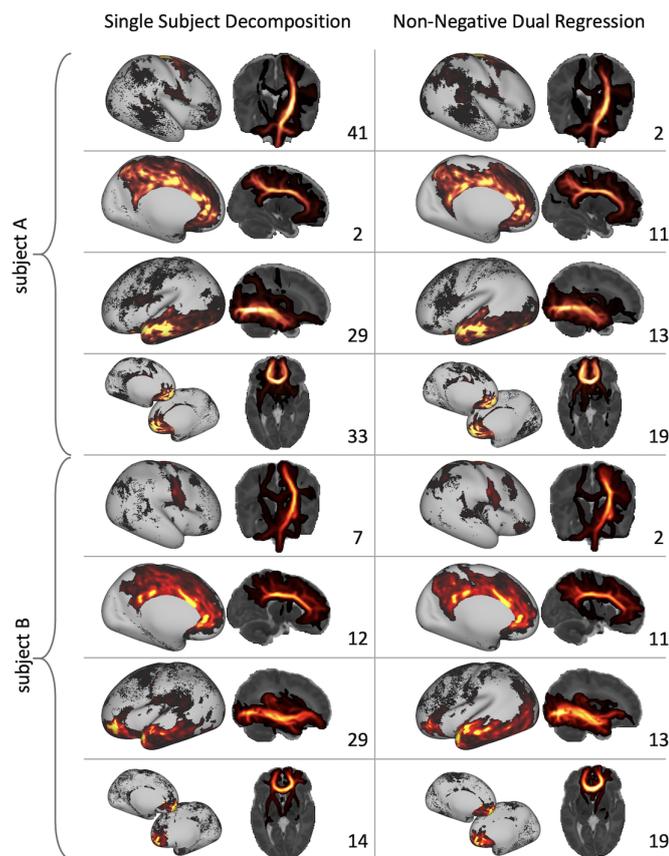


Figure 6.8: Selected components from two representative example subjects from the cohort, Subject A (median correlation = 0.89, yellow star in figure 6.7), and Subject B (median correlation = 0.83, red star in figure 6.7). The component numbers are included in the lower right corner of each box, to show how dual regression provides consistent component ordering across subjects.

subject level parcellations, with the lowest and highest average (across parcels) Dice coefficient, respectively. Most parcels are relatively stable across subjects (average Dice coefficient > 0.7).

Split-Half Analysis

Figure 6.12 further quantifies the similarity between the group parcellations by showing the distributions of Dice coefficients of matched parcels between the parcellations from the two halves of the cohort. This was compared against distributions of Dice coefficients obtained from the comparison of the split 1 parcellation with 100 random Voronoi parcellations, each with the same number of parcels as the decomposition

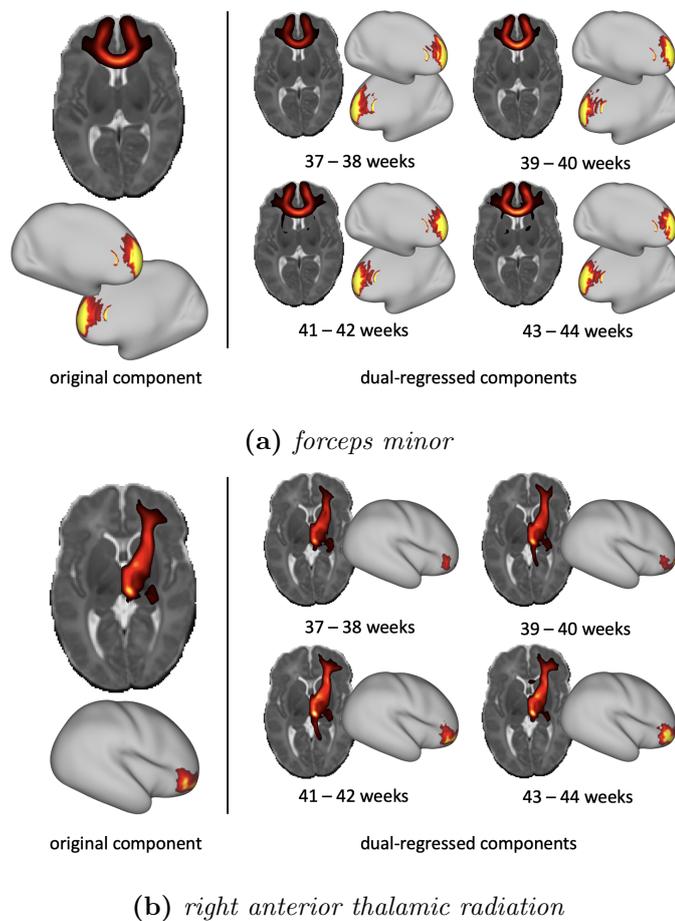


Figure 6.9: Example components from a group-level dual regression, where the $K = 50$ NMF components from split 1 of the cohort were projected onto groups of 15 subjects scanned at different ages.

used. The parcellations generated using the NMF components are significantly more consistent across subjects than the equivalent randomly generated parcellations.

Silhouette Score

To further gain insight into the quality of these parcellations, the mean Silhouette score across parcels was calculated for the NMF-based parcellations at each model order, and for each split-half of the cohort. The measure was also computed for 100 randomly generated Voronoi parcellations with the same number of parcels, as a comparison. A Silhouette score measures the similarity of the data within a parcel, relative to their dissimilarity to data in other parcels. Figure 6.13 shows that the

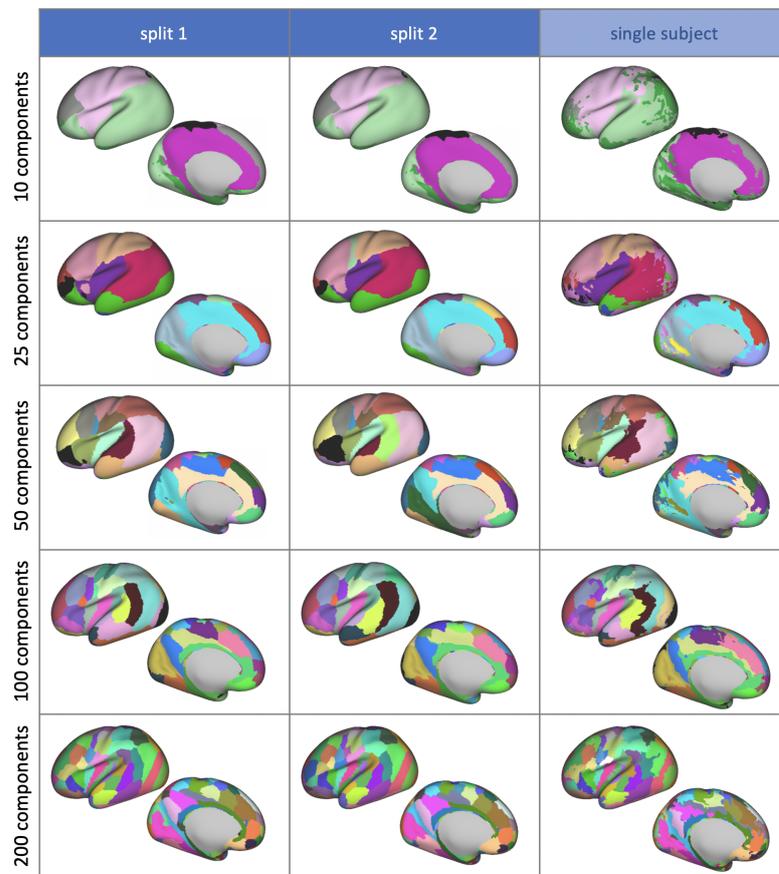


Figure 6.10: *Hard parcellations of the cortical surface from NMF, from each split-half of the cohort and from dual regression of the group-level results onto a single subject. The left hemisphere displayed only. Parcels are colour matched according to the correlation values between the original grey matter components.*

mean Silhouette score across parcels for the data-driven parcellations is consistently higher than for the equivalent random parcellations. Furthermore, it shows that the ability of the data-driven parcellations to identify meaningful clusters in the data increases with increasing numbers of parcels. On the contrary, for random parcellations, the Silhouette score peaks at $K = 25$, and then decreases for greater values of K . This indicates that the data-driven parcellations provide a better clustering of the data than random parcellations, even when the random parcellation has spatial contiguity enforced.

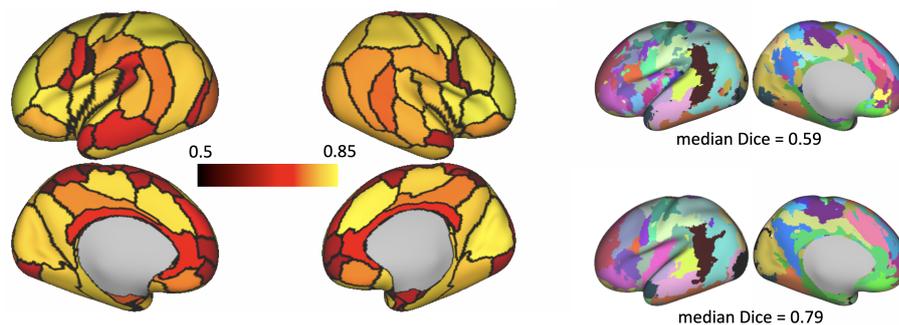


Figure 6.11: *Left: variability of the $K = 100$ parcellation borders, colour coded according to the average Dice coefficient between the group level parcellations with the subject level parcellations from split 1 of the cohort (dark red: small overlap of parcel across subjects, bright yellow: large overlap of parcel across subjects). Right: examples of subject-level parcellations with low and high average Dice coefficient with the group parcellation.*

6.5 Discussion

6.5.1 Reproducibility

Different analyses were performed to test how reproducible the data-driven components were across subject groups. A split-half analysis showed that similar patterns of connectivity could be obtained for large groups of subjects, as shown in figure 6.3. Even for groups of subjects as small as 5, patterns can be obtained that strongly resemble those from the larger groups (over 100 subjects). This shows that this method can be applied to smaller datasets and still provide reliable information.

A new non-negative dual regression method was developed to obtain subject-level representations of components obtained at the group level. Non-negative least squares was used to perform the regression, instead of the pseudoinverse as is normally used (Beckmann et al., 2009; Nickerson et al., 2017). This provides subject-level components that are sparse and non-negative like their group-level counterparts (figure 6.6). Subject-level maps derived from non-negative dual regression were shown to be spatially similar to the results from NMF applied directly to the connectivity matrices of individual subjects. Therefore, using non-negative dual regression provides connectivity patterns that characterise the white matter connections of the individual subject, with the advantage of providing a consistently

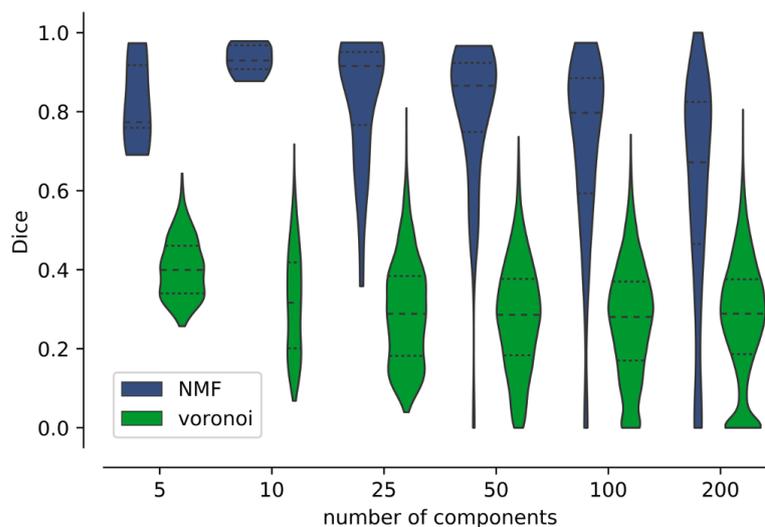


Figure 6.12: *Dice coefficients of matching parcels across the split-half analysis. For comparison, the Dice coefficient was also calculated between one of the split’s NMF parcellations and 100 randomly generated Voronoi parcellations of the same model order.*

ordered set of bundles across subjects for comparison. In addition, non-negative dual regression can be used to project components from one group to another, as well as from groups to individuals. This is demonstrated in figure 6.9, in which group-level components have been projected onto different subgroups of the cohort with different age ranges.

6.5.2 Parcellations

The grey matter maps of the NMF components were used to generate a connectivity-based cortical parcellation scheme. Specifically, each vertex on the cortical mesh was labelled according to the component with the strongest weighting at each point. This leads to a parcellation in which clusters share similar patterns of structural connectivity to the rest of the brain. Depending on the model order of the decomposition, the parcellation can be coarse or more fine-grained (see figure 6.10). An advantage of this approach is that it is entirely data-driven, so the parcellations are not biased by any subjective measures. It can also be used to generate subject specific parcellations, by using the subject-level grey matter maps from dual regression.

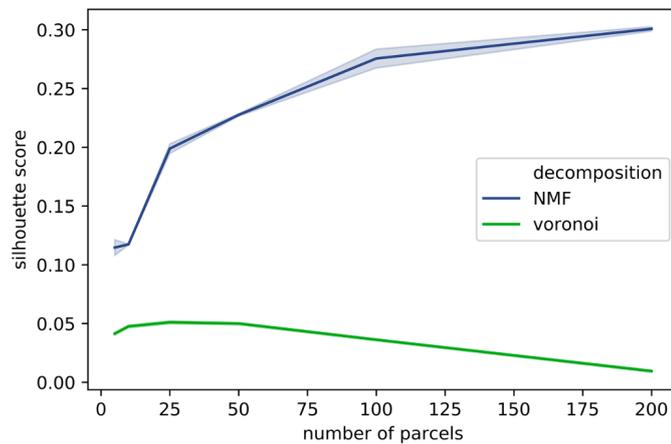


Figure 6.13: Mean Silhouette score across clusters for NMF and Voronoi parcellations with model orders of 5, 10, 25, 50, 100 and 200.

The reproducibility of the parcellations was assessed with a split-half analysis, using Dice coefficient as a similarity measure. The results were compared with a benchmark of the Dice coefficient between one split and a set of randomly generated Voronoi parcellations. For all model orders, the data-driven parcellations were more consistently similar than random parcellations. In addition, Silhouette score was used as a measure of the parcel validity, and the performance of the NMF-based parcellations was again compared against 100 random Voronoi parcellations. Silhouette score measures the similarity of the connectivity profile of a given grey matter vertex to others in its parcel, relative to the connectivity of vertices in other parcels. The data-driven parcellations consistently scored higher on this measure than the random parcellations (see figure 6.13).

Despite this evidence, validating a cortical parcellation is extremely challenging. Existing schemes for the neonatal brain have been derived from manual segmentation of high-resolution data (B. Alexander et al., 2017; B. Alexander et al., 2019), or compared against expert manual segmentations (Adamson et al., 2020; Oishi et al., 2011). While these are extremely useful pieces of work that stem from traditional invasive parcellation approaches, they are based on gyral and sulcal landmarks. These landmarks may not necessarily coincide with functional boundaries (see (Van

Essen and Glasser, 2018) for a recent review). The hope is that connectivity patterns can provide additional information that is closely linked to non-invasive functional delineation, as shown in (Glasser et al., 2016). The NMF framework presented here may be extremely useful for providing another connectivity-based modality, in addition for instance to functional connectivity approaches, and further augment multi-modal parcellations.

Impact of Gyral Bias

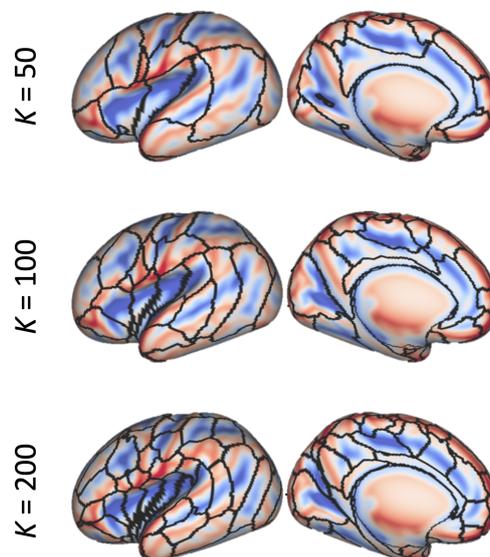


Figure 6.14: *Our group-level NMF parcellation overlaid on the dHCP’s 40-week PMA sulcal depth template (Bozek et al., 2018) (blue for sulcal fundi, red for gyral crowns). There is no consistent pattern for parcellation boundaries to follow sulci or gyri, which indicates that our parcellation is not driven (at least to a large degree) by the gyral bias.*

It has been shown that tractography streamlines are biased towards terminations in the gyri rather than the sulci (Schilling et al., 2018; Van Essen et al., 2013a), although the effects of this “gyral bias” can be minimised by seeding from the cortical surface rather than the whole brain (Donahue et al., 2016; Schilling et al., 2018), as done here. In Chapter 4 it was shown that the effects of gyral bias are less prevalent in neonates than in adults due to the less developed cortical folding. Therefore, less direct influence of such biases in the NMF results in the neonatal

brain would be expected. In fact, the parcellation borders did not show a consistent overlap with sulcal fundi or gyral crowns, as shown in figure 6.14.

6.5.3 Limitations and Future Work

The decomposition framework presented here uses whole-brain tractography data and its performance can therefore be challenged by tractography limitations, which are important to keep in mind when interpreting results. Tractography is an indirect measure of anatomy that is prone to identifying false positive connections (Maier-Hein et al., 2017). False positives in tractography can be demonstrated in two ways: first, in a noisy fashion, causing false paths that are inconsistent either spatially or across subjects; these are less likely to be major drivers of data-driven decompositions. Secondly, in a biased fashion, i.e. consistent false positives that have a certain spatial extent and are reproducible across subjects. These can form the basis of extracted components in NMF, even at the group level. A future application of this framework might be to denoise structural connectivity matrices by removing these components, in a similar manner to ICA-based denoising methods for fMRI (Salimi-Khorshidi et al., 2014). However, a number of indirect validations have been performed in chapters 5 and 6 to gain confidence in the validity and interpretability of the results. NMF decompositions, without any constraints or anatomical knowledge imposed, identified patterns that resembled constrained tractography results in white matter and patterns obtained from an independent modality (rfMRI) in grey matter, and allowed whole-brain connectivity-based parcellations that were reproducible across subjects.

Given the similarity of the structural grey matter networks to those from fMRI, and the common use of data-driven methods to analyse functional data, the work presented here provides a convenient framework for multi-modal analysis of fMRI and dMRI. This could take the form of joint analysis of structural and functional components, or data-fusion approaches in which NMF is applied to concatenated

structural and functional data. In future this could be used to better understand the links between the development of structural and functional connectivity in the brain.

Another route for future work would be the comparison of data-driven components across age-groups and species. An ICA-based data-driven framework has already been used to identify structural connections in the chimpanzee (Mars et al., 2019), where a lack of prior knowledge about brain anatomy makes standard tractography approaches difficult. Using data-driven methods would provide a common reference space for both grey matter and white matter connectivity patterns that can be used to compare between groups, without any bias or constraints from prior knowledge. These ideas are explored further in the next chapter.

6.6 Concluding Remarks

This chapter extends the work in the previous chapter, in which a data-driven framework was developed for mapping white matter bundles and grey matter networks in the neonatal brain. The resultant patterns of connectivity were found to be robust across different subject groups, even for smaller numbers of subjects. A non-negative approach for dual regression was then demonstrated, which provides subject-level representations of the group-level components, without compromising the advantageous non-negative properties of NMF. The results from non-negative dual regression were shown to accurately characterise the individual subjects' connectivity patterns. Finally, the grey matter NMF components were used to generate a connectivity-based parcellation of the neonatal cortex. This parcellation was also shown to be robust and to provide a meaningful clustering of the grey matter according to structural connectivity.

Chapter 7

Conclusions and Future Directions

Contents

7.1	Summary and Conclusions	155
7.2	Future Directions	157
7.2.1	Cross-Modal Connectivity Analysis	157
7.2.2	Applications of Non-Negative Dual Regression	158
7.2.3	Exploring the Ontogeny and Phylogeny of Brain Connections	160
7.2.4	Connectivity-Based Alignment and Parcellation	165
7.3	Final Remarks	166

7.1 Summary and Conclusions

This thesis considers different methods for mapping white matter connections in the neonatal brain, using dMRI tractography. During the neonatal period, the structural connections in the brain grow and mature rapidly. These processes are sensitive to preterm birth and environmental factors (Deoni et al., 2013; Batalle et al., 2017; Lautarescu et al., 2020), and can influence later outcomes (Counsell et al., 2008; Ball et al., 2015; Girault et al., 2019). It is therefore important to have analysis tools that are robust and sensitive to subtle changes in white matter development, if we are to disentangle these complex relationships.

In the Original Research section, different approaches were presented that provide new ways to map white matter connections in the neonatal brain, using methods that are specifically tailored to the challenges of this unique population. High-quality data were used from the developing Human Connectome Project, alongside specially developed preprocessing schemes that are optimised for this age group (Bastiani et al., 2019). This allows us to map white matter connections in the neonatal brain with unprecedented detail and accuracy.

There is a lack of standardised protocols that can be used across different studies to map a comprehensive set of white matter tracts in neonates. To address this problem, automated tractography protocols were developed in chapter 4 to map 42 white matter tracts in the neonatal brain. A key advantage of these protocols is that they have been developed to be compatible with tractography protocols for the same tracts in the adult human and macaque brains. The protocols were robust across different ages and data quality, and were used to generate probabilistic tract atlases from 489 subjects.

The second part of chapter 4 explored the impact of the gyral bias, a methodological confound that causes the tractography streamlines to preferentially terminate at the gyral crowns, rather than the sulcal walls or fundi. It was found that

the gyral bias has less impact on the neonatal connectome than it does in adults, due to the lower level of gyrification in this age group. We also found that its effects on whole-brain connectivity matrices can be mitigated by seeding streamlines from the white matter/grey matter boundary (WGB), rather than throughout the brain volume.

These whole-brain tractography matrices formed the basis of the work presented in chapter 5, in which a framework was presented for decomposing whole brain tractography data into components describing white matter connections and their associated cortical terminations and networks. This data-driven approach provides an alternative to the manually-defined tractography protocols described in the previous chapter and opens new, exciting ways to map connectivity in the brain. We explored and compared two methods for the decomposition, Independent Component Analysis (ICA) and Non-Negative Matrix Factorisation (NMF). NMF was found to better characterise the non-negative signals in simulated data, similar to those in the structural connectome. When applied to real data, the NMF white matter components were shown to resemble a large number of white matter bundles, while the corresponding grey matter components showed similarity to resting-state functional networks, as measured with fMRI.

In chapter 6, this data-driven approach was further augmented by a non-negative dual regression method, which allows the components from a group of subjects to be projected onto data from an individual subject, or another group of subjects, while retaining the advantageous non-negative properties of the original decomposition. Furthermore, the stability of the components was assessed across subjects. Results indicated that the patterns are robust even in small groups of as few as ten subjects. As further evidence of the anatomical relevance of these components, we used them to produce the first connectivity-based atlas of the neonatal cortex, which was shown to be stable across subjects and groups.

In summary, this thesis has developed novel approaches for mapping brain or-

ganisation, which can provide new insight into the development of white matter connections. There is a broad scope for applying these methods to further populations and applications, some of which will be overviewed in the following section.

7.2 Future Directions

The analysis methods that have been developed in the preceding chapters enable new, exciting opportunities for further exploration of brain connectivity and organisation. Not only can these be applied to further our understanding of early neurodevelopment, but also in a wide range of other contexts; whether this be understanding the evolutionary context of the connectome, or developing more accurate subject alignment and cortical parcellation schemes. The following section outlines the main opportunities.

7.2.1 Cross-Modal Connectivity Analysis: Fusing Diffusion MRI and Resting-State Functional MRI

The new data-driven approach developed here for structural connectivity and diffusion MRI data mirrors a common analysis pipeline for resting-state functional MRI data, in which ICA is applied to matrix of functional time-series to yield a set of grey matter components and their associated time courses (McKeown et al., 1998). These resting-state networks are consistent across subjects and correspond to patterns of task-related functional activation (S. M. Smith et al., 2009). In neonates, resting state networks have also been identified that correspond to those found in adults (Doria et al., 2010; Fransson et al., 2007). The connection strength within these networks increases with age during the first years of life (Lin et al., 2008; Gao et al., 2011)

Using a parallel analysis technique for structural connectivity data provides opportunities to explore the links between structure and function (during development

and beyond). The structural components from our NMF approach are obtained in the same grey matter space as the functional networks, which provides a convenient route for joint analysis and detailed comparison of structural and functional components in the shared grey matter space. Figure 5.11 shows an initial comparison of the grey matter components across the two modalities, however, more detailed analysis could explore how the structure-function relationship changes with age or between different groups. Networks from one modality could be used to constrain the analysis of another, by looking at the structural connections between the nodes of a functional network, or vice versa.

Importantly, the data-driven approach presented here can form the basis for devising a data fusion approach to access the joint information in the two modalities. For example, joint decomposition of the concatenated structural connectomes and functional time courses would yield a set of joint structural and functional components. The structural and functional components could also be used as features for other methods, such as linked-ICA (Groves et al., 2011), which aim to derive shared components across modalities that explain inter-subject variation in the data. Such data fusion approaches may provide greater sensitivity to inter-subject variation than post-hoc comparison of separate analyses. As dMRI and fMRI have complementary information on long-range connectivity (but have independent sources of error), such an approach has the potential to identify patterns in a more robust way than independent unimodal analyses.

7.2.2 Applications of Non-Negative Dual Regression

In chapter 6, a non-negative dual regression framework was developed to back project the data-driven structural components onto the data of individual subjects, or of different subgroups. This provides a subject-level (or different group-level) representation of the original group-level components. For example, figure 7.1a shows one of the group-level components dual regressed onto groups containing subjects

of different ages. Visual comparison of these components shows differences in the components across the age range. However, statistical analysis provides a more quantitative route to assess these changes. The spatial maps obtained from our approach can be compared across groups of subjects to look for group differences, using a general linear model (GLM). This approach models the measured data as a linear combination of the explanatory variables (EVs), which are encoded in a design matrix. Such an approach has been followed before for fMRI ICA components (Beckmann et al., 2009; Nickerson et al., 2017). Using non-negative dual regression would provide a similar analysis paradigm for structural components for the first time, as illustrated in figure 7.1.

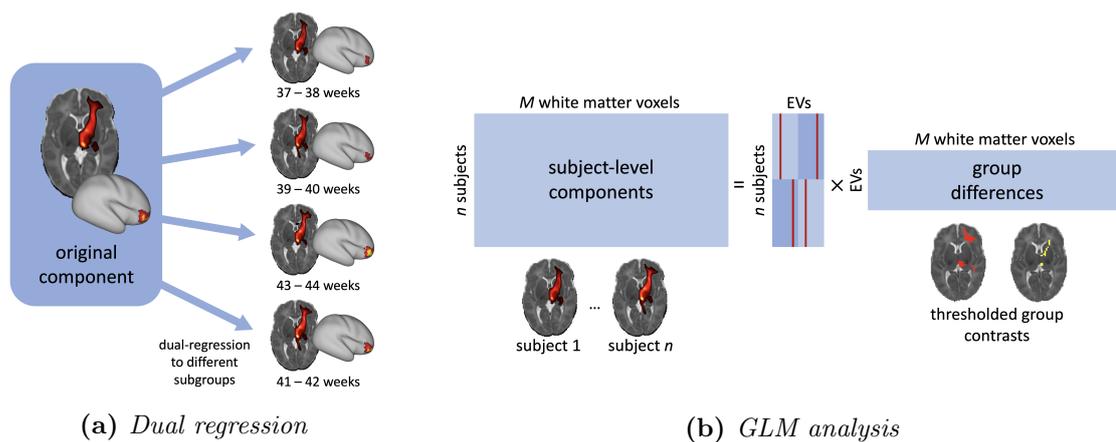


Figure 7.1: Schematic diagram illustrating the statistical analysis of dual-regressed structural components, with a general linear model.

Although the example in figure 7.1 is demonstrated on the white matter spatial maps, it could also be applied to investigate the inter-subject variability in the grey matter structural networks. In the context of neonatal development, this statistical framework could be used to examine changes in structural connectivity with age, or as a result of preterm birth. Beyond this, it could be used to measure the alterations in brain connectivity brought about by aging and disease.

7.2.3 Exploring the Ontogeny and Phylogeny of Brain Connections

The methods presented in chapters 4 - 6 not only facilitate the mapping of white matter connections in the neonatal brain, but also provide routes for joint analysis of connectivity across different age groups and across different primate species. Comparative anatomy studies with non-human primates provide an insight into the phylogeny, or evolutionary development of white matter connections, whereas mapping white matter bundles across the the lifespan allows us to study their ontogeny, in development and aging.

Connectivity Fingerprinting

How are brain connections linked to brain function? And why are they good probes of functional specialisation? A basic tenet in neuroscience is that the pattern of extrinsic connections of a brain region to the rest of the brain is tightly linked to its functional role in the brain network (Passingham et al., 2002). Connections govern the way information flows in and out of a region, therefore areas that have distinct functionality will have a distinct pattern of connections and vice versa. Connectivity patterns can even be used to predict the location and shape of functional activation in individuals, for example in response to faces (Saygin et al., 2012). This idea of unique connectivity signatures or fingerprints underpins brain connectivity research studies to a large extent.

The frameworks and methodology presented here provide novel ways to map and utilise connectivity fingerprints. A corollary of the above idea is that if we can identify similar connection patterns, we can then establish relationships between functionally equivalent regions, even across diverse brains, such as from humans and non-human primates (Mars et al., 2018). The baby-XTRACT protocols described in chapter 4 help in this direction because they are consistent with tractography

protocols for the adult human and macaque brain (Warrington et al., 2020). The protocols have been designed to map tracts that are homologous across species and age-groups. These results can also be used to generate grey matter “connectivity fingerprints”, the patterns of connections (in terms of white matter tracts) belonging to each cortical area.

The process of extracting such connectivity patterns is shown in figure 7.2. The same tracts can be identified in each diverse group (e.g. neonates, adults, macaques), using standardised protocols that are as consistent as possible between species. Tractography results are unwrapped to 1-D, yielding a (whole brain \times tracts) matrix. Taking a (whole brain \times cortex) connectivity matrix (i.e. the connectomes used for data-driven tractography in chapter 5), a dot product of the two matrices will yield a (tracts \times cortex) connectivity matrix, as shown in figure 7.2. The rows of this matrix give the cortical termination patterns of each tract, whereas the columns provide the connectivity fingerprints of each of the grey matter locations.

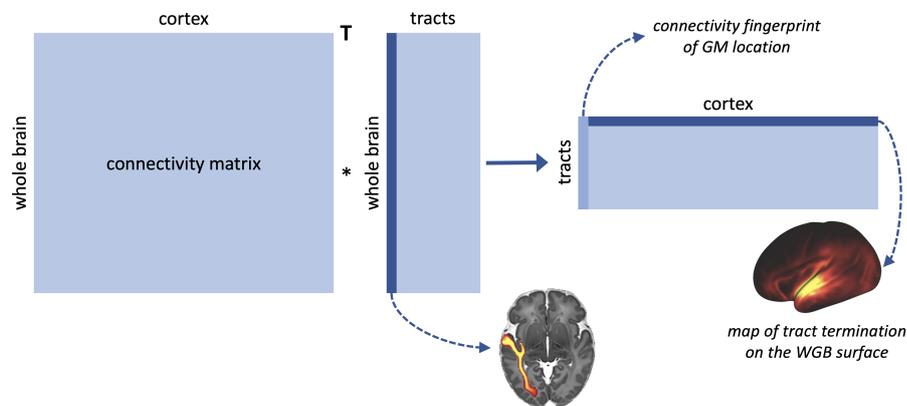


Figure 7.2: *Connectivity fingerprints are calculated by taking the dot product of a cortex-to-whole brain connectivity matrix with a matrix of tractography maps, unwrapped to 1D.*

A selection of the tractography results and their associated cortical termination patterns are shown in figure 7.3, alongside those from the adult human and macaque. These results show how the XTRACT protocols can be used to map consistent patterns of connectivity in these diverse populations, in both the grey and white matter.

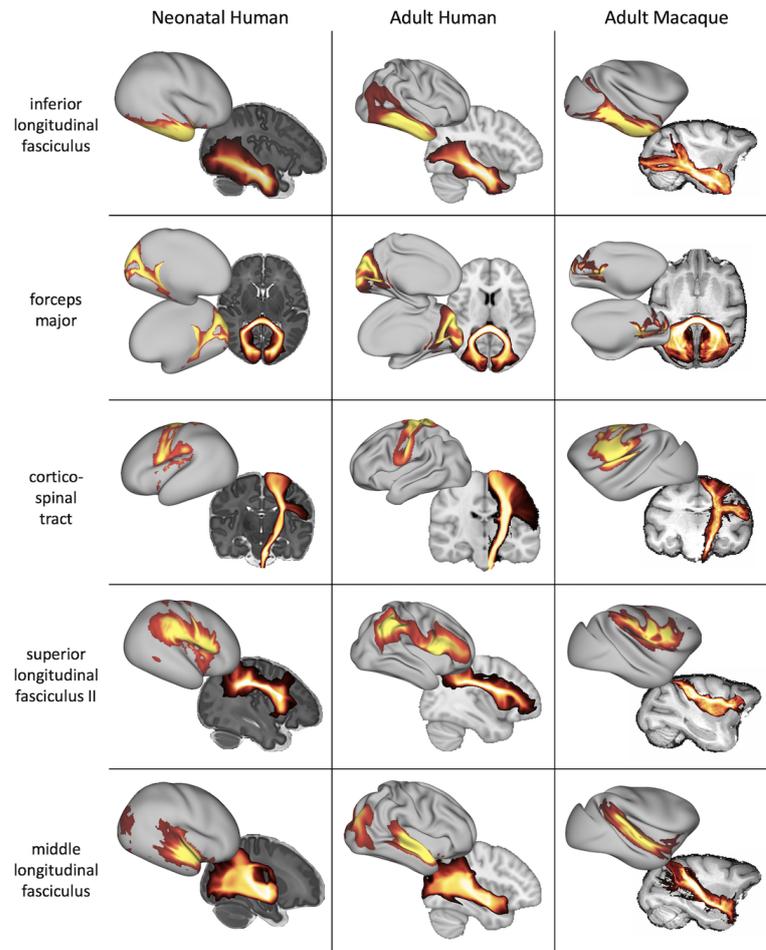


Figure 7.3: Cross-species comparison of tracts and cortical termination patterns from the XTRACT protocols. Adult human and macaque results were obtained from (Warrington *et al.*, 2020).

The connectivity fingerprints can be thought of as discrete probability distributions and can be compared using standard metrics of distribution similarity. For instance, Kullback-Leibler divergence (KL divergence) (Kullback and Leibler, 1951) can be used as a measure of the relative entropy between two probability distributions. Normalising the connectivity fingerprint of each grey matter location to sum to one yields a probability distribution describing the likelihood of each tract’s contribution to its connectivity profile. KL divergence can then be used to compare each vertex’s connectivity profile with that of each vertex from the other group. A schematic representation of this process is shown in figure 7.4, comparing the connectivity fingerprint matrix from neonatal subjects N , with that of adult subjects A .

Mars et al. (2018) used this approach to identify homologue areas between human and non-human primates. The same idea can be used to study neurodevelopment and correspondences between the neonatal and the adult brain.

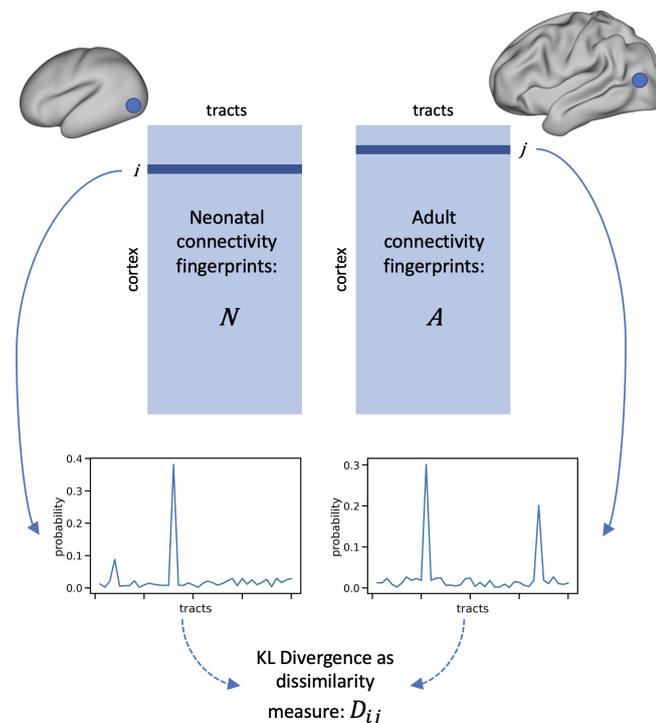


Figure 7.4: Schematic diagram illustrating how Kullback-Leibler divergence can be used to compare connectivity profiles of different groups. The connectivity profiles of different vertices are expressed in terms of the contributions of the different tracts, which can then be compared using KL divergence.

Data-Driven Components

In the previous section the connectivity fingerprints were described in terms of manually-defined landmarks: white matter tracts obtained using functionally equivalent tractography protocols (i.e. baby-XTRACT and XTRACT protocols). This has two limitations: firstly, a set of such optimised protocols needs to exist, which may not always be the case; secondly, these protocols have a limited number of tracts, which may not be adequate for all patterns and regions that can depend on other tracts that are not represented in the protocols.

An alternative for exploration of connectivity patterns across species and age-

groups would be to use data-driven white matter components as reference landmarks. Our NMF framework provides an ideal avenue for this type of analysis, as it generates both white matter tracts and their matched grey matter cortical terminations, without the need for manually-defined protocols. Furthermore, the data-driven components are derived from the fully sampled connectome, so describe both long-range and short-range fibre bundles that may not characterised by standard protocols, depending on the dimensionality of the decomposition. Early results comparing NMF components from neonates and adults indicate that the NMF components will provide a valuable basis for connectivity analysis across the lifespan. A selection of matched components, shown in figure 7.5, show very strong similarity across these two distinct groups.

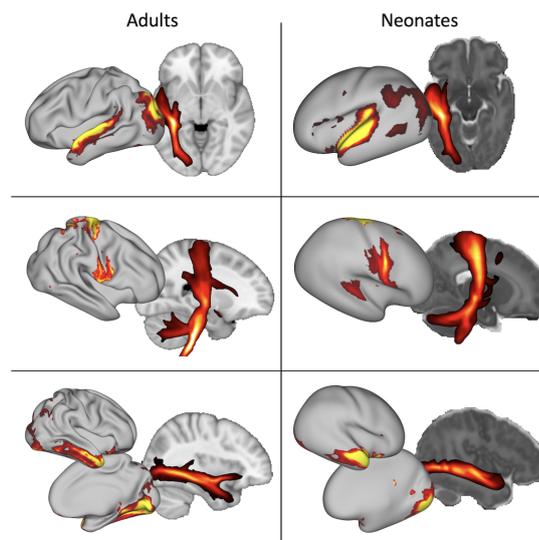


Figure 7.5: Comparison of data-driven NMF components from adult and neonatal subjects. Adult data were acquired from the HCP.

In addition, data-driven methods are advantageous for studying connectivity in species where our knowledge of the brain anatomy is incomplete. For example, ICA decomposition of connectivity matrices has been used to study white matter connections in the chimpanzee, a species for which there are no modern atlases of the neocortex (Mars et al., 2019).

7.2.4 Connectivity-Based Alignment and Parcellation

Connectivity-based features have recently been used to improve methods for between-subject cortical alignment and cortical parcellation. Surface-based registration provides a more natural approach for alignment of features on the cortex than volumetric registration, but previous methods have relied on cortical folding to drive the registration. Including connectivity information in addition to geometric features has been shown to improve the alignment of functionally homologous brain regions between subjects (Robinson et al., 2014; Robinson et al., 2018). However, to date only fMRI have been used for this purpose on a large scale.

Another domain in which multiple modalities can provide complementary information is brain parcellation. Parcellating the cortex into neurobiologically distinct regions has been a major goal of neuroscience for over a century (Van Essen and Glasser, 2018). A multi-modal parcellation developed by Glasser et al. (2016) was a ground-breaking effort in this direction, incorporating functional connectivity data alongside information about myelination and cortical thickness to generate a parcellation that is robust and provides neuroanatomically precise delineation of 180 areas in each hemisphere. As described above, areas of functional specialisation can be identified through their structural connectivity fingerprint, and so NMF features may provide an additional insight into brain organisation that can build upon these methods.

The new NMF framework presented in chapter 5 provides data-driven structural connectivity information on the cortex that could be incorporated into these schemes in future, to provide additional information and improve accuracy. Figure 7.6 shows how structural connectivity components from NMF could be combined with information from other modalities to drive both parcellation and surface-based registration.

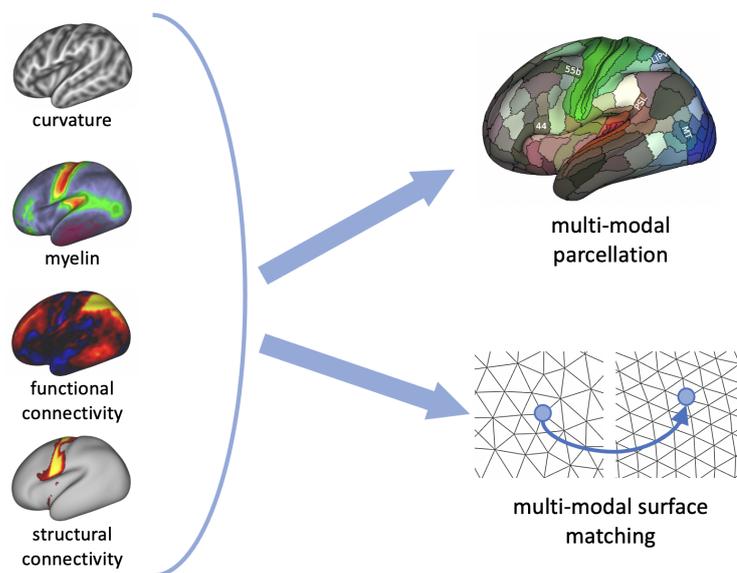


Figure 7.6: Schematic diagram illustrating how complementary information from different modalities can be used to inform cortical parcellation schemes or surface-based registration. Multi-modal parcellation taken from (Glasser et al., 2016).

7.3 Final Remarks

In this thesis, analytical techniques have been developed to map white matter connectivity in a flexible and robust manner, based on the unique challenges of neonatal brain development. In this final chapter we have discussed how these methods can be applied beyond the field of neonatal imaging, to map both the ontogeny and phylogeny of brain connections, as well as to explore the links between brain structure and function.

The proliferation of “Big Data” projects provide a wealth of high quality data for the research community, to improve our understanding of typical brain function throughout the lifespan and also how brain health deviates from this trajectory in disease. New analysis approaches are needed to make the most of these data and to combine the complementary information provided by different modalities. Computational advances mean that data-driven methods and large datasets can now be used widely to probe new questions that would not be accessible with traditional

approaches. This thesis contributes in this direction, by extending a framework of reproducible tractography protocols, and by introducing new data-driven methods for connectivity mapping.

Bibliography

- [1] C. L. Adamson, B. Alexander, G. Ball, R. Beare, J. L. Y. Cheong, A. J. Spittle, L. W. Doyle, P. J. Anderson, M. L. Seal, and D. K. Thompson. “Parcellation of the neonatal cortex using Surface-based Melbourne Children’s Regional Infant Brain atlases (M-CRIB-S)”. In: *Scientific Reports* 10.1 (2020), p. 4359.
- [2] A. M. Aertsen, G. L. Gerstein, M. K. Habib, and G. Palm. “Dynamics of neuronal firing correlation: modulation of “effective connectivity””. In: *Journal of Neurophysiology* 61.5 (1989), pp. 900–917.
- [3] I. Aganj, C. Lenglet, G. Sapiro, E. Yacoub, K. Ugurbil, and N. Harel. “Reconstruction of the orientation distribution function in single- and multiple-shell q-ball imaging within constant solid angle”. In: *Magnetic Resonance in Medicine* 64.2 (2010), pp. 554–566.
- [4] K. Akazawa, L. Chang, R. Yamakawa, S. Hayama, S. Buchthal, D. Alicata, T. Andres, D. Castillo, K. Oishi, J. Skranes, T. Ernst, and K. Oishi. “Probabilistic maps of the white matter tracts with known associated functions on the neonatal brain atlas: Application to evaluate longitudinal developmental trajectories in term-born and preterm-born infants”. In: *NeuroImage* 128 (2016), pp. 167–179.
- [5] B. Alexander, W. Y. Loh, L. G. Matthews, A. L. Murray, C. Adamson, R. Beare, J. Chen, C. E. Kelly, P. J. Anderson, L. W. Doyle, A. J. Spittle, J. L. Y. Cheong, M. L. Seal, and D. K. Thompson. “Desikan-Killiany-Tourville Atlas Compatible Version of M-CRIB Neonatal Parcellated Whole Brain Atlas: The M-CRIB 2.0”. In: *Frontiers in Neuroscience* 13 (2019), p. 34.
- [6] B. Alexander, A. L. Murray, W. Y. Loh, L. G. Matthews, C. Adamson, R. Beare, J. Chen, C. E. Kelly, S. Rees, S. K. Warfield, P. J. Anderson, L. W. Doyle, A. J. Spittle, J. L. Y. Cheong, M. L. Seal, and D. K. Thompson. “A new neonatal cortical and subcortical brain atlas: the Melbourne Children’s Regional Infant Brain (M-CRIB) atlas”. In: *NeuroImage* 147.Supplement C (2017), pp. 841–851.
- [7] D. C. Alexander. “An Introduction to Computational Diffusion MRI: the Diffusion Tensor and Beyond”. In: *Mathematics and Visualization*. Ed. by J. Weickert and H. Hagen. Springer, 2006, pp. 83–106.
- [8] D. C. Alexander, T. B. Dyrby, M. Nilsson, and H. Zhang. “Imaging brain microstructure with diffusion MRI: practicality and applications”. In: *NMR in Biomedicine* 32.4 (2019), e3841.

- [9] D. Anblagan, M. E. Bastin, S. Sparrow, C. Piyasena, R. Pataky, E. J. Moore, A. Serag, A. G. Wilkinson, J. D. Clayden, S. I. Semple, and J. P. Boardman. “Tract shape modeling detects changes associated with preterm birth and neuroprotective treatment effects”. In: *NeuroImage: Clinical* 8 (2015), pp. 51–58.
- [10] A. W. Anderson. “Measurement of fiber orientation distributions using high angular resolution diffusion imaging”. In: *Magnetic Resonance in Medicine* 54.5 (2005), pp. 1194–1206.
- [11] J. L. Andersson and S. N. Sotiropoulos. “An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging”. In: *NeuroImage* 125 (2016), pp. 1063–1078.
- [12] J. L. Andersson, M. S. Graham, I. Drobnjak, H. Zhang, N. Filippini, and M. Bastiani. “Towards a comprehensive framework for movement and distortion correction of diffusion MR images: Within volume movement”. In: *NeuroImage* 152 (2017), pp. 450–466.
- [13] M. Anjari, L. Srinivasan, J. M. Allsop, J. V. Hajnal, M. A. Rutherford, A. D. Edwards, and S. J. Counsell. “Diffusion tensor imaging with tract-based spatial statistics reveals local white matter abnormalities in preterm infants”. In: *NeuroImage* 35.3 (2007), pp. 1021–1027.
- [14] Y. Assaf and P. J. Basser. “Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain”. In: *NeuroImage* 27.1 (2005), pp. 48–58.
- [15] F. Aurenhammer. “Voronoi diagrams—a survey of a fundamental geometric data structure”. In: *ACM Computing Surveys (CSUR)* 23.3 (1991), pp. 345–405.
- [16] B. B. Avants, N. J. Tustison, G. Song, P. A. Cook, A. Klein, and J. C. Gee. “A reproducible evaluation of ANTs similarity metric performance in brain image registration”. In: *NeuroImage* 54.3 (2011), pp. 2033–2044.
- [17] G. Ball, P. Aljabar, S. Zebari, N. Tusor, T. Arichi, N. Merchant, E. C. Robinson, E. Ogunidipe, D. Rueckert, A. D. Edwards, and S. J. Counsell. “Rich-club organization of the newborn human brain”. In: *Proceedings of the National Academy of Sciences* 111.20 (2014), pp. 7456–7461.
- [18] G. Ball, R. Beare, and M. L. Seal. “Charting shared developmental trajectories of cortical thickness and structural connectivity in childhood and adolescence”. In: *Human Brain Mapping* 40.16 (2019), pp. 4630–4644.
- [19] G. Ball, J. P. Boardman, P. Aljabar, A. Pandit, T. Arichi, N. Merchant, D. Rueckert, A. D. Edwards, and S. J. Counsell. “The influence of preterm birth on the developing thalamocortical connectome”. In: *Cortex* 49.6 (2013), pp. 1711–1721.
- [20] G. Ball, S. J. Counsell, M. Anjari, N. Merchant, T. Arichi, V. Doria, M. A. Rutherford, A. D. Edwards, D. Rueckert, and J. P. Boardman. “An optimised tract-based spatial statistics protocol for neonates: Applications to prematurity and chronic lung disease”. In: *NeuroImage* 53.1 (2010), pp. 94–102.

- [21] G. Ball, L. Pazderova, A. Chew, N. Tusor, N. Merchant, T. Arichi, J. M. Allsop, F. M. Cowan, A. D. Edwards, and S. J. Counsell. “Thalamocortical connectivity predicts cognition in children born preterm”. In: *Cerebral Cortex* 25.11 (2015), pp. 4310–4318.
- [22] M. J. Barkovich, Y. Li, R. S. Desikan, A. J. Barkovich, and D. Xu. “Challenges in pediatric neuroimaging”. In: *NeuroImage* 185 (2019), pp. 793–801.
- [23] P. J. Basser, J. Mattiello, and D. LeBihan. “MR diffusion tensor spectroscopy and imaging”. In: *Biophysical Journal* 66.1 (1994), pp. 259–267.
- [24] P. J. Basser. “Inferring microstructural features and the physiological state of tissues from diffusion-weighted images”. In: *NMR in Biomedicine* 8.7 (1995), pp. 333–344.
- [25] P. J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. “In vivo fiber tractography using DT-MRI data”. In: *Magnetic Resonance in Medicine* 44.4 (2000), pp. 625–632.
- [26] P. J. Basser and C. Pierpaoli. “Microstructural and Physiological Features of Tissues Elucidated by Quantitative-Diffusion-Tensor MRI”. In: *Journal of Magnetic Resonance, Series B* 111.3 (1996), pp. 209–219.
- [27] D. S. Bassett and E. Bullmore. “Small-World Brain Networks”. In: *The Neuroscientist* 12.6 (2006), pp. 512–523.
- [28] M. Bastiani, J. L. Andersson, L. Cordero-Grande, M. Murgasova, J. Hutter, A. N. Price, A. Makropoulos, S. P. Fitzgibbon, E. Hughes, D. Rueckert, S. Victor, M. Rutherford, A. D. Edwards, S. M. Smith, J.-D. Tournier, J. V. Hajnal, S. Jbabdi, and S. N. Sotiropoulos. “Automated processing pipeline for neonatal diffusion MRI in the developing Human Connectome Project”. In: *NeuroImage* 185 (2019), pp. 750–763.
- [29] M. Bastiani, M. Cottaar, K. Dikranian, A. Ghosh, H. Zhang, D. C. Alexander, T. E. Behrens, S. Jbabdi, and S. N. Sotiropoulos. “Improved tractography using asymmetric fibre orientation distributions”. In: *NeuroImage* 158 (2017), pp. 205–218.
- [30] D. Batalle, A. D. Edwards, and J. O’Muircheartaigh. “Annual Research Review: Not just a small adult brain: understanding later neurodevelopment through imaging the neonatal brain”. In: *Journal of Child Psychology and Psychiatry* 59.4 (2018), pp. 350–371.
- [31] D. Batalle, E. J. Hughes, H. Zhang, J.-D. Tournier, N. Tusor, P. Aljabar, L. Wali, D. C. Alexander, J. V. Hajnal, C. Nosarti, A. D. Edwards, and S. J. Counsell. “Early development of structural networks and the impact of prematurity on brain connectivity”. In: *NeuroImage* 149.Supplement C (2017), pp. 379–392.
- [32] D. Batalle, J. O’Muircheartaigh, A. Makropoulos, C. J. Kelly, R. Dimitrova, E. J. Hughes, J. V. Hajnal, H. Zhang, D. C. Alexander, A. D. Edwards, and S. J. Counsell. “Different patterns of cortical maturation before and after 38 weeks gestational age demonstrated by diffusion MRI in vivo”. In: *NeuroImage* 185 (2019), pp. 764–775.

- [33] C. Beaulieu. “The basis of anisotropic water diffusion in the nervous system - a technical review”. In: *NMR in Biomedicine* 15.7-8 (2002), pp. 435–455.
- [34] C. F. Beckmann, C. E. Mackay, N. Filippini, and S. M. Smith. “Group comparison of resting-state fMRI data using multi-subject ICA and dual regression”. In: *NeuroImage* 47 (2009), S148.
- [35] C. F. Beckmann and S. M. Smith. “Probabilistic Independent Component Analysis for Functional Magnetic Resonance Imaging”. In: *IEEE Transactions on Medical Imaging* 23.2 (2004), pp. 137–152.
- [36] T. E. J. Behrens, H. J. Berg, S. Jbabdi, M. F. S. Rushworth, and M. W. Woolrich. “Probabilistic diffusion tractography with multiple fibre orientations: What can we gain?” In: *NeuroImage* 34.1 (2007), pp. 144–155.
- [37] T. E. J. Behrens, M. W. Woolrich, M. Jenkinson, H. Johansen-Berg, R. G. Nunes, S. Clare, P. M. Matthews, J. M. Brady, and S. M. Smith. “Characterization and propagation of uncertainty in diffusion-weighted MR imaging”. In: *Magnetic Resonance in Medicine* 50.5 (2003), pp. 1077–1088.
- [38] T. E. J. Behrens, S. N. Sotiropoulos, and S. Jbabdi. “MR Diffusion Tractography”. In: *Diffusion MRI*. Elsevier, 2014, pp. 429–451.
- [39] B. Biswal, F. Zerrin Yetkin, V. M. Haughton, and J. S. Hyde. “Functional connectivity in the motor cortex of resting human brain using echo-planar MRI”. In: *Magnetic Resonance in Medicine* 34.4 (1995), pp. 537–541.
- [40] F. Bloch. “Nuclear Induction”. In: *Physical Review* 70.7-8 (1946), pp. 460–474.
- [41] N. Bloembergen, E. M. Purcell, and R. V. Pound. “Relaxation Effects in Nuclear Magnetic Resonance Absorption”. In: *Physical Review* 73.7 (1948), pp. 679–712.
- [42] C. Boutsidis and E. Gallopoulos. “SVD based initialization: A head start for nonnegative matrix factorization”. In: *Pattern Recognition* 41.4 (2008), pp. 1350–1362.
- [43] J. Bozek, A. Makropoulos, A. Schuh, S. Fitzgibbon, R. Wright, M. F. Glasser, T. S. Coalson, J. O’Muircheartaigh, J. Hutter, A. N. Price, L. Cordero-Grande, R. P. A. Teixeira, E. Hughes, N. Tustor, K. P. Baruteau, M. A. Rutherford, A. D. Edwards, J. V. Hajnal, S. M. Smith, D. Rueckert, M. Jenkinson, and E. C. Robinson. “Construction of a neonatal cortical surface atlas using Multimodal Surface Matching in the Developing Human Connectome Project”. In: *NeuroImage* 179 (2018), pp. 11–29.
- [44] B. A. Brody, H. C. Kinney, A. S. Kloman, and F. H. Gilles. “Sequence of central nervous system myelination in human infancy. I. An autopsy study of myelination”. In: *Journal of Neuropathology and Experimental Neurology* 46.3 (1987), pp. 283–301.
- [45] C. J. Brown, S. P. Miller, B. G. Booth, S. Andrews, V. Chau, K. J. Poskitt, and G. Hamarneh. “Structural network analysis of brain development in young preterm neonates”. In: *NeuroImage* 101 (2014), pp. 667–680.

- [46] E. Bullmore and O. Sporns. “Complex brain networks: graph theoretical analysis of structural and functional systems”. In: *Nature Reviews Neuroscience* 10.3 (2009), pp. 186–198.
- [47] R. B. Buxton. “The physics of functional magnetic resonance imaging (fMRI)”. In: *Reports on progress in physics. Physical Society (Great Britain)* 76.9 (2013), p. 96601.
- [48] M. Catani and M. Schotten Thiebaut De. “A diffusion tensor imaging tractography atlas for virtual in vivo dissections”. In: *Cortex* 44.8 (2008), pp. 1105–1132.
- [49] M. Catani, F. Dell’Acqua, F. Vergani, F. Malik, H. Hodge, P. Roy, R. Valabregue, and M. Thiebaut de Schotten. “Short frontal lobe connections of the human brain”. In: *Cortex* 48.2 (2012), pp. 273–291.
- [50] M. Catani, F. Dell’Acqua, and M. Thiebaut de Schotten. “A revised limbic system model for memory, emotion and behaviour”. In: *Neuroscience & Biobehavioral Reviews* 37.8 (2013), pp. 1724–1737.
- [51] M. Catani, R. J. Howard, S. Pajevic, and D. K. Jones. “Virtual in Vivo interactive dissection of white matter fasciculi in the human brain”. In: *NeuroImage* 17.1 (2002), pp. 77–94.
- [52] M. Catani, M. Thiebaut de Schotten, D. Slater, and F. Dell’Acqua. “Connectomic approaches before the connectome”. In: *NeuroImage* 80 (2013), pp. 2–13.
- [53] H. Chen, T. Zhang, L. Guo, K. Li, X. Yu, L. Li, X. Hu, J. Han, X. Hu, and T. Liu. “Coevolution of gyral folding and structural connection patterns in primate brains”. In: *Cerebral Cortex* 23.5 (2013), pp. 1208–1217.
- [54] A. Cichocki and A.-H. Phan. “Fast Local Algorithms for Large Scale Non-negative Matrix and Tensor Factorizations”. In: *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences* E92-A.3 (2009), pp. 708–721.
- [55] A. H. Cohen, R. Wang, M. Wilkinson, P. MacDonald, A. R. Lim, and E. Takahashi. “Development of human white matter fiber pathways: From newborn to adult ages”. In: *International Journal of Developmental Neuroscience* 50.1 (2016), pp. 26–38.
- [56] T. E. Conturo, N. F. Lori, T. S. Cull, E. Akbudak, A. Z. Snyder, J. S. Shimony, R. C. McKinstry, H. Burton, and M. E. Raichle. “Tracking neuronal fiber pathways in the living human brain”. In: *Proceedings of the National Academy of Sciences* 96.18 (1999), pp. 10422–10427.
- [57] L. Cordero-Grande, E. J. Hughes, J. Hutter, A. N. Price, and J. V. Hajnal. “Three-dimensional motion corrected sensitivity encoding reconstruction for multi-shot multi-slice MRI: Application to neonatal brain imaging”. In: *Magnetic Resonance in Medicine* 79.3 (2018), pp. 1365–1376.

- [58] S. J. Counsell, A. D. Edwards, A. T. Chew, M. Anjari, L. E. Dyet, L. Srinivasan, J. P. Boardman, J. M. Allsop, J. V. Hajnal, M. A. Rutherford, and F. M. Cowan. “Specific relations between neurodevelopmental abilities and white matter microstructure in children born preterm”. In: *Brain* 131.12 (2008), pp. 3201–3208.
- [59] M. De Groot, M. W. Vernooij, S. Klein, M. A. Ikram, F. M. Vos, S. M. Smith, W. J. Niessen, and J. L. R. Andersson. “Improving alignment in Tract-based spatial statistics: Evaluation and optimization of image registration”. In: *NeuroImage* 76 (2013), pp. 400–411.
- [60] L. De Lathauwer, B. De Moor, and J. Vandewalle. “A Multilinear Singular Value Decomposition”. In: *SIAM Journal on Matrix Analysis and Applications* 21.4 (2000), pp. 1253–1278.
- [61] D. C. Dean, E. M. Planalp, W. Wooten, N. Adluru, S. R. Keckskemeti, C. Frye, C. K. Schmidt, N. L. Schmidt, M. A. Styner, H. H. Goldsmith, R. J. Davidson, and A. L. Alexander. “Mapping White Matter Microstructure in the One Month Human Brain”. In: *Scientific Reports* 7.1 (2017), p. 9759.
- [62] F. Dell’Acqua, G. Rizzo, P. Scifo, R. A. Clarke, G. Scotti, and F. Fazio. “A model-based deconvolution approach to solve fiber crossing in diffusion-weighted MR imaging”. In: *IEEE Transactions on Biomedical Engineering* 54.3 (2007), pp. 462–472.
- [63] F. Dell’Acqua, P. Scifo, G. Rizzo, M. Catani, A. Simmons, G. Scotti, and F. Fazio. “A modified damped Richardson–Lucy algorithm to reduce isotropic background effects in spherical deconvolution”. In: *NeuroImage* 49.2 (2010), pp. 1446–1458.
- [64] F. Dell’Acqua and J. D. Tournier. “Modelling white matter with spherical deconvolution: How and why?” In: *NMR in Biomedicine* 32.4 (2019), e3945.
- [65] S. C. Deoni, D. C. Dean, I. Piryatinsky, J. O’Muircheartaigh, N. Waskiewicz, K. Lehman, M. Han, and H. Dirks. “Breastfeeding and early white matter development: A cross-sectional study”. In: *NeuroImage* 82 (2013), pp. 77–86.
- [66] M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche. “Regularized, fast, and robust analytical Q-ball imaging”. In: *Magnetic Resonance in Medicine* 58.3 (2007), pp. 497–510.
- [67] Z. Ding, J. C. Gore, and A. W. Anderson. “Classification and quantification of neuronal fiber pathways using diffusion tensor MRI”. In: *Magnetic Resonance in Medicine* 49.4 (2003), pp. 716–721.
- [68] C. J. Donahue, S. N. Sotiropoulos, S. Jbabdi, M. Hernandez-Fernandez, T. E. Behrens, T. B. Dyrby, T. Coalson, H. Kennedy, K. Knoblauch, D. C. Van Essen, and M. F. Glasser. “Using diffusion tractography to predict cortical connection strength and distance: A quantitative comparison with tracers in the monkey”. In: *Journal of Neuroscience* 36.25 (2016), pp. 6758–6770.

- [69] V. Doria, C. F. Beckmann, T. Arichi, N. Merchant, M. Groppo, F. E. Turkheimer, S. J. Counsell, M. Murgasova, P. Aljabar, R. G. Nunes, D. J. Larkman, G. Rees, and A. D. Edwards. “Emergence of resting state networks in the preterm human brain”. In: *Proceedings of the National Academy of Sciences* 107.46 (2010), pp. 20015–20020.
- [70] J. Dubois, M. Benders, A. Cachia, F. Lazeyras, R. Ha-Vinh Leuchter, S. V. Sizonenko, C. Borradori-Tolsa, J. F. Mangin, and P. S. Hüppi. “Mapping the Early Cortical Folding Process in the Preterm Newborn Brain”. In: *Cerebral Cortex* 18.6 (2008), pp. 1444–1454.
- [71] J. Dubois, G. Dehaene-Lambertz, S. Kulikova, C. Poupon, P. Hüppi, and L. Hertz-Pannier. “The early development of brain white matter: A review of imaging studies in fetuses, newborns and infants”. In: *Neuroscience* 276 (2014), pp. 48–71.
- [72] J. Dubois, L. Hertz-Pannier, A. Cachia, J. F. Mangin, D. Le Bihan, and G. Dehaene-Lambertz. “Structural Asymmetries in the Infant Language and Sensori-Motor Networks”. In: *Cerebral Cortex* 19.2 (2009), pp. 414–423.
- [73] J. Dubois, G. Dehaene-Lambertz, M. Perrin, J. F. Mangin, Y. Cointepas, E. Duchesnay, D. Le Bihan, and L. Hertz-Pannier. “Asynchrony of the early maturation of white matter bundles in healthy infants: Quantitative landmarks revealed noninvasively by diffusion tensor imaging”. In: *Human Brain Mapping* 29.1 (2008), pp. 14–27.
- [74] N. Eichert, L. Verhagen, D. Folloni, S. Jbabdi, A. A. Khrapitchev, N. R. Sibson, D. Mantini, J. Sallet, and R. B. Mars. “What is special about the human arcuate fasciculus? Lateralization, projections, and expansion”. In: *Cortex* 118 (2019), pp. 107–115.
- [75] A. Einstein. “Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen”. In: *Annalen der Physik* 322.8 (1905), pp. 549–560.
- [76] R. W. Emerson, W. Gao, and W. Lin. “Longitudinal study of the emerging functional connectivity asymmetry of primary language regions during infancy”. In: *Journal of Neuroscience* 36.42 (2016), pp. 10883–10892.
- [77] Y. Fan, F. Shi, J. K. Smith, W. Lin, J. H. Gilmore, and D. Shen. “Brain anatomical networks in early human brain development”. In: *NeuroImage* 54.3 (2011), pp. 1862–1871.
- [78] C. Févotte and J. Idier. “Algorithms for Nonnegative Matrix Factorization with the β -Divergence”. In: *Neural Computation* 23.9 (2011), pp. 2421–2456.
- [79] B. Fischl. “FreeSurfer”. In: *NeuroImage* 62.2 (2012), pp. 774–781.
- [80] S. Fitzgibbon et al. “The developing Human Connectome Project (dHCP) automated resting-state functional processing framework for newborn infants”. In: *bioRxiv* (2019), p. 766030.

- [81] P. Fransson, B. Skiöld, S. Horsch, A. Nordell, M. Blennow, H. Lagercrantz, and U. Åden. “Resting-state networks in the infant brain”. In: *Proceedings of the National Academy of Sciences of the United States of America* 104.39 (2007), pp. 15531–15536.
- [82] K. J. Friston, C. D. Frith, P. F. Liddle, and R. S. J. Frackowiak. “Functional Connectivity: The Principal-Component Analysis of Large (PET) Data Sets”. In: *Journal of Cerebral Blood Flow & Metabolism* 13.1 (1993), pp. 5–14.
- [83] W. Gao, J. H. Gilmore, K. S. Giovanello, J. K. Smith, D. Shen, H. Zhu, and W. Lin. “Temporal and spatial evolution of brain network topology during the first two years of life”. In: *PLoS ONE* 6.9 (2011), e25278.
- [84] E. Garyfallidis, M. Brett, M. M. Correia, G. B. Williams, and I. Nimmo-Smith. “QuickBundles, a Method for Tractography Simplification”. In: *Frontiers in Neuroscience* 6 (2012), p. 175.
- [85] X. Geng, S. Gouttard, A. Sharma, H. Gu, M. Styner, W. Lin, G. Gerig, and J. H. Gilmore. “Quantitative tract-based white matter development from birth to age 2years”. In: *NeuroImage* 61.3 (2012), pp. 542–557.
- [86] J. B. Girault, B. C. Munsell, D. Puechmaille, B. D. Goldman, J. C. Prieto, M. Styner, and J. H. Gilmore. “White matter connectomes at birth accurately predict cognitive abilities at age 2”. In: *NeuroImage* 192 (2019), pp. 145–155.
- [87] M. F. Glasser, T. S. Coalson, E. C. Robinson, C. D. Hacker, J. Harwell, E. Yacoub, K. Ugurbil, J. Andersson, C. F. Beckmann, M. Jenkinson, S. M. Smith, and D. C. Van Essen. “A multi-modal parcellation of human cerebral cortex”. In: *Nature* 536.7615 (2016), pp. 171–178.
- [88] M. F. Glasser, S. N. Sotiropoulos, J. A. Wilson, T. S. Coalson, B. Fischl, J. L. Andersson, J. Xu, S. Jbabdi, M. Webster, J. R. Polimeni, D. C. Van Essen, and M. Jenkinson. “The minimal preprocessing pipelines for the Human Connectome Project”. In: *NeuroImage* 80 (2013), pp. 105–124.
- [89] M. Glickstein. “Golgi and Cajal: The neuron doctrine and the 100th anniversary of the 1906 Nobel Prize”. In: *Current Biology* 16.5 (2006), R147–R151.
- [90] I. S. Gousias, A. D. Edwards, M. A. Rutherford, S. J. Counsell, J. V. Hajnal, D. Rueckert, and A. Hammers. “Magnetic resonance imaging of the newborn brain: Manual segmentation of labelled atlases in term-born and preterm infants”. In: *NeuroImage* 62.3 (2012), pp. 1499–1509.
- [91] A. R. Groves, C. F. Beckmann, S. M. Smith, and M. W. Woolrich. “Linked independent component analysis for multimodal data fusion”. In: *NeuroImage* 54.3 (2011), pp. 2198–2217.
- [92] A. Gupta, M. Poe, M. Styner, A. Panigrahy, and M. Escolar. “Regional differences in fiber tractography predict neurodevelopmental outcomes in neonates with infantile Krabbe disease”. In: *NeuroImage: Clinical* 7 (2015), pp. 792–798.

- [93] P. Hagmann, L. Cammoun, X. Gigandet, R. Meuli, C. J. Honey, V. J. Wedeen, and O. Sporns. “Mapping the Structural Core of Human Cerebral Cortex”. In: *PLoS Biology* 6.7 (2008), e159.
- [94] E. L. Hahn. “Spin Echoes”. In: *Physical Review* 80.4 (1950), pp. 580–594.
- [95] M. N. Hallquist and F. G. Hillary. “Graph theory approaches to functional network organization in brain disorders: A critique for a brave new small-world”. In: *Network Neuroscience* 3.1 (2018), pp. 1–26.
- [96] S. R. Heilbronner and S. N. Haber. “Frontal Cortical and Subcortical Projections Provide a Basis for Segmenting the Cingulum Bundle: Implications for Neuroimaging and Psychiatric Disorders”. In: *Journal of Neuroscience* 34.30 (2014), pp. 10041–10054.
- [97] S. Herculano-Houzel. “The human brain in numbers: a linearly scaled-up primate brain”. In: *Frontiers in Human Neuroscience* 3 (2009), p. 31.
- [98] M. Hernandez-Fernandez, G. D. Guerrero, J. M. Cecilia, J. M. García, A. Inuggi, S. Jbabdi, T. E. J. Behrens, and S. N. Sotiropoulos. “Accelerating Fibre Orientation Estimation from Diffusion Weighted Magnetic Resonance Imaging Using GPUs”. In: *PLoS ONE* 8.4 (2013), e61892.
- [99] M. Hernandez-Fernandez, I. Reguly, S. Jbabdi, M. Giles, S. Smith, and S. N. Sotiropoulos. “Using GPUs to accelerate computational diffusion MRI: From microstructure estimation to tractography and connectomes”. In: *NeuroImage* 188 (2019), pp. 598–615.
- [100] C. P. Hess, P. Mukherjee, E. T. Han, D. Xu, and D. B. Vigneron. “Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic basis”. In: *Magnetic Resonance in Medicine* 56.1 (2006), pp. 104–117.
- [101] M. P. van den Heuvel and H. E. Hulshoff Pol. “Exploring the brain network: A review on resting-state fMRI functional connectivity”. In: *European Neuropsychopharmacology* 20.8 (2010), pp. 519–534.
- [102] M. P. van den Heuvel, K. J. Kersbergen, M. A. de Reus, K. Keunen, R. S. Kahn, F. Groenendaal, L. S. de Vries, and M. J. N. L. Benders. “The Neonatal Connectome During Preterm Brain Development”. In: *Cerebral Cortex* 25.9 (2015), pp. 3000–3013.
- [103] D. Holland, L. Chang, T. M. Ernst, M. Curran, S. D. Buchthal, D. Alicata, J. Skranes, H. Johansen, A. Hernandez, R. Yamakawa, J. M. Kuperman, and A. M. Dale. “Structural growth trajectories and rates of change in the first 3 months of infant brain development”. In: *JAMA Neurology* 71.10 (2014), pp. 1266–1274.
- [104] T. P. Hosey, S. G. Harding, T. A. Carpenter, R. E. Ansorge, and G. B. Williams. “Application of a probabilistic double-fibre structure model to diffusion-weighted MR images of the human brain”. In: *Magnetic Resonance Imaging* 26.2 (2008), pp. 236–245.

- [105] B. R. Howell, M. A. Styner, W. Gao, P. T. Yap, L. Wang, K. Baluyot, E. Yacoub, G. Chen, T. Potts, A. Salzwedel, G. Li, J. H. Gilmore, J. Piven, J. K. Smith, D. Shen, K. Ugurbil, H. Zhu, W. Lin, and J. T. Ellison. “The UNC/UMN Baby Connectome Project (BCP): An overview of the study design and protocol development”. In: *NeuroImage* 185 (2019), pp. 891–905.
- [106] P. O. Hoyer. “Non-negative matrix factorization with sparseness constraints”. In: *Journal of Machine Learning Research* 5 (2004), pp. 1457–1469.
- [107] H. Huang, J. Zhang, S. Wakana, W. Zhang, T. Ren, L. J. Richards, P. Yarowsky, P. Donohue, E. Graham, P. C. van Zijl, and S. Mori. “White and gray matter development in human fetal, newborn and pediatric brains”. In: *NeuroImage* 33.1 (2006), pp. 27–38.
- [108] E. J. Hughes, T. Winchman, F. Padormo, R. Teixeira, J. Wurie, M. Sharma, M. Fox, J. Hutter, L. Cordero-Grande, A. N. Price, J. Allsop, J. Bueno-Conde, N. Tusor, T. Arichi, A. D. Edwards, M. A. Rutherford, S. J. Counsell, and J. V. Hajnal. “A dedicated neonatal brain imaging system”. In: *Magnetic Resonance in Medicine* 78.2 (2017), pp. 794–804.
- [109] P. S. Hüppi and J. Dubois. “Diffusion tensor imaging of brain development”. In: *Seminars in Fetal and Neonatal Medicine* 11.6 (2006), pp. 489–497.
- [110] P. R. Huttenlocher and A. S. Dabholkar. “Regional differences in synaptogenesis in human cerebral cortex”. In: *Journal of Comparative Neurology* 387.2 (1997), pp. 167–178.
- [111] J. Hutter, A. N. Price, L. Cordero-Grande, S. Malik, G. Ferrazzi, A. Gaspar, E. J. Hughes, D. Christiaens, L. McCabe, T. Schneider, M. A. Rutherford, and J. V. Hajnal. “Quiet echo planar imaging for functional and diffusion MRI”. In: *Magnetic Resonance in Medicine* 79.3 (2018), pp. 1447–1459.
- [112] J. Hutter, J. D. Tournier, A. N. Price, L. Cordero-Grande, E. J. Hughes, S. Malik, J. Steinweg, M. Bastiani, S. N. Sotiropoulos, S. Jbabdi, J. Andersson, A. D. Edwards, and J. V. Hajnal. “Time-efficient and flexible design of optimized multishell HARDI diffusion”. In: *Magnetic Resonance in Medicine* 79.3 (2018), pp. 1276–1292.
- [113] A. Hyvärinen and E. Oja. “Independent component analysis: Algorithms and applications”. In: *Neural Networks* 13.4-5 (2000), pp. 411–430.
- [114] K. M. Jansons and D. C. Alexander. “Persistent angular structure: new insights from diffusion magnetic resonance imaging data”. In: *Inverse Problems* 19.5 (2003), pp. 1031–1046.
- [115] S. Jbabdi, S. N. Sotiropoulos, S. N. Haber, D. C. Van Essen, and T. E. Behrens. “Measuring macroscopic brain connections in vivo”. In: *Nature Neuroscience* 18.11 (2015), pp. 1546–1555.
- [116] S. Jbabdi, S. N. Sotiropoulos, A. M. Savio, M. Graña, and T. E. Behrens. “Model-based analysis of multishell diffusion MR data for tractography: How to get over fitting problems”. In: *Magnetic Resonance in Medicine* 68.6 (2012), pp. 1846–1855.

- [117] M. Jenkinson, C. F. Beckmann, T. E. J. Behrens, M. W. Woolrich, and S. M. Smith. “FSL”. In: *NeuroImage* 62.2 (2012), pp. 782–790.
- [118] B. Jeurissen, M. Descoteaux, S. Mori, and A. Leemans. “Diffusion MRI fiber tractography of the brain”. In: *NMR in Biomedicine* 32.4 (2019), e3785.
- [119] B. Jeurissen, A. Leemans, D. K. Jones, J.-D. Tournier, and J. Sijbers. “Probabilistic fiber tracking using the residual bootstrap with constrained spherical deconvolution”. In: *Human Brain Mapping* 32.3 (2011), pp. 461–479.
- [120] B. Jeurissen, A. Leemans, J. D. Tournier, D. K. Jones, and J. Sijbers. “Investigating the prevalence of complex fiber configurations in white matter tissue with diffusion magnetic resonance imaging”. In: *Human Brain Mapping* 34.11 (2013), pp. 2747–2766.
- [121] B. Jeurissen, J.-D. Tournier, T. Dhollander, A. Connelly, and J. Sijbers. “Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data”. In: *NeuroImage* 103 (2014), pp. 411–426.
- [122] B. Jian and B. Vemuri. “A Unified Computational Framework for Deconvolution to Reconstruct Multiple Fibers From Diffusion Weighted MRI”. In: *IEEE Transactions on Medical Imaging* 26.11 (2007), pp. 1464–1471.
- [123] B. Jian, B. C. Vemuri, E. Özarslan, P. R. Carney, and T. H. Mareci. “A novel tensor distribution model for the diffusion-weighted MR signal”. In: *NeuroImage* 37.1 (2007), pp. 164–176.
- [124] S. Johnson and N. Marlow. “Growing up after extremely preterm birth: Lifespan mental health outcomes”. In: *Seminars in Fetal and Neonatal Medicine* 19.2 (2014), pp. 97–104.
- [125] D. K. Jones, T. R. Knösche, and R. Turner. “White matter integrity, fiber count, and other fallacies: The do’s and don’ts of diffusion MRI”. In: *NeuroImage* 73.Supplement C (2013), pp. 239–254.
- [126] E. Kaden, T. R. Knösche, and A. Anwander. “Parametric spherical deconvolution: Inferring anatomical connectivity using diffusion MR imaging”. In: *NeuroImage* 37.2 (2007), pp. 474–488.
- [127] A. P. Kansagra, M. C. Mabray, D. M. Ferriero, A. J. Barkovich, D. Xu, and C. P. Hess. “Microstructural maturation of white matter tracts in encephalopathic neonates”. In: *Clinical Imaging* 40.5 (2016), pp. 1009–1013.
- [128] S. Kaur, S. Powell, L. He, C. R. Pierson, and N. A. Parikh. “Reliability and Repeatability of Quantitative Tractography Methods for Mapping Structural White Matter Connectivity in Preterm and Term Infants at Term-Equivalent Age”. In: *PLoS ONE* 9.1 (2014), e85807.
- [129] H. C. Kinney, B. A. Brody, A. S. Kloman, and F. H. Gilles. “Sequence of Central Nervous System Myelination in Human Infancy II. Patterns of Myelination in Autopsied Infants”. In: *Journal of Neuropathology and Experimental Neurology* 47.3 (1988), pp. 217–234.

- [130] C. G. Koay, J. D. Carew, A. L. Alexander, P. J. Basser, and M. E. Meyerand. “Investigation of anomalous estimates of tensor-derived quantities in diffusion tensor imaging”. In: *Magnetic Resonance in Medicine* 55.4 (2006), pp. 930–936.
- [131] I. Kostović, G. Sedmak, and M. Judaš. “Neural histology and neurogenesis of the human fetal and infant brain”. In: *NeuroImage* 188 (2019), pp. 743–773.
- [132] I. Kostović and N. Jovanov-Milošević. “The development of cerebral connections during the first 20–45 weeks’ gestation”. In: *Seminars in Fetal and Neonatal Medicine* 11.6 (2006), pp. 415–422.
- [133] M. Kuklisova-Murgasova, G. Quaghebeur, M. A. Rutherford, J. V. Hajnal, and J. A. Schnabel. “Reconstruction of fetal brain MRI with intensity matching and complete outlier removal”. In: *Medical Image Analysis* 16.8 (2012), pp. 1550–1564.
- [134] S. Kullback and R. A. Leibler. “On Information and Sufficiency”. In: *The Annals of Mathematical Statistics* 22.1 (1951), pp. 79–86.
- [135] A. Lautarescu, D. Pecheva, C. Nosarti, J. Nihouarn, H. Zhang, S. Victor, M. Craig, A. D. Edwards, and S. J. Counsell. “Maternal Prenatal Stress Is Associated With Altered Uncinate Fasciculus Microstructure in Premature Neonates”. In: *Biological Psychiatry* 87.6 (2020), pp. 559–569.
- [136] P. C. Lauterbur. “Magnetic resonance zeugmatography”. In: *Pure and Applied Chemistry* 40.1-2 (1974), pp. 149–157.
- [137] D. Le Bihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis, and M. Laval-Jeantet. “MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders.” In: *Radiology* 161.2 (1986), pp. 401–407.
- [138] D. D. Lee and H. S. Seung. “Learning the parts of objects by non-negative matrix factorization”. In: *Nature* 401.6755 (1999), pp. 788–791.
- [139] S. J. Lee, R. J. Steiner, Y. Yu, S. J. Short, M. C. Neale, M. A. Styner, H. Zhu, and J. H. Gilmore. “Common and heritable components of white matter microstructure predict cognitive function at 1 and 2 y”. In: *Proceedings of the National Academy of Sciences* 114.1 (2017), pp. 148–153.
- [140] J. D. Lewis, A. C. Evans, J. R. Pruett, K. Botteron, L. Zwaigenbaum, A. Estes, G. Gerig, L. Collins, P. Kostopoulos, R. McKinstry, S. Dager, S. Paterson, R. T. Schultz, M. Styner, H. Hazlett, and J. Piven. “Network inefficiencies in autism spectrum disorder at 24 months”. In: *Translational Psychiatry* 4 (2014), e388.
- [141] W. Lin, Q. Zhu, W. Gao, Y. Chen, C.-H. Toh, M. Styner, G. Gerig, J. Smith, B. Biswal, and J. Gilmore. “Functional Connectivity MR Imaging Reveals Cortical Functional Connectivity in the Developing Brain”. In: *American Journal of Neuroradiology* 29.10 (2008), pp. 1883–1889.

- [142] R. F. Ling, C. L. Lawson, and R. J. Hanson. “Solving Least Squares Problems.” In: *Journal of the American Statistical Association* 72.360 (1977), p. 930.
- [143] Y. Liu, D. Balériaux, M. Kavec, T. Metens, J. Absil, V. Denolin, A. Pardou, F. Avni, P. Van Bogaert, and A. Aeby. “Structural asymmetries in motor and language networks in a population of healthy preterm neonates at term equivalent age: A diffusion tensor imaging and probabilistic tractography study”. In: *NeuroImage* 51.2 (2010), pp. 783–788.
- [144] N. Lopez Rios, A. Foias, G. Lodygensky, M. Dehaes, and J. Cohen-Adad. “Size-adaptable 13-channel receive array for brain MRI in human neonates at 3 T”. In: *NMR in Biomedicine* 31.8 (2018), e3944.
- [145] K. H. Maier-Hein et al. “The challenge of mapping the human connectome based on diffusion tractography”. In: *Nature Communications* 8.1 (2017), p. 1349.
- [146] N. Makris, M. G. Preti, D. Wassermann, Y. Rathi, G. M. Papadimitriou, C. Yergatian, B. C. Dickerson, M. E. Shenton, and M. Kubicki. “Human middle longitudinal fascicle: segregation and behavioral-clinical implications of two distinct fiber connections linking temporal pole and superior temporal gyrus with the angular gyrus or superior parietal lobule using multi-tensor tractography”. In: *Brain Imaging and Behavior* 7.3 (2013), pp. 335–352.
- [147] N. Makris, G. M. Papadimitriou, J. R. Kaiser, S. Sorg, D. N. Kennedy, and D. N. Pandya. “Delineation of the Middle Longitudinal Fascicle in Humans: A Quantitative, In Vivo, DT-MRI Study”. In: *Cerebral Cortex* 19.4 (2008), pp. 777–785.
- [148] A. Makropoulos et al. “The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction”. In: *NeuroImage* 173 (2018), pp. 88–112.
- [149] P. Mansfield and A. A. Maudsley. “Medical imaging by NMR”. In: *The British Journal of Radiology* 50.591 (1977), pp. 188–194.
- [150] D. S. Marcus, J. Harwell, T. Olsen, M. Hodge, M. F. Glasser, F. Prior, M. Jenkinson, T. Laumann, S. W. Curtiss, and D. C. Van Essen. “Informatics and data mining tools and strategies for the human connectome project”. In: *Frontiers in Neuroinformatics* 5 (2011), p. 4.
- [151] R. B. Mars, S. Foxley, L. Verhagen, S. Jbabdi, J. Sallet, M. P. Noonan, F.-X. Neubert, J. L. Andersson, P. L. Croxson, R. I. M. Dunbar, A. A. Khrapitchev, N. R. Sibson, K. L. Miller, and M. F. S. Rushworth. “The extreme capsule fiber complex in humans and macaque monkeys: a comparative diffusion MRI tractography study”. In: *Brain Structure and Function* 221.8 (2016), pp. 4059–4071.
- [152] R. B. Mars, J. O’Muircheartaigh, D. Folloni, L. Li, M. F. Glasser, S. Jbabdi, and K. L. Bryant. “Concurrent analysis of white matter bundles and grey matter networks in the chimpanzee”. In: *Brain Structure and Function* 224.3 (2019), pp. 1021–1033.

- [153] R. B. Mars, S. N. Sotiropoulos, R. E. Passingham, J. Sallet, L. Verhagen, A. A. Khrapitchev, N. Sibson, and S. Jbabdi. “Whole brain comparative anatomy using connectivity blueprints”. In: *eLife* 7 (2018), p. 245209.
- [154] M. J. Mckeown, S. Makeig, G. G. Brown, T.-P. Jung, S. S. Kindermann, A. J. Bell, and T. J. Sejnowski. “Analysis of fMRI data by blind separation into independent spatial components”. In: *Human Brain Mapping* 6.3 (1998), pp. 160–188.
- [155] K. L. Miller, F. Alfaro-Almagro, N. K. Bangerter, D. L. Thomas, E. Yacoub, J. Xu, A. J. Bartsch, S. Jbabdi, S. N. Sotiropoulos, J. L. Andersson, L. Griffanti, G. Douaud, T. W. Okell, P. Weale, I. Dragonu, S. Garratt, S. Hudson, R. Collins, M. Jenkinson, P. M. Matthews, and S. M. Smith. “Multimodal population brain imaging in the UK Biobank prospective epidemiological study”. In: *Nature Neuroscience* 19.11 (2016), pp. 1523–1536.
- [156] R. J. Morecraft, G. Ugolini, J. L. Lanciego, F. G. G. Wouterlood, and D. N. Pandya. “Classic and Contemporary Neural Tract-Tracing Techniques”. In: *Diffusion MRI: From Quantitative Measurement to In vivo Neuroanatomy: Second Edition*. Academic Press, 2013, pp. 359–399.
- [157] S. Mori, B. J. Crain, V. P. Chacko, and P. C. M. Van Zijl. “Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging”. In: *Annals of Neurology* 45.2 (1999), pp. 265–269.
- [158] L. D. Nickerson, S. M. Smith, D. Öngür, and C. F. Beckmann. “Using Dual Regression to Investigate Network Shape and Amplitude in Functional Connectivity Analyses”. In: *Frontiers in Neuroscience* 11 (2017), p. 115.
- [159] J. Nie et al. “Axonal fiber terminations concentrate on gyri”. In: *Cerebral Cortex* 22.12 (2012), pp. 2831–2839.
- [160] L. O’Donnell and C.-F. Westin. “Automatic Tractography Segmentation Using a High-Dimensional White Matter Atlas”. In: *IEEE Transactions on Medical Imaging* 26.11 (2007), pp. 1562–1575.
- [161] J. O’Muircheartaigh and S. Jbabdi. “Concurrent white matter bundles and grey matter networks using independent component analysis”. In: *NeuroImage* 170 (2017), pp. 296–306.
- [162] J. O’Muircheartaigh, C. Vollmar, C. Traynor, G. J. Barker, V. Kumari, M. R. Symms, P. Thompson, J. S. Duncan, M. J. Koepp, and M. P. Richardson. “Clustering probabilistic tractograms using independent component analysis applied to the thalamus.” In: *NeuroImage* 54.3 (2011), pp. 2020–2032.
- [163] S. Ogawa, T. M. Lee, A. R. Kay, and D. W. Tank. “Brain magnetic resonance imaging with contrast dependent on blood oxygenation.” In: *Proceedings of the National Academy of Sciences of the United States of America* 87.24 (1990), pp. 9868–9872.
- [164] K. Oishi, S. Mori, P. K. Donohue, T. Ernst, L. Anderson, S. Buchthal, A. Faria, H. Jiang, X. Li, M. I. Miller, P. C. van Zijl, and L. Chang. “Multi-contrast human neonatal brain atlas: Application to normal neonate development analysis”. In: *NeuroImage* 56.1 (2011), pp. 8–20.

- [165] K. Oishi, K. Zilles, K. Amunts, A. Faria, H. Jiang, X. Li, K. Akhter, K. Hua, R. Woods, A. W. Toga, G. B. Pike, P. Rosa-Neto, A. Evans, J. Zhang, H. Huang, M. I. Miller, P. C. van Zijl, J. Mazziotta, and S. Mori. “Human brain white matter atlas: Identification and assignment of common anatomical structures in superficial white matter”. In: *NeuroImage* 43.3 (2008), pp. 447–457.
- [166] E. St-Onge, A. Daducci, G. Girard, and M. Descoteaux. “Surface-enhanced tractography (SET)”. In: *NeuroImage* 169 (2018), pp. 524–539.
- [167] M. Ouyang, J. Dubois, Q. Yu, P. Mukherjee, and H. Huang. “Delineation of early brain development from fetuses to infants with diffusion MRI and beyond”. In: *NeuroImage* 185 (2019), pp. 836–850.
- [168] E. Özarslan, T. M. Shepherd, B. C. Vemuri, S. J. Blackband, and T. H. Mareci. “Resolution of complex tissue microarchitecture using the diffusion orientation transform (DOT)”. In: *NeuroImage* 31.3 (2006), pp. 1086–1103.
- [169] E. Panagiotaki, T. Schneider, B. Siow, M. G. Hall, M. F. Lythgoe, and D. C. Alexander. “Compartment models of the diffusion MR signal in brain white matter: A taxonomy and comparison”. In: *NeuroImage* 59.3 (2012), pp. 2241–2254.
- [170] K. Pannek, S. M. Scheck, P. B. Colditz, R. N. Boyd, and S. E. Rose. “Magnetic resonance diffusion tractography of the preterm infant brain: a systematic review”. In: *Developmental Medicine & Child Neurology* 56.2 (2014), pp. 113–124.
- [171] S. C. Partridge, P. Mukherjee, R. G. Henry, S. P. Miller, J. I. Berman, H. Jin, Y. Lu, O. A. Glenn, D. M. Ferriero, A. J. Barkovich, and D. B. Vigneron. “Diffusion tensor imaging: Serial quantitation of white matter tract maturity in premature newborns”. In: *NeuroImage* 22.3 (2004), pp. 1302–1314.
- [172] R. E. Passingham, K. E. Stephan, and R. Kötter. “The anatomical basis of functional localization in the cortex”. In: *Nature Reviews Neuroscience* 3.8 (2002), pp. 606–616.
- [173] D. Pecheva, P. Yushkevich, D. Batalle, E. Hughes, P. Aljabar, J. Wurie, J. V. Hajnal, A. D. Edwards, D. C. Alexander, S. J. Counsell, and H. Zhang. “A tract-specific approach to assessing white matter in preterm infants”. In: *NeuroImage* 157.Supplement C (2017), pp. 675–694.
- [174] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. “Scikit-learn: Machine learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.
- [175] M. Petrides, F. Tomaiuolo, E. H. Yeterian, and D. N. Pandya. “The prefrontal cortex: Comparative architectonic organization in the human and the macaque monkey brains”. In: *Cortex* 48.1 (2012), pp. 46–57.
- [176] C. Pierpaoli and P. J. Basser. “Toward a quantitative assessment of diffusion anisotropy”. In: *Magnetic Resonance in Medicine* 36.6 (1996), pp. 893–906.

- [177] E. M. Purcell, H. C. Torrey, and R. V. Pound. “Resonance Absorption by Nuclear Magnetic Moments in a Solid”. In: *Physical Review* 69.1-2 (1946), pp. 37–38.
- [178] M. Radoš, M. Judaš, and I. Kostović. “In vitro MRI of brain development”. In: *European Journal of Radiology* 57.2 (2006), pp. 187–198.
- [179] J. M. Rasmussen, F. Kruggel, J. H. Gilmore, M. Styner, S. Entringer, K. N. Conring, S. G. Potkin, P. D. Wadhwa, and C. Buss. “A novel maturation index based on neonatal diffusion tensor imaging reflects typical perinatal white matter development in humans”. In: *International Journal of Developmental Neuroscience* 56.1 (2017), pp. 42–51.
- [180] N. Ratnarajah and A. Qiu. “Multi-label segmentation of white matter structures: Application to neonatal brains”. In: *NeuroImage* 102 (2014), pp. 913–922.
- [181] N. Ratnarajah, A. Rifkin-Graboi, M. V. Fortier, Y. S. Chong, K. Kwek, S.-M. Saw, K. M. Godfrey, P. D. Gluckman, M. J. Meaney, and A. Qiu. “Structural connectivity asymmetry in the neonatal brain”. In: *NeuroImage* 75 (2013), pp. 187–194.
- [182] M. Reisert, E. Kellner, and V. G. Kiselev. “About the Geometry of Asymmetric Fiber Orientation Distributions”. In: *IEEE Transactions on Medical Imaging* 31.6 (2012), pp. 1240–1249.
- [183] M. Reisert and V. G. Kiselev. “Fiber Continuity: An Anisotropic Prior for ODF Estimation”. In: *IEEE Transactions on Medical Imaging* 30.6 (2011), pp. 1274–1283.
- [184] C. Reveley, A. K. Seth, C. Pierpaoli, A. C. Silva, D. Yu, R. C. Saunders, D. A. Leopold, and F. Q. Ye. “Superficial white matter fiber systems impede detection of long-range cortical connections in diffusion MR tractography”. In: *Proceedings of the National Academy of Sciences of the United States of America* 112.21 (2015), E2820–E2828.
- [185] A. Rifkin-Graboi, M. J. Meaney, H. Chen, J. Bai, W. B. Hameed, M. T. Tint, B. F. P. Broekman, Y.-S. Chong, P. D. Gluckman, M. V. Fortier, and A. Qiu. “Antenatal Maternal Anxiety Predicts Variations in Neural Structures Implicated in Anxiety Disorders in Newborns”. In: *Journal of the American Academy of Child & Adolescent Psychiatry* 54.4 (2015), 313–321.e2.
- [186] E. C. Robinson, K. Garcia, M. F. Glasser, Z. Chen, T. S. Coalson, A. Makropoulos, J. Bozek, R. Wright, A. Schuh, M. Webster, J. Hutter, A. Price, L. Cordero Grande, E. Hughes, N. Tusor, P. V. Bayly, D. C. Van Essen, S. M. Smith, A. D. Edwards, J. Hajnal, M. Jenkinson, B. Glocker, and D. Rueckert. “Multimodal surface matching with higher-order smoothness constraints”. In: *NeuroImage* 167 (2018), pp. 453–465.
- [187] E. C. Robinson, S. Jbabdi, M. F. Glasser, J. Andersson, G. C. Burgess, M. P. Harms, S. M. Smith, D. C. Van Essen, and M. Jenkinson. “MSM: A new flexible framework for multimodal surface matching”. In: *NeuroImage* 100 (2014), pp. 414–426.

- [188] A. Romano, G. D’Andrea, G. Minniti, L. Mastronardi, L. Ferrante, L. M. Fantozzi, and A. Bozzao. “Pre-surgical planning and MR-tractography utility in brain tumour resection”. In: *European Radiology* 19.12 (2009), pp. 2798–2808.
- [189] P. J. Rousseeuw. “Silhouettes: A graphical aid to the interpretation and validation of cluster analysis”. In: *Journal of Computational and Applied Mathematics* 20 (1987), pp. 53–65.
- [190] J. Sa de Almeida, L. Lordier, B. Zollinger, N. Kunz, M. Bastiani, L. Gui, A. Adam-Darque, C. Borradori-Tolsa, F. Lazeyras, and P. S. Hüppi. “Music enhances structural maturation of emotional processing neural pathways in very preterm infants”. In: *NeuroImage* 207 (2020), p. 116391.
- [191] N. Sadeghi, J. H. Gilmore, and G. Gerig. “Twin-singleton developmental study of brain white matter anatomy”. In: *Human Brain Mapping* 38.2 (2017), pp. 1009–1024.
- [192] N. Saito, B. M. Larson, and B. Benichou. “Sparsity vs. statistical independence from a best-basis viewpoint”. In: *Wavelet Applications in Signal and Image Processing VIII*. Vol. 4119. 2000, p. 474.
- [193] G. Salimi-Khorshidi, G. Douaud, C. F. Beckmann, M. F. Glasser, L. Griffanti, and S. M. Smith. “Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical fusion of classifiers”. In: *NeuroImage* 90 (2014), pp. 449–468.
- [194] Z. M. Saygin, D. E. Osher, K. Koldewyn, G. Reynolds, J. D. Gabrieli, and R. R. Saxe. “Anatomical connectivity patterns predict face selectivity in the fusiform gyrus”. In: *Nature Neuroscience* 15.2 (2012), pp. 321–327.
- [195] B. Scherrer, A. Schwartzman, M. Taquet, M. Sahin, S. P. Prabhu, and S. K. Warfield. “Characterizing brain tissue by assessment of the distribution of anisotropic microstructural environments in diffusion-compartment imaging (DIAMOND)”. In: *Magnetic Resonance in Medicine* 76.3 (2016), pp. 963–977.
- [196] K. Schilling, Y. Gao, V. Janve, I. Stepniewska, B. A. Landman, and A. W. Anderson. “Confirmation of a gyral bias in diffusion MRI fiber tractography”. In: *Human Brain Mapping* 39.3 (2018), pp. 1449–1466.
- [197] A. Schuh, A. Makropoulos, E. Robinson, L. Cordero-Grande, E. Hughes, J. Hutter, A. Price, M. Murgasova, R. P. Teixeira, N. Tusor, J. Steinweg, S. Victor, M. Rutherford, J. Hajnal, A. D. Edwards, and D. Rueckert. “Unbiased construction of a temporally consistent morphological atlas of neonatal brain development”. In: *bioRxiv* (2018), p. 251512.
- [198] A. Serag, P. Aljabar, G. Ball, S. J. Counsell, J. P. Boardman, M. A. Rutherford, A. D. Edwards, J. V. Hajnal, and D. Rueckert. “Construction of a consistent high-definition spatio-temporal atlas of the developing brain using adaptive kernel regression”. In: *NeuroImage* 59.3 (2012), pp. 2255–2265.

- [199] V. Siless, K. Chang, B. Fischl, and A. Yendiki. “AnatomyCuts: Hierarchical clustering of tractography streamlines based on anatomical similarity”. In: *NeuroImage* 166 (2018), pp. 32–45.
- [200] R. E. Smith, J. D. Tournier, F. Calamante, and A. Connelly. “Anatomically-constrained tractography: Improved diffusion MRI streamlines tractography through effective use of anatomical information”. In: *NeuroImage* 62.3 (2012), pp. 1924–1938.
- [201] S. M. Smith, M. Jenkinson, H. Johansen-Berg, D. Rueckert, T. E. Nichols, C. E. Mackay, K. E. Watkins, O. Ciccarelli, M. Z. Cader, P. M. Matthews, and T. E. Behrens. “Tract-based spatial statistics: Voxelwise analysis of multi-subject diffusion data”. In: *NeuroImage* 31.4 (2006), pp. 1487–1505.
- [202] S. M. Smith, P. M. T. Fox, K. L. Miller, D. C. Glahn, P. M. T. Fox, C. E. Mackay, N. Filippini, K. E. Watkins, R. Toro, A. R. Laird, and C. F. Beckmann. “Correspondence of the brain’s functional architecture during activation and rest.” In: *Proceedings of the National Academy of Sciences of the United States of America* 106.31 (2009), pp. 13040–13045.
- [203] S. M. Smith, A. Hyvärinen, G. Varoquaux, K. L. Miller, and C. F. Beckmann. “Group-PCA for very large fMRI datasets”. In: *NeuroImage* 101 (2014), pp. 738–749.
- [204] J. W. Song, P. D. Mitchell, J. Kolasinski, P. Ellen Grant, A. M. Galaburda, and E. Takahashi. “Asymmetry of White Matter Pathways in Developing Human Brains”. In: *Cerebral Cortex* 25.9 (2015), pp. 2883–2893.
- [205] S.-K. Song, S.-W. Sun, M. J. Ramsbottom, C. Chang, J. Russell, and A. H. Cross. “Dysmyelination Revealed through MRI as Increased Radial (but Unchanged Axial) Diffusion of Water”. In: *NeuroImage* 17.3 (2002), pp. 1429–1436.
- [206] A. Sotiras, S. M. Resnick, and C. Davatzikos. “Finding imaging patterns of structural covariance via Non-Negative Matrix Factorization”. In: *NeuroImage* 108 (2015), pp. 1–16.
- [207] A. Sotiras, J. B. Toledo, R. E. Gur, R. C. Gur, T. D. Satterthwaite, and C. Davatzikos. “Patterns of coordinated cortical remodeling during adolescence and their associations with functional specialization and evolutionary expansion”. In: *Proceedings of the National Academy of Sciences of the United States of America* 114.13 (2017), pp. 3527–3532.
- [208] S. N. Sotiropoulos, T. E. Behrens, and S. Jbabdi. “Ball and rackets: Inferring fiber fanning from diffusion-weighted MRI”. In: *NeuroImage* 60.2 (2012), pp. 1412–1425.
- [209] S. N. Sotiropoulos, S. Jbabdi, J. Xu, J. L. Andersson, S. Moeller, E. J. Auerbach, M. F. Glasser, M. Hernandez, G. Sapiro, M. Jenkinson, D. A. Feinberg, E. Yacoub, C. Lenglet, D. C. Van Essen, K. Ugurbil, and T. E. Behrens. “Advances in diffusion MRI acquisition and processing in the Human Connectome Project”. In: *NeuroImage* 80 (2013), pp. 125–143.

- [210] S. N. Sotiropoulos, M. Hernández-Fernández, A. T. Vu, J. L. Andersson, S. Moeller, E. Yacoub, C. Lenglet, K. Ugurbil, T. E. J. Behrens, and S. Jbabdi. “Fusion in diffusion MRI for improved fibre orientation estimation: An application to the 3T and 7T data of the Human Connectome Project”. In: *NeuroImage* 134 (2016), pp. 396–409.
- [211] S. N. Sotiropoulos and A. Zalesky. “Building connectomes using diffusion MRI: why, how and but”. In: *NMR in Biomedicine* 32.4 (2019), e3752.
- [212] O. Sporns, G. Tononi, and R. Kötter. “The Human Connectome: A Structural Description of the Human Brain”. In: *PLoS Computational Biology* 1.4 (2005), e42.
- [213] E. O. Stejskal and J. E. Tanner. “Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient”. In: *The Journal of Chemical Physics* 42.1 (1965), pp. 288–292.
- [214] M. R. Swanson, J. J. Wolff, M. D. Shen, M. Styner, A. Estes, G. Gerig, R. C. McKinstry, K. N. Botteron, J. Piven, and H. C. Hazlett. “Development of White Matter Circuitry in Infants With Fragile X Syndrome”. In: *JAMA Psychiatry* 75.5 (2018), p. 505.
- [215] M. Takahashi, D. B. Hackney, G. Zhang, S. L. Wehrli, A. C. Wright, W. T. O’Brien, H. Uematsu, F. W. Wehrli, and M. E. Selzer. “Magnetic resonance microimaging of intraaxonal water diffusion in live excised lamprey spinal cord”. In: *Proceedings of the National Academy of Sciences* 99.25 (2002), pp. 16192–16196.
- [216] H. Takemura, F. Pestilli, K. S. Weiner, G. A. Keliris, S. M. Landi, J. Sliwa, F. Q. Ye, M. A. Barnett, D. A. Leopold, W. A. Freiwald, N. K. Logothetis, and B. A. Wandell. “Occipital White Matter Tracts in Human and Macaque”. In: *Cerebral Cortex* 27.6 (2017), pp. 3346–3359.
- [217] E. W. Y. Tam, V. Chau, A. J. Barkovich, D. M. Ferriero, S. P. Miller, E. E. Rogers, R. E. Grunau, A. R. Synnes, D. Xu, J. Foong, R. Brant, and S. M. Innis. “Early postnatal docosahexaenoic acid levels and improved preterm brain development”. In: *Pediatric Research* 79.5 (2016), pp. 723–730.
- [218] C. M. Tax, B. Jeurissen, S. B. Vos, M. A. Viergever, and A. Leemans. “Recursive calibration of the fiber response function for spherical deconvolution of diffusion MRI data”. In: *NeuroImage* 86 (2014), pp. 67–80.
- [219] M. Thiebaut de Schotten, F. Dell’Acqua, S. J. Forkel, A. Simmons, F. Vergani, D. G. M. Murphy, and M. Catani. “A lateralized brain network for visuospatial attention.” In: *Nature neuroscience* 14.10 (2011), pp. 1245–6.
- [220] M. Thiebaut de Schotten, D. H. Ffytche, A. Bizzi, F. Dell’Acqua, M. Allin, M. Walshe, R. Murray, S. C. Williams, D. G. M. Murphy, and M. Catani. “Atlasing location, asymmetry and inter-subject variability of white matter tracts in the human brain with MR diffusion tractography”. In: *NeuroImage* 54.1 (2011), pp. 49–59.

- [221] J. D. Tournier, F. Calamante, and A. Connelly. “MRtrix: Diffusion tractography in crossing fiber regions”. In: *International Journal of Imaging Systems and Technology* 22.1 (2012), pp. 53–66.
- [222] J. D. Tournier, F. Calamante, and A. Connelly. “Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution”. In: *NeuroImage* 35.4 (2007), pp. 1459–1472.
- [223] J. D. Tournier, F. Calamante, D. G. Gadian, and A. Connelly. “Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution”. In: *NeuroImage* 23.3 (2004), pp. 1176–1185.
- [224] D. S. Tuch. “Q-ball imaging”. In: *Magnetic Resonance in Medicine* 52.6 (2004), pp. 1358–1372.
- [225] D. S. Tuch, T. G. Reese, M. R. Wiegell, N. Makris, J. W. Belliveau, and V. J. Wedeen. “High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity”. In: *Magnetic Resonance in Medicine* 48.4 (2002), pp. 577–582.
- [226] O. Tymofiyeva, C. P. Hess, E. Ziv, P. N. Lee, H. C. Glass, D. M. Ferriero, A. J. Barkovich, and D. Xu. “A DTI-Based Template-Free Cortical Connectome Study of Brain Maturation”. In: *PLoS ONE* 8.5 (2013), e63310.
- [227] D. C. Van Essen and M. F. Glasser. “Parcellating Cerebral Cortex: How Invasive Animal Studies Inform Noninvasive Mapmaking in Humans”. In: *Neuron* 99.4 (2018), pp. 640–663.
- [228] D. C. Van Essen, S. Jbabdi, S. N. Sotiropoulos, C. Chen, K. Dikranian, T. Coalson, J. Harwell, T. E. Behrens, and M. F. Glasser. “Mapping Connections in Humans and Non-Human Primates. Aspirations and Challenges for Diffusion Imaging.” In: *Diffusion MRI: From Quantitative Measurement to In vivo Neuroanatomy: Second Edition*. Academic Press, 2013, pp. 337–358.
- [229] D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. J. Behrens, E. Yacoub, K. Ugurbil, and WU-Minn HCP Consortium. “The WU-Minn Human Connectome Project: An overview”. In: *NeuroImage* 80 (2013), pp. 62–79.
- [230] A. R. Verde, F. Budin, J.-B. Berger, A. Gupta, M. Farzinfar, A. Kaiser, M. Ahn, H. Johnson, J. Matsui, H. C. Hazlett, A. Sharma, C. Goodlett, Y. Shi, S. Gouttard, C. Vachet, J. Piven, H. Zhu, G. Gerig, and M. Styner. “UNC-Utah NA-MIC framework for DTI fiber tract analysis”. In: *Frontiers in Neuroinformatics* 7 (2014).
- [231] S. Wakana, A. Caprihan, M. M. Panzenboeck, J. H. Fallon, M. Perry, R. L. Gollub, K. Hua, J. Zhang, H. Jiang, P. Dubey, A. Blitz, P. van Zijl, and S. Mori. “Reproducibility of quantitative tractography methods applied to cerebral white matter”. In: *NeuroImage* 36.3 (2007), pp. 630–644.
- [232] S. Wakana, H. Jiang, L. M. Nagae-Poetscher, P. C. M. van Zijl, and S. Mori. “Fiber Tract-based Atlas of Human White Matter Anatomy”. In: *Radiology* 230.1 (2004), pp. 77–87.

- [233] K. Walhovd, H. Johansen-Berg, and R. Káradóttir. “Unraveling the secrets of white matter – Bridging the gap between cellular, animal and human imaging studies”. In: *Neuroscience* 276 (2014), pp. 2–13.
- [234] S. Warrington, K. L. Bryant, A. A. Khrapitchev, J. Sallet, M. Charquero-Ballester, G. Douaud, S. Jbabdi, R. B. Mars, and S. N. Sotiropoulos. “XTRACT - Standardised protocols for automated tractography in the human and macaque brain”. In: *NeuroImage* 217 (2020), p. 116923.
- [235] V. J. Wedeen, P. Hagmann, W.-Y. I. Tseng, T. G. Reese, and R. M. Weisskoff. “Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging”. In: *Magnetic Resonance in Medicine* 54.6 (2005), pp. 1377–1386.
- [236] S. Wold, K. Esbensen, and P. Geladi. “Principal component analysis”. In: *Chemometrics and Intelligent Laboratory Systems* 2.1-3 (1987), pp. 37–52.
- [237] L. Wu, V. D. Calhoun, R. E. Jung, and A. Caprihan. “Connectivity-based whole brain dual parcellation by group ICA reveals tract structures and decreased connectivity in schizophrenia”. In: *Human Brain Mapping* 36.11 (2015), pp. 4681–4701.
- [238] Y. Wu, Y. Hong, Y. Feng, D. Shen, and P.-T. Yap. “Mitigating gyral bias in cortical tractography via asymmetric fiber orientation distributions”. In: *Medical Image Analysis* 59 (2020), p. 101543.
- [239] P. T. Yap, Y. Fan, Y. Chen, J. H. Gilmore, W. Lin, and D. Shen. “Development trends of white matter connectivity in the first years of life”. In: *PLoS ONE* 6.9 (2011), e24678.
- [240] F.-C. Yeh, T. D. Verstynen, Y. Wang, J. C. Fernández-Miranda, and W.-Y. I. Tseng. “Deterministic Diffusion Fiber Tracking Improved by Quantitative Anisotropy”. In: *PLoS ONE* 8.11 (2013), e80713.
- [241] A. Yendiki. “Automated probabilistic reconstruction of white-matter pathways in health and disease using an atlas of the underlying anatomy”. In: *Frontiers in Neuroinformatics* 5 (2011), p. 23.
- [242] H. Zhang, T. Schneider, C. A. Wheeler-Kingshott, and D. C. Alexander. “NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain”. In: *NeuroImage* 61.4 (2012), pp. 1000–1016.
- [243] L. Zöllei, C. Jaimés, E. Saliba, P. E. Grant, and A. Yendiki. “TRActs constrained by UnderLying INfant anatomy (TRACULInA): An automated probabilistic tractography tool with anatomical priors for use in the newborn brain”. In: *NeuroImage* 199 (2019), pp. 1–17.

Acronyms

BCP Baby Connectome Project. 50

BOLD blood oxygenation level-dependent. 11

CSD constrained spherical deconvolution. 25

dHCP developing Human Connectome Project. 49, 51, 57, 60, 68, 72, 73, 76, 86, 96, 100, 107

dMRI diffusion MRI. i, 3, 5, 10, 12, 24, 28, 30, 35, 38, 40, 44, 47, 50, 72, 73, 74, 86, 93, 98, 134, 152, 155

dODF diffusion orientation distribution function. 24

DSI Diffusion Spectrum Imaging. 24

DTI Diffusion Tensor Imaging. 23, 40, 42, 43, 86

FA fractional anisotropy. 21, 22, 42, 43, 72, 79

fMRI functional MRI. 11, 35, 93, 95, 96, 98, 107, 108, 124, 130, 135, 152, 156, 159, 165

fODF fibre orientation distribution function. 23, 25, 26, 27, 30, 31, 33, 88

HARDI High Angular Resolution Diffusion Imaging. 23, 24, 40, 49

HCP Human Connectome Project. 49, 50, 72, 164

ICA Independent Component Analysis. 6, 35, 50, 95, 96, 98, 99, 102, 104, 105, 106, 109, 111, 113, 120, 124, 126, 128, 129, 130, 131, 134, 135, 136, 138, 140, 141, 156, 157, 159, 164

MD mean diffusivity. 22, 42, 72, 79

MIGP Melodic's Incremental Group PCA. 108

MRI Magnetic Resonance Imaging. 3, 12, 14, 38, 49, 68, 72, 73

-
- NMF** Non-Negative Matrix Factorisation. xiii, xiv, 6, 95, 96, 98, 99, 102, 104, 105, 106, 107, 108, 109, 111, 113, 120, 121, 123, 124, 126, 127, 128, 129, 130, 131, 134, 135, 136, 137, 138, 140, 141, 142, 143, 144, 146, 148, 151, 152, 153, 156, 158, 164, 165
- NMR** nuclear magnetic resonance. 12
- NNLS** non-negative least squares. 136
- NODDI** Neurite Orientation Dispersion and Density Imaging. 43
- ODF** orientation distribution function. 24, 25, 26, 33
- PCA** Principal Component Analysis. 98, 102, 105, 106, 108
- PMA** post-menstrual age. 38, 40, 45, 49, 60, 69, 70, 72, 80, 101, 107
- ROI** region of interest. 40, 41, 42, 56, 79, 91
- SNR** signal-to-noise ratio. 47, 49, 87
- SVD** singular value decomposition. 98, 99
- TBSS** tract-based spatial statistics. 43
- uODF** uncertainty orientation distribution function. 33
- WGB** white matter/grey matter boundary. xii, 58, 72, 74, 75, 80, 81, 82, 83, 85, 87, 88, 89, 97, 100, 101, 108, 156