

# Brain connectivity mapping with diffusion MRI across individuals and species

Shaun Warrington

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

March 2021

### Abstract

The human brain is a highly complex organ that integrates functionally specialised subunits. Underpinning this complexity and functional specialisation is a network of structural connections, which may be probed using diffusion tractography, a unique, powerful and non-invasive MRI technique. Estimates of brain connectivity derived through diffusion tractography allow for explorations of how the brain's functional subunits are inter-linked to subsequently produce experiences and behaviour.

This thesis develops new diffusion tractography methodology for mapping brain connectivity, both across individuals and also across species; and explores frameworks for discovering associations of such brain connectivity features with behavioural traits. We build upon the hypothesis that connectional patterns can probe regions of functional equivalence across brains. To test this hypothesis we develop standardised and automated frameworks for mapping these patterns in very diverse brains, such as from human and non-human primates. We develop protocols to extract homologous fibre bundles across two species (human and macaque monkeys). We demonstrate robustness and generalisability of these protocols, but also their ability to capture individual variability. We also present investigations into how structural connectivity profiles may be used to inform us of how functionally-related features can be linked across different brains. Further, we explore how fully data-driven tractography techniques may be utilised for similar purposes, opening the door for future work on data-driven connectivity mapping.

Subsequently, we explore how such individual variability in features that probe brain organisation are associated with differences in human behaviour. One approach to performing such explorations is the use of powerful multivariate statistical techniques, such as canonical correlation analysis (CCA). After identifying issues in out-of-sample replication using multi-modal connectivity information, we perform comprehensive explorations into the robustness of such techniques and devise a generative model for forward predictions, demonstrating significant challenges and limitations in their current applications. Specifically, we predict that the stability and generalisability of these techniques requires an order of magnitude more subjects than typically used to avoid overfitting and mis-interpretation of results. Using population-level data from the UK Biobank and confirmations from independent imaging modalities from the Human Connectome Project, we validate this prediction and demonstrate the direct link of CCA stability and generalisability with the number of subjects used per considered feature.

### Acknowledgements

This PhD was supported by a Medical Research Council (MRC IMPACT DTP) PhD Studentship (MR/N013913/1). An additional grant supported a visit to Yale University to undertake a collaborative project described in Chapter 6: University of Nottingham Graduate School Medical Research Council student flexible funding (MR/K501360/1). I would like to thank the MRC IMPACT DTP leads, particularly Karen Robinson at Nottingham, and our Cohort Training Officer, Vikki Harrison, who have delivered a brilliant and supportive course, and continue to support me as I transition to becoming an independent researcher.

Much of the computations described in this thesis were performed using the University of Nottingham's Augusta HPC service and the Precision Imaging Beacon Cluster, which provide High Performance Computing service to the University's research community. If these systems (specifically the GPUs on the HPC) were not in place, tractography for the HCP cohort would still be running. As such, I am very grateful to the IT department and cluster support staff at the University of Nottingham. MRI data were provided in part by a) The Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University and b) The UK Biobank Resource under Application Number 43822.

I would like to thank my supervisors Dr Stam Sotiropoulos and Prof. Theo Kypraios. Stam, your support, guidance and advice at every stage of my PhD has been invaluable. I am deeply grateful for the opportunities you have helped me explore, and for your patience. I would also like to thank Dr Paul Morgan, who trained me in MRI physics, for guiding me through the very early stages of my career. Thank you to all of my collaborators who contributed their expertise and time to various aspects of my thesis. I would like to express my gratitude towards collaborators in Oxford University, in particular Katherine Bryant, Rogier Mars and Saad Jababdi, who contributed significantly to the development of XTRACT and its tractography protocols - the anatomical expertise provided by Katherine Bryant and Rogier Mars underpins XTRACT.

I would like to thank Alan Anticevic at Yale University who hosted me for a twomonth research placement. The time I spent in Yale had a significant impact on my training, gave me fresh perspectives on neuroscience research and led to successful collaborative projects. This gratitude extends to the whole of the Anticevic Lab, and to John Murray and Markus Helmer for their collaborative efforts in Chapter 6.

Thank you to everyone in the office who I've shared so many coffee breaks with and who have endured my grumpy face. A special thanks goes to fellow "Stamite", Dr Ellie Thompson, who paved the way as the lab's first PhD student, setting a high bar in the process, and has always been available for the occasional venting of PhD stress.

Most of all, I would like to thank my family: my Mum and Dad, my brother and sister Alex and Lyndsey, my wife Dr Kayleigh Warrington, and my cat Monty. Mum and Dad, you have always been supportive and helped whenever you can - thank you. Monty, thanks for the fluff, head bumps and awkward, but much appreciated, mid-work cuddles during lockdown. To Kayleigh - I wouldn't and couldn't have done this without you - thank you for being with me every step of the way, from A-levels to PhD, infinity and beyond.

### List of Publications

### Journals

- Warrington, S., Bryant, K. L., Khrapitchev, A. A., Sallet, J., Charquero-Ballester, M., Douaud, G., Jbabdi, S., Mars, R. B., and Sotiropoulos, S. N. "XTRACT Standardised protocols for automated tractography in the human and macaque brain", NeuroImage, 217:116923, 2020
- Helmer, M., <u>Warrington, S.</u>, Ji, L.J., Howell, A., Rosand, B., Anticevic, A., Sotiropoulos, S.N., Murray, J.D. "On stability of Canonical Correlation Analysis and Partial Least Squares with application to brain-behavior associations", bioRxiv, 2020 DOI: 10.1101/2020.08.25.265546v1
- Schilling, K.G., Rheault, F., Petit, L., Hansen, C.B., Nath, V., Yeh, F.-C., Girard, G., Barakovic, M., Rafael-Patino, J., Yu, T., Fischi-Gomez, E., Pizzolato, M., Ocampo-Pineda, M., Schiavi, S., Canales-Rodriguez, E.J., Daducci, A., Granziera, C., Innocenti, G., Thiran, J.-P., Mancini, L., Wastling, S., Cocozza, S., Petracca, M., Pontillo, G., Mancini, M., Vos, S.B., Vakharia, V.N., Duncan, J.S., Melero, H., Manzanedo, L., Sanz-Morales, E., Pena-Melian, A., Calamante, F., Attye, A., Cabeen, R.P., Korobova, L., Toga, A.W., Ambili Vijayakumari, A., Parker, D., Verma, R., Radwan, A., Sunaert, S., Emsell, L., De Luca, A., Leemans, A., Bajada, C.J., Haroon, H., Azadbakht, H., Chamberland, M., Genc, S., Tax, C.M.W., Yeh, P.-H., Srikanchana, R., Mcknight, C., Yang, J.Y.M., Chen, J., Kelly, C.E., Yeh, C.-H., Cochereau, J., Maller, J.J., Welton, T., Almairac, F., Seunarine, K.K., Clark, C.A., Zhang, F., Makris, N., Golby, A., Rathi, Y., O'Donnell, L.J., Xia, Y., Baran Aydogan, D., Shi, Y., Fernandes, F.G., Raemaekers, M., Warrington, S., Michielse, S., Ramirez-Manzanares, A., Concha, L., Aranda, R., Rivera Meraz, M., Lerma-Usabiaga, G., Roitman, L., Fekonja, L.S., Calarco, N., Joseph, M., Nakua, H., Voineskos, A.N., Karan, P., Grenier, G., Haitz Legarreta, J., Adluru, N., Nair, V.A., Prabhakaran, V., Alexander, A.L., Kamagata, K., Saito, Y., Uchida, W., Andica, C., Masahiro, A., Bayrak, R.G., Gandini Wheeler-Kingshott, C.A.M., D'Angelo, E., Palesi, F., Savini, G., Rolandi, N., Guevara, P., Houenou, J., Lopez-Lopez, N., Mangin, J.-F., Poupon, C., Roman, C., Vazquez, A., Maffei, C., Arantes, M., Andrade, J.P., Silva, S.M., Raja, R., Calhoun, V.D., Caverzasi, E., Sacco, S., Lauricella, M., Pestilli, F., Bullock, D., Zhan, Y., Brignoni-Perez, E., Lebel, C., Reynolds, J.E., Nestrasil, I., Labounek, R., Lenglet, C., Paulson, A., Aulicka, S., Heilbronner, S., Heuer, K., Anderson,

A.W., Landman, B.A., Descoteaux, M. "Tractography dissection variability: what happens when 42 groups dissect 14 white matter bundles on the same dataset?", bioRxiv, 2020 DOI: 10.1101/2020.10.07.321083v1

 Tewarie, P., Meier, J., Prasse, B., Mandke, K., <u>Warrington, S.</u>, Stam, C.J., Brookes, M.J., Van Mieghem, P., Sotiropoulos, S.N., Hillebrand, A. "Predicting time-resolved electrophysiological brain networks from structural eigenmodes", Under review

#### Conferences

- Warrington, S., Helmer, M., Ji, L.J., Mohammadi-Nejad, A.-R., Anticevic, A., Murray, J.D. Sotiropoulos, S.N. "Exploring the stability of canonical correlation analysis between imaging and non-imaging datasets", Organization for Human Brain Mapping Annual Meeting, June 2020.
- Rafipoor, H., <u>Warrington, S.</u>, Bryant, K., Sotiropoulos, S.N., Cottaar, M., Mars, R.B., Jbabdi, S. "Cross-species parcellation of the corpus callosum using joint embedding of connectivity blueprints", Organization for Human Brain Mapping, June 2020
- Helmer, M., <u>Warrington, S.</u>, Ji, L.J., Howell, A., Rosand, B., Anticevic, A., Sotiropoulos, S.N., Murray, J.D. "Reliable estimation of Canonical Correlation Analysis (CCA) and Partial Least Squares (PLS) with application to brain-behavior associations", Organization for Human Brain Mapping Annual Meeting, June 2020.
- Howell, A., <u>Warrington, S.</u>, Ji, J.L., Kolobaric, A., Adkinson, B., Fonteneau, C., Sotiropoulos, S.N., Murray, J.D., Anticevic, A. "Exploring the topography of structure-function mappings across cortico-thalamic systems", Organization for Human Brain Mapping, June 2020
- Rafipoor, H., <u>Warrington, S.</u>, Bryant, K. L., Sotiropoulos, S. N., Cottarr M., Mars, R. B., Jbabdi, S. "Cross-species parcellation of the corpus callosum using joint embedding of connectivity blueprints", Organization for Human Brain Mapping Annual Meeting, June 2020.
- Howell, A., <u>Warrington, S.</u>, Ji, J.L., Adkinson, B., Sotiropoulos, S.N., Murray, J.D., Anticevic, A. "Defining Targeted Projection Patterns in Thalamus Using Diffusion Weighted Imaging", Biological Psychiatry 87(9):S380, DOI: 10.1016/j.biopsych.2020.02.973, 2020
- <u>Warrington, S.</u>, Bryant, K. L., Khrapitchev, A. A., Sallet, J., Charquero-Ballester, M., Douaud, G., Jbabdi, S., Mars, R. B., and Sotiropoulos, S. N.

"Standardised protocols for automated tractography and connectivity blueprints", MRC IMPACT DTP Student Symposium, 2019 - Awarded 1st Place in the poster competition

- Warrington, S., Bryant, K. L., Khrapitchev, A. A., Sallet, J., Charquero-Ballester, M., Douaud, G., Jbabdi, S., Mars, R. B., and Sotiropoulos, S. N. "Standardised protocols for automated tractography and connectivity blueprints", Organization for Human Brain Mapping Annual Meeting, June 2019.
- Howell, A.M., <u>Warrington, S.</u>, Ji, J.L., Adkinson, B., Sotiropoulos, S.N., Murray, J.D., Anticevic, A. "Characterizing higher-order thalamo-cortical projection patterns in humans", Society for Neuroscience, October 2019
- Ganjgahi, H., Bijsterbosch, J., Daw, E.W., Donohue, B., Fieremans, E., Gahn, D., Glasser, M., Harms, M., Hodge, M., Jbabdi, S., Kochunov, P., Marchini, J., Novikov, D., Smith, S., Sotiropoulos, S., Van Essen, D., Veraart, J., <u>Warrington, S.</u>, Winkler, A., Nichols, T. "Genetics of the Young Adult Human Connectome Project", Organization for Human Brain Mapping, June 2018

### Publicly-released Software

- XTRACT, along with periphery tools xtract\_viewer and xtract\_stats, have been freely and publically released in FMRIB's Software Library (FSL) https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/XTRACT
- WM tract atlases for the human (HCP) have been released in FSL, and equivalent atlases for the macaque brain and derived from the UK Biobank are available via GitHub (https://github.com/SPMIC-UoN/XTRACT\_atlases). Atlases of the connectivity blueprints for the human and macaque brain are also available under the same location.
- Microstructural diffusion atlases, "HCP1065 standard-space DTI templates", derived from the "WU-Minn" Human Connectome Project, have also been released is FSL. https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/Atlases
- gemmr Generative Modeling of Multivariate Relationships (led by Markus Helmer). A python-based package to estimate the required sample sizes for Canonical Correlation Analysis (CCA) and Partial Least Squares (PLS), generate synthetic datasets and perform CCA and PLS analyses. https://github.com/murraylab/gemmr

## Contents

A	bstra	$\mathbf{ct}$		i			
A	cknov	wledge	ments	ii			
Li	st of	Publi	cations	iv			
1	Intr	Introduction					
	1.1	Overv	iew	2			
	1.2	Layou	t of This Thesis	7			
	1.3	Public	cly-released Software	9			
	1.4	Comp	uting Infrastructure and Software	10			
<b>2</b>	Bac	kgrou	nd	11			
	2.1	Introd	luction to MRI Physics	12			
		2.1.1	Spins, Precession and Relaxation	12			
		2.1.2	Imaging and k-space	15			
		2.1.3	Spin Echo and Echo Planar Imaging	21			
		2.1.4	Summary	23			
	2.2	Introd	luction to Diffusion MRI	24			
		2.2.1	The Basics of dMRI	24			
		2.2.2	Diffusion Modelling - The Diffusion Tensor	27			
		2.2.3	Beyond The Diffusion Tensor	31			
		2.2.4	Tractography	37			
	2.3	Introd	luction to Resting-State Functional MRI	43			
		2.3.1	The Basics of fMRI	43			
		2.3.2	Estimating Functional Connectivity Using fMRI	44			
	2.4	Summ	ary and Thesis Context	48			
3	$\operatorname{Big}$	Data	and Large-Scale Neuroimaging Projects	49			
	3.1	Introd	luction	50			
	3.2	The H	Iuman Connectome Project	51			
		3.2.1	Cohort Overview	51			

		3.2.2	Non-Imaging Data	52
		3.2.3	Imaging Data	54
		3.2.4	The HCP Minimal Preprocessing Pipeline (MPP)	56
	3.3	The U	K Biobank	59
		3.3.1	Cohort Overview	59
		3.3.2	Non-Imaging Data	60
		3.3.3	Imaging Data	61
		3.3.4	Preprocessing Pipeline	62
	3.4	HCP :	Structural Connectivity Features and Atlases	64
		3.4.1	White Matter Microstructural Atlases	65
		3.4.2	Structural Connectomes	67
	3.5	Summ	ary and Thesis Context	70
4	Star tivi	ndardi ty Feat	sed and Automated Tractography for Deriving Connec- tures	71
	4.1	Introd	luction	73
	4.2	Metho	ds	76
		4.2.1	XTRACT and the Tractography Protocols	76
		4.2.2	Data and Preprocessing	88
		4.2.3	Fibre Orientation Estimation and Tractography	90
		4.2.4	Connectivity Blueprints	90
		4.2.5	Atlas Generation	91
		4.2.6	Protocol Evaluation and Validation Methods	92
		4.2.7	Capturing Inter-Subject Variability	95
	4.3	Result	S	96
		4.3.1	Atlases	96
		4.3.2	The Effect of Sample Size on Atlas Generation	98
		4.3.3	Robustness Against Datasets	99
		4.3.4	Reflecting Known Anatomy	101
		4.3.5	Capturing Inter-Subject Variability	107
	4.4	Discus	ssion	108
	4.5	Concl	uding Remarks	114

4.6 Appendix			ndix	. 116
		4.6.1	Summary of the Tractography Protocols Previously Defined	
			in the Literature	. 116
		4.6.2	XTRACT User-Guide	. 118
		4.6.3	Summary of the Tract Lateralisation Literature	. 122
		4.6.4	The Effect of Resolution on Connectivity Blueprints	. 123
<b>5</b>	Bui	lding a	and Mapping Onto Common Connectivity Spaces	124
	5.1	Introd	luction	. 126
	5.2	Cross-	-species Mapping onto Connectivity Spaces	. 128
		5.2.1	Connectivity Blueprints	. 129
		5.2.2	Joint-Spectral Embedding	. 129
		5.2.3	Projecting Myelin Maps Across Species	. 131
		5.2.4	Data	. 133
		5.2.5	Results	. 134
	5.3	Buildi	ing Connectivity Spaces using Data-Driven Methods	. 135
		5.3.1	Myelin Map Predictions	. 140
		5.3.2	Data	. 141
		5.3.3	Results	. 141
	5.4	Discus	ssion	. 143
	5.5	5.3.3 Results 14   Discussion 14   Concluding Remarks 14		
	5.6	Apper	ndix	. 148
		5.6.1	Landmark regions-of-interest	. 148
		5.6.2	Landmark connectivity	. 149
6	Mu	ltivaria	ate Statistical Approaches for Brain-Behaviour Associ	a-
	tion	IS		150
	6.1	Introd	luction	. 152
	6.2	CCA	Theory	. 154
	6.3	Prior	Work using CCA	. 158
	6.4	Evalu	ating CCA Stability Against Multi-modal Data	. 160
		6.4.1	Feature Extraction: Imaging Features	. 163
		6.4.2	Feature Extraction: Non-Imaging Features	. 165

A	crony	vms		<b>235</b>
Bi	bliog	graphy		208
		7.1.3	Brain-Behaviour Associations	. 207
		7.1.2	Using Connectivity to Map Brain Diversity	. 206
		7.1.1	The Standardisation of Tractography	. 205
	7.1	Future	e Directions	. 205
7	Con	clusio	ns and Future Directions	203
		6.9.4	UK Biobank CCA top Non-imaging Loadings	. 200
		6.9.3	Behavioural Measures used in CCA - UK Biobank Analyses	. 196
		6.9.2	Behavioural Measures used in CCA - HCP Analyses	. 194
		6.9.1	Double-dipping and CCA	. 192
	6.9	Appendix		. 192
	6.8	Conclu	uding Remarks	. 190
		6.7.1	Recommendations	. 186
	6.7	Discus	sion $\ldots$	. 184
		6.6.6	Summary	. 183
		6.6.5	Results: Correspondence Between Modes	. 177
		6.6.4	Results: Overall Stability	. 176
		6.6.3	Analysis	. 176
		6.6.2	Feature Extraction: Non-Imaging Features	. 175
		6.6.1	Feature Extraction: Imaging Features	. 174
	6.6	CCA A	Applied to Large Samples	. 174
		6.5.3	Summary	. 172
		6.5.2	Results	. 171
		6.5.1	Simulation Framework	. 169
	6.5	CCA v	with Known Ground-Truths	. 169
		6.4.5	Summary	. 167
		6.4.4	Results	. 166
		6.4.3	Analysis	. 166

## List of Tables

3.1	Summary demographics of the HCP cohort	52
3.2	Summary subject health and family history of the HCP cohort	54
3.3	MRI acquisition parameters for the HCP	55
3.4	Summary demographics of the UK Biobank cohort $\hfill \ldots \ldots \ldots$ .	60
3.5	MRI acquisition parameters for the UK Biobank	62
4.1	WM tracts supported in XTRACT	76
4.2	A review of the tractography protocols previously defined in the lit-	
	erature	117
4.3	WM tract lateralisation results	122
6.1	Summary of the CCA literature	161

# List of Figures

1.1	Neurons as sketched by Cajal and imaged via electron microscopy	3
2.1	The simple MRI experiment	16
2.2	The spatial localisation of signal in MRI and the gradient echo sequence	20
2.3	The spin echo and echo planar imaging sequences $\ldots \ldots \ldots \ldots$	22
2.4	Schematics of the neuron and examples of diffusion in the three main	
	brain tissue types	25
2.5	The Stejskal-Tanner pulsed gradient spin echo sequence	26
2.6	Examples of diffusion weighted images of the brain	29
2.7	Microstructural feature maps derived from diffusion MRI	32
2.8	Examples of diffusion tensor ellipsoids	32
2.9	The issue of multiple fibre bundles	33
2.10	A schematic of tractography	39
2.11	A simple comparison of deterministic and probabilistic tractography .	40
2.12	The BOLD effect	44
3.1	The HCP cohort-averaged FA map	66
3.2	An enhanced view of the HCP group-averaged first eigenvector map .	67
3.3	Cohort-averaged orientation dispersion index and intra-cellular vol-	
	ume fraction maps	68
3.4	The group averaged structural connectomes derived from the HCP	69
4.1	The XTRACT processing pipeline	78
4.2	Schematic of the construction of connectivity blueprints	92
4.3	Schematic of the calculation of the KL divergence between hemispheres	95
4.4	WM tract atlases for the human and macaque brain	97
4.5	Connectivity blueprint atlases for the human and macaque brain	98
4.6	Comparison of atlases for varying sample sizes	99
4.7	Summary of inter-cohort robustness	100
4.8	Comparisons of the tract volumes across cohorts	101

4.9	Summary of the WM tract lateralisation
4.10	The minimum KL divergence between hemispheres
4.11	Examples of the WM tract contribution of the KL divergence 104
4.12	Parcel-averaged tract contributions to a subset of the cortex 106
4.13	Comparisons of tract similarity considering familial structure 107
4.14	Examples of XTRACT in cases of atypical anatomy
4.15	The XTRACT command-line user interface
4.16	Comparisons of the connectivity blueprints across resolutions 123
5.1	Summaries of the myelin map prediction approaches
5.2	The Kullback-Leibler divergence between the macaque and human
	cortical termination maps
5.3	The joint-similarity matrix and embedded components
5.4	Predictions of the macaque myelin map
5.5	Absolute difference between the predicted and measured macaque
	myelin maps
5.6	The extraction of WM and GM components through NMF 139
5.7	Examples of the NMF components
5.8	Predicting myelin across subjects using data-driven tractography 143 $$
5.9	The homologous landmarks used in deriving the joint-embedded com-
	ponents
5.10	The tract contribution to each of the homologous landmarks 149
6.1	Schematic of the CCA equation
6.2	Schematic of a typical permutation testing scheme
6.3	Schematic of the CCA subsampling scheme
6.4	Schematic of the imaging-data preprocessing for CCA
6.5	Plots of the CCA stability using the HCP data
6.6	Schematic of the CCA simulation framework
6.7	Simulation framework sanity checks
6.8	The stability of the canonical correlations and canonical weights with
	known ground-truths

6.9	The stability of CCA at the population-level
6.10	The stability of canonical correlations considering the top 10 compo-
	nents using the UK Biobank cohort
6.11	Correlation matrices showing the correspondence between canonical
	modes derived from the UK Biobank CCA analysis $\ . \ . \ . \ . \ . \ . \ . \ . \ . \ $
6.12	The top non-imaging contributors to each canonical mode for the UK
	Biobank CCA
6.13	CCAs reported in the population neuroimaging literature might often
	be unstable
6.14	Sample sizes required for sparse CCA compared to classic CCA 189
6.15	The consequence of double-dipping in CCA
6.16	The top non-imaging contributors to each canonical mode for the UK
	Biobank CCA with 20k subjects

# Chapter 1

## Introduction

### Contents

1.1	Overview	<b>2</b>
1.2	Layout of This Thesis	7
1.3	Publicly-released Software	9
1.4	Computing Infrastructure and Software	10

#### 1.1 Overview

Understanding the human brain and how it gives rise to behaviour and eventually cognition is one the most compelling challenges of our time. Early investigations into the anatomy and organisation of the brain, and its relationship to function, were speculative and often founded on misguided conceptions of what the brain actually is. In the 17th century, the Danish anatomist Nicolaus Steno argued for a move towards a more methodical approach to understanding the brain. Considering the brain as a machine, Steno argued that not only should we describe the component parts of the brain, we should seek to understand the global consequence of those components and their individual function (Steno, 1669; Cobb, 2020). Of course, this is not a simple task and the general principles that Steno argued for still guide current neuroscience research.

Through the 17th-19th centuries, developments in our understanding of electricity, the finding that nerve fibres transmit signals, the consensus that all biological tissues are composed of discrete cells, and the development of histological cell staining led to the conception of the neuron doctrine. This defined the nervous system as a collection of individual cells based on the works of Santiago Ramón y Cajal and others (Figure 1.1a). Although in line with the consensus on other biological tissues and winning the Nobel Prize for Physiology or Medicine in 1906, the neuron doctrine was disputed by some until the definitive observations made via electron microscopy in the 1950s which depicted neuronal cells separated by the synapse (Figure 1.1b) (Palade, 1954).

Still to this day, our understanding of how the brain works is in general terms - we can measure activity and relate it to external events and stimuli; we can define subdivisions of the brain based on the differences in anatomical structure; and we can precisely measure the electrical activity of neurons, yet, these concepts, and others, lack a comprehensive description and unification in how matter gives rise to



Figure 1.1: a. An example of Cajal's drawings of pyramidal cells. b. Early electron microscope images of neurons, showing clear boundaries - the synapse - between two neurons. The arrows indicate the synapse with the neurons labelled **a** and **b**. Edited from López-Muñoz et al. (2006).

human behaviour. We are not yet at the stage of unifying models of coarse brainbehaviour relationships and neuronal activity. To progress towards this end-goal, further developments in our understanding of all aspects of the brain are required. Two approaches to understanding the brain are of particular interest in this work, both of which consider the coarse, global features of the brain. The first relates to the structural architecture of the brain and how mapping this architecture may inform our understanding of brain function, even across non-human primate species. The second is towards a better understanding of the relationships between the architecture of the brain and human behaviour. Such investigations have become possible in recent decades through the rapid development of medical imaging techniques, computational power and large-scale neuroscience efforts.

The advances and frameworks presented in this thesis are based on magnetic resonance imaging (MRI) - a widely used clinical and research imaging modality. Since its clinical adoption, MRI has reformed diagnostic imaging and led to significant advancements in our understanding of the normal and abnormal structure and function of the human body (Haacke et al., 1999; McRobbie et al., 2005; Edelman, 2014; Lui et al., 2016; Lerch et al., 2017; Beek et al., 2019; Sarubbo and Petit, 2019). Most notably, MRI has propelled forward the field of neuroscience, which previously relied on poor tissue contrast methods such as x-ray imaging techniques or field-of-view limited and time consuming post-mortem methods, such as lightimaging techniques, microscopy, anatomical tracing and histology. MRI allows for safe, non-invasive, high-soft-tissue-contrast imaging of the human brain with useful spatial resolution and full brain coverage. This makes MRI ideal for investigating the underlying nature of the brain, albeit at the macroscale. More specifically, in neuroscience, MRI has advanced our understanding of large-scale brain anatomy, the functional specialisation and organisation of the brain, and has afforded us the opportunity the measure brain changes during disease progression non-invasively.

Several MRI based techniques - modalities - have been developed which are capable of providing unique information with regard to the structure and physiology of the brain. Two MRI modalities are of focus in this research: diffusion MRI (dMRI) and functional MRI (fMRI). dMRI and fMRI seek to understand the structural and functional architecture of the brain respectively.

Structural architecture may be explored at a range of scales. High-spatial resolution imaging methods, such as electron microscopy, aim to understand the brain at a ultra-structural/neuronal level (DeFelipe, 2010) by probing brain tissue structure at a microscale level. These methods are time consuming and limited in their fieldof-view. MRI-based methods are capable of probing brain connectivity on a global level by seeking to reconstruct macroscale connections in relatively short acquisition times. However, measuring brain connectivity through MRI is limited by the fact that measurements are indirect. Modelling is used to estimate quantities of interest and, in part, overcome some of the limitations associated with MRI (Jbabdi et al., 2015). Although limited, MRI is capable of providing in vivo semi-quantitative measures of brain connectivity on a scale and level of detail not achieved by any other imaging modality (Jbabdi et al., 2015).

Brain connectivity may be summarised by constructing connectivity adjacency matrices. In analogy to the term used in the study of the genome - genomics the term connectomics has been coined to encapsulate the study of brain connectivity matrices, or "connectomes" (Sporns et al., 2005). A connectivity matrix, or connectome, is derived by mapping the functional co-activation of brain regions or the structural connections (fibre pathways) between brain regions (Rubinov and Sporns, 2010). Such data may be acquired using an array of multi-modal techniques such as functional and diffusion MRI, electroencephalography (EEG) and magnetoencephalography (MEG). The connectome is defined by the connections (edges) between pairs of regions/vertices (nodes) and aims to mathematically represent complex real-world problems (Rubinov and Sporns, 2010). These connectomes allow us to study brain connectivity on a local and global level through a number of summary measures (for instance node degree and strength, community structure and modularity, network efficiency and topology) (Rubinov and Sporns, 2010) or regional patterns of connections ("connectivity fingerprints") (Passingham et al., 2002; Jbabdi et al., 2015; Mars et al., 2018b).

The two types of connectome, functional and structural, are closely related. The functional specialisation of brain regions is well-documented through history (Finger, 2009) with recent work dividing the brain into hundreds of specialised regions (Glasser et al., 2016). Underpinning this functional localisation and specialisation is the extremely complex network of axonal fibres, which carry signals between brain regions. It is the architecture of this vast network that constrains brain function.

It is the patterns of connections that define functional territories; they comprise a signature of brain organisation (Jbabdi and Behrens, 2012; Passingham et al., 2002). As such, it is the unique variation in structural architecture that underpins the individual mind. Therefore, we can gain insight into the brain and its function/dysfunction by understanding its structural architecture.

"A mechanism may operate before us without our recognizing the exact relation between its function and its architecture. But, on the other hand, if we are acquainted with the principles upon which this mechanism operates, we may infer its function from its structure, regarding the former as the natural outcome of the latter." - Meynert (1885), pg. 156

In the era of modern - 21st century - neuroscience, there is an abundance of data that can assist in understanding the brain (Fan et al., 2014; Sejnowski et al., 2014; Vu et al., 2018; Smith and Nichols, 2018; Bzdok et al., 2019). Neuroscience has learnt from and progressed beyond observations of individual brains and now relies on large-scale and population-based studies. Recent neuroimaging projects, unprecedented in scale, make it possible to study the links between the structural architecture of the brain and function at the population-level. At the time of writing, the UK Biobank has recruited half a million subjects, approximately forty-thousand of whom have brain MRI data available. Further, significant advancements in acquisition technology and techniques have led to a vast pool of superb quality data. The "WU-Minn" Human Connectome Project (HCP) pushed the boundaries of data acquisition and quality, producing the highest-quality multi-modal MRI data available for a large cohort of subjects. These advancements are promising in seeking to unravel the associations between structure and function. However, the methodology to utilise this data lags behind and new techniques are required.

In this thesis, we present approaches that are linked to these two challenges. First we develop methodology for mapping brain connectivity from diffusion MRI and extract patterns of connectivity in a way that is generalisable even in the absence of geometric similarities, across human and non-human primates. We then demonstrate how these patterns can used to probe functional organisation and similarities, even across very diverse brains. Second, we use inter-individual variability of connectivity features to explore links with behavioural traits. We demonstrate challenges and issues of existing techniques and propose solutions for robust data-driven brain-behaviour associations.

#### 1.2 Layout of This Thesis

This thesis focuses on the development of tools for robustly mapping structural connectivity features across individuals and species and relating those measures to functionally-relevant features, ultimately human behaviour. As such, Chapter 2 introduces the basic concepts behind magnetic resonance imaging (MRI), focusing on diffusion and touching on functional MRI. A brief overview of the physics of MRI is provided, followed up by the background to diffusion MRI data, modelling and tractography, specific to the development of automated tractography pipelines, in Chapter 4. In addition, an overview of the physics governing fMRI and how such techniques may be used to estimate functional connectivity are also presented. Chapter 3 provides overviews of the HCP and UK Biobank cohorts, which are the main sources of data for this thesis, and summarises the extraction of key connectivity feature maps used throughout this thesis or in side-projects associated with this thesis.

Chapter 4 presents the development of a standardised and automated crossspecies tractography toolbox (XTRACT). This includes comprehensive descriptions of a library of standardised tractography protocols, the automated tractography pipeline, and the validation of said protocols and pipeline. The utility of the toolbox is demonstrated by reconstructing major white matter bundles in the human brain using varying quality data, and their homologous bundles in the non-human primate (macaque) brain. For the human brain, these bundles are used to explore interhemispheric differences in structural connectivity relating to functional specialisation and lateralisation. Chapter 5 furthers the work presented in Chapter 4 to explore how such methods can be used to extract connectivity patterns and define a common connectivity space for probing functional similarity, even across very diverse brains, such as those from different primate species. We compare two approaches for identifying correspondence between human and macaque brains, using the tractography patterns derived in Chapter 4 and employing a) statistical similarity of connection patterns, b) joint-embedding of connectivity. Furthermore, we extend these ideas and explore how fully data-driven approaches for mapping connection patterns (as opposed to using a predefined and limited set of manually-devised tractography protocols) can be also used for similar purposes.

Chapter 6 focuses on the application of multivariate statistical techniques in exploring the associations between brain connectivity and behavioural traits, as captured by non-imaging measures, such as lifestyle, physical health and cognition. Multivariate statistical techniques, such as canonical correlation analysis (CCA) and partial least squares (PLS), are powerful techniques which, recently, have been used to reveal latent associations in large dimensionality data, often with few subjects. Chapter 6 performs comprehensive investigations into the robustness of such techniques, focusing on CCA, and demonstrates significant challenges and limitations in their application when not enough data are available. Issues with regard to the number of subjects compared to the number of features are identified, which lead to significant overfitting and the potential mis-interpretation of results. These results challenge the definition of "large" datasets in the context of multidimensional brainbehaviour association studies. Using data from the UK Biobank, we demonstrate that, in order to obtain robust patterns, tens of thousands of subjects are needed when 50-100 imaging/non-imaging features are explored, suggesting that an order of magnitude more subjects are needed than used in similar (under-powered) studies before.

Finally, conclusions and future directions are presented in Chapter 7. The works

presented in this thesis are discussed in terms of their limitations, utility, and their impact and wider context.

#### 1.3 Publicly-released Software

This thesis develops, or contributes to, several publicly released software. The first of which is XTRACT. This is a tool for the standardised and automated extraction of major white matter fibre bundles in the human and non-human primate brain (see Chapter 4 for full details). It was publicly released in September 2019 as a part of one of the most commonly used neuroimage processing software packages, FMRIB's Software Library (FSL, v6.0.2).

A set of white matter tract atlases and a set of diffusion microstructural atlases, both derived from the HCP cohort, have also been publicly released in FSL (see Chapter 4 for full details). The white matter tract atlases were derived by applying XTRACT to each HCP subject and averaging across the cohort. Equivalent atlases for the non-human primate (macaque) brain and from the UK Biobank, and atlases of the connectivity blueprints for both the human and macaque brain have also been made available via GitHub (github.com/SPMIC-UoN/XTRACT\_atlases).

Chapter 6, which is concerned with investigations into the use of multivariate statistical techniques, was a collaborative project with colleagues from the Murray Lab and Anticevic Lab at Yale University. Led by Markus Helmer, a python-based tool to calculate the required sample sizes for such techniques was developed: *GEM-MER* - Generative Modeling of Multivariate Relationships. The details of this package may be found in the associated manuscript (Helmer et al., 2020) and through GitHub (github.com/murraylab/gemmr).

#### 1.4 Computing Infrastructure and Software

The primary processing system in the early stages of this work was an 80 core computing server (Intel Xeon CPU E5-2698 v4, 2.20 GHz) CentOS Linux v7 (630 GB total memory). Importantly, the server has four graphical processing units (GPUs) available (NVIDIA Tesla K80, 11.5 GB memory per GPU). GPUs allow for extensive and rapid data processing by employing parallel processing on a large scale (Hernandez-Fernandez et al., 2019).

From year two of this work and onwards, the primary processing system was the "Imaging Beacon" partition of the University of Nottingham High-Performance Computer (HPC). This partition consists of 600 cores with memory specifications ranging from 192 GB to 1536 GB depending on the node and 10 high-end GPUs (6 Nvidia Tesla V100 and 4 Nvidia Tesla P100).

Statistical analysis were performed using the Python programming language (version 3.7.2) (Python Software Foundation, https://www.python.org) along with several community-developed python libraries (mainly NumPy, SciPy, pandas, Ni-Babel and Matplotlib) and MATLAB 2018a (The MathWorks, Inc., Natick, Massachusetts, United States). Visualisation and processing of surface data was completed using the Connectome Workbench (version 1.3.2; Van Essen Laboratory, Washington University). General image processing were completed using the FM-RIB software library, FSL (version 5.0.11 and beyond) (Jenkinson et al., 2012).

# Chapter 2

# Background

### Contents

<b>2.1</b>	Intro	duction to MRI Physics
	2.1.1	Spins, Precession and Relaxation
	2.1.2	Imaging and k-space
	2.1.3	Spin Echo and Echo Planar Imaging
	2.1.4	Summary
2.2	Intro	duction to Diffusion MRI
	2.2.1	The Basics of dMRI 24
	2.2.2	Diffusion Modelling - The Diffusion Tensor
	2.2.3	Beyond The Diffusion Tensor
	2.2.4	Tractography
<b>2.3</b>	Intro	duction to Resting-State Functional MRI 43
	2.3.1	The Basics of fMRI 43
	2.3.2	Estimating Functional Connectivity Using fMRI 44
<b>2.4</b>	Sum	nary and Thesis Context

#### 2.1 Introduction to MRI Physics

#### 2.1.1 Spins, Precession and Relaxation

Magnetic resonance imaging (MRI) is an imaging technique that uses powerful magnetic fields and radio frequency (RF) fields to measure distributions of protons in biological samples. Most commonly, these protons are bound in water and fatty tissues. Protons, as with all subatomic particles, possess the quantum mechanical property known as spin, which describes the particle's angular momentum. The spin (i.e. spin quantum number) of a proton is s = 1/2, in dimensionless units, corresponding to the quotient of the spin angular momentum  $(\mathbf{S}, \text{ units of } \text{kgm}^2 \text{s}^{-1})$  and the reduced Planck constant  $\hbar$ . Given this non-zero spin and the magnetic and mass properties of protons, we can deduce the spin magnetic moment  $\mu_s$  as  $\mu_s = \gamma S$ , where  $S = \frac{\hbar}{2}$  for a proton and  $\gamma$  is the gyromagnetic ratio. The gyromagnetic ratio (with units of  $rads^{-1}T^{-1}$  but often reported as  $MHzT^{-1}$ ) is fundamental in the understanding and workings of MRI. It describes the ratio between the magnetic moment and the angular momentum of a particle. The proton has a gyromagnetic ratio of  $42.57 \text{ MHzT}^{-1}$ . Importantly, from the gyromagnetic ratio, we can see that the angular momentum of a proton is not only dependent on its own properties, but also on any external magnetic field.

When an external static magnetic field  $\mathbf{B}_0$  is introduced, as in MRI, protons within that field experience a torque, the size of which is dependent on the strength of the magnetic field. This torque forces the protons towards alignment/anti-alignment (known as spin-up and spin-down) with the magnetic field. Given a distribution of protons, the ratio of the number of protons in the spin-up state compared to the spin-down state is given by Boltzmann's distribution,  $\frac{N_{up}}{N_{down}} \approx 1 + \frac{\gamma \hbar \mathbf{B}_0}{k_B T}$ , where  $k_B$ is the Boltzmann constant and T is the temperature of the sample (e.g. in the field of MRI, this is most commonly body temperature).

On average, in the typical case where the temperature of a biological sample is approximately 300 Kelvin, there is a greater number of protons in the spin-up state compared to the spin-down state by a few parts per million. This slightly greater favouring to the spin-up state leads to a non-zero net magnetisation  $\mathbf{M}_0$ caused by the averaging of the magnetic fields carried by the protons. It is this net magnetisation that MRI seeks to quantify. However, due to the tiny magnitude of the net magnetisation compared to the main magnetic field  $\mathbf{B}_0$ , to be able to observe it, we must rotate its direction with respect to the external magnetic field of the net magnetisation. A simple choice, and one in which the observable signal is maximised, is to rotate the axis of the net magnetisation  $\mathbf{M}_0$  such that the net magnetisation and the main magnetic field  $\mathbf{B}_0$  are orthogonal, i.e.  $\mathbf{M}_0 \perp \mathbf{B}_0$ . In fact, the axis of the net magnetisation may be rotated by any angle, however, the observable signal is limited to the component of the vector of the net magnetisation along the plane orthogonal to the main magnetic field. From this we may define a set of axes x, y, z. We will consider the main magnetic field  $\mathbf{B}_0$  and, initially, the net magnetisation  $\mathbf{M}_0$  as being along the z-axis. We can refer to this net magnetisation as  $M_z = |\mathbf{M}_0|$  and net magnetisation in the x - y plane as  $M_{xy}$  which, initially, is zero. Before discussing the effect of rotating the plane of the net magnetisation, we must understand an important aspect of the interaction between the protons and the main magnetic field.

Due to the spin angular momentum of the proton, protons exposed to an external magnetic field cannot align exactly to the magnetic field. Instead, their axis is offset. Therefore, they start to precess about the main axis of the magnetic field. The frequency of this precession is known as the Larmor frequency and is determined by the magnetic field strength and the gyromagnetic ratio, given by Equation 2.1.

$$\omega_0 = \gamma |\mathbf{B}_0| \tag{2.1}$$

where  $\mathbf{B}_0$  (units of Tesla, T) is used to denote the magnetic field strength of the main magnetic field in MRI.

By applying an RF electromagnetic field pulse (with a magnetic field strength of  $\mathbf{B}_1$ ) at the resonant/Larmor frequency, and thus meeting the resonance condition, it is possible to rotate the axis about which the protons precess, thus altering the direction of the net magnetisation with reference to the main magnetic field's axis (the z-axis). Typically, the RF field is applied such that the net magnetisation axis is rotated 90° compared to the main axis. This may be thought as providing enough energy at the resonant frequency such that the distribution of proton spin states becomes equalised, i.e. there is no net magnetisation in the axis of the main magnetic field. Mathematically, this may be summarised as  $M_z = |\mathbf{M}_0|, M_{xy} = 0 \stackrel{\alpha}{\Rightarrow}$  $M_z = 0, M_{xy} = |\mathbf{M}_0|$  where  $\alpha = 90^\circ$  RF pulse. This concept may be extended to considering various angles of rotation. For example, a 180° rotation may be considered as flipping the spin state of each proton such that the net magnetisation is now anti-parallel to the main magnetic field.

In addition to the effect of rotating the axis about which the protons precess, the RF field also enforces phase coherence on to the collection of protons. Whereas before, when precessing about the main magnetic field's axis, the protons were out of phase with each other, and thus there was zero net magnetisation about any axis other than the main magnetic fields axis (i.e.  $M_x = M_y = 0$ ), there is now a relatively strong, phase coherent net magnetisation in the plane of the new **B**<sub>1</sub> field which is still precessing about the main magnetic field **B**<sub>0</sub>. We refer to this new plane, when a 90° RF pulse has been applied, as the transverse plane (the x - yplane).

Using a wire coil<sup>1</sup> tuned to the Larmor frequency, via electromagnetic induction described by the Maxwell-Faraday equation, it is possible to measure the electrical

<sup>&</sup>lt;sup>1</sup>In fact, two coils are used in quadrature, recording two signals: a real ( $\Re$ ) and an imaginary ( $\Im$ ) signal, that are 90° out of phase. The subsequent averaging of these signals ( $\sqrt{\Re^2 + \Im^2}$ ) boosts the signal-to-noise ratio (SNR)

voltage induced in the coil by the oscillating/precessing transverse plane net magnetisation. Over time, the protons lose phase coherence and return to their original spin-state via proton-proton and proton-environment interactions respectively. In other words, the net magnetisation in the transverse plane reduces back to zero and the net magnetisation in the main magnetic field axis returns to its original magnitude (i.e.  $M_z \stackrel{T1}{\Rightarrow} |\mathbf{M}_0|, M_{xy} \stackrel{T2}{\Rightarrow} 0$ , where T1 and T2 are the decay times described below). The time taken for these processes to occur is governed by the frequency with which the proton-proton and proton-environment interactions (and additional interactions not covered here) occur.

Commonly, the process of proton de-phasing is referred to as transverse relaxation and is described by the T2 time constant. Additional relaxation, caused by local inhomogeneities in the main magnetic field, also occurs, giving rise to the  $T2^*$ time constant. This is due to the Larmor frequency being dependent on the magnetic field strength. Imperfections in the external magnetic field and local changes in sample susceptibility result in small changes in the Larmor frequency of protons. Therefore, relative to neighbouring protons, protons affected by local inhomogeneities will precess at a slightly different rate, thus causing loss of phase coherence at a greater rate.

The process of the orientation of the net magnetisation returning to the axis of the main magnetic field (the z-axis) is referred to as longitudinal relaxation and is described by the T1 time constant. It is during this relaxation, or decay, that we may measure the voltage, or signal, produced by the proton distribution. The processes discussed above are the basic principles of the generation and measurement of signal in MRI and are summarised in Figure 2.1.

#### 2.1.2 Imaging and k-space

Further to the precession and relaxation processes, we may refine the signal we observe by embedding spatial information. We have mentioned that the Larmor



**Figure 2.1:** Schematics of the basic processes included in producing an observable signal in a simple MRI experiment, including the application of external magnetic and radiofrequency fields to a pool of protons, the process of proton relaxation, in which protons return to their equilibrium states following radiofrequency excitation, and the overall observed signal observed as a combination of the two main types of proton relaxation and spin precession. **1**. A pool of protons at equilibrium. The net magnetisation  $\mathbf{M}_0$  produced by the protons is approximately zero. **2**. An external magnetic field  $\mathbf{B}_0$  is applied to the pool of protons along the z-axis, causing the protons to move to parallel or anti-parallel orientations, in accordance with the Boltzmann distribution. There is now a non-zero net magnetisation. **3**. The protons precess about the z-axis (as indicated by the dashed blue line) at the Larmor frequency  $\omega_0$ with a  $\mathbf{M}_0$  parallel to the  $\mathbf{B}_0$  field. **4**. An RF electromagnetic field (carrying the associated magnetic field  $\mathbf{B}_1$ ) is applied to the system, flipping the  $\mathbf{B}_0$  from the z-axis into the x - y plane and forcing phase coherence on the protons, denoted as  $M_{xy}$ . Caption continued on the following page.

Figure 2.1: 5. Following the RF electromagnetic field pulse, the protons are initially in-phase with a maximum possible  $M_{xy} = |\mathbf{M}_0|$  amplitude. With time, governed by the transverse relaxation Bloch equation with time constant  $T_2$ , the protons begin to lose phase, thus decreasing  $M_{xy}$ , until completed loss of coherence where  $M_{xy} = 0$ . 6. Also immediately following the RF electromagnetic field pulse, the  $M_z$  net magnetisation has been flipped into the x-y plane (as shown in panel 4.), leaving zero net magnetisation in along the z-axis (i.e.  $M_z = 0$ ) at the end of the 90° RF pulse. Over time, governed by the longitudinal relaxation Bloch equation with time constant T1, the protons begin to return to their pre-RF pulse equilibrium and  $M_z$  increases until equaling  $|\mathbf{M}_0|$ . 7. The exponential decay and growth of  $M_{xy}$ and  $M_z$  respectively following a 90° RF pulse (top-right). The decay of  $M_{xy}$ , i.e. de-phasing of protons, combined with the precession of the spins induces a sinusoidal voltage in a wire coil (the 'Signal' axis) decreasing in amplitude over time. This signal is known as the free induction decay (FID). The envelope of the FID, shown by the green line corresponds to the  $T_2$  decay curve. Note: panels 1-3. are visualised from a stationary frame of reference. As is common in the MRI community, panel 4. and beyond are viewed within a rotating frame of reference in order to ease visualisation, indicated by the blue curved arrow in panel 4. (the frame of reference is rotating at the Larmor frequency).

frequency is dependent on the external magnetic field (Equation 2.1), giving rise to spin precession and, due to inhomogeneities, loss of phase coherence. This concept may be applied in a controlled manner to take advantage of these effects for spatial localisation.

Using an additional external spatially-varying magnetic field, we can adjust the Larmor frequency of the protons in a sample. Then, by applying an RF pulse at the adjusted Larmor frequency, we can excite protons in, and observe signal from, a given region. For example, we may apply the additional magnetic field  $G_z = f(z)$ , where f(z) is a linear function of z, such that there is a gradient in the magnetic field strength along the z-axis. In turn, this gives us a gradient of Larmor frequencies along the z-axis. During the application of the additional magnetic field, referred to as gradient magnetic fields, the RF pulse excites a spatially localised collection of the protons. This principle may be extended to three dimensions through the use three orthogonal gradient magnetic fields  $G_x, G_y, G_z$ . These fields may be applied in combination to extract signal for a specific three-dimensional confined space, i.e. a volume element - voxel (typically around  $1 - 10 \text{ mm}^3$  volume). Typically, the  $G_z$  gradient is used to restrict proton excitement to a given slab/slice of the sample (Figure 2.2 - panels 1-2). Hence, it is referred to as the slice selection gradient  $(G_{SS})$ . The location (z) of the slice along the  $G_{SS}$  gradient, and therefore along the sample, is controlled by the central frequency (referred to as  $\omega$  here) of the applied RF pulse, and slice thickness  $(\Delta z)$  is controlled by the bandwidth  $(\Delta \omega)$  (Equations 2.2-2.3). The  $G_x$  and  $G_y$  gradients are then used to sample the select slice.

$$\omega = \gamma (B_0 + z \cdot G_{SS}) \tag{2.2}$$

$$\Delta \omega = \gamma \cdot G_{SS} \cdot \Delta z \tag{2.3}$$

In reality, MRI acquisitions happen in what is referred to as k-space, also known as the spatial frequency domain: the Fourier transform of the spatial domain. Here, we acquire information regarding the spatial frequencies of the signal in a two dimensional plane  $k_x$ ,  $k_y$ , which are traversed by the  $G_x$  and  $G_y$  gradients. As discussed, a slice of the sample is selected using the slice selection gradient  $G_{SS}$  along with an initial RF pulse. The  $G_x$  (referred to as the frequency-encoding gradient  $G_{FE}$ ) and  $G_y$  (referred to a the phase-encoding gradient  $G_{PE}$ ) are then used to restrict the signal sampling to a location in k-space (each 'location' in k-space represents a spatial frequency). More specifically, the phase-encoding gradient  $G_{PE}$  controls the which row  $k_y$  of k-space is being sampled (k is the temporal integral of the given gradient, e.g.  $k_y = \frac{\gamma}{2\pi} \int_0^t G_{PE}(t') dt'$  and the frequency-encoding gradient  $G_{FE}$  is used to move along that row (i.e. along  $k_x$ ). By selecting a row of k-space  $k_y$  and sampling along  $k_x$ , a single line of k-space is acquired. The process is repeated whilst adjusting the  $G_{PE}$  gradient to move through k-space along the  $k_y$  axis, each time applying  $G_{FE}$  to collect data along the columns  $k_x$ . The combination of the frequency- and phase-encoded measurements allows for the extraction of the contribution from unique signal frequencies: multiple measurements at different phases (along  $k_y$ ) of signals with equal frequencies (along  $k_x$ ) allows for the separation of said signals, which otherwise would be represented as the sum of all signals at a given  $k_x$  location.

The central points in k-space, i.e. low frequencies, provide high contrast information but lack edge detail. The peripheral points, i.e. high frequencies, provide edge detail in the final image. This may be conceptualised by considering the high contrast information as low frequency signal undulations and the edge information as sharp, high frequency, changes in signal. Therefore, to obtain high-quality spatial information (i.e. an image with high contrast and edge detail), we must sample both the centre and periphery of k-space. Ideally, we would sample k-space infinitely, both in terms of the frequency resolution (i.e. an infinitely small difference between sampled k-space locations) and peripheral reach (i.e. sampling infinitely high frequencies). Of course, this is not possible and we must reach a compromise between the required field of view (FOV) in the spatial domain, the spatial resolution and the acquisition time. The required FOV in the spatial domain is generally well-defined and there is limited room for compromise as, for example, we must cover the whole brain. We then seek to maximise the spatial resolution whilst maintaining a feasible acquisition time. These spatial domain features (FOV and spatial resolution) relate inversely to the sampling of the frequency domain. The FOV of is controlled by the frequency resolution ( $\Delta k$ , i.e. the space between sampled locations) and the spatial resolution  $\Delta w$  (i.e. pixel width) is controlled by the range of k-space sampled ( $k_{\rm FOV}$ ) i.e. how far into the periphery we sample) (Equations 2.4-2.5).

$$FOV = \frac{1}{\Delta k} \tag{2.4}$$

$$\Delta w = \frac{1}{k_{\rm FOV}} \tag{2.5}$$

Therefore, we seek to sample k-space with a small enough  $\Delta k$  such that our FOV

covers the sample of interest, and far enough into the periphery of k-space to obtain useful spatial resolution in a feasible acquisition time. Advancements in k-space sampling schemes and image reconstruction aid in maximising spatial resolution and FOV whilst reducing acquisition time, however, details of such techniques are beyond the scope of this thesis. Once k-space has been filled (i.e. all  $k_y$  and  $k_x$  locations within the desired  $k_{\text{FOV}}$  have been sampled) for a single slice of the sample, the Fourier transform is applied, giving a two dimensional image. A three-dimensional image is obtained by repeating the k-space scheme whilst adjusting the slice selection gradient  $G_{SS}$ . The application of RF pulses and gradient magnetic fields to sample k-space may be described schematically using sequence design diagrams, the most simple of which is a gradient echo (GE) sequence. Schematics of the process of slice selection, acquisition of k-space data and an example GE sequence are provided in Figure 2.2.



**Figure 2.2:** Schematics of the use of gradient magnetic fields, acquiring data in kspace and a basic gradient echo (GE) sequence diagram. **1.** With an applied gradient magnetic field  $G_z$ , the Larmor frequency becomes dependent on spatial location. At the central point, where there is zero contribution from  $G_z$ ,  $\omega = \omega_0$  which is the carrier frequency, i.e. Larmor frequency. Caption continued on the following page.

Figure 2.2: 2. The process of applying a gradient magnetic field (in this example, in the z-axis  $G_z$ ) to extract a slice of a sample. A brain is placed in an external magnetic field  $\mathbf{B}_0$  (left). If we were to acquire signal now, there would be no spatial localisation of said signal. We would therefore have an indication of proton density for the whole sample with little useful information. By applying  $G_z$  and an RF pulse to select a slice of the brain (middle), we are able to spatially localise the signal and produce a two dimensional image (right). 3. The process of acquiring information about the spatial frequencies in k-space. The row  $k_y$  is selected using the phaseencoding gradient (indicated with the blue dashed line) and then we acquire data whilst moving along the  $k_x$  axis, using the frequency-encoding gradient (indicated with the green dashed line) (left). The process is repeated until the whole of k-space has been filled (middle). The Fourier transform of the k-space data gives a two dimensional image of the selected slice. 4. A schematic of the basic GE sequence with radio frequency (RF), slice select gradient  $G_{SS}$ , frequency-encoding gradient  $G_{FE}$  and phase-encoding gradient  $G_{PE}$  and data acquisition (Acq). For each row  $k_{y}$  of k-space, the sequence is repeated, as indicated by the orange arrow, and the amplitude of  $G_{PE}$  is adjusted (green arrow).

#### 2.1.3 Spin Echo and Echo Planar Imaging

The principles of the GE sequence may be extended to more complex sequences. The spin echo (SE) sequence (Figure 2.3 - panels 1 and 2) is a common alternative to the GE sequence. The basic GE sequence is adapted to include a second RF pulse with a flip angle of 180°. This second RF pulse flips the spin state of the protons, reversing the precession phase in the process. The protons now start to regain the phase they lost in the time between the 90° RF pulse and the second pulse, thus increasing the observable signal. These two sequences, and variations of them, form the basis for many MRI imaging protocols. Most of important of which, with respect to the work presented here, is the echo planar imaging (EPI) sequence (Mansfield, 1977; Mansfield, 2007). The EPI sequence is an extension of the simple k-space sampling scheme described in Figure 2.2. However, rather than acquiring a single line of k-space per sequence repetition, in EPI, we acquire many lines of k-space: typically the whole of the subspace being sampled is acquired per sequence repetition. Using EPI it is possible to acquire a series of slices (e.g. covering the entire brain) in a few hundred of milliseconds. This significantly reduces artefacts
caused by subject motion but comes at the cost of relatively poor signal-to-noise ratio (SNR) and increased sensitivity to artefacts caused by local changes in the sample's magnetic susceptibility. Figure 2.3 (panels 3 and 4) shows the k-space sampling and corresponding sequence diagram used in the GE EPI sequence.



Figure 2.3: Schematics of the spin echo (SE) and echo planar imaging (EPI) sequences. 1. The SE sequence diagram. The sequence is similar to the gradient echo (GE) sequence with the addition of the  $180^{\circ}$  RF pulse in combination with a second application of the  $G_{ss}$  magnetic field, restricting the RF excitation to the proton sample of interest. 2. The combination of the 90° and 180° RF pulses rephases the protons (indicated by the dotted green line), thus increasing the observable signal. Note however, that the signal does not return to its maximum amplitude due to T1 and T2<sup>\*</sup> decay. 3. The EPI sequence diagram. Typically, a relatively small flip angle is used,  $\alpha \leq 90^{\circ}$ . Initially, a slice is selected. The phase-encoding gradient is "blipped" on intermittently whilst the frequency-encoding gradient continuously flips from positive to negative. 4. The frequency- and phase-encoding gradients are initially used to move to a starting k-space location, as indicated by the orange dashed arrow. The frequency-encoding gradient moves along the  $k_x$  axis until reaching a pre-specified limit. At this point the phase-encoding gradient "blips" to move to the next  $k_u$  position. The now negative amplitude frequency-encoding gradient moves back along  $k_x$  and the process repeats until the  $k_x$ ,  $k_y$  plane is filled. The sequence is repeated to acquire many  $k_x$ ,  $k_y$  planes (as indicated by the series of k-space trajectories on the right), thus forming a three dimensional image.

## 2.1.4 Summary

These descriptions provide the basic underpinnings of MRI leading to the specific techniques used in this thesis, which is concerned with the extraction of brain connectivity. Most commonly in MRI, connectivity is probed using functional MRI or diffusion MRI data (both of which use EPI), representing the functional co-activation of brain regions and the structural links between those regions respectively. Typically, functional connectivity is derived by observing the co-activation of brain regions: regions that activate together, or their activation correlates, during a given cognitive task are said to belong to a functional network. Structural connectivity is derived by tracking the connections - axons - between brain regions. Although limited in their indirect natures, each method provides useful and unique information regarding how the brain works and how it is organised, allows us to probe the relationships between the brain's function and structure, and ultimately between brain organisation and behaviour. The following sections provide overviews of the basics of how dMRI and fMRI work, with a focus on dMRI, as this is the main MRI modality this thesis is concerned with.

## 2.2 Introduction to Diffusion MRI

## 2.2.1 The Basics of dMRI

Diffusion MRI (dMRI) is a magnetic resonance technique designed to be particularly susceptible to a molecular transport process known as diffusion, more specifically, to the diffusion of water molecules. Referred to as Brownian motion, and first observed by Robert Brown in 1827 (Brown, 1827) and later described in detail by Albert Einstein in 1905 (Einstein, 1905), the translational motion, or diffusion, of molecules (water molecules in this case) is random when in a free environment with the diffusion distance being zero on average, but the variance of the diffusion displacements being given by

$$\langle r^2 \rangle = 2D\Delta \tag{2.6}$$

where  $\langle r^2 \rangle$  is the mean-squared diffusion distance, D is the diffusion coefficient and  $\Delta$  is the diffusion time. In a barrier-free environment, this random motion is isotropic with equal variance of these displacements in all orientations. In a non-free environment, such as a biological tissue sample, molecular motion is dependent on its environment's structure. Within the brain, diffusion of water is influenced by the presence of various microstructures, such as membranes and myelin sheaths, parts of neuronal (Figure 2.4a) and/or glial cells, which pose barriers to free motion (Le Bihan, 2014).

These barriers can cause different diffusion scatter patterns within different tissue types, dependent on their overall microstructures. In cerebrospinal fluid (CSF)filled areas, the diffusion of water molecules is almost similar to that in a barrier-free medium, it is therefore isotropic. The grey matter (GM), forming the cortex around the outside of the brain and the subcortical structures, consists primarily of the cell bodies and axon terminals. Diffusion of water molecules in GM is hindered, but equally likely to be so in all directions due to the lack of systematic structure, resulting in low magnitude, roughly isotropic diffusion profiles. The white matter (WM), forming the majority of the volume of the brain, consists of the axonal fibres. Diffusion in WM is more free along the long axis of a bundle of axons compared to across their cross-section. In other words, the diffusion of water molecules along axons, and therefore in WM, is anisotropic as there is preference for diffusion towards certain orientations. The neuron, forming the basis of these tissues, and the diffusion of water molecules in these tissues is summarised in Figure 2.4b.



Figure 2.4: A schematic of the neuron and of the diffusion of water molecules in the brain tissues: cerebrospinal fluid (CSF), white matter (WM), and grey matter (GM). **a.** The neuron, the basic functional unit of the brain, consists of the cell body, axon and axon terminals, each supporting the building and transmission of action potentials. **b.** The diffusion of water molecules is dependent on the structure of the tissue. In WM, diffusion prefers the long axis of the axon, i.e. is anisotropic, whereas in GM and CSF, diffusion is isotropic with the magnitude modulated by the relative hindrance of the structure. Created with BioRender.com

This difference in the observed diffusion patterns for different underlying microstructure forms the basis of dMRI, which uses the measured diffusion pattern to indirectly inform us about the microstructure. The Stejskal-Tanner pulsed gradient spin echo (PGSE) MRI sequence (Figure 2.5) aims to measure signal changes due to water diffusion (Stejskal and Tanner, 1965). A  $90^{\circ} - 180^{\circ}$  spin-echo RF pair with two high and equal amplitude diffusion-sensitising gradients either side of the  $180^{\circ}$ RF pulse is used probe diffusion and introduce relevant contrast to the MRI signal.

The first diffusion-sensitiving gradient offsets the Larmor frequency  $\omega(t) =$ 



Figure 2.5: A schematic of the diffusion-weighting section of a Stejskal-Tanner pulsed gradient spin echo (PGSE) sequence diagram (not including spatial localisation gradients). The top line represents the applied radiofrequency fields, the middle line represents the diffusion-sensitising magnetic fields and the bottom line represents the observed signal during sampling.  $\delta$  is the duration of the gradient magnetic field (pulse width) in milliseconds; **G** is the gradient magnetic field (GMF) amplitude in mT/m and  $\Delta$  is the diffusion "mixing time". These factors control the b-value as shown in Equation 2.11.

 $\gamma \mathbf{G}(t)$ , where  $\mathbf{G}(t)$  is the gradient. For a given spin at position  $\mathbf{x}_0$ , this results in a phase shift according to Equation 2.7

$$\phi_1 = -\int_0^\infty \omega(t)dt = -\gamma \mathbf{x_0} \int_0^\infty \mathbf{G}(t)dt$$
(2.7)

The net phase shift considering each gradient pulse and the phase reversing effect of the  $-180^{\circ}$  RF pulse may be given by

$$\phi = \phi_2 - \phi_1 = -\mathbf{q} \cdot (\mathbf{x_1} - \mathbf{x_0}) \tag{2.8}$$

where  $\mathbf{x_1}$  is the position of the spin at the application of the second diffusionsensitising gradient and  $\mathbf{q} = \gamma \int_0^\infty \mathbf{G}(t) dt$  is the "diffusion weighting". In the case of a non-diffusing molecule/spin, or where  $\Delta \to 0$  ( $\Delta$  is the "mixing time"), there is no spin displacement (i.e.  $\mathbf{x_1} = \mathbf{x_0}$ ) and, therefore, the encoded phase will be reversed and nulled by the application of the  $-180^{\circ}$  RF pulse and the second diffusionsensitising gradient, i.e.  $\phi = -\mathbf{q} \cdot (\mathbf{x_1} - \mathbf{x_0}) = 0$ . In the case of diffusing spins, spin displacement during  $\Delta$  results in a net phase shift given by Equation 2.8.

The net phase shift is dependent on the orientation of diffusion relative to the orientation of the applied diffusion gradients, the average diffusion path length and the sequence parameters. If the diffusion-sensitising gradient's orientation matches the preferred diffusion orientation (e.g. along the axon) there will be maximal phase shift as the displacement of spins is maximal (i.e.  $\mathbf{x_1} - \mathbf{x_0} > 0$ ). If, for example, the gradient orientation is perpendicular to the preferred orientation of water diffusion, no diffusion-induced phase shift occurs. The signal loss caused by these phase shifts may be described by Equation 2.9.

$$E(\mathbf{q},\tau) = \int P(\mathbf{r},\tau) \cdot \exp(-i\mathbf{q}\cdot\mathbf{r})d\mathbf{r}$$
(2.9)

This equation is reached by considering the signal attenuation as a function of  $\mathbf{q}$ :  $E(\mathbf{q}) = S(\mathbf{q})/S(0)$ , where  $S(\mathbf{q})$  is the signal for a given  $\mathbf{q}$  and S(0) is the signal if no diffusion sensitising gradient were applied (i.e.  $\mathbf{q} = 0$ ). Through Fourier transform, this signal attenuation equation may then be related to the diffusion propagator  $P(\mathbf{r}, \tau)$ , which describes the probability of a spin having a displacement of  $\mathbf{r}$  due to diffusion over time  $\tau$  (Stejskal and Tanner, 1965). It should also be noted that in these descriptions, it is assumed that the application time of the diffusion sensitising gradients is negligible compared to the the mixing time (i.e.  $\delta \ll \Delta$ ). A full derivation of Equation 2.9 is available elsewhere (Kuchel et al., 2012).

## 2.2.2 Diffusion Modelling - The Diffusion Tensor

Using the PGSE sequence, the signal loss observed with the application of diffusionsensitising gradients can be measured. Considering a simple diffusion experiment in which molecules are free to diffuse and with a constant diffusion coefficient D, it is the case that the diffusion propagator P takes the form of a zero-mean Gaussian profile (according to Einstein's equation). The signal attenuation (Equation 2.9) may then be simplified and, by substituting this simplified expression into the signal attenuation as a function of **q**, Equation 2.10 is obtained (Stejskal and Tanner, 1965).

$$S(\mathbf{q},\tau) = S(0) \cdot \exp(-q^2 D\tau) \tag{2.10}$$

This is typically further simplified and re-written to take the form of Equation 2.11 where S(b) is the measured signal and b is the "b-value" which describes the diffusion weighting introduced by the experimental parameters relating to the diffusion sensitising gradients in units of s/mm<sup>2</sup> (Stejskal and Tanner, 1965). b = 0 indicates no diffusion weighting.

$$S(b) = S(0) \cdot \exp(-bD), \quad where \quad b = \gamma^2 G^2 \delta^2 \left(\Delta - \frac{\delta}{3}\right). \tag{2.11}$$

First applications of the PGSE experiment and the calculation of voxel-wise diffusion coefficients in the human brain were reported in 1986 (Le Bihan et al., 1986). The term "apparent diffusion coefficient" (ADC) was introduced; this accounts for the fact that our measurements of the self-diffusion coefficient of water are perturbed by the presence of microstructures in the tissue which hinder and restrict diffusion. The ADC is not only dependent on the tissue environment but also on the acquisition parameters, namely the diffusion time  $\Delta$ .

From Equation 2.11, it is clear that by varying the sequence parameters one may obtain a series of images with varying diffusion weighting. By repeating the PGSE experiment and iteratively adjusting the orientation of the diffusion-sensitising gradients, it is possible to build a picture of the signal loss related to the directionality of diffusion in three dimensions. A series of diffusion-weighted MR images are exampled in Figure 2.6, highlighting the gradient direction dependence and the use of b-values to enhance the effect of signal loss due to diffusion. This enhanced signal loss associated with higher b-values improves the angular contrast in the estimation of diffusion orientation but suffers from poor SNR.



**Figure 2.6:** Examples of the directionally and b-value dependent signal loss in diffusion weighted imaging (DWI). The first image represents the anatomy with no diffusion-weighting, i.e. b = 0. A series of b = 2000 images with varying gradient directions, as indicated by the vec(x, y, z) values and the white arrow, demonstrate how signal loss is gradient direction dependent. From these examples, it is possible to begin to understand how certain features may be extracted from this information. For example, the splenium of the corpus callosum (indicated by the yellow arrow) becomes more or less visible depending on the gradient direction. The final two images (bottom-right) provide examples of how the overall signal changes with respect to the selected b-value. Lower b-values (b = 1000) provide less diffusion weighting but offer greater signal-to-noise ratio. The higher b-value b = 3000 offers highly diffusion-weighted information but suffers from a relatively low of signal-to-noise ratio.

#### The Diffusion Tensor

The concept of a scalar (isotropic) diffusion coefficient (e.g. the ADC) was extended to incorporate multiple orientations (therefore diffusion anisotropy) and a new diffusion-based imaging method was introduced (Basser et al., 1994): diffusion tensor imaging (DTI). DTI is a method for estimating the diffusion tensor  $\mathbf{D}$ , which describes the average water displacement profile, and therefore WM tract orientation, in three dimensional space. In practical terms, the diffusion tensor is an array of diffusion coefficients, each representing the diffusion in a specified orientation. Typically, the diffusion tensor is a symmetric  $3 \times 3$  array, with six unique elements, represented in a Cartesian (x, y, z) coordinate system:

$$\mathbf{D} = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{bmatrix}$$
(2.12)

where  $D_{x,y,z}$  are the diffusion coefficients in the (x, y, z) orientations respectively. The tensor is derived by solving a modified version of Equation 2.11 whilst considering sufficient signal measurements at various gradient directions. It is typically estimated by log-transforming the observed signal attenuation and then performing linear regression (Basser et al., 1994). As we need only consider the six unique elements of the symmetric tensor, through log-transformation, the measured signal is related to the tensor elements as  $\ln(\frac{S(b)}{S(0)}) = -\sum_{i=1}^{3} \sum_{j=1}^{3} b_{ij} D_{ij} =$  $-(b_{xx}D_{xx} + 2b_{xy}D_{xy} + 2b_{xz}D_{xz} + b_{yy}D_{yy} + 2b_{yz}D_{yz} + b_{zz}D_{zz})$ . By measuring S(0)and S(b) and defining each b element, we may calculate each element of the tensor as the gradient of the slope in a plot between  $\ln(\frac{S(b)}{S(0)})$  and each b-value: i.e. we may perform linear regression. In order to calculate each unique tensor element, we therefore require at least six unique measurements of S(b), a measurement per tensor direction (xx, xy, xz, yy, yz, zz), plus a measurement of the signal without any diffusion-induced attenuation S(0).

Through eigen-decomposition of the tensor, it is possible to derive descriptive terms such as principal diffusivities, corresponding to the eigenvalues  $\lambda$ , and principal diffusion directions, corresponding to the eigenvectors  $\mathbf{v}$  (Basser et al., 1994). These values may be used to simply describe the local diffusion profile, relative to the anatomy, and provide insight into microstructure. The eigenvalues and eigenvectors may be used to derive further descriptive terms. For example, the translational displacement profile and scalar invariant measures of anisotropy: the trace of the diffusion tensor (closely related to the mean diffusivity (MD), Equation 2.13), the fractional anisotropy (FA) (a measure of the degree of diffusion anisotropy, defined in Equation 2.14) and a measure of the principal diffusivities distribution relative (or skewness) to their mean (Basser, 1997b).

$$MD = \frac{1}{3}(\lambda_1 + \lambda_2 + \lambda_3) \tag{2.13}$$

where  $\lambda_{1,2,3}$  are the tensor eigenvalues.

$$FA = \sqrt{\frac{(\lambda_1 - \lambda_2)^2 + (\lambda_2 - \lambda_3)^2 + (\lambda_1 - \lambda_3)^2}{2(\lambda_1^2 + \lambda_2^2 + \lambda_3^2)}}$$
(2.14)

These properties are calculated at the voxel level, giving maps of the microstructural features derived from diffusion imaging. Examples of typical FA, MD and the first eigenvector, describing the principle diffusion direction, are presented in Figure 2.7.

The tensor may also be visually depicted as a three-dimensional ellipsoid, the "diffusion ellipsoid". Typically, diffusion ellipsoids are constructed from the tensors on a voxel-by-voxel basis such that the polar axis of the ellipsoid represents the fibre orientation (Basser et al., 1994), as exampled in Figure 2.8.

## 2.2.3 Beyond The Diffusion Tensor

Although the diffusion tensor is a useful model, it is limited in its ability to represent specific features of the true underlying anatomy, in particular with respect to multiple fibre populations and complex fibre architecture. In DTI, each voxel is represented by a single tensor, i.e. a single Gaussian profile. Given the scale



**Figure 2.7:** Examples of the microstructural maps derived from the diffusion tensor, corresponding to the fractional anisotropy (FA), mean diffusivity (MD) and the first eigenvector (v1). v1 brightness is modulated by the FA map, highlighting regions of high FA, i.e. areas corresponding to fibre populations.



Figure 2.8: Examples of the voxel-wise tensor description of diffusion colour-coded in the standard diffusion MRI community way such that left-right directions are red, anterior-posterior directions are green and superior-inferior directions are blue. Three regions are highlighted demonstrating the tensor geometry for the primary brain tissues: highly directional, prolate tensors in the corpus callosum genu (i.e. in WM), low-magnitude omnidirectional tensors in GM, and omnidirectional, spherical tensors in the cerebrospinal fluid in the ventricles.

of neurons and axons (micrometer range) relative to typical MRI voxel sizes (millimeter range), this means that a single diffusion tensor represents the average over tens to hundreds of thousands of cells/axons. In highly heterogeneous regions where large bundles of parallel WM fibres are found, such as the midline of the corpus callosum, this averaging provides a useful, although limited, approximation of diffusion/fibre direction. In reality, the contents of a single voxel is complex: neurons are not perfectly packed and, therefore, "free"-space is present, voxels may average over tissue borders (e.g. between WM and CSF) and multiple fibre populations may be present, giving rise to crossing, merging, fanning and kissing fibres. Figure 2.9 examples these complications schematically.



**Figure 2.9:** Schematic of fibre structures containing two fibre bundles. In a small example region of the brain, crossing, kissing and merging/fanning fibres are observed. Examples depicting how the observed diffusion direction may comprise of multiple compartments are shown. Created with BioRender.com

These limitations prompted the development of more sophisticated diffusion modelling techniques which seek to model multiple types of diffusion within a voxel, relating the presence of differing diffusion profiles (e.g. isotropic and anisotropic) and multiple fibre bundles. Many diffusion modelling approaches exist and a full review is out of the scope of this thesis (see Alexander et al., 2017). One model of particular relevance in this work, the "ball and stick" model, was first introduced in Behrens et al. (2003b) and furthered by Behrens et al. (2007), Jbabdi et al. (2012), and Sotiropoulos et al. (2016).

The "ball and stick" model is a multi-compartment model which seeks to describe intravoxel diffusion whilst considering the consequences of partial-volume effects. In the most simple form, it consists of two compartments which separately model diffusion in the fibre orientation (e.g. diffusion within axons) and the diffusion of free water (e.g. in extracellular space or due to partial volume with GM or CSF). The anisotropic diffusion associated with the fibre bundle is defined to be perfectly anisotropic, and is therefore represented as a delta-like function, or "stick". The isotropic diffusion compartment is represented by a perfect sphere, or "ball". The ability to represent multiple fibre populations may be achieved by considering more "stick" components in the model, representing each population as a separate stick. The signal  $S(\mathbf{g})$  is then given as a function of N (the number of fibre populations being modelled) compartments (Behrens et al., 2003b; Behrens et al., 2007):

$$S(\mathbf{g}) = S(0) \left[ \left( 1 - \sum_{i=1}^{N} f_i \right) \exp(-bd) + \sum_{i=1}^{N} f_i \exp\left(-bd\left(\mathbf{g} \cdot \mathbf{v}_i\right)^2\right) \right]$$
(2.15)

where  $f_i$  is the volume fraction for the given compartment, d is the diffusivity,  $\mathbf{v}_i$  is the orientation of the *i*-th fibre population and  $\mathbf{g}$  and b define the sequence parameters: b is the b-value as before and  $\mathbf{g}$  is the gradient direction. The ball compartments, with their isotropic diffusion profiles, are described by the term in Equation 2.15 highlighted in red and the stick compartments, with their signal component dependent on the gradient direction relative to the fibre orientation, by the term highlighted in blue.

An issue now presents in determining how many intravoxel fibre orientations there are and, therefore, how many compartments should be included in the model, i.e. determining how complex the model should be. Automatic relevance determination (ARD) (Mackay, 1995) is a stochastic model selection technique which may be employed to select a data-driven number of compartments. In cases where the most simple model (e.g. a single compartment/stick) is strongly supported by the data, further complexity is not required and therefore higher compartments are forced to zero. If the data supports additional intravoxel fibre orientations, model complexity is allowed to increase. This model selection approach proved effective in determining the number of fibres in the underlying data, particularly benefiting regions of crossing fibres, aiding in the estimation of global connectivity (Behrens et al., 2007).

The ball and stick model increases the complexity of the solution compared to the diffusion tensor: we now have three unknown variables (the volume fractions  $f_i$ , the diffusivity d and the fibre orientations  $\mathbf{v}_i$ ) for each modelled compartment. Further, we wish to not only estimate the most likely solution, but to also estimate the uncertainty associated with these parameters. Thus, instead of solving our model analytically, a Bayesian framework is employed to estimate the parameters and to provide measures of uncertainty. The basic principle is to guess each parameter based on prior distributions, model the signal, and compare the modelled signal to the measured signal. If the modelled signal fits the measured signal, the parameter choices are accepted. In fitting the signal in the ball and stick model, a mathematical framework based on Markov Chain Monte Carlo (MCMC) sampling is used to estimate the parameters by preferentially sampling prior distributions in areas of high probability. Through this sampling scheme, many acceptable solutions (i.e. combinations of our unknown parameters) are found and each model parameter is described as a probability density function (PDF), capturing each of the acceptable solutions, which represents the uncertainty in dMRI data (Behrens et al., 2003b). Modelling uncertainty in this way becomes particularly useful in the use of tractography to define global connectivity later.

The ball and sticks model was furthered to take advantage of multi-shell, i.e.

multiple b-values, dMRI data in Jbabdi et al. (2012). As previously mentioned, higher b-values allow for higher angular diffusion contrast, which is particularly beneficial for crossing fibre estimation. However, an additional complexity arises. For high b-values (> 1500 s/mm<sup>2</sup>) the signal no longer exhibits mono-exponential decay as we become sensitive to, and therefore sample, the exchange of intracellular and extracellular pools of water. This sampling of mixed pools results in a more complex decay profile: a non-mono-exponential decay. Therefore, the previous mono-exponential approach to estimating diffusion (i.e. the  $\exp(-bd)$  and  $\exp(-bd(\mathbf{g}\cdot\mathbf{v}_i)^2)$  components of Equation 2.15 which describe a single diffusion coefficient d) is no longer valid and its application to data acquired with higher b-values may lead to overfitting, resulting in more fibre populations being estimated than expected. Jbabdi et al. (2012) proposes a solution to this issue, introducing a continuous distribution of diffusion coefficients to Equation 2.15 in the form of a gamma distribution with a shape parameter  $\alpha$  and scale  $\frac{1}{\beta}$  (which describe the mean and variance of the distribution of the diffusion coefficients):

$$S(b_k, \mathbf{g}) = S(0) \left[ \left( 1 - \sum_{i=1}^N f_i \right) \left( \frac{\beta}{\beta + b_k} \right)^{\alpha} + \sum_{i=1}^N f_i \left( \frac{\beta}{\beta + b_k (\mathbf{g} \cdot \mathbf{v}_i)^2)} \right)^{\alpha} \right]$$
(2.16)

This model now considers multiple fibre populations whilst accounting for effects of non-mono-exponential signal decay due to higher b-values, improving our ability to utilise multi-shell diffusion data which improves angular resolution and, therefore, crossing fibre estimation.

Returning to the issue of model selection, ARD may be employed once again to determine the number of compartments to be used. In the case of Equation 2.16, a shrinkage prior is placed on the standard deviation  $(\sqrt{\frac{\alpha}{\beta}})$  of the diffusivities which forces the distribution variance to 0 if the data does not support the model and reduces model complexity towards a single-compartment. In practice, a maximum

complexity is defined based on prior expectations (typically a maximum of three compartments).

Using this non-mono-exponential model with ARD, Jbabdi et al. (2012) reports improved performance compared to the mono-exponential form. Specifically, overfitting (i.e. the inclusion of too many compartments) is significantly reduced, particularly in the presence of partial volume effects. Although limited in its findings, Leuze et al. (2021) provides some histiological validation of this model in the estimation of fibre orientations. This model is the one used in this thesis for estimating fibre orientations and is a prerequisite of probabilistic tractography as performed in this thesis and described in the following section.

## 2.2.4 Tractography

By applying field vector calculus concepts to fibre orientations extracted from dMRI datasets it is possible to computationally produce reconstructions of fibre trajectories (i.e. to perform tractography) (Basser, 1997a). Fibre trajectories,  $\mathbf{r}(s)$ , can be defined as a curve with arc length s tangent to the vector field of voxel-wise fibre orientations,  $v_1$ :

$$\frac{d\mathbf{r}(s)}{ds} = v_1(\mathbf{r}(s)) \tag{2.17}$$

This may be solved numerically to estimate the progress of the track, i.e. to estimate  $\mathbf{r}(s_i)$ , in the most simple case, employing Euler's method (Basser et al., 2000):

$$\mathbf{r}(s_{i+1}) = \mathbf{r}(s_i) + \alpha_{step} v_1(\mathbf{r}(s_i))$$
(2.18)

where  $\alpha_{step}$  is the step length and  $\mathbf{r}(s_i)$  is the position of the track at the *i*-th step with  $\mathbf{r}(s_0)$  being the starting/seed point (Conturo et al., 1999; Mori et al., 1999).

Tracking commences from a seed point and the direction is taken to be the direction of least-hindered diffusivity, representing the major fibre orientation. The track (or "streamline") follows that direction until reaching the next position (as defined by the chosen step length) where it then follows an interpolated orientation. These interpolated orientations can be obtained using nearest-neighbour interpolation (e.g. FACT - Jones et al. (1999)) to more advanced (e.g. trilinear) interpolation approaches, reducing interpolation errors (Conturo et al., 1999). This process is repeated until an end-criterion is reached, for example, reaching a pre-specified FA threshold or a pre-defined 'stop' region. Thus, the streamline is defined as a series of positions  $\mathbf{r}(s)$  with associated orientations.

Mori et al. (1999) demonstrated this approach by estimating known fibre bundles in an ex vivo rat brain, including the corpus callosum, anterior commissure, fornix and more. These methods later became known as tractography, and, the addition of anatomical priors allowed for "in vivo interactive dissection" (Catani et al., 2002; Mori et al., 2005). The general principle of tractography (seeding a track from a given region, following the principal direction across voxels, and terminating tracking at a given criteria) is depicted in Figure 2.10.

An inherent limitation of these first tractography methods is the inability to deal with the noisy nature of diffusion data: as discussed, the echo planar imaging (EPI) sequence used in the collection of dMRI data suffers from reduced SNR compared to conventional MRI sequences. Tractography algorithms are quick to produce false positives when small changes due to noise lead to the fibre orientation being misrepresented (Basser et al., 2000; Jones, 2003). Methods discussed so far may be referred to as deterministic tractography where a single curve is reconstructed per starting point, effectively representing point-estimates of the underlying pathway.



**Figure 2.10:** A schematic of the principles of tractography. Tractography is seeded from a starting voxel and the local fibre orientations (depicted by the coloured lines and representing the first eigenvector of the diffusion tensor) are used to guide a streamline, or pathway. Tractography is terminated once the streamline reaches pre-defined criteria, which, in this case, is the "end"/termination voxel.

#### Probabilistic Tractography

An alternative to deterministic estimation of tracts is a statistical estimation, where the spatial distribution of tracts arising from a seed is estimated given uncertainties in the data and the models employed to analyse them. Probabilistic tractography approaches offer such a route. By repeating the tracking algorithm from the same starting point and stochastically sampling the modelled fibre orientation distribution, it is possible to build a spatial distribution of streamlines, representing the uncertainty in estimation, e.g. the probability with which a reconstructed fibre pathway passes through a given voxel. Whereas previously, in deterministic tractography, results were essentially a maximum likelihood pathway with no measure of uncertainty, repeated estimates allow for a greater understanding of the confidence in the results (Behrens et al., 2003b).

Probabilistic tractography was introduced, in part, in an attempt to overcome the poor performance of tractography in brain regions which typically have poorly defined fibre orientation due to, for example, low FA (Behrens et al., 2003b). Such regions are poorly represented by conventional deterministic methods, as these methods stop propagating into regions with high uncertainty in their fibre orientation. Using the ball and stick model as describe above we can build probability density functions (PDFs) which represent the uncertainty in the dMRI data. By repeatedly sampling the PDFs and seeding streamlines with each sample, the tracking process is repeated with slightly different starting parameters, representing the uncertainty in the data. That way, propagation can continue even in regions with high uncertainty, which, however, is reflected in the final results as reduced confidence in going through certain regions. Figure 2.11 provides a schematic comparison of deterministic and probabilistic tractography.



**Figure 2.11:** A simple schematic of the differences between deterministic (left) and probabilistic (middle) tractography. Four seed points are used, indicated by the circles in the anterior-right portion of the frontal lobe. In each case, streamlines are commenced from their respective seeds and follow the primary direction of diffusion, indicated by a color coded map of the principal eigenvector. In the deterministic case, there is no preferred pathway whereas in the probabilistic case the genu of the corpus callosum is preferred. Right: An example of the spatial distribution of the reconstructed streamlines from probabilistic tractography. Brighter regions indicate greater confidence in the presence of a fibre bundle.

#### The Utility of Tractography

Although the voxel-by-voxel values reported in probabilistic tractography are not necessarily an absolute measure of the probability that a given fibre passes through a voxel, it is a useful metric when interpreting tractography results as it allows us to assign relative confidence to the connectivity between two regions via the reconstructed curve (Behrens and Johansen-Berg, 2005; Jbabdi and Johansen-Berg, 2011). As such, tractography has been applied widely to investigating brain circuitry and connectivity. The application of tractography may be divided into three main applications depending on the primary aim of the study: the segmentation of GM structures, the construction of connectivity matrices/networks, and the segmentation of WM.

Probably the most common application of tractography, and one of primary interests in this thesis, is to perform "virtual dissection" of WM fibre bundles (Catani et al., 2002; Mori et al., 2005). Here, the focus is on the main body of the WM fibre bundles, i.e. how the axons traverse through the deep WM. This application has been utilised widely in diseased cohorts (Johansen-Berg and Behrens, 2006; Yamada et al., 2009; Lipp et al., 2020), for surgical planning (Johansen-Berg and Behrens, 2006; Berman, 2009; Panesar et al., 2019; Ashmore et al., 2020), to better understand cross-species differences in brain connectivity (as in this thesis) (Thiebaut de Schotten et al., 2012; Hecht et al., 2013; Mars et al., 2018c; Eichert et al., 2020), and in healthy cohorts to better understand the relationships between connectivity and behaviour (as in this thesis) (Thiebaut de Schotten et al., 2011b; Thiebaut de Schotten et al., 2011c; Howells et al., 2018; Madden et al., 2012; Wang et al., 2018c; Reginold et al., 2016).

Tractography may be used to define connectivity profiles between GM structures and to define functional boundaries within GM by exploring changes in these connectivity profiles. In other words, tractography may be used to perform brain segmentation/parcellation. Such studies may be focused on specific brain regions or cover the whole-brain GM. The segmentation of subcortical structures is exampled in early tractography work where global connectivity maps were applied to estimate thalamo-cortical connectivity allowing for connectivity based segmentation of the thalamic nuclei (Behrens et al., 2003a). This segmentation was later reproduced and extended following similar approaches (Traynor et al., 2010; Plantinga et al., 2018; Patriat et al., 2018; Akram et al., 2018) and by using functional and anatomical approaches (Johansen-Berg et al., 2005), demonstrating that subcortical segmentations based on structural connectivity can be robust and align across modalities. Changes in structural connectivity profiles may also be used to identify cortical boundaries (Johansen-Berg et al., 2004; Beckmann et al., 2009; Mars et al., 2012; Thiebaut de Schotten et al., 2014b). In these examples, the structural connectivity between a given GM region and the rest of the brain (or some other region) is estimated. Spatial gradients in this connectivity matrix are then used to define sub-divisions of that region which are often reported to closely relate to functionally specialised regions.

Connectivity matrices/networks may be built at the regional or whole-brain level by seeding from each constituent voxel and, for example, counting the number of streamlines passing through/to the rest of the region of interest. This allows for the construction of whole-brain connectivity matrices, or connectomes, for example (Hagmann et al., 2007; Tymofiyeva et al., 2014; Sotiropoulos and Zalesky, 2017). Further, network representation allows for the derivation of simple network connectivity metrics, such as node strength and community structure which may relate to functional connectivity, measures of cognition and disease (Bassett and Bullmore, 2009; Rubinov and Sporns, 2010; Heuvel and Sporns, 2013; Wang et al., 2018b; Bohlken et al., 2014; Yeh et al., 2020).

## 2.3 Introduction to Resting-State Functional MRI

## 2.3.1 The Basics of fMRI

While diffusion MRI probes the structure and microstructure of brain tissue, functional MRI (fMRI) probes neuronal activity. fMRI is sensitive to the inherent magnetic properties of blood flow during physiological processes. More specifically, fMRI is sensitive to the variation of blood-oxygenation levels, accompany neuronal firing, and can therefore be used as an indirect proxy for brain activity. Oxyhaemoglobin is diamagnetic, whereas deoxyhaemoglobin is paramagnetic (Pauling, 1977). As is well documented, the magnetic properties of media may heavily affect the signal obtained during MR image acquisition (Morelli et al., 2011). Diamagnetic materials have little consequence on the final signal. However, ferromagnetic and, to a lesser extent, paramagnetic materials have a significant effect on the final signal. These effects are commonly referred to as the BOLD (blood oxygen level dependent) effect and may be capitalised on to probe neuronal activation (Ogawa et al., 1992).

During neuronal activation, there is an increased demand for oxygen. Neuronal oxygen uptake increases and oxyhaemoglobin levels in the passing blood decrease. This results in a greater concentration of deoxyhaemoglobin in the blood vessel exiting the activated region and, therefore, increased paramagnetism. This leads to increased local magnetic field distortions and reduced MR signal. In response to the reduced oxyhaemogloin concentration and a remaining demand for oxygen, heart rate increases and, more importantly, vasodilation occurs, leading to increased cerebral blood flow (CBF) and cerebral blood volume (CBV). Increases in CBF and CBV are much greater than the increase in oxygen demand, resulting in an "overshoot": much more oxyhaemoglobin is delivered than is required. These neuronal activation-induced vascular changes are termed the haemodynamic response and may be mathematically described by the haemodynamic response function (HRF) (Buxton et al., 2004).

The excess in oxyhaemoglobin alters the blood oxygen concentration, leaving a smaller fraction of deoxyhaemoglobin. This concentration shift reduces the overall paramagnetic nature of blood, reduces the local magnetic field distortions (and hence the local T2 and T2\* relaxation times) and increases local MR signal. A schematic representation is presented in Figure 2.12. These principles form the basis of fMRI.



Figure 2.12: A schematic representation of the variation in oxy/deoxy-haemoglobin during neuronal activation and the canonical signal response function (or HRF). **a**. Top: increased oxygen demand results in a reduced oxyhaemoglobin concentration in veins beyond the neurons, reducing the MR signal. Bottom: following the haemodynamic response, the concentration of oxyhaemoglobin in the veins increases due to an overshoot effect, increasing the MR signal. **b**. An example of a haemodynamic response function (HRF) showing typical signal changes with time. The stimulus is presented and is followed by a short lag, an initial signal dip and then a the signal increase due to the overshoot effect. Created in part with BioRender.com

By rapidly, and repeatedly, measuring the fMRI signal using EPI techniques (Mansfield, 1977; Mansfield, 2007), it is possible to observe signal variance related to neuronal activity (Biswal et al., 1995) with relatively high temporal resolution. From such a series of images, time courses of the signal fluctuations may be derived which may then be used to identify regions of (co-)activation.

## 2.3.2 Estimating Functional Connectivity Using fMRI

The BOLD effect may be applied in the case of task-induced neuronal activation to extract information regarding the localisation of activation and how it relates to an external stimulus or task, thereby revealing regional brain specialisation. For example, the regional specialisation of the production and/or comprehension of language may be probed by measuring fMRI signal fluctuations during stimulus paradigms (Agarwal et al., 2019). Such studies have allowed for the non-invasive localisation of a range of functions, including memory (Henson, 2005; Rajah and D'Esposito, 2005), vision (Kashou, 2012; DeYoe and Raut, 2014), object recognition (Grill-Spector, 2009), motor control (Toma and Nakai, 2002; Biswal et al., 1995) and much more (Yarkoni et al., 2011).

These methods are referred to as task-fMRI and are not the focus of this thesis. However, similar image acquisition and analysis methods may be applied during task-free periods to observe the constant low frequency fluctuations of neuronal activity while the brain is at "rest" (Biswal et al., 1995). Here, rest refers to the subject remaining inactive, i.e. not actively engaged in any specific task. Whilst at rest, spontaneous co-activation between functionally related brain regions occurs (Smith et al., 2013b). This method is known as resting-state fMRI (RS-fMRI).

By comparing (e.g. correlating) voxel-wise, or regional, time courses of signal fluctuations, these functionally related brain regions may be identified. The "functional" connectivity may be estimated by measures of synchrony between pairs of fMRI time-courses, providing another indirect measure of "brain connections", complementary to those derived from dMRI. At a whole-brain level, this can be represented as functional networks or as matrices (functional connectivity matrix), comprised of nodes and edges where nodes represent each voxel, or region, and the edges represent the correlation, or connectivity, between two nodes. In the case of RS-fMRI, these networks are referred to resting state networks (RSNs): the most common example, i.e. that which is most readily identifiable and reported/used extensively in the literature, of which is the default mode network (DMN), first described using positron-emission tomography (PET) (Raichle et al., 2001) and later using fMRI (Greicius et al., 2003). The DMN, and several other RSNs, have since become well-established characteristics of brain function and organisation (Cole et al., 2010; Lee et al., 2013).

#### **Estimating Co-activation**

The most typical measures of co-activation for defining functional connectivity are correlational measures, e.g. Pearson's correlation coefficient or partial correlation (Smith et al., 2013a). Partial correlation aims to represent direct connections more accurately by considering the association of two nodes after first regressing out the recorded activity of all other nodes. For example, in a three-node system, the partial correlation between two nodes is estimated by first regressing out the third node, thus reducing indirect correlations, and therefore indirect connections. Seedbased fMRI analysis, which includes the pre-selection of a region of interest (ROI), or "seed", may be used to estimate the functional connectivity between the given seed and the rest of the brain (Biswal et al., 1995). Since its introduction, seedbased analysis has been widely implemented as it allows for relatively quick analysis. However, the requirement of a priori assumptions about the data and the carefully considered selection of seeds limits its application (Lee et al., 2013).

More recently, data-driven methods for extracting functional co-activation maps have proved beneficial. In fMRI, we measure signal fluctuations at the voxel level and seek to understand how said signals relate to the task at hand. However, these signals are composed of multiple sources of variation, including task-related activation, physiological signal changes and artefactual signal changes. Therefore, we wish to separate unwanted signal variation (e.g. from head motion) from signal variation of interest (that due to neuronal activation). Independent component analysis (ICA) is an exploratory multivariate analysis method which is capable of estimating multiple signal sources (independent components) in a mixed-signal dataset by maximising statistical independence between said components (Jutten and Herault, 1991; Comon, 1994). The mixed-signal dataset  $\mathbf{X}$ , containing the measured fMRI signal, is decomposed into source signals (i.e. independent components)  $\mathbf{S}$ , representing various sources of signal fluctuations, through the estimation of mixing coefficients  $\mathbf{W}$  such that  $\mathbf{X} = \mathbf{WS}$ .  $\mathbf{W}$  is estimated whilst seeking to maximise the statistical independence of the estimated components, under the assumption that the measured signal is a linear combination of the source signals. It may be considered as seeking to find the primary axes of a dataset, i.e. low-dimensional projections of the data, which, in the case of ICA are defined such that each projection is statistically independent.

ICA was first applied to fMRI with the aim of separating task-fMRI data into independent components (McKeown et al., 1998). These components included taskrelated activation, physiological signal changes and artefactual signal changes. Component separation allows for the exclusion of unwanted signal (e.g. head motion) and therefore improves interprebility (McKeown et al., 1998). ICA may also be applied to investigating cerebral vasculator related signal changes (Kiviniemi et al., 2003), demonstrating the capability of ICA to separate resting-state neuronal activity from other sources. ICA is limited by the need for user-selection of the independent components. As mentioned, ICA provides an output of independent components, which could be interesting signal changes due to neuronal activation or uninteresting signal changes due to motion, physiological noise or MRI artefacts, but it does not specify which components are which. Aspects of this process may be automated but require sufficient data quality and quantity (Salimi-Khorshidi et al., 2014; Griffanti et al., 2014; Pruim et al., 2015).

## 2.4 Summary and Thesis Context

This chapter introduces the basic concepts of MRI and the two key MRI techniques used in this thesis: diffusion MRI and functional MRI. As discussed, both dMRI and fMRI are unique tools for probing brain connectivity in vivo and non-invasively and they provide complementary information. fMRI allows us to estimate functional activation in the brain relating to stimuli or at when at rest, and subsequently, infer regions of co-activation and functional networks. dMRI allows us to estimate the anatomical connections through white matter between brain regions. Combined, these modalities provide unique insight into how the brain is organised.

Both modalities are used in this thesis, although dMRI is the focus. Chapter 4 capitalises on the developments in diffusion MRI and tractography to introduce a new, cross-species, automated and standardised tractography toolkit. Chapter 5 builds on this to explore how measures of structural connectivity carry wholebrain architectural information that can be used to project functionally-relevant information across diverse brains. Finally, Chapter 6 uses multi-modal connectivity measures to explore brain-behaviour associations and, specifically, to explore the stability of typically used approaches in the field. The following chapter introduces large neuroimaging efforts which provide the data used in this thesis. They acquire both dMRI and fMRI data for large cohorts of subjects and allow for the estimation of brain connectivity, either using high-quality data or data at the population-level.

## Chapter 3

# Big Data and Large-Scale Neuroimaging Projects

## Contents

3.1	Introduction				
<b>3.2</b>	The Human Connectome Project				
	3.2.1 Cohort Overview	1			
	3.2.2 Non-Imaging Data	2			
	3.2.3 Imaging Data	4			
	3.2.4 The HCP Minimal Preprocessing Pipeline (MPP) 5	6			
<b>3.3</b>	The UK Biobank	9			
	3.3.1 Cohort Overview	9			
	3.3.2 Non-Imaging Data	0			
	3.3.3 Imaging Data	1			
	3.3.4 Preprocessing Pipeline	2			
<b>3.4</b>	HCP Structural Connectivity Features and Atlases 6	4			
	3.4.1 White Matter Microstructural Atlases 6	5			
	3.4.2 Structural Connectomes	7			
<b>3.5</b>	Summary and Thesis Context	0			

## 3.1 Introduction

Significant advancements in the field of neuroscience have been made in recent decades, driven in part by MRI studies. As in other modern sciences (Sagiroglu and Sinanc, 2013; Hansen et al., 2014; Leonelli, 2019; Pal et al., 2020), neuroscience and neuroimaging have entered the era of "big data". This transition is necessary if, one day, we are to fully understand the intricate detail of the brain and how it varies across the population and gives rise to cognition. "Big data" is a broad term relating generally to the richness of data available. Specific to neuroimaging, big data encapsulates boundary pushing resolution (Edlow et al., 2019; Wang et al., 2020a), both temporal and spatial, and number of samples, for example, subjects (Jack et al., 2008; Thompson et al., 2014; Miller et al., 2016; Casey et al., 2018) or brain coverage. A prime example of big data in neuroscience is in the development of the most complete brain connectome to date: that of the adult Drosophila central brain, which represents 25,000 neurons and 20 million synapses (Scheffer et al., 2020). We are not yet at the stage of such mapping in the human brain. However, several groundbreaking neuroimaging efforts have focused on acquiring high-quality brain MRI data from large cohorts of subjects. This thesis capitalises on data from two such efforts: the "WU-Minn" Human Connectome Project (HCP) (Van Essen et al., 2012; Van Essen et al., 2013) and the UK Biobank imaging study (Miller et al., 2016; Alfaro-Almagro et al., 2018; Sudlow et al., 2015). The following sections provide overviews of these cohorts, highlighting the diffusion MRI, functional MRI and behavioural data available. We also describe our early efforts in high-throughput analysis of such datasets that led to some diffusion MRI microstructure atlases, now publicly released in FMRIB's Software Library (FSL).

## 3.2 The Human Connectome Project

The "WU-Minn" Human Connectome Project (HCP)<sup>1</sup> was a ground-breaking research project with the aim of pushing image acquisition and processing technology to probe the human connectome in unprecedented ways. The collected datasets comprise of high-quality anatomical, functional and diffusion MRI data for a relatively large (1,200 subjects) family-structured (twins and non-twin siblings) healthy cohort. In addition to MRI data, the HCP also provides a broad range of behavioural and demographic data, genetic data and, for a sub-cohort, magnetoencephalography (MEG) data. MRI data have been quality controlled, preprocessed (as discussed in the following sections) and all MRI, behavioural and genetic data have been made publicly available. The HCP data forms the basis of this research. This section summarises the data available via the HCP and describes briefly the HCP Minimal Preprocessing Pipeline (MPP).

## 3.2.1 Cohort Overview

In total, the HCP recruited 1,206 subjects aged 21-35 years, including males and females and a range of ethnic backgrounds (to reflect the ethnic diversity of the United States population). The HCP subject cohort is based on a family and twin structure - families with at least four siblings or two twins (both monozygotic and dizygotic twins) within the age range were eligible to be recruited. The purpose of this is to boost power with regard to genetic associations, to aid in analysing heritability and to distinguish between genetic and environmental influences (Van Essen et al., 2013). In total, 457 families were recruited. Behavioural data was acquired from all subjects; genetic data was acquired for 1,142 subjects; a total of 1,113 subjects were scanned at 3-Tesla (T) (1,065 subjects with diffusion MRI data), 184 at 7 T and 95 subjects were studied using MEG. Cutting-edge MRI scanning

<sup>&</sup>lt;sup>1</sup>https://www.humanconnectome.org/

protocols were devised to push spatial and temporal resolution (Uğurbil et al., 2013).

Subject demographics are centered around the young-healthy adult United States population and are summarised in terms of age, race, ethnicity, handedness, employment and education in Table 3.1 (considering only the 1,065 subjects with dMRI data). Further descriptors include twin status, family structure, data completeness and the release in which the data was included (accounting for slight variance in data acquisition and processing between releases).

Demographic Measure	Summary		
Gender	575 (54%) female, 490 (46%) male		
Age in Years	Mean $= 28.75$ (standard deviation $= 3.67$ )		
Race	803 (75%, White), 150 (14%, Black or African Amer-		
	ican), 65 (6.1%, Asian/Native American, Hawai-		
	ian/Other Pacific Island), 27 (2.5%, More than one),		
	20 (1.9%, Unknown or Not Reported)		
Handedness	Mean $= 66$ (standard deviation $= 44.01$ )		
Employment Status	155 (14.5%, Not Working), 187 (17.6%, Part-time Em-		
	ployment), 722 (67.8%, Full-time Employment)		
Years of Education	Mean $= 14.94$ (standard deviation $= 1.79$ )		

**Table 3.1:** Summary demographics of the HCP subjects. Handedness scale is -100 (left) to +100 (right). For Year of Education, if the value is less than 11 or more than 17, it is reported as 11 and 17 respectively. For comparison, US population demographics: 51% male, 49% female (aged 24-54); 75% White, 12.7% Black or African American, 6.4% Asian/Native American, 3.1% More than one; 3.7-6.2% unemployed. Statistics taken from wikipedia.org/wiki/Demographics\_of\_the\_United\_States#Race\_and\_ethnicity

## 3.2.2 Non-Imaging Data

A broad range of behavioural and cognitive tests and questionnaires were completed and demographic and lifestyle data collected. In total, 476 measures of behaviour, cognition, family history, demographics, physical health and lifestyle were collected. Many of these measures are derived from tests included in the NIH Toolbox (for the Assessment of Neurological and Behavioural Function) (Gershon et al., 2013) (see http://www.nihtoolbox.org) and the "Penn Neuropsychological Battery" (PNB) (Gur et al., 2001; Gur et al., 2010). A brief description of the relevant tests and measures is provided here. For a full descriptions, see the HCP data dictionary<sup>2</sup>.

Health and family history were captured through a series of physical body measures and questionnaires. Physical body measures include the subject's height, weight, body mass index (BMI), self-reported heaviest historical BMI and results of blood test measurements of hematocrit, thyroid stimulating hormone (TSH) and hemoglobin A1C (i.e. percentage glycated hemoglobin). Physical health and motor skills were also measured by assessing cardiovascular endurance (2-minute walk test and gait speed), manual dexterity (9-hole pegboard test) and grip strength. Self-reported history of hyper/hypo-thyroidism, any other endocrine problems and, where applicable, regularity of menstrual cycle and any birth control are given. Family history was self-reported and covers a range of psychiatric disorders for both the subject's mother and father. Summaries of these measures of subject health and family history are given in Table 3.2.

Further behavioural measures capture cognitive ability, personality, mental health, social well-being and more. More specifically, measures span from: cognitive tasks; psychological well-being; measures of psychiatric and life functioning; illicit substance, alcohol and tobacco use; and blood samples for DNA sequencing. Cognition was assessed using an array of established tests, including the mini-mental state examination (MMSE), probing overall cognitive status, (Folstein et al., 1975) and a reduced version of Raven's Progressive Matrices (Raven, 2000), capturing fluid intelligence. Further, the NIH toolbox and the PNB capture several aspects of cognitive ability (e.g. overall intelligence, memory), psychological well-being (e.g. emotion processing) and personality. Psychiatric health, life functioning and illicit substance, alcohol and tobacco use were assessed through self-report questionnaires. Illicit substance, alcohol and tobacco use at the point of data acquisition was further assessed through screening testing. Full descriptions of these behavioural measures

 $<sup>^{2}</sup>$ Full list available at:

 $<sup>\</sup>label{eq:https://wiki.humanconnectome.org/display/PublicData/HCP+Data+Dictionary+Public+Updated+for+the+1200+Subject+Release} the the set of the set of the set of the the set of t$ 

Measure	Summary		
Height (Inches)	Mean $= 67.5$ (standard deviation $= 3.9$ )		
Weight (Pounds)	Mean = $171.7$ (standard deviation = $38.9$ )		
BMI	Mean = $26.4$ (standard deviation = $5.1$ )		
Blood Pressure	Mean Systolic = $123.4$ (standard deviation =		
	13.8), Mean Diastolic = $76.5$ (standard devia-		
	tion = 10.5)		
Family History (Mental Disor-	279 reported		
ders)			
Family History (Drugs and Al-	176 reported		
cohol)			
Family History (Neurological	19 reported		
Conditions)			
Family History (None)	912 (Mother), 852 (Father)		

**Table 3.2:** Summary subject health and family history of the HCP subjects. Healthy BMI is 18.5–24.9, average BMI in the United States is 26.6 for men and 26.5 for women (https://www.cdc.gov/nchs/data/nhanes/databriefs/adultweight.pdf). Healthy systolic and diastolic blood pressure is below 120 and below 80 respectively. For family history, "Mental Disorders" includes schizophrenia, depression, bipolar disorder and anxiety; "Neurological Conditions" includes Alzheimer's disease and Parkinson's disease; and "None" indicates that the subject reported no family history of the above conditions for their Mother or Father. For "Mental Disorders", "Drugs and Alcohol" and "Neurological Conditions", the total is given across conditions and parents.

are available in Barch et al. (2013). A full list of the behavioural measures used in this thesis is available in the Appendix of Chapter 6 (Section 6.9.2).

## 3.2.3 Imaging Data

#### Scanner Details

Imaging data were acquired at 3 T, while a subset of the cohort was scanned at 7 T as well. This thesis uses the 3 T data and so will focus on these. The customised 3 T Siemens Skyra (Siemens AG, Erlanger, Germany) scanner, coined the 3 T "Connectom Skyra", at Washington University was used for all image acquisition. This scanner is customised to be able to achieve a maximum gradient strength of 100 mT/m (a significant improvement over the standard 40 mT/m) and maximum slew rate of 91 mT/ms. This is used in combination with a 32-channel Siemens head

RF receive coil and a separate Siemens body RF transmission coil. This improved gradient magnetic field system performance, which is particularly beneficial during the acquisition of dMRI data (Uğurbil et al., 2013; Sotiropoulos et al., 2013a).

#### Data

The HCP acquired anatomical, diffusion and functional MRI brain data (Uğurbil et al., 2013). The main protocol details of all MRI modalities are shown in Table 3.3. The anatomical data consist of high-resolution T1-weighted (T1w) and T2-weighted (T2w) scans.

	T1w	T2w	dMRI	RS-fMRI
Sequence	3D MPRAGE	3D T2-SPACE	Spin-Echo EPI	Gradient-Echo
				EPI
TR (ms)	2400	3200	5520	720
TE (ms)	2.14	565	89.5	52
Voxel Size	0.7 isotropic	0.7 isotropic	1.25 isotropic	2.0 isotropic
(mm)				
iPAT, MB	2,-	2,-	-,3	-,8
# Volumes			270	4800
b-values	-	-	1000, 2000,	-
$(s/mm^2)$			3000	
Acquisition	16	16	60	60
time (mins)				

**Table 3.3:** Summary of image acquisition parameters for the HCP cohort. TR = Repetition Time, TE = Echo Time, iPAT = In-Plane Acceleration, MB = MultiBand factor

fMRI (specifically resting-state) data in the HCP were acquired in four blocks, each containing 1,200 volumes over 15 minutes with flipped phase-encoding direction (left-right and right-left) and a spatial resolution of 2.0 mm, totalling to 1 hour of RS-fMRI data. Repeated acquisitions and flipped phase-encoding directions reduce signal dropout and allows for distortion correction (Smith et al., 2013b).

As previously mentioned, the HCP scanner was significantly upgraded prior to data acquisition. These upgrades had great benefit with respect to dMRI data quality. A monopolar Stejskal–Tanner diffusion-encoding scheme spin-echo EPI sequence was used in combination with multiband acceleration. Three b-value shells, each with 90 non-collinear directions were acquired twice with the phase-encoding flipped between acquisitions (allowing for distortion corrections described later) with a spatial resolution of 1.25 mm (total scan time  $\sim 60$  min per subject) (Sotiropoulos et al., 2013a).

## 3.2.4 The HCP Minimal Preprocessing Pipeline (MPP)

Prior to public release (and use in this thesis), all HCP data is preprocessed using a processing pipeline developed by the HCP consortium (Glasser et al., 2013; Sotiropoulos et al., 2013a). These minimal processing pipelines (MPP) are applied to all data (anatomical, diffusion and functional) with six key goals:

- 1. Removal of spatial artefacts and distortions
- 2. The generation of cortical surfaces, segmentations and myelin maps
- Ease data visualisation by making the data compatible with the Connectome Workbench visualisation software (Marcus et al., 2013)
- 4. To generate precise cross-modal registration
- 5. To generate spatial normalisation warp fields for volume and surface spaces
- 6. To generate and make available spatially normalised "grayordinate"<sup>3</sup> data

Anatomical MRI processing: Following image data acquisition, scan quality was graded (Marcus et al., 2013), images were anonymised (including defacing) and the images were converted from DICOM (Digital Imaging and Communications in Medicine) image file type to NIfTI file format<sup>4</sup>. Anatomical data processing includes distortion and bias field corrections, co-registration of the T1w and T2w images, and

<sup>&</sup>lt;sup>3</sup>Grayordinate: the term used to describe the grouping of cortical surface vertices and subcortical voxels

<sup>&</sup>lt;sup>4</sup>For details on the NIfTI (Neuroimaging Informatics Technology Initiative) image file format, see: https://nifti.nimh.nih.gov/.

spatial normalisation to the MNI152 co-ordinate system<sup>5</sup> (Grabner et al., 2006). Spatial normalisation to the MNI152 co-ordinate system was achieved through an initial affine registration followed by a non-linear registration (Andersson et al., 2007; Jenkinson et al., 2012), giving two sets of subject data: one in the subject's native T1w space and one in the MNI152 co-ordinate system.

FreeSurfer (Fischl, 2012) and its outputs were used to segment the anatomical images, producing pial, WM and cortical myelination surface maps, used throughout this thesis to generate connectivity features. A subset of the FreeSurfer-based brain segmentation (the subcortical volumes) were combined with a downsampled standard brain surface meshes to produce a standard grayordinate CIFTI file<sup>6</sup>, allowing for the standardised visualisation and analysis of brain connectivity data. These CIFTI files contain 91,282 grayordinates consisting of approximately 32,000 vertices per hemisphere surface mesh and approximately 30,000 voxels for the 2mm isotropic resolution subcortical grey matter (GM) volumes.

**Resting-state Functional MRI processing**: Temporal and spatial preprocessing were completed on the RS-fMRI data. Spatial preprocessing included correction for spatial distortions, motion correction and registration to standard space. Temporal preprocessing included minimal highpass frequency filtering and ICA based classification of "good" (i.e. true functional networks) and "bad" (artefacts and other unwanted signal sources) components using FSL's MELODIC (Beckmann and Smith, 2004) and FIX (Salimi-Khorshidi et al., 2014; Griffanti et al., 2014). These components are aligned to the standard grayordinate space, covering both the cortical vertices and subcortical voxels, allowing further analysis to be completed in either space - in this thesis, connectomes obtained from the surface vertices are used. At this stage the RS-fMRI data is made publicly available (as the minimally processed data) (Smith et al., 2013b). However, the option to download further processed

<sup>&</sup>lt;sup>5</sup>See http://nist.mni.mcgill.ca/?p=858 for details

<sup>&</sup>lt;sup>6</sup>For details on the CIFTI (Connectivity Informatics Technology Initiative) image file format, see: https://www.nitrc.org/projects/cifti/.
data is available.

The data used in this research (in Chapter 6) comes from the parcellated network matrices ("netmats"). A group averaged dense connectome is derived using group-PCA (Smith et al., 2014). Group-ICA is applied to the group averaged dense connectome to produce group-level parcellations (with a range of dimensionalities: 15-300). Each parcel (component) represents a functional node. Via dual-regression, nodes are then used to extract subject-specific timeseries and further dual regression provides subject-specific parcellations. Full and partial correlation connectivity matrices (netmats) are then derived (Smith et al., 2013a). The parcellations, timeseries and netmats are available for the full "S1200" HCP data release and are used in this research. The netmats are, simply put, symmetric matrices representing the functional connectivity between two brain regions. Thus, in this work, we are only interested in the unique entries corresponding to the upper (or lower) triangle of the matrix giving d(d-1)/2 (where d is the group-level parcellation dimensionality) unique functional connectivity features.

**Diffusion MRI processing**: Diffusion data preprocessing includes the correction of susceptibility-induced and eddy-current-induced distortions as well as head motion correction (Andersson et al., 2003; Andersson et al., 2016; Andersson and Sotiropoulos, 2016). Following this, spatial distortion induced by gradient nonlinearities were corrected for (Sotiropoulos et al., 2013a). In addition to the previously described spatial normalisation, the diffusion data were registered to the T1w space through a rigid-body transform and the diffusion gradient vectors were rotated accordingly. Finally, the crossing fibre model described in Jbabdi et al. (2012) was applied. This is a parametric spherical deconvolution model that accounts for the non-monoexponential decay of the dMRI signal with higher b-values. Up to three fibre orientations were estimated in each voxel along with their uncertainty. This modelling is a precursor to the tractography performed in this thesis.

# 3.3 The UK Biobank

The UK Biobank<sup>7</sup> is a large cohort, multi-modal, multi-site epidemiological study. Simply, the wealth of data available through the UK Biobank is unprecedented, with approximately half a million subjects in the cohort. Such a large cohort was achieved by inviting approximately 9.2 million individuals to take part (with 5.45% of those invited agreeing to take part) (Fry et al., 2017). Data is collected using a broad range of diagnostic tools, including MRI. One of the aims of the UK Biobank imaging study is to collect brain MRI data for 100,000 healthy subjects whilst tracking their health over the following decades. The following sections summarise the cohort demographics, non-imaging data and imaging data relevant to this thesis. Full descriptions are provided elsewhere: for brain imaging data, see Miller et al. (2016); for subject demographics, see Fry et al. (2017); for a general overview of the project, see Sudlow et al. (2015); for in-depth details on all measures, see the UK Biobank Data Showcase (https://biobank.ctsu.ox.ac.uk/crystal/index.cgi).

#### 3.3.1 Cohort Overview

In total, the UK Biobank has recruited 503,317 subjects, 37,848 of whom have fMRI data at the time of writing. Subjects are aged between 40–69 years and live within 25 miles of one of 22 assessment centers located throughout England, Wales, and Scotland. The cohort is considered generally healthy at the point of recruitment but one of the key aims of the project is to track subject health, thus affording researchers the opportunity to identify early signs of disease. Although impressive in scale, it should be noted that the UK Biobank cohort suffers significantly from recruitment biases: the cohort is healthier and less socioeconomically deprived compared to the UK population (Fry et al., 2017). Table 3.4 provides a summary of the key demographics of these subjects.

<sup>&</sup>lt;sup>7</sup>https://www.ukbiobank.ac.uk/

Demographic Mea-	Summary			
sure				
Gender	20,122 (53%) female, 17,722 (47%) male			
Age in Years	Mean $= 54.8$ (standard deviation $= 7.5$ )			
Ethnicity	34,506 (91.17%, White British), 38 (0.10%, White and Black			
	Caribbean), 271 (0.72%, Indian), 129 (0.34%, Caribbean),			
	917 (2.4%, White Irish), 21 (0.05%, White and Black			
	African), 59 (0.16%, Pakistani), 103 (0.27%, African), 1139			
	(3.01%, Any other white background), 57 (0.15%, White and)			
	Asian), 6 (0.02%, Bangladeshi), 2 (0.01%, Any other Black			
	background), 54 (0.14%, Any other mixed background), 58			
	(0.15%, Any other Asian background)			
Handedness	33,698 (right-handed), 3,538 (left-handed), 591 (both			
	equally)			
Employment Status	428 (1.13%, Unemployed), 9,159 (24.20% Retired), 26,317			
	(69.53%, Employed), 997 (2.63%, Looking after home/fam-			
	ily), 407 (1.08%, Unable to work because of sickness/disabil-			
	ity), 181 (0.48%, Unpaid/voluntary work)			
Age completed full-	Mean $= 17.2$ (standard deviation $= 2.4$ )			
time education				

**Table 3.4:** Summary demographics of the UK Biobank subjects. Notes: age is recorded at recruitment; ethnicity is reported using a hierarchical structure where participants select "White" and "British", for example. For comparison, UK population demographics: 51% female, 49% male (aged 40-69); 91.3% White, 2.4% Black or black British, 0.8% Mixed, 1.9% Indian, 1.0% Pakistani, 0.3% Bangladeshi, 0.5% Chinese, 1.0% Other Asian, 0.6% Other ethnic group Fry et al. (2017)

# 3.3.2 Non-Imaging Data

Subject data is extensive, consisting of lifestyle, family and medical history questionnaires; physical and cognitive measures; blood, urine and saliva samples; and genetic data. In the current work, a subset of these measures are used. These include self-report questionnaires on diet, lifestyle, family and medical history; summary measures of physical and cardiac health derived from non-MRI modalities (e.g. ultrasound bone densitometry, electrocardiograms) and body impedance analysis; and scores from a range of cognitive tests. The measures broadly capture information similar in nature to that provided in the HCP. Data is available locally for 39,696 subjects (under UK Biobank project number 43822): those without fMRI data have been removed, leaving 37,848 subjects. A full list of the behavioural measures used in this thesis is available in the Appendix of Chapter 6 (Section 6.9.3).

# 3.3.3 Imaging Data

Imaging data in the UK Biobank is acquired using various techniques and modalities. For a cohort of this size, this is both a benefit and hindrance: the time and cost associated with acquiring these data are huge with each minute of MRI data acquisition costing approximately £1 million (Miller et al., 2016; Alfaro-Almagro et al., 2018). Therefore, MRI data quality is more typical of a clinical setting rather than the very high quality data of the HCP.

#### Scanner Details

MRI brain data is acquired at three dedicated imaging centers using Siemens Skyra 3 T scanners with standard Siemens 32-channel RF receive head coils. Hardware and software are kept consistent across sites and software is maintained at a single version throughout the study.

#### Data

The MRI modalities of interest here are the anatomical, diffusion, and functional MRI. Although the anatomical data is not used directly in this work, as in the HCP preprocessing, it is used in the processing of the diffusion and functional data prior to public release. Table 3.5 provides a summary of the key acquisition parameters for the MRI data<sup>8</sup>. Briefly, anatomical data consists of T1w and T2w scans with a resolution of 1.0 mm.

The functional data (RS-fMRI) was acquired over 490 timepoints with an isotropic spatial resolution of 2.4 mm and a multisclice acceleration factor of 8 ( $\sim 6$  minutes of data acquisition).

<sup>&</sup>lt;sup>8</sup>See http://biobank.ctsu.ox.ac.uk/crystal/docs/brain\_mri.pdf for complete protocol details.

The diffusion data was acquired using monopolar Stejskal-Tanner diffusionencoding scheme spin-echo EPI sequence with an isotropic spatial resolution of 2.0 mm and diffusion weighting of b = 0,1000,2000. Each b > 0 shell is made up of 50 non-collinear diffusion-encoding directions, giving 100 non-collinear directions in total. Flipped phase-encoding data was also acquired to allow for the calculation of fieldmaps and to perform geometric distortion correction.

	T1w	T2w	dMRI	RS-fMRI
Sequence	3D MPRAGE	3D T2-SPACE	Spin-Echo EPI	Gradient-Echo
				EPI
TR (ms)	2000	5000	3600	735
TE (ms)	2.01	395	92	39
Voxel Size	1.0 isotropic	$1.05 \times 1.0 \times 1.0$	2.0 isotropic	2.4 isotropic
(mm)				
iPAT, MB	2,-	2,-	-,3	-,8
# Volumes			100	490
b-values	-	-	1000, 2000	-
$(s/mm^2)$				
Acquisition	5	6	7	6
time $(mins)$				

**Table 3.5:** Summary of image acquisition parameters for the UK Biobank cohort. TR = Repetition Time, TE = Echo Time, iPAT = In-Plane Acceleration, MB = MultiBand factor

#### 3.3.4 Preprocessing Pipeline

Anatomical MRI processing: The preprocessing pipeline is similar to that described for the HCP MMP (Alfaro-Almagro et al., 2018). However, one major difference is that the UK Biobank do not yet provide analysis in the previously discussed grayordinate system and, therefore, surface analysis is not possible. The general principles for the anatomical data are to perform gradient distortion correction, spatial normalisation to the standard MNI152 coordinate system (as in the HCP MPP) and brain segmentation. In addition, automated quality control (QC) is performed on the data. This is achieved by training a machine learning classifier on what constitutes a good or bad image based on the anatomical T1w scans. This pipeline results in the identification of usable data within the cohort with 81% of the datasets (across all modalities) deemed suitable for use (Alfaro-Almagro et al., 2018).

Further registration warp fields were generated for the functional and diffusion data. Specifically, warp fields for the diffusion data were generated through nonlinear alignment to the standard space (MNI152 co-ordinate system) FA FMRIB58 atlas<sup>9</sup> based on FA maps, allowing for standard space analysis (e.g. standardised tractography).

**Diffusion MRI processing**: Fieldmaps, to correct susceptibility-induced geometric distortions in EPI data, are generated based on the phase-encoding pair (AP - PA) of b = 0 diffusion data using FSL's topup (Andersson et al., 2003). Subsequently, susceptibility-induced distortions and eddy-current and head motion induced distortions are corrected for using FSL's eddy (Andersson et al., 2016; Andersson and Sotiropoulos, 2016). The UK Biobank processing pipeline extends the diffusion data processing further by; applying diffusion models including the diffusion tensor, neurite orientation dispersion and density imaging (NODDI, Zhang et al. (2012)) and the crossing fibre model (Jbabdi et al., 2012) to the data; extracting tract-based spatial statistics (TBSS, Smith et al. (2006)); and performing ROI-based tractography (Groot et al., 2013; Hernandez-Fernandez et al., 2019). These processing steps are used to generate microstructural feature maps and a set of image-derived phenotypes (IDPs), summarising structural connectivity features (not used in this thesis).

**Resting-state Functional MRI processing**: Functional data preprocessing closely follows that of the HCP volumetric preprocessing pipeline. In summary, EPI unwarping (using the previously estimated fieldmaps), gradient distortion correction, motion correction, intensity normalisation and highpass temporal filtering were applied using FSL's MELODIC (Beckmann and Smith, 2004). Structured artefacts

<sup>&</sup>lt;sup>9</sup>FA FMRIB58 atlas: https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FMRIB58\_FA

were removed using FSL's FIX (Beckmann and Smith, 2004; Salimi-Khorshidi et al., 2014; Griffanti et al., 2014): this utilises independent component analysis (ICA) and manual training to classify components as (non-)artefactual.

Resting-state networks are first defined at a group level based on the data from 4,100 subjects. This is achieved by performing group-level principle component analysis (PCA) with MIGP (MELODIC's Incremental Group-PCA) (Smith et al., 2014), giving a group-averaged dense connectome. The resultant 1,200 components were then further reduced through spatial-ICA, i.e. parcellated, at two dimensionalities (25 and 100) using FSL's MELODIC. The 100-dimensional set is used in this thesis. Of the 100 components, 45 were identified as artefactual, leaving 55 components. These group-wise components were then dual-regressed back onto each subject's timeseries, giving a timeseries per component, per subject. These data are used to extract partial correlation connectivity matrices (netmats) using FSLNets (Smith et al., 2013b). This results in a final d = 55 symmetric connectivity matrix per subject with  $\frac{d}{2}(d-1) = 1485$  unique entries.

# 3.4 HCP Structural Connectivity Features and Atlases

As part of early projects in this thesis and in setting-up high-throughput analysis pipelines, we used the HCP dMRI data and standard tools to obtain high-quality white matter atlases and connectivity features. These have led directly to publicly available data (some of them available already in FSL<sup>10</sup>) or have contributed significantly to this thesis and to projects other than those presented in this thesis and they are briefly presented here.

<sup>&</sup>lt;sup>10</sup>HCP1065 standard-space DTI templates https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/Atlases

#### 3.4.1 White Matter Microstructural Atlases

A number of high-resolution microstructure feature maps have been generated using all subjects in the HCP cohort (1,065 subjects). These include quantitative maps from the diffusion tensor model and the neurite orientation dispersion and density imaging (NODDI) model (Zhang et al., 2012).

These feature maps have been further used to generate atlases of these features. Diffusion tensors were fitted in each subject, using the  $b = 1000 \text{ s/mm}^2$ data and accounting for gradient non-linearities. Each tensor was transformed to the standard MNI152 space using the subject-specific non-linear transformations, after appropriate re-orientation of the tensor elements (Alexander et al., 2001). The standard-space tensors were averaged and the averaged diffusion tensor was then decomposed into its eigenvectors and eigenvalues, giving rise to average fractional anisotropy (FA) and mean diffusivity (MD) maps. Figure 3.1 shows the HCP cohortaveraged FA map ("FSLHCP1065\_FA") and compares this averaging approach to a simplistic map averaging (simply transforming the subject-wise FA maps to the standard space and taking the cohort mean) and the current average FA map available in FSL (the "FMRIB58\_FA" map). Vast improvements in detail are observed when comparing the new FA map (bottom row) or the simple average (middle row) to the FMRIB58 FA map (top row). Clear improvements in image sharpness are still achieved in the new map when compared to the simple average. The level of detail achieved in this map will be greatly beneficial in image registration, for example, and was particularly useful in the development of the tractography protocols developed in Chapter 4.

To truly appreciate the level of detail achieved, an enhanced region corresponding to the splenium of the corpus callosum is displayed in Figure 3.2. The direction of the first principal eigenvector of the diffusion tensor is displayed in the standard colour-coding format. Each line represents the direction of the eigenvector in a given



**Figure 3.1:** Sagittal, coronal and axial views of the group-averaged fractional anisotropy (FA) map (FSLHCP1065\_FA, bottom row) and, for comparison, the simple mean FA map (middle row) and the FMRIB58\_FA map (top row).

voxel with an isotropic spatial resolution of 1 mm and the brightness modulated by the underlying FA values. These microstructural atlases are freely available within FMRIB's Software Library<sup>11</sup>.

Further cohort-averaged microstructural feature maps were obtained from NODDI<sup>12</sup>

 $<sup>^{11}\</sup>mathrm{See}$  "HCP1065 standard-space DTI templates" at https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/Atlases for details

<sup>&</sup>lt;sup>12</sup>NODDI modelling was completed using the CUDA Diffusion Modelling Toolbox (cuDIMOT) toolbox: https://users.fmrib.ox.ac.uk/ moisesf/cudimot/index.html (Hernandez-Fernandez et al., 2019).



**Figure 3.2:** The group-averaged first eigenvector map (modulated by FA) with the standard colour coding. A exemplar region of the brain (the splenium of the corpus callosum), highlighting the superb quality achieved by averaging over a large cohort of high-quality data.

modelled data, using all 3 shells of the HCP data, shown in Figure 3.3. In the case of the orientation dispersion map (top row), lower intensities indicate regions in which the angular variation of neurites is lower, i.e. diffusion is strongly directional. The intra-cellular volume fraction map (bottom row) describes the fraction of the voxel bound by neurite membranes.

# 3.4.2 Structural Connectomes

Multiple "flavours" of structural connectomes have also been calculated at the subject-level and subsequently averaged. One of which describes the connectivity between the whole-brain grey matter (GM) (actually the white-grey matter boundary (WGB)) and whole-brain white matter (WM) used in the calculation of connectivity blueprints in Chapter 4 (shown in Figure 3.4a). Another is the whole-brain GM to GM connectivity matrix. Here, tractography is seeded from each WGB surface vertex (approximately 60,000 vertices). The same WGB surface is taken as a termination mask and streamlines are allowed to pass through whole-brain WM. In addition, the streamlines are allowed to enter and propagate through the subcortical structures but terminate upon exiting. In addition, in order to avoid gyral jumping (falsely passing through the boundaries between neighbouring gyri), the pial sur-



**Figure 3.3:** Examples of the microstructure measures obtained from the groupaveraged HCP diffusion NODDI modelled data. Top row: orientation dispersion map. Bottom row: intra-cellular volume fraction map.

face is used as a termination mask. The seeding approach used here is preferred over seeding at each voxel as reductions in inherent biases are observed (Donahue et al., 2016). The number of streamlines reaching each WGB vertex is recorded and normalised by the total number of valid streamlines. This produces a dense connectivity matrix  $60k \times 60k$  describing the connectivity between a given cortical vertex and the rest of the cortex, i.e. the GM to GM connectome (Figure 3.4b).



Figure 3.4: Group average structural connectomes derived from the HCP cohort. a. the white-grey matter boundary (WGB, with  $N \approx 60k$  vertices) to white matter (WM, with  $M \approx 120k$  voxels) connectome averaged across 100 subjects. b. the grey matter (GM) to GM connectome, represented by N surface vertices averaged across the whole HCP cohort (1022 subjects).

# **3.5** Summary and Thesis Context

This chapter introduces the primary datasets used in this thesis. As described, the "WU-Minn" Human Connectome Project provides superb quality MRI data for a relatively large cohort of subjects with substantial behavioural data. The UK Biobank provides population-level MRI data of quality more typical of a goodquality data in a clinical setting but makes this available for forty thousands subjects.

This chapter also described the derivation of structural connectomes and atlases of diffusion microstructural features. The diffusion microstructural atlases are available in FSL (version 6.0.1 and later). The structural connectomes provide a valuable resource for exploring the associations between structural brain connectivity and other, imaging or non-imaging, features in future work - these have also been utilised in side-projects (Tewarie et al., 2020).

The diffusion MRI data from both datasets is utilised in the development of XTRACT in Chapter 4 and for its applications shown in Chapter 5. The unique family structure of the HCP allowed for the investigation of WM tract similarity within twin and non-twin groups, providing evidence that our protocols reflect the underlying anatomical variation across subjects. The scale and quality of the HCP data also allowed for the calculation of high quality WM tract atlases, as well as the diffusion microstructural feature map atlases described in Section 3.4. The availability of a dataset more similar in quality to that of a clinical scenario (the UK Biobank) allowed us to investigate the generalisability of XTRACT across data quality. This is particularly useful in ensuring that, as a publicly released tool, tract reconstructions are representative of the true anatomy despite significant differences in data quality.

Finally, both the imaging and non-imaging data from the HCP (diffusion and functional MRI) and UK Biobank (functional MRI) are vital in the investigations presented in Chapter 6. The availability of brain imaging data and the wide array of behavioural data for these large cohorts allows for the exploration of the stability of brain-behaviour association techniques in a robust manner.

# Chapter 4

# Standardised and Automated Tractography for Deriving Connectivity Features

## Contents

4.1	Introduction				
4.2	Methods				
	4.2.1	XTRACT and the Tractography Protocols			
	4.2.2	Data and Preprocessing			
	4.2.3	Fibre Orientation Estimation and Tractography 90			
	4.2.4	Connectivity Blueprints			
	4.2.5	Atlas Generation			
	4.2.6	Protocol Evaluation and Validation Methods			
	4.2.7	Capturing Inter-Subject Variability			
4.3	$\mathbf{Resul}$	$\operatorname{ts}$			
	4.3.1	Atlases			
	4.3.2	The Effect of Sample Size on Atlas Generation 98			
	4.3.3	Robustness Against Datasets			
	4.3.4	Reflecting Known Anatomy			
	4.3.5	Capturing Inter-Subject Variability			
4.4	Discu	ussion			
4.5	Conc	Concluding Remarks			
4.6	Appendix				
	4.6.1	Summary of the Tractography Protocols Previously De-			
		fined in the Literature $\ldots \ldots 116$			
	4.6.2	XTRACT User-Guide 118			
	4.6.3	Summary of the Tract Lateralisation Literature 122			
	4.6.4	The Effect of Resolution on Connectivity Blueprints 123			
4.6	<b>Appe</b> 4.6.1 4.6.2 4.6.3 4.6.4	ndix 11   Summary of the Tractography Protocols Previously De- 11   fined in the Literature 11   XTRACT User-Guide 11   Summary of the Tract Lateralisation Literature 12   The Effect of Resolution on Connectivity Blueprints 12			

# Foreword

Tractography is a powerful and unique technique but, traditionally, is a timeconsuming and subjective process. It is used by the neuroscience community to localise and reconstruct white matter tracts which connect brain regions. These representations may be used to derive features, to construct networks and for in vivo neuroanatomy explorations, potentially across large groups of individuals. The application of tractography for these purposes requires standardisation and automation.

This chapter describes the development of a toolbox for the standardised and automated extraction of 42 major white matter bundles. A special feature of our approach is that the devised tractography protocols are generalisable not only across individuals, but also across non-human species (macaque). The toolbox, called XTRACT (cross-species 'X' tractography 'TRACT'), has been released in the FM-RIB Software Library (FSL).

In the following sections, the background to standardised and automated tractography is presented and the development of the toolbox and tractography protocols are discussed. The toolbox is applied to the HCP, UK Biobank and to a macaque dataset, and white matter tract atlases are generated, along with representations of the cortical termination points of the tracts. The approach is evaluated in a number of ways. We indirectly validate: a) through the assessment, and comparison to the literature, of the hemispheric lateralisation of the resultant tracts, both in volume space and on the cortical surface and b) by demonstrating generalisability across datasets of different data quality and demographics. In addition, despite the standardisation of the process, we illustrate that individual variability in reconstructing tracts is preserved by exploring the similarity of tracts in the context of family structure (twins and siblings), as well as in cases of structural brain abnormalities.

# 4.1 Introduction

Diffusion tractography is a unique tool for extracting white matter (WM) pathways non-invasively and in vivo. The virtual dissection of major WM tracts enables the study of brain organisation (Catani et al., 2013; Jbabdi et al., 2015) and offers a probe to brain development (Huppi and Dubois, 2006) and WM pathology (Ciccarelli et al., 2008; Griffa et al., 2013). It further allows explorations of individual variations (Assaf et al., 2017) and cross-species variations (Mars et al., 2018c) in anatomy and connectivity. This information has functional relevance, as the pattern of extrinsic WM connections of each functional brain subunit to the rest of the brain are unique (Passingham et al., 2002; Mars et al., 2018b).

To be able to reliably study individual variability in WM pathways, tractography approaches often utilise protocols to extract a pre-defined set of WM tracts. Such protocols typically comprise of regions of interest (ROIs) and respective rules on using them (for instance as seed, exclusion and termination masks). They reflect prior anatomical knowledge used to guide and constrain curve propagation, reducing the chance of false positives (Catani et al., 2002; Wakana et al., 2004). Tractography protocols must be robust and reproducible, allowing reconstruction of WM tracts in a consistent manner across subjects, whilst respecting the underlying anatomical variation and individual differences. One approach that may be used is to define subject-specific tractography protocols (Conturo et al., 1999), considering the specific variations in individual anatomy. However, defining masks on a subject-wise basis is time-consuming and subjective (Jones, 2008; Nucifora et al., 2012), while for large cohorts these limitations become prohibitive. The alternative to this manual approach is to define a set of standardised masks in template space. which are then registered to the individual geometry and used in a consistent and automated manner for each subject.

These automated region of interest (ROI) -based tractography approaches have proven powerful in the extraction of a range of tracts (Thiebaut de Schotten et al., 2011b; Catani et al., 2012; Makris et al., 2013; Menjot de Champfleur et al., 2013; Hecht et al., 2015; Nowell et al., 2016; Hau et al., 2016; Zhao et al., 2016; Takemura et al., 2017; Howells et al., 2018; Maffei et al., 2019; Eichert et al., 2019). They have further allowed for the application to large cohorts for development of tract-specific atlases (Archer et al., 2018; Chenot et al., 2019) and extraction of tract-specific features (Miller et al., 2016).

In this chapter, a new approach is presented. We devise a library of standardised tractography protocols for the automated extraction of WM tracts, both in the human and the macaque brain. Prior anatomical knowledge for the human and non-human primate brain is translated to a set of species-equivalent tractography rules in order to extract homologous reconstructions for a range of WM tracts. Such equivalence can provide unique possibilities for comparative neuroanatomy and the identification of functionally equivalent cortical regions between different species, in the absence of any geometrical similarity (Mars et al., 2018c).

This work builds on and extends previous efforts that developed libraries of tractography protocols in the human brain (Wakana et al., 2007; Catani and Thiebaut de Schotten, 2008; Zhang et al., 2008; Hua et al., 2008; Thiebaut de Schotten et al., 2011b; Groot et al., 2013; Wassermann et al., 2016) (see Table 4.2 in the Appendix for a summary). A set of standard-space masks for the extraction of 20 tracts was developed by Wakana et al. (2007). They reported high inter- and intra-rater reproducibility and suggested that some tracts may display left-right asymmetry. In Hua et al. (2008), the authors extended this work to generate probabilistic tract atlases for 22 tracts (11 left/right) and exampled their use by investigating tract-specific abnormalities in multiple sclerosis. Zhang et al. (2008) applied these standardised masks to 10 subjects and reported high agreement between automated and manual tract segmentation approaches.

Similarly, Catani and Thiebaut de Schotten (2008) defined standard-space masks for the reconstruction of 19 tracts (7 left/right, 5 commissural) and assessed the reproducibility of their protocols. This work was furthered by Thiebaut de Schotten et al. (2011b) through the extension of the tractography protocols to 31 tracts (14 left/right, 3 commissural), where good correspondence between their automated tractography technique and histological atlases was reported.

In Groot et al. (2013), standardised protocols for probabilistic tractography

were used to reconstruct 27 tracts (12 left/right, 3 commissural) in two datasets with varying quality. More recently, Wassermann et al. (2016) proposed a framework for describing WM anatomy and tracts which uses subject-wise anatomical segmentation, clustering and a query language (the white matter query language -WMQL) to extract 57 (25 left/right, 7 commissural) tracts from whole-brain tractography and a grey matter (GM) parcellation. This approach reduces the definition of tracts to sets of logical rules with reference to the tract position and termination, the tract route and its spatial relationship to given brain parcels.

Despite the great potential of all previous developments, none of these have targeted multiple species, which is the aim of this study. The presented approach is also complementary to non-ROI-based methods for WM tract reconstruction, such as unsupervised clustering-based, e.g. O'Donnell and Westin (2007), Garyfallidis et al. (2012), and Siless et al. (2018), or supervised methods (Wasserthal et al., 2018). The former are "data-driven", whereas in the current study, prior anatomical knowledge is imposed in order to reduce false positives in tractography. The latter can be benefitted during training from new approaches, such as the one presented here, that aim to reconstruct anatomically known tracts in a consistent and reproducible manner.

The work in this chapter extends these previous efforts in devising an extended set of ROI-based tractography protocols and an automated tractography toolbox. The contribution of this chapter is as follows: 1) tractography protocols for 42 tracts are designed and their robustness against data quality is illustrated, using high-resolution data from the HCP (Sotiropoulos et al., 2013b) and more typical data from the UK Biobank (Miller et al., 2016) (see Sections 3.2 and 3.3 for more details), 2) the generalisability of the tractography protocols to the macaque brain is illustrated, 3) high-quality tract atlases are derived using these protocols both for the human brain (1000 HCP subjects) and the macaque brain (high-resolution, ex vivo datasets from 6 animals), 4) indirect validation is performed by assessing lateralisation of the extracted tracts (in humans), 5) it is illustrated that, despite being template-driven, reconstructed tracts preserve individual variability as assessed via twinship analysis, and 6) an open-source flexible framework for publicly exchanging tractography protocols is made available within FSL. New standard space WM tractography protocols may be defined and "plugged into" the toolbox, allowing for further expansion and tract exchange, contributing to open science and reproducibility of results.

# 4.2 Methods

# 4.2.1 XTRACT and the Tractography Protocols

Tractography protocols for 42 WM tracts (19 bilateral and 4 commissural) were devised in a generalisable manner that allows equivalent mask definitions to apply to both the human and the macaque brain. The full list of tracts that are currently supported is presented in Table 4.1. Further, a new cross-species tractography (XTRACT) toolbox, capable of reading the standard space tractography protocols and performing probabilistic tractography (Behrens et al., 2007), with the option of GPU acceleration (Hernandez-Fernandez et al., 2019) was developed and implemented.

	Tract	Abbreviation	Left/Right?	Seeding Strategy
	Arcuate Fasciculus	AF	Yes	Reverse-seeding
	Frontal Aslant Tract	FA	Yes	Single-ROI
	Inferior Longitudinal Fasciculus	ILF	Yes	Reverse-seeding
Association Fibres	Inferior Fronto-Occipital fasciculus	IFO	Yes	Reverse-seeding
	Middle Longitudinal Fasciculus	MdLF	Yes	Reverse-seeding
	Superior Longitudinal Fasciculus I	SLF1	Yes	Single-ROI
	Superior Longitudinal Fasciculus II	SLF2	Yes	Single-ROI
	Superior Longitudinal Fasciculus III	SLF3	Yes	Single-ROI
	Uncinate Fasciculus	UF	Yes	Single-ROI
	Vertical Occipital Fasciculus	VOF	Yes	Reverse-seeding
Commissural	Anterior Commissure	AC	No	Reverse-seeding
	Forceps Major	FMA	No	Reverse-seeding
Fibres	Forceps Minor	FMI	No	Reverse-seeding
	Middle Cerebellar Peduncle	MCP	No	Reverse-seeding
Limbic	Cingulum subsection: Dorsal	CBD	Yes	Single-ROI
	Cingulum subsection: Perigenual	CBP	Yes	Single-ROI
Fibres	Cingulum subsection: Temporal	CBT	Yes	Single-ROI
	Fornix	FX	Yes	Single-ROI
	Acoustic Radiation	AR	Yes	Reverse-seeding
Projection	Anterior Thalamic Radiation	ATR	Yes	Single-ROI
	Corticospinal Tract	CST	Yes	Single-ROI
Fibres	Optic Radiation	OR	Yes	Reverse-seeding
	Superior Thalamic Radiation	STR	Yes	Single-ROI

**Table 4.1:** The list of reconstructed WM tracts, their corresponding tract type and the seeding strategy used.

XTRACT is effectively a wrapper command-line tool which requires users, in

the most basic implementation, to provide a single subject's diffusion modelled data and their registration fields between the native diffusion space and the standard space for the given species. In addition, users may develop and apply their own protocols using XTRACT. XTRACT is built on pre-existing FSL tools, primarily utilising the registration tools (applywarp) and the probabilistic tractography tool (probtrackx). Figure 4.15 illustrates the main stages in XTRACT to reconstruct a given tract:

Step 1.) each tract is reconstructed using a unique combination of binarised masks (i.e. tractography protocols), defined in standard space (MNI152 for humans and F99 for macaques) using a combination of hand-drawn regions and atlas-based masks developed and agreed upon by multiple experts<sup>1</sup>. The masks include seeds, targets/waypoints, exclusions and stop/termination regions. Seed masks provide the starting points of the tractography streamlines. Target/waypoint masks are regions through which a streamline should pass in order to be valid. Exclusion masks serve to reject any streamline running through them and stop/termination masks stop any streamline running through them. Step 2.) using a non-linear registration warp field defining the transform between the standard space and the subject's native space, XTRACT warps the protocols to the subject's native space, and step 3.) performs probabilistic tractography in the subject's native space. XTRACT has the option to store the tractography results in either the subject's native diffusion space or, steps 4-5.) in the standard space in which the protocols are stored (its default and most commonly used behaviour). In this default usage, following tractography, results are directly resampled to the standard space using the inverse non-linear registration warp field, allowing between-subject geometrical correspondence, necessary in certain contexts (e.g. atlasing). Appendix 4.6.2 provides further details on the requirements, usage and output of XTRACT.

<sup>&</sup>lt;sup>1</sup>Protocol masks developed by Rogier Mars, Katherine Bryant (led the macaque protocols), Shaun Warrington, Gwenaëlle Douaud and Stamatios Sotiropoulos (led the human protocols).



Figure 4.1: Schematic of the stages for automated tractography, as implemented in XTRACT, with an example of the left arcuate fasciculus (AF) for the human brain. 1.) Tractography protocol masks are defined in standard space with seed (green), exclusion (black), waypoint (blue) and termination (orange) masks. 2.) The protocol masks are warped to the subject's native space using the subject-specific non-linear warp fields. 3.) Probabilistic tractography is performed in the subject's native space using the crossing fibre modelled diffusion data. 4.) Notice that results are mapped directly into standard space, leading to a single interpolation step. 5.) The resultant tract stored in standard space, overlaid on the FSL\_HCP1065 FA atlas.

#### **Protocol Definitions**

The sections below describe in detail the protocol for each tract in consideration in the case of human tractography and any adjustments for the macaque brain. With the exception of the brainstem and commissural tracts all protocols include the midline sagittal plane as an exclusion mask to restrict fibres to the ipsilateral hemisphere.

XTRACT supports two seeding strategies: a) a standard single-ROI seed and b) a "reverse-seeding" approach, where a pair of seed-target masks exchange roles and the final path distributions are added. In cases where the reverse-seeding strategy is used, this is to improve the robustness of the reconstruction and to ensure that the whole bundle is represented in the final reconstruction. Where possible, in order to reduce processing time, the single-ROI seed approach is used. Details of seeding strategy for each protocol are provided in Table 4.1.

In general, protocols were defined/guided based on anatomical priors from the literature. Adjustments and additional masks, compared to those described in the literature, were included to increase tract reconstruction success rate, improve tract coverage and reduce false positives (here, 'false positives' refers to the inclusion of streamlines that do not belong to a given fibre bundle) based on trial and error. Details of this trial and error approach are not provided in full. In brief, a protocol was tested on a small subset of subjects (tens of subjects in the human data) and tract quality was assessed considering the tract volume in each subject's tract reconstruction and by visual assessment of cohort average tracts. For each protocol, an example visualisation of the resultant tract is provided for the human and macaque brain.

#### Association Fibres

Superior Longitudinal Fasciculus (SLF) 1/2/3: The three branches of the superior longitudinal fasciculus 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 final coronal exclusion mask through subcortical areas posterior to the caudal end of the genu of the corpus callosum prevented leaking into ventral longitudinal tracts.



<u>Arcuate Fasciculus (AF)</u>: The arcuate fasciculus forms part of the system of dorsal longitudinal fibres, but in the human brain is distinguished by its posterior curve ventrally into the temporal cortex. The human AF was reconstructed with a seed in the supramarginal gyrus (SMG), a temporal target mask was in the WM 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. Following the observation in the macaque that this tract runs along the fundus of the circular insular sulcus (Petrides et al., 2012), a seed mask was placed there, just posterior to the level of the central sulcus. An axial target mask was placed in the parietal-temporal WM posterior to the caudal end of the Sylvian fissure. An additional axial plane was placed in the IFG. This protocol was validated by Eichert et al. (2019).



Middle/Inferior Longitudinal Fasciculus (MdLF, ILF): Three tracts that course along the temporal lobe were reconstructed (MdLF, ILF, IFO). The middle and inferior longitudinal tracts stay within the lateral posterior cortex. The MdLF was seeded in the anterior part of the superior frontal gyrus (SFG) (Makris et al., 2009); the ILF in the middle and inferior temporal gyri to account for the expansion of the temporal cortex in the human brain compared to the macaque (Latini et al., 2017; Roumazeilles et al., 2020). For the MdLF, large axial and coronal planes covering the WM in the temporo-parietal-occipital junction were used as targets, based on anatomical descriptions from Makris et al. (2009). For the 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 the 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 WM 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., 2016a), 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.



<u>Uncinate Fasciculus (UF)</u>: The bottom part of the extreme capsule contains fibres belonging to the uncinate fasciculus, curving from the inferior frontal cortex to the anterior temporal cortex. The tract was reconstructed using a seed in the STG at the first location where 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 accidental leaking into the fibres running longitudinally through the temporal lobe.



Frontal Aslant Tract (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 WM of the IFG, the target axially in that of the SFG. A posterior coronal exclusion mask prevented leakage into longitudinal fibres.



<u>Vertical Occipital Fasciculus (VOF)</u>: The VOF runs in a predominantly dorsalventral orientation in the occipital lobe. We used an adapted version of the recipe described by Takemura et al. (2017). An axial seed mask was placed in the lateral part of the ventral occipital WM 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.

Human

Macaque



#### **Commissural Fibres**

Middle Cerebellar Peduncle (MCP): The MCP was seeded in the cerebellar WM with a target in the opposite hemisphere (and their inverses). Exclusion masks were placed sagitally along the cerebellar midline and axially through the thalamus.

Human

Macaque



Corpus Callosum Splenium (FMA) & Genu (FMI): We reconstructed callosal con-

nections to the occipital lobe 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 though the inferior fronto-occipital WM 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.



<u>Anterior Commissure (AC)</u>: 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 WM lateral to the globus pallidae. 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.



#### Limbic Fibres

Cingulum subsections (CBT, CBP, CBD): Recently, Heilbronner and Haber (2014) proposed a segmentation of the cingulum bundle into distinct sections based on the presence of fibres connecting specific cingulate, non-cingulate frontal, and subcortical targets. Therefore, protocols for three distinct subsections of the cingulum bundle were created. 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-genual part of the cingulum bundle (CBP) was seeded anteriorly above the corpus callosum and a target placed below the sub-genual 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., 2013). 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. 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. We should point out that due to the relatively small size of the stria terminalis and its close proximity to the fornix, the fornix tracking may leak into the stria terminalis. This is a common issue in diffusion tractography and is yet to be overcome using approaches in line with those used in the current study (Mori and Aggarwal, 2014; Kamali et al., 2015; Mori et al., 2017; Pascalau et al., 2018).



#### **Projections Fibres**

<u>Corticospinal Tract (CST)</u>: 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 cerebral 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. 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.



<u>Acoustic Radiation (AR)</u>: 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 genic-

ulate 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 of 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 leaking into longitudinal fibres.



#### Adjustments for the Macaque Brain

Although the protocols described above are such that they allow for equivalent definitions in the macaque brain, some adjustments were required to ensure anatomical accuracy. For all macaque protocols, the reverse-seeding method was used, as this was found to increase robustness in the resulting tracts. In addition, the AF and MdLF protocols were adjusted to reflect the macaque brain. In the case of the AF, a seed is placed in the caudal STG, a target directly above the principal sulcus extending posterior to 8Ad (based on the tract-tracing data of Schmahmann and Pandya (2006)). In addition, a target placed in the caudal STG, immediately inferior and posterior to the seed ensured tracking occurred via caudal end of the lateral fissure. For the MdLF, a single axial plane in the posterior part of the STG was used as a target.

#### 4.2.2 Data and Preprocessing

To assess robustness across varying data quality, very high-quality diffusion MRI data from the HCP (Sotiropoulos et al., 2013a; Van Essen et al., 2013) (multi-shell data, 1.25 mm isotropic spatial resolution, 570 dMRI volumes, 1 hour of scanning) and data from the UK Biobank (Miller et al., 2016), which have overall quality

(2 mm isotropic resolution data, 100 dMRI volumes, 7 mins of scanning) closer to that typically available through clinical scanners, were utilised. For the HCP data 44 subjects with identified anatomical abnormalities were removed from the statistical comparisons and group atlases (see the HCP quality control website for details<sup>2</sup>), leaving a total of 1,021 subjects from the 1,065 subjects with diffusion data. For both the HCP and UK Biobank, the pre-processed dMRI data was used (see Sections 3.2 and 3.3 for details). Non-linear registration warp fields between the subject's diffusion space and the MNI152 standard space are provided by the HCP and UK Biobank for each subject. Details on their derivation are provided in Sections 3.2.4 and 3.3.4.

To ensure generalisability of the protocols across species, we also utilised diffusion MRI data from the macaque brain. This data consisted of an extended set of animals used in Eichert et al. (2020) and Mars et al. (2018b). Post-mortem data were acquired locally on a 7 T magnet with an Agilent DirectDrive console (Agilent Technologies, Santa Clara, CA, USA) using a 2D diffusion-weighted spin-echo protocol with single line readout (DW-SEMS, TE/TR: 25 ms/10 s; matrix size: 128 × 128; resolution:  $0.6 \times 0.6$  mm; number of slices: 128; slice thickness: 0.6 mm; diffusion data were acquired over the course of 53 hours). 16 non-diffusion-weighted (b = 0s/mm<sup>2</sup>) and 128 diffusion-weighted (b = 4000 s/mm<sup>2</sup>) volumes were acquired with diffusion directions distributed over the whole sphere. The brains were soaked in PBS before scanning and placed in fomblin or fluorinert during the scan. These data are available via PRIME-DE (Milham et al., 2018). Using FSL's FNIRT (Andersson et al., 2007; Jenkinson et al., 2012), estimations of the nonlinear transformations to standard space (F99) (Van Essen, 2002) were obtained based on the fractional anisotropy (FA) maps<sup>3</sup>.

<sup>&</sup>lt;sup>2</sup>HCP quality control website -

https://wiki.humanconnectome.org/pages/viewpage.action?pageId=88901591

<sup>&</sup>lt;sup>3</sup>Macaque registrations to the F99 atlas are performed using a custom configuration file. The F99 atlas volumes and surfaces are distributed within XTRACT, along with the custom configuration file.

# 4.2.3 Fibre Orientation Estimation and Tractography

Prior to XTRACT, the crossing fibre model described in Jbabdi et al. (2012), and previously discussed in Section 2.2.3, was applied to the diffusion data and used to estimate orientations to inform tractography. Up to three fibre orientations were estimated in each voxel along with their uncertainty. The XTRACT toolbox read the standard space tractography protocols and performed probabilistic tractography (Behrens et al., 2007; Hernandez-Fernandez et al., 2019). As discussed before, tractography protocols were defined for each bundle using a unique combination of seed, target, exclusion and stop masks, along with a seeding strategy (see Table 4.1).

A number of default tractography termination criteria were also used in all protocols (curvature threshold:  $\pm 80$  degrees, max streamline steps: 2000, subsidiary fibre volume threshold: 1%, randomly sampled initial fibres in case of fibre crossings in a seed location, no minimum length constraint, loop-checking and termination) (Behrens et al., 2003b; Behrens et al., 2007). A step size of 0.5 mm and 0.2 mm were used for human and macaque tractography respectively. The protocols described above have been developed using these tractography criteria and, therefore, deviation from these criteria may result in non-optimal performance. As shown in Figure 4.15, the masks were warped to the subject's native space and, after tractography, the results are directly resampled to standard space. The resultant distributions are normalised with respect to the total number of valid streamlines generated (i.e. streamlines that have not been rejected by inclusion/exclusion mask criteria).

# 4.2.4 Connectivity Blueprints

Tractography estimates discussed in the previous section focus mostly on the route of fibre bundles through deep WM, i.e. the main body of the tracts. A number of potential approaches exist for representing the cortical connectivity mediated by these bundles using structural connectomes (Sotiropoulos and Zalesky, 2017). One method of particular interest is the "connectivity blueprint", that allows us to estimate the "cortical termination"<sup>4</sup> map of each tract (Mars et al., 2018c).

This is achieved by first building a whole brain connectivity matrix, seeding from every white-grey matter boundary (WGB) location (with 2,000 streamlines initiated per seed) and counting the number of visitations to the whole WM<sup>5</sup>, giving a  $WGB \times WM$  connectivity  $\mathbf{C}_1$  matrix (exampled in Figure 3.4a). The tracts obtained using the tractography protocols (see Table 4.1)<sup>6</sup> were vectorised and concatenated into a single  $WM \times tracts \mathbf{C}_2$  matrix. Multiplying the two matrices provides the connectivity blueprint, i.e. a  $\mathbf{CB} = \mathbf{C}_1 \times \mathbf{C}_2$  ( $WGB \times tracts$ ) matrix. This process is shown schematically in Figure 4.2. Columns of this matrix represent the termination points of the corresponding tract on the WGB surface, while rows illustrate the connectivity pattern of each cortical location (i.e. how each tract contributes to the overall connectivity of each cortical location). This process was performed for the human (HCP) and macaque subjects, providing termination maps for all tracts across both species.

#### 4.2.5 Atlas Generation

For each HCP (1,065), UK Biobank (1,000) and macaque (6) subject, tractography as described above was used to extract estimates of 42 the WM tracts. For each cohort, WM tract atlases are produced. This was achieved by binarising the normalised WM tract reconstructions at a threshold value. These binarised tract masks were then cohort-averaged to produce, per tract, a distribution of the probability of a given tract being present for each voxel. In addition to the WM tract atlases, an atlas representing the average connectivity blueprint for the HCP and macaque cohort were generated. For both human cohorts, the effect of sample size on the generation of WM tract atlases was investigated. To do this, tract atlases were generated using 10, 100, 200, 500 and 1,000 subjects. In each case, each tract atlas was

<sup>&</sup>lt;sup>4</sup>Here, when referring to the cortical surface, this actually refers to the boundary surface between white matter and grey matter (as depicted in Figure 4.2). This is to overcome the limitations of diffusion MRI when used in grey matter.

 $<sup>^5\</sup>mathrm{The}$  WM volumetric data was downsampled to a resolution of 3 mm, see Appendix 4.6.4 for details.

<sup>&</sup>lt;sup>6</sup>The middle cerebellar peduncle (MCP) is excluded from this process as the cerebellum is not usually included in the cortical surface representations.



**Figure 4.2:** A schematic of the construction of connectivity blueprints. Tractography is seeded from the white-grey matter boundary (WGB), represented by the purple outline, and then counting the number of visitations to the whole white matter (WM), giving  $\mathbf{C}_1$  ( $N \sim 60k$  by  $M \sim 58k$ ). Columns of this matrix represent vertex-wise connectivity profiles. Next, the tractography reconstructions produced using XTRACT are vectorised and stacked to give a WM by tracts matrix,  $\mathbf{C}_2$  (M by T = 41). Multiplying the two matrices gives the connectivity blueprint,  $\mathbf{CB}$  (N by T). Columns represent maps of tract termination on the WGB surface; rows represent WGB connection profiles and reflect the contribution of each tract to the connection pattern of each WGB vertex.

correlated to the corresponding 1,000 subject atlas to asses how well lower sample sizes capture the population variance.

# 4.2.6 Protocol Evaluation and Validation Methods

#### **Robustness Against Datasets**

One of the aims of these standardised tractography protocols is to be generalisable across datasets. Section 4.2 describes how WM tract reconstructions are obtained from the HCP and UK Biobank. To explore robustness against varying data quality, reconstructions across the datasets were compared. First, the population percent tract atlases were compared by correlating tracts across datasets. Next, within- and across-cohort comparisons were made at the subject level by correlating random subject pairs. To avoid possible family structure-induced bias in the HCP, we restricted our subjects to the 339 unrelated subjects. To match this, the same number of gender-matched UK Biobank subjects were randomly chosen. Within-cohort correlations were performed by correlating a random HCP subject with another random HCP subject. The same was repeated for the UK Biobank cohort. Across-cohort comparisons were made by correlating a random HCP subject with a random UK Biobank subject.

#### Reflecting the Known Anatomy - Hemispheric Lateralisation

It is well known that certain functional brain regions are represented to a greater or lesser extent compared to their corresponding contralateral region, i.e. they are lateralised. For example, functional regions associated with language are typically lateralised (Hiscock and Kinsbourne, 2008). Given this functional lateralisation, it may be expected that the underlying structural architecture also be lateralised. Indeed, this is the case (Thiebaut de Schotten et al., 2011a; Menjot de Champfleur et al., 2013; Hecht et al., 2015; Nowell et al., 2016; Hau et al., 2016; Panesar et al., 2018; Howells et al., 2018) (see Table 4.3 for a summary).

#### WM Tract Lateralisation

Therefore, in order to demonstrate whether our protocols produce tracts representative of the anatomical expectations, tract lateralisation using a large number of subjects was investigated. We assessed tract lateralisation using tract volume. Specifically, lateralisation (L) was calculated as the relative right-left volume ( $V_R$ and  $V_L$ ) difference, after binarising the normalised tracts at 0.5% and taking the voxel count (Equation 4.1), in line with the literature on calculating tract lateralisation (O'Donnell et al., 2010; Propper et al., 2010; Thiebaut de Schotten et al., 2011b).

$$L = \frac{V_R - V_L}{V_R + V_L} \tag{4.1}$$
where  $V_R$  and  $V_L$  are the volumes for the right and left tracts respectively.

### Inter-Hemispheric Differences Using Connectivity Blueprints

Further to the WM tract lateralisation investigations, inter-hemispheric differences in the cortical termination maps were also explored. A connectivity blueprint  $\mathbf{CB}_L$ , including only the left-hemisphere tracts/columns and  $\mathbf{CB}_R$ , including only the right-hemisphere tracts/columns, were obtained. Both matrices were row-normalised so that the sum of all elements in each row was equal to 1. Subsequently, following similar approach to Mars et al. (2018c), the symmetric Kullback-Leibler divergence (KL divergence, Equation 4.2) between every pair of  $\mathbf{CB}_R$  and  $\mathbf{CB}_L$  rows was calculated (Kullback and Leibler, 1951).

$$\mathbf{KL}_{ij} = \sum_{t} \mathbf{CB}_{Rit} \log_2 \frac{\mathbf{CB}_{Rit}}{\mathbf{CB}_{Lit}} + \sum_{t} \mathbf{CB}_{Rjt} \log_2 \frac{\mathbf{CB}_{Rjt}}{\mathbf{CB}_{Lit}}$$
(4.2)

where  $\mathbf{KL}_{ij}$  is the KL divergence matrix, *i* indexes the vertices on the right cortex, *j* indexes the vertices on the left cortex and *t* indexes the tracts. KL divergence, or relative entropy, is a measure of statistical dissimilarity between discrete probability distributions where a value of zero indicates that two probabilities are identical and larger values indicate relative divergence. Using this approach allows us to probe the similarity in connectivity patterns to all fibre bundles between the connectivity blueprints. The right-left similarity in connectivity patterns in every hemispheric location *i* was obtained between all possible pairs using the minimum KL divergence value, i.e. min( $\mathbf{CB}_{Ri}, \mathbf{CB}_{Lj}$ ), with j spanning all *N* WGB locations. This process is shown schematically in Figure 4.3.

By selecting a vertex on a given hemisphere, the minimum KL divergence and the corresponding contralateral hemisphere vertex may be identified. It is expected that regions that are similar across hemispheres, i.e. have a relatively small minimum KL divergence, will have a contralateral vertex close, in distance, to the vertex of interest. To investigate this, vertices in regions of varying dissimilarity on the right hemisphere were selected and the contribution of the WM tracts to that vertex were compared to the corresponding contralateral (left hemisphere) minimum KL



**Figure 4.3:** Schematic of the calculation of Kullback-Leibler divergence. The connectivity blueprint is split into their respective matrices to give  $(\mathbf{CB}_L, \mathbf{CB}_R)$ , containing their respective tracts and the commissural tracts T with N vertices. i and j represent the vertices in two blueprints  $(\mathbf{CB}_R, \mathbf{CB}_L)$  respectively and t represents the tract number. i.e. the probability that a vertex i in  $\mathbf{CB}_R$  is connected to a tract t is given by the value in  $\mathbf{CB}_{Ri}$ . Entropy between each distribution  $(\mathbf{CB}_{Ri}, \mathbf{CB}_{Lj})$  is calculated for each element (i, j). The KL divergence equation provides a symmetric dissimilarity measure and dissimilarity may be represented as the vertex-wise minimum (along i for example).

divergence vertex.

### 4.2.7 Capturing Inter-Subject Variability

### **Respecting Similarities Stemming from Twinship**

Whilst we aimed for the automated tractography protocols to be robust against data quality, be reproducible and generalisable between species, it was further tested whether they could respect features stemming from the inherent individual variability in WM anatomy across subjects. To demonstrate this, the similarity of tract reconstructions within twin and non-twin sub-groups in the HCP cohort was explored. It is anticipated that monozygotic twin pairs will illustrate larger similarities than dizygotic twins and non-twin siblings, and subsequently than unrelated subject pairs, in line with the literature on the heritability of structural connections (Bohlken et al., 2014; Jansen et al., 2015) and as may be expected

from the literature on sulcal similarities in twinship (Amiez et al., 2019). Using the 72 pairs of monozygotic twins available in the HCP cohort, 72 randomly chosen pairs of dizygotic twins, 72 randomly chosen pairs of non-twin siblings and 72 randomly chosen pairs of unrelated subjects, correlations between tract reconstructions were performed and used to assess whether the automated protocols respect the underlying tract variability across individuals. We subsequently compared these distributions between the different groups.

### **Respecting Individual Differences due to Atypical Anatomy**

A subset of subjects with gross anatomical abnormalities were identified using the HCP quality control manual. These subjects were removed from statistical analysis and atlas generation, however, they afford the opportunity to investigate how XTRACT performs in the presence of small anatomical abnormalities. XTRACT was applied in the same manner as for all other subjects and the individually generated tracts were visually compared to cohort-averaged tracts. The cohort-averaged tracts used here are simple averages across subjects, i.e. not the population percent atlases - this allows for a consistent probability threshold to be applied in the individual and averaged tracts for visualisation.

# 4.3 Results

### 4.3.1 Atlases

#### Human and Macaque WM Tract Atlases

As described in Section 4.2.5, tract-specific atlases were generated for the HCP, UK Biobank and macaque subjects. Figure 4.4 provides maximum intensity projection representations of the WM tract atlases for each cohort.



(c) The macaque subjects (threshold = 0.5%)

(d) The macaque subjects (threshold = 0.1%, display range = 30%-100%)



### **Tract Termination Atlases**

Connectivity blueprints were also derived for each HCP and macaque subject and averaged to obtain atlases for the tract terminations. Examples of columns of the average connectivity blueprints are shown in Figure 4.5, representing atlases of termination points of each tract on the WGB surface.



**Figure 4.5:** white-grey matter boundary (WGB) endpoints for a subset of tracts (i.e. columns of the average connectivity blueprint) derived from the HCP and macaque cohorts.

### 4.3.2 The Effect of Sample Size on Atlas Generation

The effect of sample size on atlas creation was investigated in both human cohorts (restricted to 1,000 subjects in each case). Figure 4.6 shows the distributions of the tract-wise correlations (i.e. the average correlation across subjects per tract) for each of the sample size atlases. The top plot includes an atlas set with a sample size of 10 subjects which, whilst already showing high correlations (> 0.9), performs relatively poorly compared to using a sample size of 10 or greater (> 0.98), i.e. the average correlation obtained from a sample size of 10 is considerably lower than that obtained from a sample size of 100 or greater.



**Figure 4.6:** Summary of atlas sample size comparisons. The top plot includes the 10-subject atlas and the bottom plot shows the distributions without the 10-subject atlas to allow for improved visualisation of the other distributions.  $\sigma$  is the mean of the correlations across tracts and subject pairs. p is the p-value obtained from the Mann-Whitney U test, summarising the differences across comparisons.

### 4.3.3 Robustness Against Datasets

To explore robustness against varying data quality, tract atlases and inter-subject variability of the tract reconstructions within and across cohorts were compared. To compare atlases, each tract from the HCP atlas set was cross correlated with its corresponding UK Biobank tract atlas (population threshold of 30% applied to each tract atlas). The average correlation across tracts was 0.80 (standard deviation = 0.07).

Inter-subject correlations were obtained by cross correlating random subject pairs tract-wise (i.e. correlating the normalised path distributions in MNI space for each tract) and averaging the correlation across tracts for each subject pair. This was



**Figure 4.7:** Summary of inter-cohort robustness. Plots of the correlations between 339 subject pairs within and across cohorts.  $\sigma$  is the mean of the correlations across tracts and subject pairs. p is the p-value obtained from the Mann-Whitney U test, summarising the differences across comparisons.

repeated for many subject pairs within and across cohorts using the 339 unrelated HCP subjects and the gender-matched UK Biobank subjects. The across-cohort comparison was made by correlating a random subject in the unrelated HCP subject pool with a random UK Biobank subject, giving a distribution of 339 correlations per tract. Within-cohort comparisons gave average correlation values of 0.51 (standard deviation = 0.03) and 0.54 (standard deviation = 0.03) for the HCP and UK Biobank respectively ( $p = 8 \times 10^{-46}$ ) (Figure 4.7). Across-cohort comparison gave an average correlation of 0.41 (standard deviation = 0.03) which is lower than within-cohort comparisons ( $p = 1 \times 10^{-110}$  and  $2 \times 10^{-110}$ ), yet it is comparable enough, particularly given the differences in data quality and the age difference of subjects in the two cohorts (HCP: 22-35 years old, UK Biobank: 40-69 years old, mean age for our chosen subjects in HCP = 28.6 (standard deviation = 3.7) and UK Biobank = 62.6 (standard deviation = 7.5)).

Indeed, we have found that tract volumes are, on average, larger in the HCP cohort compared to the UK Biobank cohort  $(6.5\%\pm11.3\%$  larger in the HCP cohort), which is in line with the literature on age-related changes in WM volume (Westlye et al., 2010; Lebel et al., 2012; Rathee et al., 2016). Figure 4.8 shows the distributions of tract volumes for each tract reconstructed for the HCP and UK Biobank cohorts.



**Figure 4.8:** Comparisons of the cohort-averaged volumes of each tract for the HCP (green) and UK Biobank (orange). Volume is taken as the sum of non-zero voxels following the binarisation of the waytotal normalised tract density maps (threshold of 0.5%). Percentages indicate the percent difference between the average tract volume in the HCP cohort compared to the UK Biobank cohort, relative to the HCP cohort, i.e.  $100 \times \frac{V_{HCP}-V_{Bio}}{V_{HCP}}$ .

### 4.3.4 Reflecting Known Anatomy

### WM Tract Lateralisation

Lateralisation was assessed using both of the human cohorts. As shown in Figure 4.9, the AF is left-lateralised and IFO, MdLF and SLF3 are right-lateralised in both the HCP and the UK Biobank cohorts. SLF1 is symmetric in the HCP cohort but reaches rightward significance in the UK Biobank cohort. SLF2 is also variable across cohorts with left-lateralisation in the HCP and right-lateralisation in the UK Biobank.

### Inter-Hemispheric Differences Using Connectivity Blueprints

In addition to volume-based measures of lateralisation, inter-hemispheric differences on connectivity patterns were also assessed on the WGB surface using the tracts-



Figure 4.9: Summary of WM tract lateralisation for the arcuate fasciculus (AF), inferior fronto-occipital fasciculus (IFO), middle longitudinal fasciculus (MdLF) and the superior longitudinal fasciculi (SLFs) using the HCP (top) and UK Biobank (bottom) data. L is the cohort median WM tract lateralisation, p is the p-value obtained from the Mann-Whitney U test and  $\sigma$  is the variance for the given WM tract lateralisation. A threshold value of 0.5% has been used to binarise tracts and obtain their volume. Corrected p-value is 0.05/12 = 0.0042.

derived connectivity blueprints. As explained in Section 4.2.6, KL divergence was calculated to explore connectivity similarity between the two hemispheres of the human brain. For every location on the right hemisphere surface, the minimum KL divergence value assesses the most similar connectivity pattern on the left hemisphere. In doing so, we can probe cortical locations that demonstrate dissimilar connection patterns between left and right hemispheres and assess which tracts are contributing to these dissimilarities. In areas of high minimum KL divergence, we would expect to observe differences in the tract contribution profiles between the corresponding vertices. Figure 4.10a shows the minimum KL divergence values obtained for all WGB surface locations, overlaid with a subset of the Glasser parcellation (Glasser et al., 2016). From Figure 4.10a we can see that regions of the temporal junction and prefrontal/frontal cortices are dissimilar across hemispheres. Many of the regions highlighted are associated with language, a function known to be associated with lateralised brain regions (Hiscock and Kinsbourne, 2008).



Figure 4.10: a.) The minimum Kullback-Leibler divergence between the right and left hemispheres plotted on the surface of the right hemisphere. A subset of the Glasser parcellation (Glasser et al., 2016) is used to highlight regions of dissimilarity. The display range has been set to highlight regions with low- to mid-range dissimilarity. b.) A threshold of 0.3 has been applied to highlight regions of higher dissimilarity in the frontal cortex.

Figure 4.11 gives examples of the tract contribution to the connection pattern of specific white-grey matter boundary locations on the right hemisphere and how these compare with the connection patterns of the best matching location on the left hemisphere. In regions of high dissimilarity, the corresponding minimum KL divergence vertices drift apart and the underlying tract contributions diverge. For example, there is very little difference between the right and left profiles of the tract contribution to a vertex in the 4th visual area (V4) resulting in a very low minimum KL divergence.



**Figure 4.11:** Examples of the contribution of the WM tracts to vertex connectivity across the two hemispheres. The plots show the normalised tract contributions to the vertex on the right (blue line) and left (black line) hemispheres for the vertices indicated by the white and black circles respectively. The white circle shows the selected right hemisphere vertex. The black circle shows the corresponding minimum KL divergence left hemisphere vertex.

A small dissimilarity is also observed in a vertex in the 2nd visual area (V2), primarily driven by the IFO. An example vertex in the temporo-parietal-occopital junction area 1 (TPOJ1), an area associated with high-order functions such as language, visuo-spatial recognition and more (De Benedictis et al., 2014), highlights small differences in the contributions from AF and SLF3. The relatively low dissimilarity in a vertex in the anterior cingulate motor cortex (24dd), an area associated mostly with autonomic functions and motor control (Paus, 2001; Jumah and Dossani, 2019), is driven by minor differences in the contributions of the CBD, FA and STR.

Interhemispheric differences, primarily in the IFO, leads to a relatively small dissimilarity in an example vertex in the medial orbital frontal cortex (OFC): a relatively poorly understood area thought to be functionally involved in monitoring, learning and memory of the reward value of reinforcers (Kringelbach, 2005). The anterior ventral insular area (AVI), an area associated with affective processes (Uddin et al., 2017), shows greater dissimilarity driven by differences in the ATR, IFO, SLF3 and UF.

The inferior frontal sulcus (IFSa), an area associated with language processing and speech production (forming a part of Broca's area), and surrounding areas (45, p45r, IFSp, p9-46v and FOP5) show strong interhemispheric dissimilarity. An example vertex in the IFSa reveals that differences in the AF, FA and SLFs drive this dissimilarity. These examples highlight that, in general, in cases where dissimilarity is observed, the tracts contributing to said dissimilarity are those known to relate to lateralised functions.

Figure 4.10b presents the minimum KL divergence surface map after applying a threshold, highlighting the greater dissimilarity mostly confined the to frontal cortex. To further investigate the dissimilarity in the frontal cortex, the cohort averaged connectivity blueprint and minimum KL divergence map were parcellated using the Glasser parcellation (Glasser et al., 2016): in each case, the mean and standard deviation for a given parcel was calculated. Next, for each parcel in the frontal cortex, the mean minimum KL divergence is reported along with plots of the tract contribution to the given parcel. Figure 4.12 presents these results, demonstrating that interhemispheric dissimilarity is primarily driven by tracts known to be lateralised, for example, the right-lateralised SLF3.



**Figure 4.12:** Tract contributions for the regions highlighted in Figure 4.10b for the right (blue) and left (black) hemispheres, including regions in the inferior frontal sulcus (IFSa and IFSp), a central portion of the dorsolateral prefrontal cortex (pg-46v), area 45, the region posterior to area 47 (p47r), and an area in the frontal operculum (FOP5) parcels.

### 4.3.5 Capturing Inter-Subject Variability

#### **Respecting Similarities Stemming from Twinship**

To explore whether the tractography protocols preserved individual variability in WM anatomy, we compared tract reconstructions for 72 pairs of monozygotic twins, dizygotic twins, non-twin siblings and unrelated subjects from the HCP. Figure 4.13 shows the distributions of the average tract-wise correlations (i.e. average across tracts for each subject pair) for each twin/non-twin group. As shown, monozygotic twin pairs, on average (median), have a higher correlation (0.588, standard deviation = 0.036) with their corresponding twin compared to dizygotic twin pairs (0.545, standard deviation = 0.031), non-twin sibling pairs (0.543, standard deviation = 0.034), and unrelated subject pairs (0.507, standard deviation = 0.029). A Kruskal-Wallis test demonstrates statistically significant differences between subgroup medians:  $\chi^2 = 122.3$ ,  $p = 2 \times 10^{-26}$ . These results agree with a priori expectations stemming from the heritability of WM structure (Bohlken et al., 2014; Shen et al., 2014) and demonstrate that our protocols can preserve biological variability.



**Figure 4.13:** Twin/non-twin WM tract similarity using 72 subject pairs per group. Correlations are performed on normalised tract density maps with a threshold of 0.5%.  $\mu$  is the group median across tracts and subjects and  $\sigma$  is the standard deviation. A Kruskal-Wallis test is used to determine whether the groups come from the same median:  $\chi^2 = 122.3$ ,  $p = 2 \times 10^{-26}$ , mean ranks = 222.6 (monozygotic), 145.1 (dizygotic), 134.4 (non-twin siblings) and 69.9 (unrelated).

### Respecting Individual Differences due to Atypical Anatomy

In addition to exploring the similarity of tract reconstructions in twins, we further investigated whether the automated tractography respected individual variability in the case of anatomical abnormalities. A subset of subjects with gross anatomical abnormalities were identified using the HCP quality control. Figure 4.14 gives examples of anatomically atypical subjects and highlights the difference between the cohort-averaged tracts and the individual subject tractography results, which reflect the presence of cavernomas, cysts and developmental venous anomalies in WM. Although the abnormalities are small, the individual tractography results clearly deviate from the cohort-average results, suggesting that XTRACT does respect these abnormalities.

# 4.4 Discussion

This chapter presents a new toolbox (XTRACT) for automated probabilistic tractography along with standardised protocols for extracting white matter bundles in the human and the macaque brain. We have demonstrated that the protocols are robust when applied to data of varying image quality and to data from a non-human primate species. Human WM tract atlases using an order of magnitude more data than previous efforts, as well as macaque atlases using a small number of, however high-quality ex vivo, datasets have been generated. Indirect validation based on tract lateralisation has been performed, illustrating that reconstructed tracts are left/right asymmetric, when they are expected to be based on prior literature. It has also been demonstrated that despite automatically generating tracts using standard-space protocols, the protocols respect the underlying individual variability, as reflected in twinship-induced similarities and in respecting anatomical abnormalities. The toolbox, tractography protocols and atlases are freely and openly available as a part of FMRIB's software library (FSL) (version 6.0.2 and later).

A current issue in the field of tractography is that protocol definitions in the literature often lack detail or are designed without data-sharing in mind. Here,



(a) 1.) a cyst in the right parietal lobe affecting the MdLF. 2.) a small cavernoma in the brain stem affecting the MCP. 3.) a small cavernoma in the right parietal lobe affecting the CST. 4.) a developmental venous anomaly (DVA) in the right frontal lobe affecting the ATR. 5.) a cavernoma in the left occipital lobe affecting the OR.



(b) 6.) a small cavernoma in the right parietal lobe (same subject as Figure 4.14a, panel 2.) affecting the MdLF. 7.) a DVA in the right parietal lobe affecting the MdLF. 8.) a cavernoma in the left occipital lobe (same subject as Figure 4.14a, panel 5.) affecting the FMA and IFO.

Figure 4.14: Examples of tractography results for a subset of the subjects found to have anatomical abnormalities, demonstrating that cohort-averaged tracts do not respect underlying anatomical abnormalities. In each case the anatomical abnormality (as described on the HCP quality control website) is given and indicated by the white arrows, and affected tracts are presented. Both the individual subject's extracted tract (indicated by "S") and the corresponding cohort-averaged tract (indicated by "A") are overlaid on a zoomed in individual T1-weighted scan (the orange boxes show the zoomed regions). Tracts are displayed with a threshold of 0.1%.

we offer a platform for direct sharing of standardised protocol masks and tract atlases, allowing for the standardisation of tractography protocols across studies and aiding reproducibility. Moreover, we made the protocol definitions generalisable across species to directly facilitate comparative anatomy studies. Anatomists/tractographers may exchange their tract definitions/methods and hopefully converge on consensus protocols, as is the focus of a current multi-centre consortium on defining and standardising WM tractography (cf. Schilling et al. (2020b)). We believe that XTRACT can contribute to these efforts.

A platform for tractography protocol definitions is also presented in Wassermann et al. (2016), where a white matter query language (WMQL) is devised. Our approach uses a similar logic to WMQL, in the sense that it also relies on masks and Boolean operations on streamlines going through those masks. However, our protocols are generalisable and their utility in both the human and non-human primate brain as has been demonstrated here. Also, a main conceptual difference is that we don't rely on an automated grey matter (GM) parcellation, such as WMQL, but instead on hand-defined WM masks. WMQL relies on seeding in WM and defining cortical endpoints through subject-wise brain parcellations. Using this approach for tracking into GM has challenges; bottlenecks exist in the cortical gyrus, after which fibres fan out, and most tractography algorithms have issues resolving this fanning (Maier-Hein et al., 2017). To mitigate these issues, our protocols are not dependent on GM masks. We instead focus on reconstructing the main body of the tracts of interest using ROIs mostly in WM. In order to obtain cortical termination points of WM tracts, we take the opposite approach of tracking from the GM surface towards the WM, thus following the direction in which the fibres are expected to merge, rather than to fan out. We then multiply the surface-to-WM tractrogram with that of the body of the tract to create the WM by GM surface projections of the tract (connectivity blueprint). This avoids some of the major problems associated with tracking towards the surface (Mars et al., 2018c; Eichert et al., 2020).

We reconstructed tracts using imaging datasets of different quality and we generated atlases for both the HCP and the UK Biobank cohorts. Comparisons of tract reconstructions within and across the human cohorts demonstrate that the method and protocols are robust across subjects and against data quality. The HCP and UK Biobank cohorts provide examples of high-quality data and more typical quality data respectively. Within cohort comparisons reveal similar intersubject tract correlations across the varying quality data, with greater inter-subject tract correlations observed in the UK Biobank. This may reflect a reduced level of detail in the lower resolution UK Biobank data compared to the HCP data, but also differences in the mean age of subjects in the two cohorts. In addition, we have generated atlases using a smaller cohort of macaques. To compensate for the small number of subjects, we used high-quality and high-resolution ex vivo data. The respective results demonstrate the generalisability of our method to the macaque brain. Recent efforts to obtain macaque data from larger cohorts (HCP-style protocols) are ongoing (Milham et al., 2018; Autio et al., 2019) and the XTRACT toolbox will be a useful resource for these new initiatives for the non-human primate brain (Thiebaut de Schotten et al., 2019).

As a means to indirectly validate the presented results, left-right tract lateralisation was investigated. Lateralisation results were compared to a priori knowledge from the literature. For both human cohorts (HCP and UK Biobank), it was found that reconstructed AF is strongly left-lateralised, while SLF3, IFO and MdLF were right-lateralised, as expected from the literature (Thiebaut de Schotten et al., 2011a; Hecht et al., 2015; Nowell et al., 2016; Hau et al., 2016; Zhao et al., 2016; Panesar et al., 2018; Eichert et al., 2019). Results were less clear-cut for SLF1 and SLF2, where prior studies (with much fewer numbers of subjects) are inconclusive (Thiebaut de Schotten et al., 2011a; Hecht et al., 2015; Howells et al., 2018). This may be due to the large variance observed (in the case of SLF2), perhaps reflecting some underlying interaction, such as handedness (Howells et al., 2018).

Further sanity checks were performed by investigating lateralisation using the connectivity blueprint obtained from the reconstructed tracts. By using the KL divergence between connectivity patterns to assess inter-hemispheric dissimilarity, we identified that regions associated with lateralised functions, such as language (Hiscock and Kinsbourne, 2008), have dissimilar connectivity patterns across the two hemispheres (Figures 4.10 and 4.11).

Whilst being a robust automated method for the consistent reconstruction of tracts, XTRACT also respects the underlying anatomical variation. This was demonstrated by assessing inter-subject tract similarity in monozygotic twins, dizygotic twins, non-twin siblings and unrelated subject pairs. The results show greater similarity in twin pairs compared to unrelated pairs, as would be expected from the heritability literature (Bohlken et al., 2014; Shen et al., 2014).

We further demonstrated that the automated method respects underlying anatomical variation by exploring how tractography results differ from the cohort-averaged results in the case of subjects with anatomical abnormalities. In the presented cases, we show that the toolbox is capable of respecting atypical anatomy. These example cases are of course not exhaustive but offer insight into how the toolbox performs in the presence of relatively small anatomical abnormalities, induced by pathology. We should point out however that our approach is registration-based, therefore tracking performance in the presence of geometrical malformations is likely to depend on the extent and location of the abnormality and the influence it has to registration (Zhang et al., 2008). Smaller malformations, such as focal WM abnormalities/hyperintensities, are less likely to reduce reliability, particularly given the use of inclusive ROIs in our protocol definitions, as demonstrated in Figure 4.14 and in agreement with similar findings in recent studies that used tract-derived features (Ressel et al., 2018; Horbruegger et al., 2019). For larger malformations (such as tumours, oedema), even if some compensation can be achieved by performing conditional registration (i.e. by masking out large malformations in the computation of the warp fields), reductions in tracking accuracy may occur and case-specific alternatives may need to be considered (Fekonja et al., 2019). Nevertheless, a range of pathologies do not induce geometrical malformations (for instance the spectrum of psychiatric/neurodevelopmental/mental health disorders) and we expect our approach for delineating major WM tracts to be robust in such cases.

As mentioned, XTRACT and its success is reliant on image registration, specifically on the accuracy of registration warp fields between a subject's native diffusion space and the standard (i.e. protocol) space. That is, in order to ensure successful tract reconstruction, the anatomical priors must be accurately translated to the subject's native space. Non-linear registration between the native and standard space is therefore essential to ensure the matching of anatomy at the local level (e.g. matching gyri geometry). Such registration techniques are mostly accurate (Klein et al., 2009; Groot et al., 2013; Ou et al., 2014; Andersson et al., 2019), particularly when considering the use of inclusive ROIs in our protocol definitions, which alleviates minor inaccuracies in registration. Although we do not provide an evaluation of XTRACT tract reconstruction success in relation to registration accuracy, XTRACT has been applied to thousands of datasets and appears robust in the absence of significant registration errors.

# 4.5 Concluding Remarks

The development and validation of a set of robust and standardised tractography protocols for the automated cross-species delineation of white matter bundles, along with a platform to use them, has been presented. The demonstrated toolbox (XTRACT) is freely available along with the tractography protocols and human tract atlases as a part of FMRIB's software library (FSL version 6.0.2 and later). Given the benefits with regards to data and protocol sharing, we expect that this toolbox will aid reproducibility in the field of tractography and facilitate comparative neuroanatomy studies.

The framework and connectivity features presented here are utilised in Chapters 5 and 6. In Chapter 5, the XTRACT tracts and connectivity blueprints from the HCP and macaque subjects are used to explore cross-species prediction of myelination maps, linking closely to functional specialisation. These are then compared to a newly developed framework for extracting homologous connectivity patterns. In Chapter 6, the connectivity blueprints derived from the HCP cohort (Figure 4.5) are used as structural connectivity features in explorations of the stability of multivariate statistical techniques.

# Contributions

This chapter describes the development and validation of the XTRACT toolbox along with a library of standardised protocols for the human and macaque brain. The works presented were a collaborative effort that I led. I would like to extend my gratitude to all contributors. The XTRACT toolbox was written and developed by me, Stamatios Sotiropoulos and Saad Jbabdi. Macaque data were acquired by Jerome Sallet, Rogier Mars and Alexandr A. Khrapitchev. Protocol masks were developed by me, Gwenaëlle Douaud and Stamatios Sotiropoulos (led the human protocols), and Rogier Mars, Katherine Bryant (led the macaque protocols). My contributions cover all aspects of data processing, analyses and visualisations and further protocol refinement and assessment (spanning both species).

# Software

XTRACT and the associated tractography protocols were developed using FMRIB's Software Library (FSL, v5.0.11 onwards) (Jenkinson et al., 2012). Statistical analyses and plotting were performed in MATLAB 2018a (The MathWorks, Inc., Natick, Massachusetts, United States) and Python (version 3.7.2) (Python Software Foundation, https://www.python.org). Volume-space and surface-space visualisations were performed using FSL's FSLeyes and Connectome Workbench (version 1.3.2; Van Essen Laboratory, Washington University) respectively.

# 4.6 Appendix

# 4.6.1 Summary of the Tractography Protocols Previously Defined in the Literature

	Wakana <sup>1</sup>	Catani <sup>2</sup>	Hua <sup>3</sup>	Zhang <sup>4</sup>	Thiebaut de	Yendiki <sup>6</sup>	de	Wassermann <sup>8</sup>
					Schotten <sup>5</sup>		Groot <sup>7</sup>	
	N=4	N=12	N=28	N=10	N=40	N=67	N=60	N=97
Acoustic Radiation							~	
Anterior Commis-		~			1			
sure								
Anterior Thalamic	✓		√	✓		✓	√	
Radiation								
Arcuate Fasciculus		~			✓*			~
Cingulum subsec-	✓		$\checkmark$	~		1	~	
tion: Dorsal		$\checkmark$			~			$\checkmark$
Cingulum subsec-	✓		~	~		~	√	
tion: Peri-genual								
Cingulum subsec-								
tion: Temporal								
Corpus Callosum								√§
Corpus Callosum:	$\checkmark$	$\checkmark$	$\checkmark$	✓	<b>↓</b>	$\checkmark$	$\checkmark$	$\checkmark$
Forceps Major								
Corpus Callosum:	$\checkmark$		$\checkmark$	√		~	$\checkmark$	$\checkmark$
Forceps Minor								
Cortico-Ponto-					~			
Cerebellar								
Corticospinal Tract	√		$\checkmark$	~	~	✓	√	✓
Extreme Capsule								✓
Fornix		✓			~			
Inferior Cerebellar		~						
Peduncle								
Inferior Fronto-	$\checkmark$	$\checkmark$	$\checkmark$	~	~		$\checkmark$	$\checkmark$
Occipital Fascicu-								
lus								
Inferior Longitudi-	$\checkmark$	$\checkmark$	$\checkmark$	√	~	$\checkmark$	$\checkmark$	$\checkmark$
nal Fasciculus								
Internal Capsule/		$\checkmark$			✓			
Corona Radiata								
Medial Lemniscus							√	
Middle Cerebellar		$\checkmark$					$\checkmark$	
Peduncle								
Middle Longitudi-								$\checkmark$
nal Fasciculus								
Optic Radiation/					✓		<b>√</b>	$\checkmark$
Posterior Thalamic								
Radiation								
Spino-Cerebellar					✓			
Tract								
Striato-Fronto-								$\checkmark$
Orbital								

Striato-Prefrontal								$\checkmark$
Striato-Premotor								$\checkmark$
Striato-Precentral								$\checkmark$
Striato-Postcentral								$\checkmark$
Striato-Parietal								$\checkmark$
Striato-Occipital								~
Superior Cerebellar		$\checkmark$			✓			
Peduncle								
Superior Longitudi-								1
nal Fasciculus I	$\checkmark^+$		$\checkmark^+$	✓+		$\checkmark^+$	$\checkmark$	
Superior Longitudi-								~
nal Fasciculus II								
Superior Longitudi-								~
nal Fasciculus III								
Superior Thalamic							$\checkmark$	
Radiation								
Thalamo-								×
Prefrontal								
Thalamo-Premotor								$\checkmark$
Thalamo-								✓
Precentral								
Thalamo-								×
Postcentral								
Thalamo-Parietal								✓
Thalamo-Occipital								✓
Uncinate Fasciculus	~	~	~	~	~	1	~	1
Total:	20	19	22	22	31	18	27	57

Table 4.2: A brief review of the protocols previously defined in the literature. \* - Arcuate and anterior, long and posterior segments. <sup>§</sup> - Rostrum, rostral body, anterior midbody, posterior midbody, isthmus.  $^+$  - SLF and the temporal component of the SLF.

<sup>1</sup>Wakana, S., et al., 2004. Fiber Tract–based Atlas of Human White Matter Anatomy. Radiology. <sup>2</sup>Catani, M. & Thiebaut de Schotten, M., 2008. A diffusion tensor imaging tractography atlas for virtual in vivo dissections. Cortex.

<sup>3</sup>Hua, K., et al., 2008. Tract probability maps in stereotaxic spaces: Analysis of white matter anatomy and tract-specific quantification. NeuroImage.

 $^4{\rm Zhang},$  W., et al., 2008. Automated fiber tracking of human brain white matter using diffusion imaging. NeuroImage.

<sup>5</sup>Thiebaut de Schotten, M., et al., 2011. Atlasing location, asymmetry and inter-subject variability of white matter tracts in the human brain with MR diffusion tractography. NeuroImage.

<sup>6</sup>Yendiki, A., Panneck, P., Srinivasan, P., Stevens, A., Zöllei, L., Augustinack, J., et al., 2011. Automated probabilistic reconstruction of white-matter pathways in health and disease using an atlas of the underlying anatomy. Front. Neuroinform.

<sup>7</sup>de Groot, M., et al., 2013. Improving alignment in Tract-based spatial statistics: Evaluation and optimization of image registration. NeuroImage.

<sup>8</sup>Wassermann, D., et al., 2016. The white matter query language: a novel approach for describing human white matter anatomy. Brain Struct. Funct.

### 4.6.2 XTRACT User-Guide

XTRACT is a command line tool available within FSL which can be used to automatically extract a set of carefully dissected tracts in humans and macaques. It can also be used to define one's own tractography protocols where all the user needs to do is to define a set of masks in standard space (e.g. MNI152). XTRACT reads the standard space protocols and performs probabilistic tractography (probtrackx2) in the subject's native space. Resultant tracts may be stored in either the subject's native space or in standard space. The user must provide the crossing fibres fitted data (bedpostx) and diffusion to standard space registration warp fields (and their inverse).

Jsage: xtract -bpx <bedpostx_dir> -out <out xtract -bpx <bedpostx_dir> -out <out xtract -bpx <bedpostx_dir> -out <out< td=""><td>putDir&gt; -str <structuresfile> -p <protocolsfolder> [options] putDir&gt; -species HUMAN [options] putDir&gt; -species MACAQUE [options]</protocolsfolder></structuresfile></td></out<></bedpostx_dir></out </bedpostx_dir></out </bedpostx_dir>	putDir> -str <structuresfile> -p <protocolsfolder> [options] putDir&gt; -species HUMAN [options] putDir&gt; -species MACAQUE [options]</protocolsfolder></structuresfile>
Compulsory arguments:	
-bpx <folder> -out <folder></folder></folder>	Path to bedpostx folder Path to output folder
And EITHER: -str <file> -p <folder></folder></file>	Structures file Protocols folder
Or: -species <species></species>	One of HUMAN or MACAQUE
Optional arguments:	
-stdwarp <std2diff> <diff2std> -gpu -native -res <mm></mm></diff2std></std2diff>	Standard2diff and Diff2standard transforms Use GPU version Run tractography in native (diffusion) space Output resolution
-ptx_options <options.txt></options.txt>	Pass extra probtrackx2 options

Figure 4.15: The XTRACT command-line user interface.

### Running XTRACT

In the most basic implementation of XTRACT, you must provide a subject's crossing fibres fitted data diffusion and the standard space registration warp fields (and their inverse). In the terminal window:

xtract -bpx dMRI.bedpostX -stdwarp std2diff.nii.gz diff2std.nii.gz -out xtract -species HUMAN

This command call would run XTRACT using its default settings for all 42 human brain tractography protocols (Warrington et al., 2020). The option of GPU acceleration may be used by adding the "-gpu" flag. Additional tractography options (probtrackx options) and a specified tract resolution may also be supplied to XTRACT using the "-ptx\_options" and "-res" arguments. In the most recent release, tracts may be stored in any reference space by providing the warp fields between the protocol space, the reference space and the diffusion space using the "-stdwarp <std2diff> <diff2std>" and "-ref <refimage> <diff2ref> <ref2diff>" arguments.

### **Outputs of XTRACT**

XTRACT will produce a subdirectory for each tract containing the resultant tract image and some other useful information. They key files are summarised below. Output directory structure:

- "logs" directory containing the probtrackx log files
- "tracts" directory containing tractography results
  - "<tractName>" directory per tract, each containing:
    - \* "waytotal" txt file containing the number of valid streamlines
    - \* "densityNorm.nii.gz" NIFTI file containing the waytotal normalised fibre probability distribution. This is the main output of XTRACT.

### Pre-processing

Prior to running XTRACT, you must complete the FDT processing pipeline (see FSL wiki page: FDT User Guide), including

- 1. Brain extraction using BET
- 2. Susceptibility distortion correction using topup (only if spin-echo fieldmaps have been acquired if you don't have these, skip to step 3)
- 3. Eddy current distortion and motion correction using eddy
- 4. Fit the crossing fibre model using bedpostx
- 5. Non-linear registration (FNIRT) to standard space (MNI152 for human data)

### Adding your own tracts

Users may define their own tractography protocols following a similar design to those provided in XTRACT and use XTRACT to reconstruct said tract. Protocols consist of a combination of seed, target/waypoint, stop and exclusion masks. Seed masks provide the starting points of the tractography streamlines. Target/waypoint masks are regions through which a streamline should pass in order to be valid. Exclusion masks serve to reject any streamline running through them and stop/termination masks stop any streamline running through them. The minimum requirement in the construction of a protocol is to define a seed mask, however, if the intention is to reconstruct a well-defined fibre bundle, it is more than likely that the protocol will need to contain at least one target/waypoint mask and exclusion mask to prevent leakage into neighbouring fibre bundles.

Suppose you want to create an automated protocol for a tract called 'mytrack'. First you need to create a folder called 'mytrack' which you can add e.g. in the protocols folder. Then create the following NIFTI files (with this exact naming) and copy them into 'mytrack':

### Compulsory:

• seed.nii.gz : a seed mask

### Optional:

- stop.nii.gz : a stop mask
- exclude.nii.gz : an exclusion mask
- one of the following:
  - target.nii.gz : a single target mask
  - target1.nii.gz, target2.nii.gz, etc. : a number of targets, in which case streamlines will be kept if they cross all of them
- invert : an empty file to indicate that a seed->target and target->seed run (i.e. the inverse seeding strategy) will be added and combined. If such an option is required, a single "target.nii.gz" file is also expected.

The latest version of XTRACT allows the user to define these protocols in any space, as long as the registration warp field is to the same space. Next, make a structure file using the format <tractName> <nsamples> per line and call XTRACT using -species <SPECIES> -str <file> -p <folder>, pointing to your new protocols folder 'mytrack'.

Tract	Thiebaut de Schotten <sup>1</sup>	Menjot de Champfleur <sup>2</sup>	Hecht <sup>3</sup>	Nowell <sup>4</sup>	Hau <sup>5</sup>	Panesar <sup>6</sup>	Howells <sup>7</sup>
	N=20	N=4	N=64	N=57	N=60	N = 30	N=51
AF	-	Left (FA), Right (#s) <sup>§</sup>	-	Left (vol)*	-	-	-
ILF	-	Left (FA, #s) <sup>§</sup>	-	-	-	Left (vol)	-
IFO	-	Left (FA), Right (#s) <sup>§</sup>	-	-	Right (vol)	-	-
MdLF	-	Left (FA), Right (#s) <sup>§</sup>	-	-	-	-	-
SLF1	Sym (vol)	-	-	Sym (vol)	-	-	Left for right-handed, right for left-handed (vol)
SLF2	Right trend (vol)*	Right trend (vol)*	-	-	-	-	Right for right-handed*, Sym for left-handed (vol)
SLF3	Right (vol)	Right (vol)	-	-	-	-	Right for right- and left-handed (vol)

### 4.6.3 Summary of the Tract Lateralisation Literature

**Table 4.3:** Summary of WM tract lateralisation as reported in the literature. Sym = symmetric, #s = total number of streamlines generated, vol = voxel count following binarisation, FA = mean fractional anisotropy. \* reported trend is not significant. § significance is not reported.

<sup>1</sup>Thiebaut de Schotten, M., et al., 2011. Atlasing location, asymmetry and inter-subject variability of white matter tracts in the human brain with MR diffusion tractography. NeuroImage.

<sup>2</sup>Menjot de Champfleur, N., et al., 2013. Middle longitudinal fasciculus delineation within language pathways: a diffusion tensor imaging study in human. Eur. J. Radiol. 82, 151–157.

<sup>3</sup>Hecht, E.E., et al., 2015. Virtual dissection and comparative connectivity of the superior longitudinal fasciculus in chimpanzees and humans. Neuroimage 108, 124–137.

<sup>4</sup>Nowell, M., et al., 2016. Meyer's loop asymmetry and language lateralisation in epilepsy. J. Neurol. Neurosurg. Psychiatry 87, 836–842.

<sup>5</sup>Hua, K., et al., 2008. Tract probability maps in stereotaxic spaces: Analysis of white matter anatomy and tract-specific quantification. NeuroImage.

<sup>6</sup>Panesar, S.S., et al., 2018. A quantitative tractography study into the connectivity, segmentation and laterality of the human inferior longitudinal fasciculus. Front. Neuroanat. 12.

<sup>7</sup>Howells, H., et al., 2018. Frontoparietal tracts linked to lateralized hand preference and manual specialization. Cerebr. Cortex 28, 2482–2494.

### 4.6.4 The Effect of Resolution on Connectivity Blueprints

Connectivity blueprints were reconstructed for a subset of 5 HCP subjects using varying WM mask resolutions: 2 mm, 3 mm, and 4 mm. Figure 4.16 shows examples of minor differences across these resolutions in a subset of tracts for one of those subjects. The connectivity blueprint columns (i.e. tracts) were compared, via Pearson's correlation, across resolutions to the 2 mm case. The average correlation (across tracts and subjects) between the 3 mm and 2 mm results was  $0.993 \pm 0.0008$  and between the 4 mm and the 2 mm results was  $0.986 \pm 0.0019$ . In other words, the resolution, within these ranges, makes little difference to the final result, however, a lower resolution to 3 mm saves 10 minutes of processing time and 4800 MB of storage space per subject compared to the 2 mm case. Over the entire HCP cohort, this saves approximately a week of processing time and 5 TB of storage space.



**Figure 4.16:** Comparisons of the connectivity blueprints across WM mask resolutions for a subset of tracts, including the acoustic radiation (AR) (left column), inferior fronto-occipital fasciculus (IFO) (middle column) and superior longitudinal fasciculus (SLF) 3 (right column) for three different resolutions: 2 mm (top row), 3 mm (middle row), and 4 mm (bottom row).

# Chapter 5

# Building and Mapping Onto Common Connectivity Spaces

### Contents

5.1	Intro	duction
5.2	Cross	s-species Mapping onto Connectivity Spaces 128
	5.2.1	Connectivity Blueprints
	5.2.2	Joint-Spectral Embedding
	5.2.3	Projecting Myelin Maps Across Species
	5.2.4	Data
	5.2.5	Results
<b>5.3</b>	Build	ing Connectivity Spaces using Data-Driven Methods135
	5.3.1	Myelin Map Predictions
	5.3.2	Data
	5.3.3	Results
<b>5.4</b>	Discu	ussion
<b>5.5</b>	Conc	luding Remarks 146
<b>5.6</b>	Appe	endix
	5.6.1	Landmark regions-of-interest
	5.6.2	Landmark connectivity

# Foreword

There is a strong yet complex relationship between the structural and functional architecture of the brain, and the precise nature of this relationship is yet to be resolved. However, the idea of connectivity fingerprinting has been a basic tenet in neuroscience and computational neuroanatomy; the pattern of extrinsic connections of a brain region is linked to its functional specialisation and is therefore unique. Here we extend and test this idea for probing regions of functional equivalence across diverse brains using connections, building upon the cross-species developments from the previous chapter.

Firstly, we use the XTRACT protocols developed in the previous chapter to define connectivity patterns of cortical areas to white matter bundles. And we use similarity in these patterns across diverse brains to probe functional similarity. Given that the white matter bundles are defined analogously in both humans and monkeys, we treat them as a "common reference space" and we anticipate that functionallyhomologous areas across diverse brains will carry similar connection patterns to these bundles, i.e. will exhibit similarity in that common space. Specifically, we explore and compare the efficacy of two recently-introduced approaches for utilising such common spaces at a whole-brain level. We map diverse brains onto the spaces and use them to link brains by projecting scalar feature maps between them.

Secondly, we explore alternative ways for building such a common connectivity space. Alternatively to manually-defined tractography protocols for a limited set of major white matter bundles (as in XTRACT), more recent data-driven approaches can be used for a more complete representation of connectivity patterns. As such, we apply non-negative matrix factorisation to tractography data to extract structural connectivity components and use these to define connectivity patterns and probe functionally-similar features.

# 5.1 Introduction

The functional specialisation and organisation of the brain is underpinned by a complex network architecture, mediated by long-range connections through white matter (WM) (Passingham et al., 2002; Jbabdi and Behrens, 2012; Mars et al., 2018b). Diffusion tractography is a powerful and unique technique which may be used to probe these structural connections in vivo and non-invasively. With the application of well-informed priors, tractography can be robust and anatomically accurate (Wakana et al., 2007; Groot et al., 2013; Maier-Hein et al., 2017; Schilling et al., 2020a). Additionally, connectivity patterns derived from diffusion tractography can be relevant for probing the brain's functional organisation.

The set of extrinsic connections of a given cortical region to other brain regions is unique and directly linked to the functional role and specialisation of that region (Passingham et al., 2002; Jbabdi and Behrens, 2012). Often referred to as the connectivity fingerprint, it has been demonstrated to be predictive of functional boundaries between cortical regions (Johansen-Berg et al., 2004; Beckmann et al., 2009; Mars et al., 2011; Cerliani et al., 2012; Thiebaut de Schotten et al., 2014b) and subcortical volumes (Behrens et al., 2003a; Draganski et al., 2008; Tziortzi et al., 2014). These patterns have been also shown to be predictive of the shape and location of stimulus-induced (task) functional activation (Saygin et al., 2012; Osher et al., 2016; Saygin et al., 2016).

These principles hold and can be extended even in the case of very diverse brains, such as across species (Vincent et al., 2007; Sallet et al., 2013; Neubert et al., 2014; Neubert et al., 2015; Mars et al., 2016b; Mars et al., 2018c; Mars et al., 2018a; Eichert et al., 2020; Benn et al., 2020). Mars et al. (2018c) introduced the idea of an abstract space of structural connectivity patterns (through WM), which can be used to project, link and compare different brains and grey matter (GM) regions. A key benefit of this "common connectivity space" approach is that similarity in brain organisation can be probed even in the absence of any geometrical similarity. In the Mars et al. (2018c) study, the pattern of connections were defined through homologous WM fibre bundles across two species (humans and macaques), allowing functional resemblance to be expressed through these patterns, i.e. using these homologous WM fibre bundles as reference. Statistical measures can then be used to assess similarity, e.g. the Kullback-Leibler divergence. A number of more recent studies have mapped different brains to that common space of connectivity patterns and have used similarity in these patterns as a proxy of functional similarity (Mars et al., 2018c; Mars et al., 2018a; Eichert et al., 2020; Benn et al., 2020). Extending this to the whole-brain provides a way to project and link functionally-relevant features across very diverse brains, for instance, projecting scalar cortical myelin maps between mammalian brains using connections (Mars et al., 2018c; Mars et al., 2018a; Eichert et al., 2020; Benn et al., 2020).

In this chapter we evaluate the XTRACT protocols and demonstrate their ability to be used for mapping human and monkey brains into the common space, defined by Mars et al. (2018c). We also explore and compare to a different recent approach for mapping diverse brains, based on joint-spectral embedding. Xu et al. (2020) applied this approach to cross-species resting-state fMRI data. They defined functional connectivity in relation to homologous landmarks in GM (i.e. a finite set of cross-species homologous cortical regions) and extracted homologous patterns of functional connectivity in a common space, allowing for comparison. These were then used to perform multi-modal surface matching (MSM) (Robinson et al., 2014) and identify warp fields driven by (and therefore aligning) these joint components. They then used these MSM warps to project the human cortical myelin map to the macaque cortex with strong predictive power. Here, we apply this idea to structural patterns (as derived from XTRACT) and compare it with the Mars et al. (2018c) approach.

Further, apart from exploring options for mapping into this common space, we also devise new ways for building such a space. To overcome the limitations arising from the limited set of manually-defined homologous WM tracts or GM landmarks as in the above approaches, we explore a recent fully data-driven approach (Thompson et al., 2020) for mapping structural connectivity components and use those to define connectivity patterns. This approach uses non-negative matrix factorisation (NMF) to decompose whole-brain tractography data to GM networks and their underlying WM connections. It provides an alternative to a limited set of manually-defined tractography protocols and may therefore provide a more complete representation of connectivity patterns. As a proof of concept, we use NMF to extract components of structural connectivity and demonstrate that these measures may also be used across brains from the HCP to define the common connectivity space of Mars et al. (2018c) in a data-driven manner.

To summarise, in this chapter we devise and explore ways for building and mapping connectivity spaces. We first define a connectivity space using a finite set of reference bundles (XTRACT protocols) and further project these to a new common space (Xu et al., 2020), and compare two approaches (Mars et al., 2018c; Xu et al., 2020) for mapping diverse brains into these spaces. As an exemplar, we utilise very different brains from different primate species (humans and macaque monkeys). Subsequently, we demonstrate an alternative way of building such connectivity patterns in a completely data-driven way, using NMF rather than specific tractography protocols. We show that the more complete representation of patterns obtained through the NMF approach offers improved predictive power of functionally-relevant features (cortical myelin maps) from structural connections, providing a different way for defining connectivity spaces.

# 5.2 Cross-species Mapping onto Connectivity Spaces

We explore and assess two approaches for constructing and mapping onto connectivity spaces across two species: humans and macaque monkeys. We use the white matter (WM) bundles defined homologously in the previous chapter to define connectivity blueprints (patterns of connections through WM for each cortical location) for different brains. We then use two different approaches for linking brains through these features, a) a statistical similarity of the connectivity patterns, as described in (Mars et al., 2018c) and b) a joint-embedding (i.e. non-linear similarity) approach (Xu et al., 2020), driven by the connectivity blueprints.

We assess the efficacy of these two approaches in functionally aligning diverse brains by their success of projecting scalar features maps across species. As an example, we demonstrate projections of cortical myelin maps (Glasser and Van Essen, 2011; Glasser et al., 2014) from humans to macaques through the connectivity spaces. Myelin maps are proxies of cortical hierarchy (Demirtaş et al., 2019), as heavily myelinated areas correspond to sensory areas and less myelinated areas to higher-order association regions. Below, we present the approaches utilised here and compared: 1) the connectivity blueprints derived through XTRACT (Mars et al., 2018c) and 2) the joint-embedding approach introduced by Xu et al. (2020), driven by the connectivity blueprints.

### 5.2.1 Connectivity Blueprints

Connectivity blueprints were constructed independently in each species by first building whole-brain connectivity matrices, seeding from every white-grey matter boundary (WGB) location and counting the number of visitations to the whole WM, giving the  $WGB \times WM$  connectivity matrix for each species. XTRACT tracts (*T* number of tracts, see Table 4.1) were estimated for each species and subsequently vectorised and concatenated into a single  $WM \times tracts$  matrix. The connectivity blueprint ( $WGB \times tracts$ ) was then obtained through the multiplication of these matrices (see Figure 4.2 in Chapter 4). Columns of the connectivity blueprint represent the distribution of cortical terminations for each tract and each row (normalised so that row-sum is 1) represents the connectivity pattern for a given cortical location considering each of the tracts. This was performed for a number of human and macaque subjects and then averaged within species, providing connectivity blueprints for both (**H** and **M**).

### 5.2.2 Joint-Spectral Embedding

The connectivity blueprints ( $\mathbf{H}$  and  $\mathbf{M}$ ) were further mapped through joint-spectral embedding as in Xu et al. (2020). Spectral embedding seeks to remove redundancies in data by projecting the data onto a low-dimensional non-linear manifold, resulting in series of independent components. By performing joint-embedding on both species simultaneously using homologous landmarks, components remain inde-
pendent within species (i.e. they represent independent modes of variation across components) but are linked across species (i.e. corresponding components represent equivalent modes of variation across species).

In order to ensure like-for-like comparison of connectivity patterns across species, we used a set of homologous cortical landmarks (Xu et al., 2020). These consisted of l = 27 cortical regions, functionally homologous in the human and macaque brain (see Appendix 5.6.1). For each of these landmarks, the average connectivity blueprint pattern across the constituent vertices was calculated, giving a landmark connectivity matrix for each species  $\mathbf{L}_H$  and  $\mathbf{L}_M$  with l = 27 landmarks (rows) and T tracts (columns). Each row (i.e. vertex) of the connectivity blueprints ( $\mathbf{H}$  and **M**) were then compared through cosine similarity to each of the landmark profiles  $(\mathbf{L}_H \text{ and } \mathbf{L}_M)$ , giving a vertex by landmark similarity matrix for each species  $(\mathbf{H}_L)$ and  $\mathbf{M}_L$ ). These two matrices were concatenated row-wise and the joint 2N \* lmatrix was cross-correlated (via non-linear cosine similarity) to produce a 2N \* 2Njoint similarity matrix (Figure 5.1). In other words, for every vertex, the similarity with every other vertex was calculated for the human and macaque brain, within (H2H and M2M) and across species (H2M and M2H). Each quadrant of this joint-similarity matrix was globally demeaned and the joint-similarity matrix was decomposed using spectral embedding (Belkin and Nivogi, 2002), giving a set of low-dimensional components ( $\mathbf{H}_e$  and  $\mathbf{M}_e$ ). These embedding steps were performed separately for each hemisphere.

The joint-similarity matrix used is symmetric across species, with a the diagonal containing within species entries which provide the overall structure for the components, i.e. they provide the component correspondence across species. The off-diagonal entries, informed by the homologous landmarks, capture the acrossspecies coupling. Therefore, these components represent, jointly across the two species, whole-brain maps of regions that exhibit similar connectivity patterns (as captured by the blueprints) to the patterns of homologous cortical landmarks.

#### 5.2.3 Projecting Myelin Maps Across Species

In order to evaluate and example these connectivity spaces, we compared them across species by projecting myelin maps from the human brain to the macaque brain using two approaches. First, following Mars et al. (2018c), we applied Kullback-Leibler divergence (KL divergence) to the connectivity blueprints. Here, KL divergence was applied to calculate the dissimilarity between the macaque and human structural connectivity features and subsequently used to project the human myelin map to the macaque brain. The KL divergence between the connectivity blueprint of vertex i in the macaque brain and that of vertex j in the human brain was first calculated using Equation 5.1 (Kullback and Leibler, 1951):

$$\mathbf{KL}_{ij} = \sum_{t} \mathbf{M}_{it} \log_2 \frac{\mathbf{M}_{it}}{\mathbf{H}_{it}} + \sum_{k} \mathbf{M}_{jt} \log_2 \frac{\mathbf{M}_{jt}}{\mathbf{H}_{it}}$$
(5.1)

where  $\mathbf{KL}_{ij}$  is the KL divergence matrix,  $\mathbf{M}$  is the macaque connectivity blueprint with each row representing the normalised probability distribution,  $\mathbf{H}$  is the human connectivity blueprint, *i* indexes the macaque vertices, *j* indexes the human vertices and *t* indexes the tracts. This dissimilarity index was then used to transform the human myelin map  $\mathbf{h}_i$  to give a prediction of the macaque myelin map  $\mathbf{p}_j$  using Equation 5.2:

$$\mathbf{p}_j = \frac{\sum \mathbf{K} \mathbf{L}_{ji}^{\gamma} \mathbf{h}_i}{\sum \mathbf{K} \mathbf{L}_{ji}^{\gamma}} \tag{5.2}$$

with  $\gamma = -4$  (as in Mars et al. (2018c)). This predicted myelin map is then compared to the macaque myelin map.

Another way of "aligning" brains using connectivity patterns is to use MSM, a flexible framework for registering cortical surfaces (Robinson et al., 2014), to drive the alignment through the joint-embedding structural components and use the resulting warp field to register the myelin maps across species, as in Xu et al. (2020). MSM was configured to perform data variance normalisation and used 3level discrete (Pearson's correlation) alignment on the components derived above. In addition, we used the medial wall masks as additional features to avoid misalignment in the medial wall across species.

Finally, we further used MSM to warp the human myelin map, this time driving the alignment using the connectivity blueprints themselves instead of the joint embedding components. In summary, we assessed three methods for projecting myelin maps across species using these connectivity spaces: 1) KL divergence applied to the connectivity blueprints, 2) MSM applied to the connectivity blueprints, and 3) MSM applied to the joint-embedding components (summarised in Figure 5.1).



Figure 5.1: Summaries of the construction of the joint-embedding components and myelin map prediction approaches. Connectivity blueprints ( $\mathbf{H}$  and  $\mathbf{M}$ ) were constructed: example columns (i.e. tract cortical termination maps) are shown on the left for both species. Further, joint-embedding was applied to extract jointcomponents of structural connectivity (path 3). Connectivity spaces are then compared and used to project the human myelin map  $\mathbf{h}$  to predict the macaque myelin map  $\mathbf{p}$  using three different approaches. 1) By calculating the KL divergence between the two connectivity blueprints. 2) By performing multi-modal surface matching (MSM) between the human and macaque connectivity blueprints. 3) As in 2, but using the landmark-based joint-spectral embedding components to drive the alignment.

#### Assessing Prediction Performance

Myelin map predictions using these average connectivity patterns were compared to an average measured macaque myelin map through Spearman's correlation and mean absolute difference ( $E = \overline{|\mathbf{p}_j - \mathbf{m}_j|}$ , where **p** is the predicted myelin map and **m** is the measured). Calculations of the Spearman's correlation and mean absolute difference did not include the medial wall.

#### 5.2.4 Data

We used diffusion MRI data for the human and macaque brain. Human data consists of 20 randomly chosen subjects from the "WU-Minn" Human Connectome Project (HCP) cohort (Sotiropoulos et al., 2013a; Van Essen et al., 2013), as described in the previous chapters. The macaque data utilised are as in Chapter 4, from 6 animals.

Cortical surfaces for the human brain, including the pial, white-grey matter boundary (WGB) and myelin<sup>1</sup>, are provided through the HCP preprocessing pipeline (Glasser et al., 2013). Cortical surface data are provided in the standard space, obtained through cortical folding-based surface registration (Glasser and Van Essen, 2011). Further, non-linear volumetric registration warps fields between the subject's native diffusion space and the MNI152 standard space are also provided, allowing for tractography to be run in the subject's native space yet stored in the standard space, as is required for inter-subject comparison and averaging. For the macaque brain, surfaces were obtained through the non-linear registration of a single macaque brain surface model, derived through Freesurfer surface reconstruction (Dale et al., 1999; Fischl, 2012), to each other macaque. Subsequently, these were registered to the F99 standard space. All surface data, structural and myelin maps, were downsampled to approximately 2,000 vertices per hemisphere (N = 3,924 vertices in total) to reduce computational times. XTRACT was ran on all brains to reconstruct WM homologous bundles and the connectivity blueprints were estimated for each subject and subsequently averaged. Here, we pre-selected tracts such that each hemisphere

<sup>&</sup>lt;sup>1</sup>Cortical myelin maps were estimated by dividing the T1w and T2w surface maps (Glasser and Van Essen, 2011).

contains ipsilateral and commissural tracts only, i.e. we remove right tracts from the left hemisphere and visa versa, leaving two connectivity matrices for each species representing the left and right hemispheres, each with T = 22 columns/tracts (see Chapter 4: Table 4.1 for details). For the myelin map predictions and comparisons, we used averaged connectivity blueprints and myelin maps for the 20 human subjects and the 6 macaques.

#### 5.2.5 Results

The calculated vertex-wise KL divergence matrix of connectivity blueprints across human and macaque vertices is shown in Figure 5.2, demonstrating an overall diagonal trend and indicating overall similarity across species. The vertex-wise joint similarity matrix was calculated within and across species (Figure 5.3a) using the homologous cortical landmarks. Figure 5.10 in Appendix 5.6.1 provides comparisons of the average tract contribution to each of the used landmarks. From this, low-dimensional components were extracted using spectral embedding, keeping the top 10 components (99.8% of the original data variance explained based on principal component analysis). Examples of these components for the human and macaque brain are shown in Figure 5.3 (b-c), highlighting that homologous landmarks cluster across species in the embedded space.

#### **Myelin Map Predictions**

The human myelin map was projected to the macaque cortical surface using the three approaches, described in Methods. Results are presented in Figure 5.4. The prediction using KL divergence and the connectivity blueprints was strongly correlated with the measured macaque myelin map (i.e. the average over the 6 macaques) with r = 0.731 and a mean absolute difference of E = 0.074. Using MSM, the correlations between the predicted and measured macaque myelin maps were r = 0.703 and r = 0.642 and the mean absolute differences were E = 0.085 and E = 0.094 for the connectivity blueprint and joint-embedding component driven alignments respectively. Figure 5.4 shows the measured myelin map for the human and macaque



**Figure 5.2:** The Kullback-Leibler divergence between the macaque and human cortical termination maps. Each row represents the KL divergence between a given macaque vertex and all of the human vertices. Likewise, each column represents the KL divergence between a given human vertex and all of the macaque vertices. Dark banding is a consequence of the inclusion of the medial wall.

brain and the obtained predictions, as well as scatter plots between the predicted and measured maps for each approach. Figure 5.5 plots the absolute difference between the predicted and measured myelin maps for each approach.

# 5.3 Building Connectivity Spaces using Data-Driven Methods

The above approaches offer promising ways for functionally-aligning very diverse brains using structural connectivity information. Yet, they are limited in that they rely on a finite set of manually-defined landmarks, either in WM (through homologous fibre bundles) or in GM. Hence, we explored a different approach for building connectivity spaces that overcomes this limitation. We used data-driven methodology for extracting connectivity patterns, specifically through non-negative matrix factorisation (NMF) of whole-brain tractography data (Thompson et al., 2020).

NMF is an unsupervised dimensionality reduction method which, when applied to whole-brain tractograms (e.g. cortical GM to whole-brain WM connectivity), is



Figure 5.3: The joint-similarity matrix and embedded components for the right hemisphere. **a.**) The joint-similarity matrix within and across species. Each quadrant reflects a similarity matrix: the within human similarity (top-left, H2H), the within macaque similarity (bottom-right, M2M), and the cross-species similarity matrix (bottom-left, M2H) and its transpose (top-right, H2M). Prior to joint-embedding, each quadrant is globally demeaned. **b.**) Examples of the joint-embedding components derived for the human and macaque brain visualised on the cortical surfaces, demonstrating the cross-species structure. **c.**) The first two components visualised in the joint-embedding space with homologous landmarks (area 3, area 9m, area MT, the primary auditory area (A1), and the primary visual area (V1)) highlighted, demonstrating that these landmarks have been projected similarly: plus and cross markers represent the human and macaque landmarks respectively.



Figure 5.4: Myelin map predictions. a.) The measured and predicted myelin maps. From top to bottom: the human myelin map - this is transformed using the described approaches, the measured macaque myelin map, the predicted macaque myelin map based on KL divergence, the predicted macaque myelin map based using the cortical termination maps to drive MSM, and the predicted macaque myelin map using the embedded components to drive MSM. b.) Scatter plots between the predicted and measured myelin maps for each approach.



**Figure 5.5:** Absolute difference between the predicted and measured macaque myelin maps. From top to bottom, the difference between the predicted myelin map and the measured myelin map based on: KL divergence, using the cortical termination maps to drive MSM, and using the embedded components to drive MSM.

capable of extracting representations of the WM bundles and their corresponding cortical termination maps (GM networks/nodes) (Figure 5.6). Importantly, NMF results in non-negative weights, vital in the interpretation of structural connectivity, employs sparsity parameters, resulting in pseudo-independent tract-like components and allows for the extraction of WM fibre bundles without the need for pre-defined tractography protocols. The GM spatial maps derived through NMF represent the cortical terminations of the corresponding WM spatial maps and may be used similarly to define connectivity blueprints and hence a common connectivity space in a fully data-driven manner. Here, to demonstrate feasibility of the approach, we extract NMF components on human data from the HCP at the group and individual level and use the KL divergence of data-driven connectivity blueprints to predict an "unseen" subject's myelin map.

NMF is best applied to group-averaged GM  $\times$  WM ( $N \times M$ ) whole-brain matrices (Figure 5.6) which may then be dual-regressed back to the individuallevel<sup>2</sup>. As such, we performed tractography in a group of HCP subjects and then averaged the whole-brain tractography matrices across the group.

 $<sup>^2 {\</sup>rm Full}$  details of the NMF procedure may be found elsewhere: Thompson et al. (2020) and github.com/ethompson93/Data-driven-tractography.



**Figure 5.6:** The extraction of WM and GM components through NMF. The wholebrain connectivity matrix is reduced via non-negative matrix factorization (NMF), resulting in equivalent sets of WM components in the volume space and GM components on the cortical surface.

Prior to tractography, crossing fibre modelling (Jbabdi et al., 2012) was applied to the data in order estimate fibre orientation (with up to three fibres estimated per voxel) and inform tractography. Whole-brain probabilistic tractography was used to obtain connectivity matrices (Behrens et al., 2007; Hernandez-Fernandez et al., 2019) as required in order to perform NMF (Thompson et al., 2020). Briefly, for each subject, tractography was seeded from each white-grey matter boundary (WGB) surface vertex and the subcortical structures, including the amygdala, caudate, thalamus, putamen and hippocampus, with 5,000 streamlines per seed. A whole-brain mask, excluding the ventricles, was used as a target with the pial surface acting as a termination mask to prevent streamlines from crossing gryi boundaries. For each seed point, the visitation count, weighted by path length, was calculated as the number of streamlines between the given seed point and each of the whole-brain mask voxels. Volume masks (subcortical and whole-brain masks) were downsampled to 2.5 mm, reducing computation time. Prior to tractography, volume and surface masks were transformed to the subject's native space, where tractography was performed. Subsequently, the resultant fibre density maps were resampled to standard space to allow for group-averaging.

Group-averaged  $(N \times M)$  connectivity matrices were calculated by first normalising the individual connectivity matrices by the total number of valid streamlines and then averaging across the group. Next, NMF was applied to the group-averaged data, extracting the top k = 200 components. The  $N \times k$  GM spatial maps of these components (the transpose of **GM** in Figure 5.6) (i.e. each column is a component GM spatial map) were used to define data-driven connectivity patterns; in an analogous manner as in the previous section, where a  $N \times T$  tracts matrix was considered the connectivity blueprint. A similar assessment of connectivity similarity as before could then be performed using the KL divergence.

#### 5.3.1 Myelin Map Predictions

To demonstrate in principle that these data-driven features may also be used to define connectivity spaces and probe functional features (as in Section 5.2), we sought to predict an unseen subject's myelin map from the group through the connectivity space (Equation 5.2). Data-driven connectivity blueprints, derived via NMF, were extracted at the group-level as described (matrix **H** in Equation 5.1). Next, the NMF components were dual-regressed to an individual (out-of-group) subject, providing a set of components for this subject which are informed by the individual anatomy yet structured as in the group-data (i.e. component correspondence between the subject and the group-average is maintained). That formed matrix **M** in Equation 5.1. The unseen subject's myelin map was then predicted using KL divergence (Equations 5.1-5.2). Prediction performance was assessed through correlation (Spearman's) and the mean absolute difference between the predicted and measured myelin map for that subject. For comparison, the prediction was repeated using the XTRACT-derived connectivity blueprints, as performed in the previous section.

#### **NMF Based Predictions**

Due to the arbitrariness in the selection of NMF decomposition order (k), we sought to reduce the NMF components to those that represent fibre bundles with strong cortical connectivity. This removes components representing tracts of no interest, for example the middle cerebellar peduncle (MCP) which would have little/no contribution to cortical connectivity. To achieve this, we selected the top 100 components with the largest mean values (considering only the surface vertex entries) based on the group-averaged data. Prior to making predictions, we smoothed the components (using a Gaussian smoothing kernel with a standard deviation of 4 mm). This reduces the sparsity inherent in these components, which would otherwise introduce issues in the calculation of the KL divergence. The connectivity patterns were normalised so that their row-sum was 1.

#### **XTRACT Based Predictions**

Connectivity blueprints were calculated using the XTRACT-derived tracts for each subject and averaged across the group. Prior to myelin predictions, and as in the cross-species explorations, tracts were sub-selected so that the left hemisphere contains only left and commissural tracts, and visa versa for the right hemisphere. In addition, we performed a third prediction in which we smoothed the connectivity blueprints as in the NMF case.

#### 5.3.2 Data

For these explorations we use 100 randomly chosen HCP subjects: 99 of these were used to generate group-averaged structural connectivity features and the remaining subject was used as the unseen subject. A group-average myelin map was generated by taking the mean across the 99 subjects. All surface data were downsampled (after feature generation) to approximately 2,000 vertices per hemisphere as before.

#### 5.3.3 Results

#### **NMF** Components

We performed whole-brain tractography and extracted the top k = 200 NMF components. Examples of the WM and GM components for the group-averaged data are shown in Figure 5.7.



Figure 5.7: Examples of the WM and GM NMF components derived from the 100 subject group-averaged data. Example components capture known major WM fibre bundles, for example a.) cingulum bundles, b.) the uncinate fasciculus (UF), c.) the second branch of the superior longitudinal fasciculus (SLF), d.) the arcuate fasciculus (AF), e.) the middle longitudinal fasciculus (MdLF), f.) temporal cingulum/fornix fibre bundles, g.) the anterior thalamic radiation (ATR), h.) sensory-motor connections, and i-j.) commissural bundles (splenium and genu of the corpus callosum)

#### **Myelin Map Predictions**

The unseen subject's cortical myelin map was predicted through KL divergence using three sets of features: the connectivity blueprints, the top 100 NMF components, and, for comparison, the connectivity blueprints also spatially smoothed with the 4 mm kernel. The group-average, measured and predicted myelin maps are shown in Figure 5.8. The correlations and mean absolute differences between the predicted and measured myelin maps were r = 0.739 and E = 0.109 (XTRACT), r = 0.777and E = 0.104 (smoothed XTRACT) and r = 0.791 and E = 0.100 (data-driven NMF). Regions of notable difference are indicated in Figure 5.8 by the white arrows, where it can be seen that the MT area and the general 'C'-shape of the cingulum are better preserved when using the data-driven connectivity blueprints, however, the increased myelin in the hippocampal area is underestimated compared to the XTRACT-derived connectivity blueprint. Generally, the lateral frontal region is variable across predictions with the predictions using the data-driven connectivity blueprints reflecting the group-average myelin map more so.



**Figure 5.8:** The measured and predicted myelin maps. From top to bottom: the group-averaged myelin map, the measured unseen subject myelin map, the predicted myelin map using the XTRACT-based connectivity blueprints, the predicted myelin map using the smoothed XTRACT-based connectivity blueprints, and the predicted myelin map using the smoothed data-driven connectivity blueprints. Arrows mark regions of notable difference across predictions. MT: Area MT, PC: Posterior Cingulate, Hipp: Hippocampus.

# 5.4 Discussion

This chapter presents explorations into using brain connectivity to probe functional similarities and to align similar features across diverse brains. We use structural connectivity patterns to define common connectivity spaces between brains. We assess two methods for building such structural connectivity patterns (connectivity blueprints and joint-embedding) and example their use through the projection of myelin maps across species. Each space provides a set of features - tracts/columns of the connectivity blueprint or components of the joint-embedding space - which are homologous, i.e. they represent equivalent variance in connectivity for each species but respect the individual species anatomy and which may be compared across very

diverse brains in the absence of geometrical correspondence.

Three approaches to performing projections of myelin maps were compared, allowing for the comparison of the two connectivity spaces: 1) using the KL divergence between the connectivity blueprints, 2) using multi-modal surface matching (MSM) with the connectivity blueprints as features, and 3) using MSM with the joint-embedding components as features. Each approach was predictive of cortical myelination as measured by MRI, in-line with expectations from the literature (Mars et al., 2018c; Xu et al., 2020; Eichert et al., 2020), suggesting that each connectivity space captures useful variation in structural connectivity in a common space which may be used to probe functional similarities. Figure 5.4 reveals that each approach, to some extent, captures the expanded and more lateral visual area and more anterior motor region in the macaque brain compared to the human brain. Although medial wall masks were used as a feature in the MSM-based surface alignments, there still appears to be some misalignment in the medial wall, particularly in the third approach. These predictions demonstrate that comparisons of structural connectivity features, derived through standardised cross-species tractography, can drive functional alignment of different brains, as suggested by the literature (Mars et al., 2018c; Eichert et al., 2020).

The application of these methods in these data may be limited by the already low-dimensionality of the data (22 tracts), reducing our ability to capture wholebrain connectivity profiles in their entirety. Further, in the joint-embedding approach, it should be noted that, due to the downsampling of the cortical surfaces and the already relatively small size of the landmark ROIs, some of the landmark connectivity profiles were the average of very few cortical locations (i.e. vertices, Figure 5.10 in Appendix 5.6.2). However, we demonstrated that joint-embedding captures components of unique variance which are structured across species: the component maps capture homologous variance in structural connectivity (Figure 5.3b) and homologous landmark regions cluster together in the embedded space (Figure 5.3c). The ability to compare structure and probe functional organisation across brains may benefit from the inclusion of a broader array of structural connectivity features. Towards this, we adapted a recently developed data-driven feature extraction technique (Thompson et al., 2020) to the adult human brain and demonstrated that structural connectivity patterns derived through this approach can also be used to build a common space and perform predictions of cortical myelination. The NMF components extracted represent a diverse set of WM connections (Figure 5.7) whilst allowing for the extraction of features not defined in some protocol-driven methods (i.e. in XTRACT). We used KL divergence to compare group-averaged GM NMF components to those from an unseen subject, making a prediction of the unseen subject's myelin map with similar, although higher in the data-driven approach, predictive power across approaches. A limitation in this approach is that, in order to avoid issues in the calculation of the KL divergence, we performed smoothing on the data, reducing the inherent sparsity of the data. This increased the overall similarity between the group-average and unseen subject data, which, in the XTRACT-based connectivity blueprint case, led to an apparent improvement in predictive power. Future work will build upon this feasibility study.

NMF for the extraction of WM fibre bundles has previously been demonstrated to be robust and generalisable across stages of neonatal brain development (Thompson et al., 2020). Such techniques could be adapted to performing tractography in multiple species which would overcome the limitations in the traditional cross-species tractography approach, where time-consuming tractography protocols must be developed in order to extract structural homologues. Using this data-driven method, we can extract any number of components which may capture fibre bundles not yet defined in the traditional approach, thus providing a more complete representation of structural connectivity patterns.

Although data-driven approaches are attractive in their ability to capture tracts not yet defined in XTRACT, and in their flexibility in being applied to any brain in principle, they do have limitations compared to anatomically driven tractography. Anatomically driven tractography is able to robustly extract well-defined reconstructions of fibre bundles with few false positives (i.e. not including other fibre bundles). Further, by restricting our explorations to known and well-defined fibre bundles, interpretation is eased: understanding the contribution of a fibre bundle is simple, whereas, understanding the contribution of a component (potentially containing aspects of several distinct fibre bundles) is more complex. Additionally, to explore changes in connectivity profiles across diverse brains, well-defined fibre bundles (homologous in the case of cross-species studies) are particularly useful. In order to utilise data-driven tractography approaches in such cases, anatomically driven tractography provides a useful comparator against which data-driven tractography may be assessed.

# 5.5 Concluding Remarks

Investigations into the use of structural connectivity features to build common connectivity spaces have been presented. We example, on a cross-species basis, that connectivity patterns derived through protocol-driven tractography may be compared in the absence of geometrical alignment, can drive alignment of diverse brains in connectivity space, are predictive of cortical myelination and, therefore, probe functional specialisation. Further, we adapted a recently developed method for the data-driven estimation of structural connectivity patterns and demonstrated, in principle, that such features may be used to build common connectivity spaces and perform such comparisons. This opens the door to a data-driven cross-species tractography framework, potentially overcoming the limitations discussed.

# Contributions

The works presented in this chapter were conceptualised and designed by me and Stamatios Sotiropoulos. Cortical landmarks were kindly provided by Ting Xu (Xu et al., 2020). Macaque data are as in Chapter 4. I am also grateful to Ellie Thompson for her guidance and expertise in the application of NMF to tractography data. My contributions span all aspects of scripting and data processing, analyses and visualisations.

# Software

Homologous tracts were derived using FSL's XTRACT (v6.0.2) and connectivity blueprints were derived using MATLAB 2018a (The MathWorks, Inc., Natick, Massachusetts, United States). Data-driven tractography (i.e. NMF) was performed using the Data-driven-tractography toolbox (https://github.com/ethompson93/Datadriven-tractography) (Thompson et al., 2020). Statistical analyses and visualisation were performed in Python (version 3.7.2) (Python Software Foundation, https://www.python.org). Brain visualisations were performed using Connectome Workbench (version 1.3.2; Van Essen Laboratory, Washington University).

# 5.6 Appendix

#### 5.6.1 Landmark regions-of-interest



Figure 5.9: The 27 homologous landmarks used in deriving the joint-embedded components. 10r: Ventro-Medial Prefrontal Cortex, p32: Area 32, 9m: Area 9 medial, 9-46d: Area 9/46 dorsal, 8Av and FEF: Area 8A and Frontal Eye Field, PFop: Area PFop, AIP: Anterial Intraparietal, VIP: Ventral Intraparietal Area, LIPd and LIPv: Mid-Intraparietal Sulcus, PFm: Mid-Inferior Parietal Lobule, 6v and 6r: Area 6 ventral and Area 6 dorsal, IFSp and IFSa: Inferiro Frontal Sulcus, IFJp and IFJa: Inferiro Frontal Junction, 44: Area 44 ventral, 45: Area 45, 471: Area 47, 47s: Operculum Frontale, 46: Area 46, 10d: Mid-Inferior Parietal Lubule, V1: Area V1, V2d: Area V2 dorsal, V2v: Area V2 ventral, A1: Area A1 (Auditory), MT: Area MT, 3a and 3b: Area 3, 4: Area 4, Hipp: Are Hippocampus). Regions were provided by Xu et al. (2020).



### 5.6.2 Landmark connectivity

**Figure 5.10:** The average and normalised tract contributions to each of the homologous landmarks for the human (blue) and macaque (orange) brains. For each cortical landmark, the top contributors for each species are shown, defined as have a contribution of greater than 0.05 (following sum-normalisation to 1). The number of vertices included in each landmark are given.

# Chapter 6

# Multivariate Statistical Approaches for Brain-Behaviour Associations

## Contents

6.1	Introduction									
6.2	$\mathbf{CCA}$	Theory								
6.3	Prior	Work using CCA								
6.4	Evalu	Evaluating CCA Stability Against Multi-modal Data 160								
	6.4.1	Feature Extraction: Imaging Features								
	6.4.2	Feature Extraction: Non-Imaging Features								
	6.4.3	Analysis								
	6.4.4	Results								
	6.4.5	Summary								
6.5	CCA with Known Ground-Truths									
	6.5.1	Simulation Framework								
	6.5.2	Results								
	6.5.3	Summary 172								
6.6	$\mathbf{CCA}$	Applied to Large Samples								
	6.6.1	Feature Extraction: Imaging Features								
	6.6.2	Feature Extraction: Non-Imaging Features								
	6.6.3	Analysis								
	6.6.4	Results: Overall Stability								
	6.6.5	Results: Correspondence Between Modes 177								
	6.6.6	Summary								
6.7	Discussion									
	6.7.1	Recommendations								
6.8	Conc	oncluding Remarks 190								
6.9	Appendix									
	6.9.1	Double-dipping and CCA								
	6.9.2	Behavioural Measures used in CCA - HCP Analyses $\ . \ . \ . \ 194$								
	6.9.3	Behavioural Measures used in CCA - UK Biobank Analyses196								
	6.9.4	UK Biobank CCA top Non-imaging Loadings 200								

# Foreword

One of primary goals of neuroscience is to reveal the links between observable features of the brain and human behaviour. One approach to revealing these associations is discovering relationships through a large set of features that characterise this brain-behaviour parameter space. Brain features can be extracted through neuroimaging, such as using the approaches presented in the previous chapters and behavioural traits can be probed using behavioural data. Multivariate techniques can then be used to explore latent relationships. Canonical correlation analysis (CCA) is a powerful multivariate statistical technique used to reveal associations between multiple, typically large dimensionality, datasets. In recent years, CCA has been applied widely in the field of neuroimaging. However, issues with regard to the stability of CCA have been demonstrated in a general context.

This chapter performs comprehensive investigations into the stability of CCA in the context of neuroimaging, specific to brain-behaviour associations. Using data from the "WU-Minn" Human Connectome Project and UK Biobank, and through a generative simulation framework, we demonstrate that results obtained through CCA are heavily dependent on the number of subjects, the number of features and the true underlying between-set association.

Specifically, we demonstrate that linear multivariate brain-behaviour associations need about 50 samples per considered feature to achieve stability, an order of magnitude more data than that used before. We show that latent associations revealed with HCP data ( $\sim 1,000$  subjects) can be highly unstable, while associations revealed with UK Biobank data ( $\sim 20,000$  subjects) are stable. We therefore caution the interpretation of these associations when the ratio of samples per feature is small.

# 6.1 Introduction

A plethora of techniques, including those introduced in this thesis, have been developed that use MRI to estimate measures of brain connectivity. One of the primary goals in estimating brain connectivity is to better understand how human behaviour and cognition, in the healthy and diseased brain, are related to variations in brain connectivity. Studies on the associations between disease effects and brain connectivity have advanced our understanding of disease mechanisms, for example, in schizophrenia (Liu et al., 2008; Fornito et al., 2012; Collin et al., 2014; Bohlken et al., 2016), Alzheimer's disease (He et al., 2009; Lo et al., 2010; Oxtoby et al., 2017; Córdova-Palomera et al., 2017), attention deficit hyperactivity disorder (ADHD) (Oldehinkel et al., 2016; Silberstein et al., 2016), and lesions (Thiebaut de Schotten et al., 2014a; Corbetta et al., 2015; Boes et al., 2015; Vaidya et al., 2019; Thiebaut de Schotten et al., 2020). In order to reveal features capable of identifying early markers of disease or to understand normal cognition and behaviour in great detail, studies must be highly sensitive. This is achieved by utilising large cohorts, robust image processing and powerful statistical methods (Miller et al., 2016; Smith and Nichols, 2018; Marek et al., 2020; Gong et al., 2020).

Canonical correlation analysis (CCA) is one such statistical method (Hotelling, 1936). It is used widely to reveal latent associations between datasets. In the field of neuroscience, CCA has been applied to exploring the associations within and across image modalities, and between imaging and non-imaging features. The latter is the focus of this chapter, where such techniques are applied to revealing latent relationships between image-derived features, such as brain connectivity, and behavioural traits, as captured by non-imaging measures, such as lifestyle, physical and psychological health, and cognition. In other words, to reveal brain-behaviour associations. These investigations use large dimensionality datasets: typically hundreds to thousands of image-derived features and tens to hundreds of behavioural trait features. CCA is used to reveal low dimensional representations of each dataset which are maximally correlated across the modalities. These latent representations - or "modes" - may then be related back to the original feature space to allow for visualisation and interpretation.

Although widely utilised in the field of brain-behaviour associations (Smith et al., 2015; Miller et al., 2016; Drysdale et al., 2017; Bijsterbosch et al., 2018; Wang et al., 2018a; Xia et al., 2018; Mihalik et al., 2019; Li et al., 2019; Becker and Hervais-Adelman, 2019; Rodríguez-Cruces et al., 2020; Han et al., 2020; Wang et al., 2020b; Ly et al., 2020), limitations of CCA in general have been reported (Thorndike and Weiss, 1973; Barcikowski and Stevens, 1975; Thorndike, 1976; Barcikowski and Stevens, 1976; Weinberg and Darlington, 1976; Thompson, 1990; Strand and Kossman, 2000; Lee, 2007; Leach and Henson, 2014) and, less-so, specific to brain-behaviour associations (Dinga et al., 2019; Yang et al., 2019; Wang et al., 2020b). These studies typically use small sample sizes to illustrate instabilities in CCA solutions, caused by overfitting in the solution, which, in turn, leads to inflated correlations (Weinberg and Darlington, 1976; Thompson, 1990; Lee, 2007). This may lead to marked differences between the reported between-set correlation and cross-validation performance. Dinga et al. (2019) illustrates this issue with regard to a specific study from the literature (Drysdale et al., 2017). This discrepancy is also observed in Smith et al. (2015), where the reported between-set correlation is 0.87 and the out-of-sample predictive performance is 0.25, suggesting potential overfitting. Furthermore, less work has focused on the stability of the actual CCA modes. Without fully understanding the limitations of CCA in the context of brainbehaviour associations, the generalisability of the technique and its solutions are severely limited.

This chapter performs comprehensive investigations into the robustness of CCA specific to brain-behaviour studies. Data from the HCP and UK Biobank are utilised to demonstrate significant challenges and limitations in the application of multivariate techniques when not enough data are available. In the following sections, the mathematical principles of CCA and a review of the brain-behaviour literature are presented. Following this, several investigations into the stability of CCA are presented: 1) CCA is explored using structural connectivity features and behavioural traits derived from the HCP cohort, revealing issues in the application of CCA in cases where the number of subjects is comparable to the number of features; 2)

reported instabilities are confirmed not to be specific to the type of feature used by repeating the previous explorations using functional connectivity features and behavioural traits derived from the HCP cohort; 3) a simulation framework is developed in order to explore these limitations in cases where the ground-truth is known; and 4) CCA is applied to the largest neuroimaging cohort available, the UK Biobank, to demonstrate the extent to which instabilities in reported CCA results occur.

Although the work presented in this chapter is specific to CCA, equivalent analyses have been performed using a similar approach, partial least squares (PLS), and in variants of CCA (sparse CCA) (Helmer et al., 2020). The work presented in Helmer et al. (2020) was a collaborative effort, benefiting from the studies presented here. Empirical analyses presented here were performed using the *GEMMER* package<sup>1</sup>.

# 6.2 CCA Theory

The aim of classical CCA is to produce low-dimensional linear combinations of the columns of two datasets such that those low-dimensional representations are maximally correlated (Hotelling, 1936). Let us define two data matrices  $\mathbf{X} = (\mathbf{x}_1, ..., \mathbf{x}_p)$  and  $\mathbf{Y} = (\mathbf{y}_1, ..., \mathbf{y}_q)$ , each with equal number of rows representing observations (e.g. subjects) n, where p and q are the respective number of features (columns) in each dataset (e.g. imaging and behavioural features), thus  $\mathbf{x}_i$  and  $\mathbf{y}_i$  are column vectors representing the features for each subject for each dataset respectively. CCA seeks to find two unmixing matrices (canonical weights)  $\mathbf{W}_X = (\mathbf{w}_{X1}, ..., \mathbf{w}_{Xp})^T$  and  $\mathbf{W}_Y = (\mathbf{w}_{Y1}, ..., \mathbf{w}_{Yq})^T$ , where T denotes the transpose and  $\mathbf{w}_{Xi}$  and  $\mathbf{w}_{Yi}$  are row vectors representing the contribution (or weighting) of a given feature to each of k modes, which, when multiplied with  $\mathbf{X}$  and  $\mathbf{Y}$  respectively, produce the canonical variates  $\mathbf{U}$  and  $\mathbf{V}$  such that column pairs are maximally correlated (Figure 6.1), i.e.

 $\boldsymbol{\rho} = corr(\mathbf{X}\mathbf{W}_X, \mathbf{Y}\mathbf{W}_Y) = [corr(\mathbf{u}_i, \mathbf{v}_i)], i = 1...k$ (6.1)

<sup>&</sup>lt;sup>1</sup>See https://github.com/murraylab/gemmr

such that

$$\mathbf{W}_X, \mathbf{W}_Y = \arg\max corr(\mathbf{X}\mathbf{W}_X, \mathbf{Y}\mathbf{W}_Y)$$
(6.2)

where *corr* means Pearson correlation and  $\rho$  is the vector of correlations between each pair of canonical variates, known as the canonical correlations.



Figure 6.1: A schematic representation of the CCA equation. **a.**) Columns of the two modalities  $\mathbf{X}$  and  $\mathbf{Y}$ , representing the imaging and non-imaging data, are projected to low dimensional space (columns of  $\mathbf{U}$  and  $\mathbf{V}$ ) through the multiplication of the mixing matrices  $\mathbf{W}_X$  and  $\mathbf{W}_Y$ . **b.**) The mixing matrices are such that the correlation between the canonical variates  $\mathbf{u}_i$  and  $\mathbf{v}_i$  is maximised, thus revealing the latent mode of variation.

The canonical variates describe the weighted-composite features of the original data. In other words, **U** and **V** are the low dimensional projections of **X** and **Y** respectively. A pair of canonical variates (e.g.  $\mathbf{u}_1$  and  $\mathbf{v}_1$ ,  $\mathbf{u}_k$  and  $\mathbf{v}_k$ ) is known as a canonical mode, which describes how the two datasets correlate across samples for a given linear combination of features. A maximum of k = min(p, q) canonical modes may be generated under the conditions that  $\rho_1 \ge \rho_2 \ge ... \ge \rho_k$ ,  $corr(\mathbf{u}_1, \mathbf{u}_{2...k}) = 0$  and  $corr(\mathbf{v}_1, \mathbf{v}_{2..k}) = 0$ .

In practice, CCA is often performed via singular-value decomposition of the between-set covariance matrix as in Equation 6.3.

$$\mathbf{C}_{XY}^{CCA} = \mathbf{C}_{XX}^{-1/2} \mathbf{C}_{XY} \mathbf{C}_{YY}^{-1/2}$$
(6.3)

where  $\mathbf{C}_{XX} = 1/(n-1)\mathbf{X}'\mathbf{X}$   $(p \times p)$  is the covariance matrix of  $\mathbf{X}$ ,  $\mathbf{C}_{YY} =$ 

 $1/(n-1)\mathbf{Y'Y} (q \times q)$  is the covariance matrix of  $\mathbf{Y}$  and  $\mathbf{C}_{XY} = 1/(n-1)\mathbf{X'Y} (p \times q)$ is the covariance matrix of  $\mathbf{X}$ ,  $\mathbf{Y}$ . The correlations between pairs of canonical vectors (known as the canonical correlations) are the singular values of  $\mathbf{C}_{XY}^{CCA}$ :  $\mathbf{C}_{XY}^{CCA} =$  $\mathbf{U}diag(\boldsymbol{\rho})\mathbf{V}^{T}$ , where  $\boldsymbol{\rho}$  are the canonical correlations as in Equations 6.1-6.2. The corresponding weights  $\mathbf{W}_{X}$  and  $\mathbf{W}_{Y}$  are given by  $\mathbf{C}_{XX}^{-\frac{1}{2}}\mathbf{U}$  and  $\mathbf{C}_{YY}^{-\frac{1}{2}}\mathbf{V}$ .

Typically, the statistical significance of the canonical correlations is found via permutation testing (Figure 6.2). This process usually includes shuffling one of the datasets (e.g.  $\mathbf{Y}' = permute(\mathbf{Y})$ ) such that the order of subjects no longer matches between the datasets, performing CCA and repeating  $N_p$  times. This process allows for the construction of a null distribution of canonical correlations  $\boldsymbol{\rho}_{null} = [\rho_1, ..., \rho_{N_p}]$ . The true, i.e. unshuffled, canonical correlation between the first pair of canonical variates  $(\rho_1)$  may then be compared to the null to assess the probability of the correlation being spurious. This is typically performed via family-wise error correction where each mode is considered significant if the *p*-value  $= \frac{1}{N_p}(1 + \sum (\boldsymbol{\rho}_{null} \ge \rho_i))$  is less than a desired threshold ( $\rho_i$  is the correlation for a given canonical mode). The lowest *p*-value obtainable is then governed by the reciprocal of  $N_p$ , typically  $N_p = 1,000 - 100,000$ . If the first mode  $\rho_1$  is found not to be spurious, i.e. statistically significant, the subsequent modes are then compared to the same null distribution.

Following the assessment of the statistical significance of modes, visualisation and interpretation is typically performed by correlating the canonical variates  $\mathbf{u}_i$ and  $\mathbf{v}_i$ , corresponding to the significant modes, with the columns of the input data  $\mathbf{X}$  and  $\mathbf{Y}$  respectively, giving two sets of loading vectors  $\mathbf{l}_X$  and  $\mathbf{l}_Y$ . These describe the contribution of the input features to the given canonical mode in the original feature space (i.e. p or q features). In practice, this correlational analysis is often performed between the canonical variates and the original, pre-PCA, input data - the application of PCA prior to CCA is common practice and aims to reduce overfitting, i.e. improve solution stability and reduce correlation inflation, and/or overcome the issue of rank deficiency. Returning the CCA solution back to the original feature space, i.e. deriving the loading vectors, overcomes issues in interpretation, allowing for direct visualisation in the original feature space, e.g. on the cortical surface or



Figure 6.2: A schematic representation of permutation testing. a.) First, CCA is performed on the original datasets X and Y, and the canonical correlations  $\rho$  for each mode are stored. b.) Next, one of the datasets, for example the non-imaging dataset Y, is permuted (i.e. shuffled) in the subject dimension so that there is no longer correspondence between modalities, giving Y'. c.) CCA is repeated using the original X and the permuted Y' and the canonical correlation for the top mode is stored  $\rho_{null}$ . Steps b.) and c.) are repeated  $N_p$  times, thus constructing a null distribution of correlation values  $\rho_{null}$  against which the non-permuted correlation value  $\rho_i$  may be compared.

in terms of meaningful behavioural measures. Higher loadings indicate a greater contribution to the canonical mode.

Although this chapter is concerned with the application of conventional CCA to exploring brain-behaviour associations, many similar techniques and variants exist which one should be aware of. Zhuang et al. (2020) provides a comprehensive overview of the CCA variants. Common variants of CCA include constrained/sparse CCA (Witten et al., 2009). Both impose penalties/constraints (e.g. non-negativity) and/or regularisation parameters (most commonly L1-norm) on the canonical coefficients with the intention of improving model performance, e.g. generalisability. Constrained and sparse CCA aim to reduce overfitting by imposing constraints on the canonical coefficients, pushing low value entries to exactly zero through L1 regularisation in the case of sparse CCA. As well as seeking to reduce overfitting, sparse CCA also allows for CCA to be applied in cases where the number of subjects n does not exceed the number of features p + q (this condition must be met for classical CCA solutions to be mathematically sound).

A similar approach to CCA is partial least squares (PLS). Here, rather than seeking to maximise the correlation between low-dimensional projections as in CCA, PLS maximises the covariance. Mathematically, this is very similar to the description of CCA:  $CV_1 = cov(\mathbf{XW}_X, \mathbf{YW}_Y) = cov(\mathbf{u}_1, \mathbf{v}_1)$  such that  $\mathbf{W}_X, \mathbf{W}_Y =$  $\arg \max cov(\mathbf{XW}_X, \mathbf{YW}_Y)$ , where  $CV_1$  is the covariance between the first pair of modes.

## 6.3 Prior Work using CCA

Initial studies using CCA in the field of neuroimaging focused on the source separation of brain activity signals (Friston et al., 1995; Friston et al., 1996; Friman et al., 2001). In recent years, CCA has been applied to exploring the associations between imaging-derived measures ( $\mathbf{X}$ ), such as brain connectivity, and non-imaging behavioural measures ( $\mathbf{Y}$ ), such as measures of intelligence and clinical diagnostic measures. A summary of key literature is provided in Table 6.1.

A number of studies have focused on CCA between functional connectivity, as derived from resting-state functional MRI, and behavioural traits. Amongst the first, Smith et al. (2015) applied CCA in such datasets using 461 subjects from the HCP. Prior to CCA, principle component analysis (PCA) was used to reduce the imaging features from 19,800 (the unique edges of a 200 by 200 functional connectome) to the top 100 principal components and the behavioural data from approximately 160 to the top 100 principal components. Therefore, the input into CCA was two datasets **X** and **Y** with dimensions of n = 461 and p = q = 100, and, therefore, the subject-feature ratio is 2.3. A single significant canonical mode was reported, capturing a "positive-negative" mode of population variance in the behavioural data that is associated with known functional networks, for example the default mode network, and with overall connectivity.

The analysis and findings presented in Smith et al. (2015) have since been replicated (Bijsterbosch et al., 2018; Goyal et al., 2020; Han et al., 2020) and the same explorations have been performed with other fMRI-derived features (Marquand et al., 2017; Bijsterbosch et al., 2018; Li et al., 2019), measures of cortical thickness (Han et al., 2020), and magnetoencephalography (MEG)-derived functional connectivity (Becker and Hervais-Adelman, 2019) using the healthy HCP cohort. In addition, similar approaches have been applied in other healthy cohorts to explore latent associations between connectivity and cognition/behaviour (Wang et al., 2018a; Wang et al., 2020b). Further, CCA has been applied to exploring brain-behaviour associations in clinical cohorts, including in depression (Drysdale et al., 2017; Mihalik et al., 2019; Dinga et al., 2019), temporal lobe epilepsy (Rodríguez-Cruces et al., 2020), schizophrenia (Lv et al., 2020), and in a neurodevelopmental cohort (Xia et al., 2018).

It should be noted that the analysis and results presented in Drysdale et al. (2017) have since been brought into question (Dinga et al., 2019) (however, see Grosenick et al. (2019) and Mihalik et al. (2020)). In brief, Drysdale et al. (2017) suggests that subjects with depression may be sub-typed, and therefore treatment may be better targeted, based on the latent associations between functional connectivity and clinical measures revealed via CCA and subsequent clustering. Dinga et al. (2019) fails to robustly replicate these findings in an independent cohort and argues that permutation testing should be applied at the feature-selection level to avoid overfitting and instabilities. In the case of Drysdale et al. (2017), imaging features are sub-selected through a mass univariate association (MUA) where features passing a pre-defined p-value threshold are kept. Appendix 6.9.1 provides a brief investigation into the use of MUA as a feature-selection procedure prior to CCA and we argue that such "double-dipping" should be avoided.

Despite the potential suggested from the literature, previous work has highlighted issues with regard to the stability of CCA (Barcikowski and Stevens, 1975; Thorndike and Weiss, 1973; Barcikowski and Stevens, 1976; Weinberg and Darlington, 1976; Thorndike, 1976; Thompson, 1990; Strand and Kossman, 2000; Lee, 2007; Leach and Henson, 2014; Yang et al., 2019; Dinga et al., 2019; Wang et al., 2020b), which may lead to unrealistic canonical correlations ( $\rho > 0.9$ ) (Drysdale et al., 2017; Becker and Hervais-Adelman, 2019; Dinga et al., 2019; Rodríguez-Cruces et al., 2020) and potentially erroneous conclusions. Here, using data from the HCP and UK Biobank, and through a simulation framework, instabilities in canonical weights and correlations are demonstrated when using data dimensionalities typical of that found in the literature (see Table 6.1 for a summary). Specifically, the importance of the subject-feature ratio (SFR = n/(p+q)) is highlighted which, as demonstrated in Table 6.1, ranges from 0.15 to 6.1 in the literature.

# 6.4 Evaluating CCA Stability Against Multi-modal Data

In this section, the stability of CCA is assessed in terms of the estimated canonical correlations, describing the overall correlation between the latent space modes, and canonical weights, which define said latent modes. In each case presented below, a set of samples (i.e. subjects)  $n_s$  were defined with two modalities of data  $\mathbf{X}_{n_s}$  and  $\mathbf{Y}_{n_s}$ , which correspond to the imaging and non-imaging features respectively. Explorations with two sources of imaging features are performed (i.e. two separate matrices  $\mathbf{X}_{n_s}$ ) obtained from independent modalities: from dMRI data and fMRI data in the HCP. Brain connectivity features are extracted from these modalities and CCA is performed independently between each of these two neuroimaging modalities and HCP behavioural measures.

The stability of the estimated canonical weights for each modality and the canonical correlations was explored against sample size. CCA was repeated on subsets of the data ( $\mathbf{X}_n$  and  $\mathbf{Y}_n$ ) for  $n < n_s$  whilst maintaining the feature dimensions p and q, i.e. reducing subject-feature ratio (*SFR*). For each iteration, the canonical correlations  $\boldsymbol{\rho}_n$  and canonical weights  $\mathbf{W}_n$  were compared to the ones obtained from the original set (here, we use  $\mathbf{W}$  to refer to the weights for each data set  $\mathbf{X}$  and  $\mathbf{Y}$ ).  $\boldsymbol{\rho}_n$  are plotted as a function of *SFR* and inflation (percentage increase) was calculated. Weight stability/error was calculated as 1- the absolute value of cosine similarity between the subsampled  $\mathbf{W}_n$  and full-sample  $\mathbf{W}_{n_s}$  weights. This was calculated for each data set  $\mathbf{X}$  and  $\mathbf{Y}$  and the greater of the two is taken:

$$\Delta \mathbf{W}_n = \max_{j \in X, Y} (1 - |\operatorname{cossim}(\mathbf{W}_n, \mathbf{W}_{n_s}^j)|)$$
(6.4)

	Cohort	Img (X)	Non-Img (Y)	Reduction	#Subs	#Feats	SFR
				(X, Y if	(n)	(p,q)	(n/(p+q))
	TICD			different)	101	100 100	
Smith 2015	HCP	WB-func	§	PCA	461	100, 100	2.3
Drysdale	Depression	WB-func	Clinical de-	MUA (to re-	220	178, 17	1.1
2017	cohort		pression scores	duce imaging			
Manaurand	UCD	Continentui	6	Ieatures)	166	0.174	25
Marquand	пог	Corticostri-	8	Folynomiai	400	9,174	2.0
2017		connectivity		olimination			
Bijstor-	НСР	Various fMBI	8	PCA	810	100 100	4.1
bosch 2018	1101	connectomes	3	1011	015	100, 100	1.1
50501 2010		and spatial					
		configuration					
Wang 2018	Nathan Kline	WB-func	Cognitive	Sparse CCA	258	1,596,	0.15
	Institute-		"resting-state"	1		31	
	Rockland		test				
	sample						
Xia 2018	Philadelphia	WB-func	Psychiatric	Percentile	663	3,410,	0.19
	Neurode-		symptoms	thresholding		111	
	velopmen-			and Sparse			
	tal Cohort			CCA			
N / : 1 1: 1-	(PNC)	XX7b a balantin	Maria I Iraa Ith	DCA	200	05 05	C 1
Minalik	Adolescent	Whole-brain	Mental health,	PCA	306	25, 25	6.1
2019	and young	connectivity	and domo				
	sion cohort	connectivity	graphics				
Li 2019	HCP	Global signal.	8	PCA	1.094	100.100	5.5
21 2010		fMRI connec-	5	1 011	1,001	100, 100	0.0
		tivity					
Dinga 2019	Depression	WB-func	Clinical de-	MUA (to re-	187	150, 17	1.1
_	cohort		pression scores	duce imaging			
				features)			
Becker 2019	HCP	MEG func-	§	PCA	89	22, 22	2.0
		tional con-					
		nectivity, 5					
		frequency					
Dodnimuoz	Tommonol	Dands Whole brois	Comiting	Domulonicod	24	ND	ND
Roariguez-	lebo opilopey	whole-brain	clinical scores	CCA	34	INK	INK
2020	cohort	connectivity	chincar scores				
Han 2020	HCP	WB-func and	8	PCA	818	100 100	4.1
11011 2020	1101	cortical thick-	3	1 0/1	010	100, 100	1.1
		ness					
Wang 2020	Healthy co-	WB-func	Cognitive	Top 5%	178	247, 13	0.69
_	hort		scores	functional			
				features and			
				Sparse CCA			
Lv 2020	Schizophrenic	Cortical thick-	AAI and la-	$N/A \text{ and } I\overline{CA}$	322	116, 1	2.8 and
	cohort	ness and	tent cognitive/			and 4	2.7
		tractional	behaviour				
		anisotropy	scores				

**Table 6.1:** Summary of the canonical correlation analyses performed in the literature showing the types of imaging and non-imaging measures used, dimensionality reduction methods and number of subjects and features used. WB-func = Wholebrain fMRI connectivity; \$ = Cognitive, lifestyle, family history, physical health; NR = Not Reported; PCA = Principal Component Analysis; MUA = Mass Univariate Association; HCP = "WU-Minn" Human Connectome Project

where

$$\operatorname{cossim}(\mathbf{W}_{n}, \mathbf{W}_{n_{s}}) = \frac{\mathbf{W}_{n} \cdot \mathbf{W}_{n_{s}}}{\|\mathbf{W}_{n}\| \|\mathbf{W}_{n_{s}}\|}$$
(6.5)

The weight error was then plotted as a function of SFR. Further, 5-fold crossvalidation was used to estimate the out-of-sample canonical correlation at each sample size. The CCA solution was estimated with 80% of the data (training set) and the resulting canonical weights  $\mathbf{W}_{train}$  were applied to the remaining 20% of the data (test set,  $\mathbf{X}_{test}$ ) to obtain test canonical variates  $\mathbf{U}_{test}$  by  $\mathbf{U}_{test} = \mathbf{X}_{test} \times \mathbf{W}_{train}$  for the imaging modality for example. The Pearson correlation was calculated between the test variates and averaged across the 5 folds. Various cross-validation strategies are available and may be used to assess CCA performance. The 5-fold scheme used here is not only a convenient choice is also well-suited to model performance assessment (Varoquaux et al., 2017; Scheinost et al., 2019, also see Helmer et al. (2020), Supplementary Figure S7). The general framework used is presented in Figure 6.3.



**Figure 6.3:** A schematic of the CCA subsampling scheme. CCA was performed on the full-sample datasets  $\mathbf{X}_{n_s}$  and  $\mathbf{Y}_{n_s}$  and the canonical weights  $\mathbf{W}_{n_s}$  and canonical correlations  $\boldsymbol{\rho}_{n_s}$  are obtained. Next, CCA was repeated on subsets  $\mathbf{X}_{n_j}$  and  $\mathbf{Y}_{n_j}$  for  $n < n_s$ , storing the canonical weights  $\mathbf{W}_{n_j}$  and canonical correlations for comparison  $\boldsymbol{\rho}_{n_j}$ . In addition, at each sample size, 5-fold cross-validation is performed (not shown).

The following describes feature extraction and preparation and the trends ob-

served through the comparison of subsampled  $n < n_S$  CCA to the full set  $n = n_S$  CCA using the multi-modal HCP data.

#### 6.4.1 Feature Extraction: Imaging Features

Imaging features were derived using the dMRI and fMRI data from the HCP cohort and CCA was performed using each of these imaging modalities against the behavioural measures. The following sections describe both imaging feature sets and the behavioural data.

#### dMRI HCP Data

First, we used the diffusion MRI data from the HCP cohort, utilising the structural connectivity features extracted in Chapter 4 and performed CCA against behavioural features. In line with the HCP recommendations, a set of 44 subjects identified as having anatomical abnormalities were removed from the dataset. In addition, one further subject was removed during data preprocessing (see below), leaving  $n_s = 1,020$  subjects.

Here, the subject-wise connectivity blueprints described in Section 4.2.4 were used as structural connectivity features. The connectivity blueprints represent the cortical termination maps for 41 major white matter (WM) fibre bundles. The blueprint matrices (~ 60k vertices by T = 41 tracts) were parcellated along rows, i.e. each column parcellated individually, into 68 regions using the Desikan-Killany atlas (Desikan et al., 2006) with each parcel representing the median of the constituent vertices. The parcellated connectivity blueprints (68 × 41) were then vectorised and stacked subject-wise, giving the imaging dataset  $\mathbf{X}_{n_s}$  (1,020 × 2,788).

#### fMRI HCP Data

In addition, we also use measures derived from the HCP resting-state functional connectivity data. For these analyses, an extended set of the imaging features as used in Smith et al. (2015) were used. Following subject removal (see below),  $n_s = 775$  subjects remained.



**Figure 6.4:** A schematic of the imaging-data preprocessing. **1.)** XTRACT was used to reconstruct 41 major WM fibre bundles in each of  $n_s = 1,020$  HCP subjects. The tracts were vectorised and stacked tract-wise to give the matrix  $C_2$  (M voxels by T tracts). **2.)** The white-grey boundary (N vertices) to WM (M voxels) connectivity matrix is constructed for each subject, giving the  $C_1$  matrix. **3.)** The multiplication of  $C_1$  and  $C_2$  gives the connectivity blueprint **CB** with N rows, corresponding to the white-grey boundary vertices, and T rows, corresponding to tracts. **4.)** The blueprints, visualised on the cortex here, were parcellated using the Desikan-Killany atlas (Desikan et al., 2006), reducing the dimensionality from  $N \approx 60k$  to P = 68. **5.)** The parcellated blueprints were vectorised and stacked subject-wise to give a subject  $n_s$  by imaging feature  $p = P \times T = 2,788$  matrix. This matrix will be reduced via principal component analysis prior to CCA.

The data contains parcellated functional connectomes at multiple dimensionalities. These were obtained by applying group-ICA (Smith et al., 2014) to the preprocessed fMRI data described in Section 3.2.3. The parcellated connectomes with a dimensionality of 200 were used in this analysis. The upper triangle of each subject's connectome, representing the unique entries of a symmetric matrix, were then vectorised and stacked subject-wise, giving the imaging dataset  $\mathbf{X}_{n_s}$  (775 × 19, 900).

#### **Imaging Feature Preprocessing**

Imaging features in both cases were first z-scored and, separately, feature-wise normalised (by division of the absolute feature mean). These were concatenated, giving 2 × 2,788 and 2 × 19,900 features for the dMRI and fMRI data sets respectively. Subsequently, the features were de-confounded as  $\mathbf{X}_{decon} = \mathbf{X} - \mathbf{C}\boldsymbol{\beta}$  where  $\mathbf{X}$  represents each imaging data set,  $\mathbf{C}$  are the confounds (described below) and  $\boldsymbol{\beta} = \mathbf{C}^+ \mathbf{X} = (\mathbf{C}^T \mathbf{C})^{-1} \mathbf{C}^T \mathbf{X}$ . Finally, the imaging datasets were reduced via PCA to the top k principal components, giving the final input dataset  $\mathbf{X}_{n_s}$  ( $n_s \times k$ ).

#### 6.4.2 Feature Extraction: Non-Imaging Features

The set of 158 behavioural measures used in Smith et al. (2015) were used here, giving a non-imaging matrix  $\mathbf{Y}_{n_s}$  ( $n_s \times 158$ ) (for a full list of the behavioural measures, see Appendix 6.9.2). Broadly, this includes measures of cognitive ability (e.g. fluid intelligence, working memory), psychological well-being (e.g. general life satisfaction, feelings of sadness), family history (e.g. history of depression, history of drugs/alcohol), alcohol and substance use (e.g. testing positive for THC, typical alcohol consumption), tobacco and alcohol dependence, and physical measures (e.g. dexterity, use of corrective eyeglasses).

The features were transformed through a rank-based inverse normal transformation and z-scored (Beasley et al., 2009). Next, the subject-wise covariance matrix was calculated  $(n_s \times n_s)$  across features whilst only considering subject pairs for whom data was available for each subject. This was subsequently projected to the nearest positive definite matrix (Seabold and Perktold, 2010), avoiding the need for
data imputation in cases of missing entries. Finally, features were reduced via PCA to the top k principal components, giving the final input dataset  $\mathbf{Y}_{n_s}$   $(n_s \times k)$ .

### Confounds

Similarly to Smith et al. (2015), a set of confound measures were defined which were regressed out of both the imaging and non-imaging datasets. Confound measures include height, weight, intracranial volume and brain volume. In addition, volume measures were raised to the power of 1/3 and the square-terms of height, weight and motion were included in the confound matrix. In the fMRI analysis, head motion during functional MRI data acquisition and software version were used as additional confounds. Missing values were set to 0 (3% and 5% for the fMRI and dMRI analyses respectively) and features were z-scored across subjects.

### 6.4.3 Analysis

CCA was performed using  $n \leq n_s$  and the canonical correlations ( $\rho_{n_s}$  and  $\rho_n$ ) and canonical weights ( $\mathbf{W}_{n_s}$  and  $\mathbf{W}_n$ ) were compared. Weight stability, the estimated canonical correlation and the cross-validation correlation were plotted as a function of *SFR* and inflation was calculated relative to  $\rho_{n_s}$ . CCA was performed at the full cohort sample size and at 5 logarithmically spaced subsample-sizes between 202 (to avoid rank deficiency) and 80% of the full sample size.

Statistical significance was assessed via permutation testing. For the HCP cohort, the permutation matrix was defined using FSL's PALM (Winkler et al., 2015) such that the family structure inherent in the HCP is respected and maintained. For these analyses,  $N_p = 1,000$  permutations were used. In cases where permutation indices were not returned, the corresponding subjects were excluded from the analysis: 3 subjects in the fMRI analysis and 1 subject in the dMRI analysis.

# 6.4.4 Results

The CCA subsampling scheme was performed using the top k = 100 principal components in both the dMRI and fMRI cases. These components explained 98.6% of the non-imaging (**Y**) data and 94.9% and 28.9% of the dMRI and fMRI data respectively. In each case, at least one significant mode (p < 0.001) was identified at the full-sample  $n = n_s$ . For both modalities, the observed canonical correlation increases as the number of subjects, and therefore SFR, decreases with a maximum inflation of 9.65% whilst remaining statistically significant: orange curves in the left column plots in Figure 6.5. The cosine similarity between the subsampled canonical weights ( $\mathbf{W}_n$ ) and the full-sample weights ( $\mathbf{W}_{n_s}$ ) reduces as SFR decreases to as low as 0.33 whilst remaining statistically significant: right column plots in Figure 6.5. Further, the out-of-sample correlations obtained via cross-validation are considerably lower than the in-sample estimates whilst decreasing further with reducing SFR, demonstrating lack of generalisability: red curves in the left column plots in Figure 6.5. Interestingly, all these trends are very similar for both cases. Two very different and independent set of features obtained from different imaging modalities (dMRI and fMRI) return almost equivalent profiles and trends, which seem to be governed by the SFR ratio.

### 6.4.5 Summary

The trends observed demonstrate that the canonical correlation inflates, the weight stability reduces and the cross validation performance decreases as the SFR decreases. This suggests that the canonical correlation values reported in the literature may be dependent on the data dimensionality. Further, given the relationship between weight stability and SFR, it suggests that importance assigned to a given feature depends on the SFR: a particular feature may be assigned high importance in one case but may be assigned low importance in another simply by changing the number of subjects. This has further consequences when considering the original (pre-PCA) feature space loading vectors.

Overall trends suggest that CCA results are sensitive to the SFR. Further, these results confirm the trends observed are not unique to a particular modality but generalise across common modalities. We also tried the analysis with a different number of PCA components, i.e. k varying from 50 to 150, with the overall trends



**Figure 6.5:** The estimated canonical correlation and weight stability with reducing numbers of subjects (i.e. SFR) for the HCP multi-modal data: dMRI top row, fMRI bottom row. Left: the relationship between the canonical correlation and SFR. The orange line shows the increasing correlation  $\rho$  with decreasing SFR with the lightorange area showing the null-distributions. The red line indicates cross-validation estimated correlation. The circle and triangle symbols correspond to in-sample and out-of-sample canonical correlations, respectively, when the full dataset is considered. Right: the relationship between weight stability and SFR. The yellow line shows the pairwise cossine similarity between the full-sample  $n_s$  weight vectors and the subsampled weight vectors. Analysis was performed using the top k = 100 principal components.

being similar for corresponding SFRs. However, these analyses are limited in their interpretation as we are considering the  $n = n_s$  case as the ground-truth, which is not the case. The  $n = n_s$  case may already suffer from overfitting. In order to formally quantify these instabilities, the ground-truth must be known.

# 6.5 CCA with Known Ground-Truths

As discussed, trends observed in the real-data examples suggest a strong sensitivity to the *SFR*. To further explore these behaviours in a robust manner, a simulation framework was devised to generate sets of input data with known ground-truth between-set correlations.

# 6.5.1 Simulation Framework

The framework generates sets of features (**X** and **Y**) with known ground-truth canonical correlations  $\rho$  and weights  $\mathbf{W}_X$  and  $\mathbf{W}_Y$ . This was achieved by generating an  $n_s$  by 2 × k matrix **Q** of random values drawn from a Normal distribution. **Q** was then globally demeaned and variance normalised before being orthonormalised via Gramm-Schmidt orthonormalisation (Schmidt, 1907; Chen, 2020), ensuring that each column of **Q** is orthogonal to each other column. This was then split equally into two, giving two matrices with dimensions of  $n_s$  by k = p = q representing the canonical variates **U** and **V** (step 1 in Figure 6.6).

To introduce the canonical correlations, columns of  $\mathbf{V}(\mathbf{v}_i)$  were rotated with respect to columns of  $\mathbf{U}(\mathbf{u}_i)$ . Vector rotation was achieved by first calculating the required rotation angle as  $\theta_i = \cos^{-1}(\rho_i)$ , where  $\rho_i$  is the desired canonical correlation for a given canonical mode and  $\theta_i$  is the required rotation.  $\mathbf{v}_i$  was then rotated with respect to  $\mathbf{u}_i$ :  $\mathbf{v}_i = \mathbf{v}_i + (\frac{1}{\tan(\theta_i)} \times \mathbf{u}_i)$ , thus forming a canonical mode (step 2 in Figure 6.6). To ensure that the rotation has been successful, the correlation between the rotated vectors ( $\mathbf{v}_i$  and  $\mathbf{u}_i$ ) was compared to the desired  $\rho_i$ . For these purposes, the top canonical correlation  $\rho_1$  is user-defined and the canonical correlations for the following k - 1 modes were calculated as  $\rho_n = \rho_{n-1} - 0.05$ , n = 2, ..., k. This produces a set of correlations which reduce across mode number in a manner not too dissimilar to typical CCA results.

Next, the canonical weights  $\mathbf{W}_X$  and  $\mathbf{W}_Y$  were generated (step 3 in Figure 6.6). This was achieved by first drawing k samples from a Gamma distribution. This vector represents the column-wise standard deviation for a given canonical



**Figure 6.6:** Schematic of the simulation framework used to generate input data **X** and **Y** with known ground-truth CCA solutions.

weight matrix. These standard deviations were then used to construct columns of the canonical weights using  $\mathbf{w}_{Xi} = \mu_i + \sqrt{\sigma_i} \times \delta_i$ , where  $\mu_i$  represents the column mean ( $\mu = 0$ ),  $\sigma_i$  is the standard deviation and  $\delta_i$  is a "noise" parameter drawn from a Gaussian distribution.

Finally, the canonical weights ( $\mathbf{W}_X$  and  $\mathbf{W}_Y$ ) and canonical variates ( $\mathbf{U}$  and  $\mathbf{V}$ ) were used to calculate the data matrices  $\mathbf{X}$  and  $\mathbf{Y}$  as  $\mathbf{X} = \mathbf{U}\mathbf{W}_X^{-1}$  and  $\mathbf{Y} = \mathbf{V}\mathbf{W}_Y^{-1}$ (step 4 in Figure 6.6).

Prior to analysis, in order to ensure that the selected  $n_s$  is large enough and captures converged trends, the simulation framework was used to investigate canonical correlation inflation for a single predefined correlation value and various magnitudes of  $n_s$ . The results from this investigation inform the selection of  $n_s$ . Sanity checks were also performed to ensure that canonical weight properties were as intended and that canonical modes obey the CCA criteria ( $\rho_1 \ge \rho_2 \ge ... \ge \rho_k$ ,  $corr(\mathbf{u}_1, \mathbf{u}_{2...k}) = 0$ and  $corr(\mathbf{v}_1, \mathbf{v}_{2...k}) = 0$ ).

Following these sanity checks, the framework was used to generate sets of input data  $\mathbf{X}_{n_s}$  and  $\mathbf{Y}_{n_s}$  with varying ground-truth canonical correlations  $\boldsymbol{\rho}_{n_s}$ . As in the real-data example, CCA was performed on  $\mathbf{X}_n$  and  $\mathbf{Y}_n$  with  $n < n_s$  to estimate the

subsampled canonical correlation  $\rho_n$  and canonical weights  $\mathbf{W}_{Xn}$  and  $\mathbf{W}_{Yn}$  which were then compared to the ground-truth parameters  $(\rho_{n_s}, \mathbf{W}_{Xn_s}, \mathbf{W}_{Yn_s})$ .

### 6.5.2 Results

CCA input matrices **X** and **Y** were generated with k = 100 features each and  $n_s = [1 \times 10^3, 1 \times 10^4, 1 \times 10^5, 1 \times 10^6, 1 \times 10^7]$  and a  $\rho_1$  of 0.7. Figure 6.7a shows cross-correlation matrices within and across the canonical variates **U** and **V**, i.e. the column-wise correlation corresponding to the correlation between column pairs. The diagonal of the cross correlation matrix in the left-hand plot  $(corr(\mathbf{U}, \mathbf{V}))$  represents the the canonical correlation values, starting at 0.7 and descending across the modes, fulfilling the CCA criteria  $\rho_1 \ge \rho_2 \ge ... \ge \rho_k$ . The middle and right-hand plots show the correlation within the canonical variates  $(corr(\mathbf{U}, \mathbf{U}))$  and  $corr(\mathbf{V}, \mathbf{V})$  which take the form a identity matrices. In accordance with the CCA criteria, the off-diagonal entries for each of the three plots are zero, indicating that each variate (e.g.  $\mathbf{v}_1$ ) is orthogonal to all other variates except for the corresponding variate from the second modality (e.g.  $\mathbf{u}_1$ ), i.e.  $corr(\mathbf{v}_n, \mathbf{u}_n) = \rho_n$ ,  $corr(\mathbf{u}_n, \mathbf{u}_j) = 0$  and  $corr(\mathbf{v}_n, \mathbf{v}_j) = 0$  with  $n = 1...k, j \neq n$ .

Figure 6.7b depicts the column-wise mean and standard deviation of the canonical weights  $\mathbf{W}_X$  and  $\mathbf{W}_Y$ , showing that each distribution of mean values is centered around zero and the standard deviation follows a positive Gamma-like distribution. Figure 6.7c shows how the estimated canonical correlation changes with SFR for each of the generated input datasets with a ground-truth of  $\rho_1 = 0.7$ . It can be seen that in each case of  $n_s$ , the overall trend is consistent apart from when  $n_s = 1 \times 10^3$ . In other words, 1,000 subjects is insufficient to represent the converged trend in these simulations. With  $n_s \geq 1 \times 10^4$ , the trend is consistent. Based on these observations, a sample size  $n_s$  of  $1 \times 10^5$  was selected, erring on the cautious side whilst significantly reducing the computation expenditure associated with larger  $n_s$ . With k = p = q = 100, this gives a SFR of 500, orders of magnitude larger than any used in the literature.

Input data  $\mathbf{X}_{n_s}$  and  $\mathbf{Y}_{n_s}$  were then generated using first mode canonical corre-



Figure 6.7: Results from the simulation framework sanity checks. **a.**) The correlation matrices within and between the generated canonical variates (U and V). **b.**) Histograms of the mean and standard deviation of the generated canonical weights  $\mathbf{W}_X$  and  $\mathbf{W}_Y$ . **c.**) Plots of the increasing  $\rho$  with SFR for varying samples sizes  $n_s$  with a ground-truth  $\rho = 0.7$ , demonstrating overall trend correspondence when  $n_s \geq 1 \times 10^4$ .

lation values in the range of  $\rho_1 = 0.1 - 0.99$  with  $n_s = 1 \times 10^5$  and k = p = q = 100. CCA was performed using  $n \leq n_s$ . The difference between the estimated and ground-truth canonical correlations  $(\boldsymbol{\rho}_n - \boldsymbol{\rho}_{n_s})$  and the correlation between the ground-truth and estimated weights  $(corr(\mathbf{W}_{n_s}, \mathbf{W}_n))$  are plotted for the top canonical mode (Figure 6.8). Statistical significance is determined via permutation testing  $(N_p = 100,000)$  with modes surviving significance testing  $(p \leq 0.001)$  highlighted in the grey region of Figure 6.8.

### 6.5.3 Summary

In each simulated case, canonical correlation inflation and weight instabilities were observed whilst remaining statistically significant. For lower correlation values (<



**Figure 6.8:** Plots of the difference between the estimated and ground-truth canonical correlations ( $\rho$ , left) and the correlation between the ground-truth and estimated canonical weights (**W**, right), i.e. stability, with SFR (top canonical mode only) for various ground-truth  $\rho$  values.

0.2), these effects were more extreme, however, significance was lost more rapidly. This occurs due to a lack of power. The sample size was not sufficient to detect true underlying associations at the given data dimensionality, i.e. false negatives were reported. For very high correlation values (> 0.8), there is little room for inflation and the weights remain relatively stable until reaching low SFR, at which point the correlation inflates and weight instability is observed. For mid-range correlation values (0.3 - 0.6), which are likely to represent more likely scenarios in true data, inflation is in the order of 0.2 and weight stability is within a range of 0.1 - 0.3 for  $SFR \approx 4$  (typical of the literature). Lower values of SFR become underpowered and false negatives are reported. For canonical correlations of 0.3 and above, to detect statistically significant correlations and keep the weight error to less than 10%, we need at least 50 subjects per feature, corresponding to a weight stability greater than 0.82. This is a much larger SFR than that found in much of the literature. These simulations demonstrate the dependencies between the observed canonical correlation and the true underlying canonical correlation, the number of samples/subjects, and the number of features.

We may conclude from this that the application of CCA may be more appro-

priate to data with either low numbers of features or with large sample sizes. Given that CCA seeks to find low dimensional representations of high dimensional data, it's usefulness and applicability would be reduced if the number of features is very low. Therefore, applying CCA in large cohorts may be optimal. To this end, the following section describes the application of CCA to one of the largest neuroimaging cohorts available, the UK Biobank, with which we can actually achieve the SFRs suggested above and confirm their validity by exploring stability. A similar exploration to that presented in the current and previous sections is first presented, followed by a closer inspection and interpretation of the consequences of the described instabilities.

# 6.6 CCA Applied to Large Samples

In order to assess these trends at the population-level, we used the largest neuroimaging dataset available to date, the UK Biobank. The UK Biobank contains a wide range of subject measures (Sudlow et al., 2015) and multi-modality imaging data, including fMRI data, allowing for similar brain-behaviour associations previously studied to be explored on a large scale ( $\approx 40$ k subjects). Due to computational limitations, a subset of  $n_s = 20,000$  subjects with complete data (i.e. with both imaging and non-imaging data and  $\geq 50\%$  complete entries in the non-imaging data) were randomly selected from the full cohort.

# 6.6.1 Feature Extraction: Imaging Features

The pre-processed resting-state fMRI data were utilised for these analyses (Miller et al., 2016; Alfaro-Almagro et al., 2018). These data include partial correlation parcellated connectomes with a dimensionality of  $55 \times 55$ , with nodes obtained from ICA analysis. The connectomes were z-score transformed and the upper triangle vectorised to give 1,485 functional connectivity features per subject, for each of the 20,000 subjects. Following this, the imaging features were preprocessed as in the HCP analyses, including normalisation, de-confounding and dimensionality reduction via PCA to the top k = 100 components, giving the final imaging dataset  $\mathbf{X}_{n_s}$   $(p \times k = 20,000 \times 100).$ 

### 6.6.2 Feature Extraction: Non-Imaging Features

The UK Biobank contains a wide range of subject measures (Sudlow et al., 2015), including physical measures (e.g. weight, height), food and drink, cognitive phenotypes, lifestyle, early life factors and sociodemographics. A subset of 389 cognitive, lifestyle and physical measures, as well as early life factors were hand-selected whilst considering data completeness (measures that had missing values in more than 50% of subjects or measures that had identical values in at least 90% of subjects were not selected). For categorical items, negative values were replaced with 0, as in Miller et al. (2016). Such negative values encode mostly "Do not know"/"Prefer not to answer". Measures with multiple visits were then averaged across visits, reducing the number of measures to 226. A redundancy check was then performed. Specifically, if the correlation between any two measures was > 0.98, one of the two items was randomly chosen and dropped. This procedure removed 2 measures, resulting in a final set of 224 behavioural measures. Finally, confound regression and dimensionality reduction via PCA were performed giving the non-imaging dataset  $\mathbf{Y}_{ns}$  ( $p \times k = 20,000 \times 100$ ).

### Confounds

The following items were used as confounds: acquisition protocol phase (due to slight changes in acquisition protocols over time), scaling of T1 image to MNI atlas, brain volume normalized for head size (sum of GM and WM), fMRI head motion, fMRI signal-to-noise ratio, age, sex. In addition, similar to Miller et al. (2016) we used the squares of all non-categorical items (i.e. T1 to MNI scaling, brain volume, fMRI head motion, fMRI signal-to-noise ratio and age), as well as age  $\times$  sex and  $age^2 \times$  sex. Altogether, these were 14 confounds. Finally, we imputed 0 for missing values (6% of entries) and z-scored all items.

### 6.6.3 Analysis

The CCA subsampling scheme was applied as in the HCP analyses, including comparisons of the estimated canonical correlation, weight stability and 5-fold crossvalidation performance with the SFR.

Further analyses were then performed to assess the implications of the observed instabilities during interpretation of the behaviour data. This was achieved by directly comparing the non-imaging loading vectors  $\mathbf{l}_Y$  describing how each of the non-imaging features (in their original feature space) contribute to the canonical modes. CCA was performed using n = [20,000,5,000,1,000] and the top 10 modes were extracted. The loading vectors were obtained by correlating the weight vectors  $(\mathbf{W}_Y)$  with the deconfounded original 224 non-imaging features, as is common practice in the literature. The subsequent loading vectors were correlated across sample sizes, comparing each of the subsampled cases n = [5,000,1,000] to the full-sample case n = 20,000. Finally, the loading vectors were summarised using the top contributors in a positive-negative distribution and the correspondence of behavioural feature importance was assessed.

# 6.6.4 Results: Overall Stability

Both the imaging and non-imaging datasets were reduced to the top k = 100 principal components, explaining 63.3% and 96.5% of the original variance respectively. With k = p = q = 100, the maximum *SFR* in these analyses is  $\frac{20,000}{100+100} = 100$ , surpassing the previously recommended *SFR* of 50.

Figure 6.9 shows the effect of SFR on the canonical correlation, out-of-sample performance and weight stability for the UK Biobank. As expected from the HCP and simulation results, the correlation increases (by as much as 31.6%) and the weight stability decreases (to as low as 0.31) with decreasing SFR, reinforcing the conclusion that CCA becomes unstable in the case of even large-sized cohorts. We find that when approaching the full UK Biobank cohort, there is near identical correspondence between weight vectors (weight similarity  $\approx$  1). In addition, the



Figure 6.9: The canonical correlation and weight stability plotted against the number of subjects (i.e. SFR). Left: the relationship between the canonical correlation and SFR. The orange line shows the increasing correlation with decreasing SFR with the light-orange line showing the null-distributions. The red line indicates cross-validation performance with SFR. Right: the relationship between weight stability and SFR. The yellow line shows the pairwise cossine similarity between the full-sample  $n_s = 20,000$  weight vectors and the subsampled weight vectors.

correlation curve flattens, indicating that there is little/no inflation at larger sample sizes, and the cross-validation performance and correlation curves converge, suggesting that the in-sample and out-of-sample canonical modes are robust, i.e. there is no overfitting.

# 6.6.5 Results: Correspondence Between Modes

Figure 6.10 shows the observed canonical correlations for the full sample and each of the subsampled analyses for the top 10 canonical modes. In the n = 20,000 subject analysis, all ten modes are statistically significant, as assessed via permutation testing with  $N_p = 10,000$ . This reduces to 7 and 2 significant modes for the n = 5,000and n = 1,000 analyses respectively.

Mode correspondence was assessed by correlating each of the n = 20,000 loading vector sets with each of the loading vector sets from n = [20,000,5,000,1,000], building a correlation matrix, shown in Figure 6.11. In cases where the sign of the mode has been arbitrarily flipped in the CCA solution (with respect to the n = 20,000 case), the loading vector was multiplied by -1. In the within n = 20,000comparison (left correlation matrix, Figure 6.11), a diagonal is observed with low-



**Figure 6.10:** The observed canonical correlations for each of the top 10 canonical modes across sample sizes. Statistical significance is indicated by \* and the grey violin plots depict the null-distributions built via permutation testing with 10,000 permutations.

level correlation values  $(r \leq 0.3)$  in the off-diagonal, as would be expected.

When comparing the n = 20,000 and n = 5,000 loading vectors (middle correlation matrix, Figure 6.11), a general diagonal trend is observed, indicating overall correspondence between equivalent modes. However, the overall correlation (i.e. correspondence) decreases across the modes and, in some cases, modes appear to be mixed or swapped. For example, modes 3 and 4 from n = 20,000 correlate similarly to modes 3 and 4 from n = 5,000, suggesting that these modes are mixed in the n = 5,000 case. The first two modes are strongly correlated across sample sizes  $(r_1 = 0.94, r_2 = 0.84)$ .



**Figure 6.11:** Correlation matrices showing the correspondence between canonical modes derived from the n = 20,000 case and each subsampled analysis. Loading vectors from the n = 20,000 (y - axis) are correlated with each loading vector from the n = 20,000 (left), n = 5,000 (middle) and n = 1,000 (right), shown on the x - axes.

When comparing the n = 20,000 and n = 1,000 loading vectors (right correlation matrix, Figure 6.11), the diagonal is much less prominent. The correspondence between the first modes is weak with a correlation of r = 0.44 and there is apparent mixing with the second mode. The second mode from the n = 1,000 case is weakly correlated (r = 0.16) to the second mode from n = 20,000. Beyond this, a very weak diagonal may be observed, however, the n = 1,000 modes are non-significant (p > 0.1).

To assess how these instabilities effect interpretation, the loading vectors for the behaviour data are visualised. The top 15 positive<sup>2</sup> (blue) and negative (red) contributors for each of the ten canonical modes from the n = 20,000 analysis are presented in Appendix 6.9.4. Interpretation of these robust modes is challenging (see Section 6.7 "Data encoding") although we may draw some conclusions with caution. The first mode seems to capture associations of socioeconomic status and overall mental health, including measures of depression, social relationships and intelligence. The second mode describes associations primarily of mental and physical health, including measures of exercise, sleep quality and stress. Mode three also describes associations of mental health specific to happiness, regular exercise and social relationship quality. The fourth mode describes associations between mental and physical health, including body mass index (BMI), smoking history and measures of poor mental health.

For comparing across sample sizes, a minimum of the top 15 positive and 15 negative contributors were considered, matching the lower loading value threshold across subsampled analyses for a given mode (Figure 6.12). As before, the sign of the modes is matched across analyses to allow for comparison. The top contributors from the subsampled cases were then compared to the n = 20,000 case with correspondence (green-highlighting) and anti-correspondence (red-highlighting) highlighted, i.e. in the case where loading vectors match perfectly across sample sizes, each feature would be highlighted in green.

 $<sup>^{2}</sup>$ Here, the terms "positive" and "negative" simply refer to the sign of the loading for the given feature. This does not imply positive or negative aspects of behaviour.



(a) Top contributors for the first canonical mode. In the n = 20,000 case, the mode describes the relationship between measures of mental health and socioeconomic status.



(b) Top contributors for the second canonical mode. In the n = 20,000 case, the mode describes the relationship between measures of mental health and socioeconomic status.

#subjects = 20000, p-value = 0.0001	#subjects = 5000, p-value = 0.0001	#subjects = 1000, p-value = 0.148
0.37 - Felt hated by family memb	0.45 Length of mobile phone us	0.23 - Nap during day
Belittlement by partner o	Belittlement by partner o	Ever believed in un-real
Length of mobile phone us	Sleeplessness / insomnia	Ankle spacing width (left Duration of other exercis
EI5 - family relationship	Number of correct matches	Belittlement by partner o
	Job involves heavy manual	Smoked cigarette or pipe
Hand grip strength (right	FI5 : family relationship	Recent worrying too much
Difference in mobile phon	Ankle spacing width (left	Mother still alive
0.25 - Total errors traversing a	Difference in mobile phon	Hands-free device/speaker
	Fed-up feelings	Total errors traversing n
Sieepiessness / Insomnia	0 16 - Maternal smoking around b	Ever believed in an un-re
Handedness (chirality/lat	Sensitivity / hurt feelin	Frequency of tenseness /
Frequency of other exerci		Forced expiratory volume
Time spent outdoors in wi	Hand grip strength (right	0.13 - Time spent watching telev
Time spent outdoors in wi	Average weekly beer plus	0.13 - Alcohol usually taken wit
Place of birth in UK - no	Total errors traversing a	-0.14 - Friendships satisfaction
Nervous feelings	Been in a confiding relat	Impedance of whole body
Recent changes in speed/a	Number of days/week of vi	Felt loved as a child
Necent changes in speeu/a	Duration of other exercis	Qualifications
Job involves heavy manual	Guilty feelings	Diagnosed with life-threa
Recent easy annoyance or	Ever addicted to any subs	Physically abused by fami
e Age first had sexual inte	Average total bousebold i	Number of incorrect match
		Digits entered correctly
ອ 0.15 - Frequency of light DIY in	Ever had prolonged feelin	Duration to complete nume
og -0.14 - Number of symbol digit ma	Time to complete round*	Worrier / anxious feeling
Duration screen displayed	0.13 - Duration spent answering	Worry too long after emba
	-0.14 - Alcohol intake frequency	Health satisfaction
Seen doctor (GP) for nerv	Ever believed in an un-re	Prospective memory result
FI3 : word interpolation	Frequency of walking for	lob involves mainly walki
Maximum digits remembered	Alcohol intake versus 10	Snoring
	Mother's age at death	Number of full sisters
Year ended full time educ		-0.17 - Place of birth in UK - ea
Frequency of stair climbi		Morning/evening person (c
Happiness	-0.16 - Duration to first press o	Number of older siblings
0.21 Usual cide of head for me	Frequency of stair climbi	Ever heard an un-real voi
	Frequency of travelling f	Alcohol intake frequency
Distance between home and	Happiness	Witnessed sudden violent
Mother's age at death	FI5 : family relationship	Ever unenthusiastic/disin
lob involves shift work	Time spent driving	Daytime dozing / sleeping
Job involves shire work	Body mass index (BMI)	Father's age at death
Ever sought or received p	Mood swings	PM: initial answer
Hand grip strength (left)	Avoided activities or sit	FI6 : conditional arithme
Recent poor appetite or o	Avoided activities of still	Job involves heavy manual
	Recent thoughts of suicia	Time elapsed
-0.35 - Number of days/week of mo	-0.23 - Ever sought or received p	-0.22 - Number of days/week of mo

(c) Top contributors for the third canonical mode. In the n = 20,000 case, the mode primarily describes the relationship between measures of mental health and socioeconomic status.

**Figure 6.12:** The top non-imaging contributors for each subsampled analysis for the top three canonical modes. A minimum of the top 15 positive and negative contributors are shown in each case, matching the lower positive and negative threshold values across subsampled analyses for a given mode. The n = [5,000,1,000] analyses are compared to the n = 20,000 case. Green highlighting indicates that the contributor appears in the n = 20,000 case, i.e. contributor correspondence. Red highlighting indicates that the contributor appears on with a flipped sign, i.e. anticorrespondence.

For the first mode (Figure 6.12a), strong correspondence is observed between the n = 5,000 and n = 20,000 cases, as suggested by Figure 6.11. This reduces when comparing the n = 1,000 and n = 20,000 cases: many of the top positive contributors are present however the negative contributors are no longer considered as "top contributors", due to instabilities in the CCA weights. The issue of reduced overall correspondence and anti-correspondence becomes more apparent in the second and third modes (Figure 6.12b-c). In the n = 1,000 case, the sign of a non-imaging measures have been flipped, for example the "Sleepiness/insomnia", "Some to take to the doctor..." and the "Recent worrying too much..." measures in the second mode. This demonstrates that, even when statistically significant, the stability of such visualisations may lead to conclusions that do not represent true underlying relationships.

# 6.6.6 Summary

As in the previous explorations with the diffusion and functional MRI connectivity features from the HCP cohort and through the simulation framework, the limitations of CCA have been demonstrated. The canonical correlations and canonical weights are strongly dependent on the ratio between the number of subjects and the number of features. Further, investigations into how these instabilities effect the interpretation of the CCA results have been presented, highlighting how erroneous conclusions may be reached. The results suggest that in order to achieve robust solutions, high SFR values are required, orders of magnitude higher than what used before.

# 6.7 Discussion

This chapter presents investigations into the stability of the multivariate statistical technique canonical correlation analysis (CCA), a popular method used to reveal latent associations in brain-behaviour datasets (Smith et al., 2015; Drysdale et al., 2017; Marquand et al., 2017; Bijsterbosch et al., 2018; Wang et al., 2018a; Dinga et al., 2019; Li et al., 2019; Becker and Hervais-Adelman, 2019; Han et al., 2020; Wang et al., 2020b; Lv et al., 2020). The focus of these studies is to interpret the canonical weights in a biologically meaningful way by first extracting the low dimensional canonical modes, typically from dimensionality reduced input data, and then relating those modes back to the original data-space to obtain a set of loading vectors. However, we have shown that the stability of CCA, and therefore the loading vectors and their interpretation, is directly linked to the subject-feature ratio and the true canonical correlation.

We have shown that these concerns generalise across modalities, using connectivity features extracted from diffusion MRI and functional MRI; across data-quality and cohorts, using data from the HCP and UK Biobank cohorts; and in simulated data with known ground-truths. We have applied CCA to one of the largest neuroimaging cohorts available, the UK Biobank, and show that very-large cohorts are required in order to avoid overfitting when using feature dimensionalities typical of the literature. The observed correlation inflation is much higher in our UK Biobank analysis compared to the HCP cohort analyses. This is likely to reflect the already inflated estimated correlations when using the full HCP cohort.

Instabilities in the loading vectors relating to the subject-feature ratio have been demonstrated. Using the UK Biobank cohort, loading vectors were extracted and compared for a range of sample sizes. The subsampled cases provide sample sizes in the more typical range of 1,000 subjects used in previous studies. More importantly, the subject-feature ratio in these cases is larger than, or typical of, the literature (see Table 6.1). The full sample case, with 20,000 subjects (SFR = 100), is, as far as we are aware, the largest study of its kind and is demonstrably robust the observed canonical correlation converges with the cross-validation performance in Figure 6.9.

These findings indicate that sample sizes much larger than those typically used in the literature are required in order to avoid overfitting and the associated instabilities. For true canonical correlations in the order of 0.3 (believed to be typical of the true underlying canonical correlation in the literature - see Figure 6.13 and Helmer et al. (2020)), we suggest a subject-feature ratio of 50. This agrees with previous literature (not specific to neuroimaging) where it is argued that at least 10-70 samples should be used per feature to ensure that the solutions are robust (Barcikowski and Stevens, 1975; Thompson, 1990; Leach and Henson, 2014). The reviewed CCA literature, although not an exhaustive review, falls short of this. Interestingly, the effect of the subject-feature ratio is strong enough that canonical correlations reported in the literature may be estimated based on the subject-feature ratio alone (Figure 6.13 and Helmer et al. (2020)).



**Figure 6.13:** Canonical correlations and the number of samples per features are extracted from the literature (each dot is a canonical correlation reported in the literature) and overlaid on predictions from a generative model. Many studies employed a small number of samples per feature and reported a large canonical correlation. These studies fall in the top-left corner of the plot. The reported canonical correlation can be predicted from the used number of samples per feature alone using linear regression ( $R^2 = 0.83$ ). We also estimated the weight error (encoded in the colourbar) for each reported CCA. The farther away a CCA lies from the predictions for permuted data the lower the mean-estimated weight error. As can be seen, the majority of studies fall between  $0.1 < r_{true} < 0.5$ , suggesting that the true underlying canonical correlation is likely in this range. (reproduced from Helmer et al. (2020))

## 6.7.1 Recommendations

CCA and other multivariate association techniques are powerful and potentially useful tools. However, we have demonstrated limitations associated with the subjectfeature ratio. Various steps may be taken to reduce the risk of overfitting, focusing on the subject-feature ratio or on the applied technique. As mentioned, we suggest a subject-feature ratio of at least 50. There are two approaches to achieving this: increasing the number of subjects or decreasing the number of features.

#### **Data Dimensionality**

Substantially increasing the number of subjects is not easily achieved in the majority of neuroimaging studies. Large cohort, multimodal datasets may be used to maximise the number of subjects, for example the HCP (Van Essen et al., 2013), the Alzheimer's Disease Neuroimaging Initiative (ADNI) (Jack et al., 2008), the UK Biobank (Miller et al., 2016; Alfaro-Almagro et al., 2018), the Adolescent Brain Cognitive Development (ABCD) (Casey et al., 2018), and the Enhancing NeuroImaging Genetics through Meta-Analysis (ENIGMA) Consortium (Thompson et al., 2014). Recent advancements in the harmonisation of data across studies (Yamashita et al., 2019; Pinto et al., 2020; Pomponio et al., 2020; Garcia-Dias et al., 2020) supports the move towards large-scale data sharing and merging (Poline et al., 2012; Eickhoff et al., 2016; Nichols et al., 2017).

The number of features used may be reduced to circumvent the risk of overfitting. In its most simple form, this would include being highly-selective in the features of interest, however, typical multivariate neuroimaging studies are seeking associations between a broad array of features. When considering the imaging features, dimensionality is typically very large: brain connectivity is often calculated at the global level. Even with parcellation, i.e. considering the connectivity between brain regions rather than voxels/vertices, brain connectivity is typically represented by hundreds to thousands of features. In some cases, brain connectivity features may be sub-selected, for example in cases where the connectivity between given brain regions is of interest, but this is only suitable in studies with such specific questions. The non-imaging features, e.g. behavioural traits, may also be sub-selected but is again limited to the question at hand.

An alternative to sub-selecting features is to remove redundancies in the data through data dimensionality reduction techniques, such as principle component analysis (PCA). This is a commonly used approach in CCA studies. The application of PCA should be carefully considered however, as the between-set correlations may be affected if the features driving the associations are not preserved. An approach which combines PCA and CCA has been developed (Song et al., 2016). Here, PCA is performed as a pre-processing step to CCA whilst considering the between-set correlation in order to optimise the number of components extracted from the PCA step.

#### Significance Testing in CCA

The statistical significance of CCA results is typically assessed by comparison to a null-distribution built through permutation testing. However, we have demonstrated here that, even in cases where statistical significance is achieved, the stability of CCA is limited. Therefore, we recommend that statistical significance should not be used to indicate that results are robust. It should be noted that in these works we use a simple permutation testing scheme typical of the literature, however, a recent study suggests that such testing is inadequate (Winkler et al., 2020).

### **Clarity in Reporting**

Section 6.3 reviews the literature using CCA to explore brain-behaviour associations. It is often difficult to understand exactly how analysis was performed and how many subjects and features were used. In some cases, the canonical correlations are not reported. We recommend to always explicitly state the used sample size, number of features in both datasets, and observed canonical correlations.

#### Data Encoding

An issue observed in the application of CCA to the datasets used in this thesis lies in the nature of the non-imaging data. These consist of behavioural traits from a broad array of behavioural and cognitive tests, lifestyle questionnaires, medical history and more. They are a mix of variable types, including categorical, continuous and discrete variables, and use complex coding schemes. For example, in the UK Biobank data, the coding used in the "Able to pay rent/mortgage as an adult" measure follows a 0-4 scale with 0 being "Never true" and 4 being "Very often true". Here, higher values would be interpreted more favourably. However, the same coding scale is used for questions such as "Physically abused by family as a child" where, of course, higher values may indicate serious childhood problems.

Further, again in the UK Biobank, negative values in categorical variables typically indicate a lack of data or a no response preference. In one study (Miller et al., 2016), and as used here, negative categorical variables are set to zero. However, in some scales, a negative value indicates the extreme of the scale. For example, for the question "At what age did you complete your continuous full time education?", a response of "Never went to school" is coded as -2 and no response/do not know also as negative values. The use of complex behavioural trait data without careful consideration of the coding of said data and their variable type may lead to difficulties in interpretation and conclusions prone to error. Beaton et al. (2019) introduces partial least squares-correspondence analysis-regression (PLS-CA-R) which seeks to overcome the issues associated with the mixing of variable types by re-coding and transforming data. Similar methods could be adapted to CCA.

#### Alternatives to CCA

Alternatives to CCA, both univariate and multivariate, should also be considered. These include the standard variants of CCA such as sparse, constrained and kernel CCA (Zhuang et al., 2020). In some cases, these variants may reduce the subjectfeature ratio requirement, however Helmer et al. (2020) suggests that the benefits of sparse CCA are limited. In a similar exploration to those presented in this thesis, Helmer et al. (2020) explores (albeit in a limited capacity) how sparsity may reduce overfitting. They used 5-fold cross-validation to find optimal sparsity parameters and determine required sample sizes to avoid overfitting. By comparison to classical CCA, it is possible to determine when sparse CCA may be beneficial (Figure 6.14). Helmer et al. (2020) report that sparse CCA is favoured, i.e. improves solution stability, in cases of many features and low true underlying between-set correlations.



**Figure 6.14:** The relative difference between the required sample size for classical and sparse CCA for a range of between-set correlation values in simulated data. The required sample size is defined such that a set of solution validity criterion must be met. Criterion include 90% statistical power (i.e. ability to identify an existing between-set association) and at most 10% error for the between-set association strength, weight, score and loading error. Blue colours indicate sparse CCA performs better compared to classical CCA. (reproduced from Helmer et al. (2020))

Further to these findings, and in the context of the overfitting reported in this thesis, it is often the case that sparse CCA, and similar variants, are used to overcome the issue of rank deficiency in the classical CCA solution where there must be more subjects than combined features. This allows for the application of multivariate approaches in cases where the number of subjects is much less than the number of features, i.e. the SFR is much less than one, as is indicated by the literature (Table 6.1). This is likely to lead to overfitting in the solution and opens the door to the widespread application of multivariate techniques in very small cohorts.

Alternatively, univariate approaches may be used, however, they are subject to similar issues with regard to the number of samples (Button et al., 2013; Schönbrodt and Perugini, 2013; Marek et al., 2020) and, as such, should be applied with a similarly cautious approach. The combination of carefully applied dimensionality reduction, univariate association techniques, cross-validation and external validation may be a particularly useful approach (Ji et al., 2020).

# 6.8 Concluding Remarks

Investigations into the stability of CCA results as applied to reveal brain-behaviour associations, have been presented. We have demonstrated that stability is strongly dependent on the number of subjects, the number of features and the underlying between-set correlation. Real-data examples from the HCP cohort demonstrate that even cohorts considered large ( $\approx 1,000$  subjects) are insufficient when considering the number of features typically used in the literature. The development and application of a simulation framework revealed, in a robust manner, the dependencies between the observed canonical correlation and the subject-feature ratio and underlying between-set correlation. Further investigations utilising the UK Biobank cohort have demonstrated the scale required to achieve stable results in the case of real-data. These analyses inform the recommendations on how CCA should be used in practice.

# Contributions

The works presented in this chapter were a collaborative effort that I co-led with the Murray and Anticevic Labs at Yale University. As such, I would like to extend my gratitude to both the Murray and Anticevic Labs. The sub-sampling analysis framework was conceptualised and designed by me, Markus Helmer, Stamatios Sotiropoulos, John Murray and Alan Anticevic. I developed, and performed analysis using, the simulation framework presented in this thesis, which is independent from the framework presented in the associated manuscript (Helmer et al., 2020) (although both follow similar ideas and draw similar conclusions). I performed all real-data analyses presented (using the gemmr toolbox developed by Markus Helmer). I developed scripts (i.e. external to the gemmr toolbox) and performed analysis for the results presented in Section 6.10.

# Software

Real-data analyses (Sections 6.4 and 6.6) were performed in Python (version 3.7.2) (Python Software Foundation, https://www.python.org) using the gemmr toolbox (https://github.com/murraylab/gemmr). The simulation framework presented in Section 6.6 was developed, and analysis was performed, using MATLAB 2018a.

# 6.9 Appendix

# 6.9.1 Double-dipping and CCA

As briefly discussed in Section 6.3, in some cases, a feature pre-selection step prior to CCA is performed, and this can be circular if not performed with caution. An example of this is in Drysdale et al. (2017), where mass univariate association (MUA) was used to sub-select imaging features prior to feeding them to CCA. MUA was performed between the 33,154 imaging features and 17 behaviour features. Imaging features which had a statistically significant correlation (p < 0.005) with any behaviour feature were then kept for CCA between this imaging features subset and the set of all 17 behavioural features.

To explore the consequence of this pre-selection approach, we applied MUA as a feature selection technique prior to CCA. Two analyses were performed: 1) CCA between functional connectivity features and behaviour data, and 2) CCA between data drawn randomly from a Normal distribution and behaviour data. The general processing principles remain as throughout the work presented in this thesis, with the extra MUA step added in both cases. Both analyses utilise the HCP cohort and use a subset of behavioural traits (q = 39) more closely matched to the features used in Drysdale et al. (2017), including summary alcohol and substance use (SSAGA), summary psychiatric health measures (DSM-V), psychological health measures (NIH negative affect survey) and family history.

The first analysis, with functional connectivity features, uses the 200-dimensionality connectomes as described in Section 6.4.1, giving 19,900 imaging features ( $\mathbf{X}_{func}$ : n = 772, 19,900 features). In the second analysis, the "imaging" dataset consists of an array of entries drawn from a Normal distribution with the same dimensions as in the first analysis ( $\mathbf{X}_{noise}$ : n = 772, 19,900 features). The datasets are demeaned and normalised as before. Confound regression is not applied to imaging data, as to avoid any potential introduction of between-set associations.

MUA between these imaging features and the behaviour data was used to sub-select imaging features. Drysdale et al. (2017) uses a p-value threshold of < 0.005. In the current case, this threshold results in so many functional connectivity features (3,579 from  $\mathbf{X}_{func}$ ) that the CCA solution would be rank deficient (i.e. more features than samples) - this may be due to the fact that we use more behavioural trait features here compared to Drysdale et al. (2017). Therefore, a lower p-value threshold is chosen (< 0.0001), giving p = 121 functional connectivity features. In the noise-data case, this same approach extracts p = 81 features.

The final input matrices are therefore  $\mathbf{X}_{func}$  with n = 772 and p = 121,  $\mathbf{X}_{noise}$  with n = 772 and p = 81, and  $\mathbf{Y}$  with n = 772 and q = 39. Permutation testing is used to assess statistical significance with 10,000 iterations and whilst considering the family structure inherent in the HCP cohort (Winkler et al., 2015). We report results for the top mode only.

CCA revealed that the top canonical mode had a correlation of  $\rho = 0.697$  (p-value = 0.0001) for the  $\mathbf{X}_{func}$  data and  $\rho = 0.605$  (p-value = 0.0001) for the  $\mathbf{X}_{noise}$  data. These are plotted with reference to the respective null-distributions (built through the permutation scheme) in Figure 6.15. This highlights the consequence of using MUA to reduce data dimensionality as a preprocessing step in CCA. The proposed permutation scheme in Dinga et al. (2019) helps alleviate these issues by incorporating the feature sub-selection into the permutation scheme.



**Figure 6.15:** The canonical correlations, and null-distributions, obtained through CCA with functional connectivity features and "noise" data reduced via mass univariate association with behavioural trait measures as a preprocessing in CCA. **a.)** The functional connectivity  $\mathbf{X}_{func}$  case, **b.)** the noise data  $\mathbf{X}_{noise}$  case. The histogram represents the null-distribution built through permutation testing. The green line shows the canonical correlation for the top mode with its value and significance annotated.

# 6.9.2 Behavioural Measures used in CCA - HCP Analyses

Note: In many cases, raw and age and sex-adjusted scores are provided. Both are used in the analyses presented, but, for brevity, they are not listed here.

### Physical

Dexterity, Regional Taste Intensity, Electronic Visual Acuity, Mars Final Contrast Sensitivity Score, Mars Contrast Sensitivity Score, Eyeglass correction, Odor Identification

### Family History

Family history of drugs and alcohol, No Father family history of drugs and alcohol, No Mother family history of drugs and alcohol, Mother history of depression, Father history of depression

#### Cognitive

Picture Vocabulary, Mini Mental State Examination (MMSE), Fluid intelligence correct responses (Penn progressive matrices), Delay Discounting \$200, Delay Discounting \$40,000, Flanker Inhibitory Control and Attention Test, Working memory (list sorting), Dimensional Change Card Sort Test, Oral Reading Recognition Test, Short Penn Continuous Performance Test: Specificity and Sensitivity, Episodic memory (picture sequence memory), Pattern Comparison Processing Speed Test, Penn Word Memory Test: Total Number of Correct Responses, Variable Short Penn Line Orientation: Total Number Correct, Total Positions Off for All Trials and Median Reaction Time Divided by Expected Number of Clicks for Correct

### Psychological Well-Being, Health and Scores

General life satisfaction, Pittsburgh Sleep Quality Index (overall score), Pain Interference Survey, Negative Affect: Anger Aggression, Anger Hostility, Fear and Anger, Friendship Survey, Sadness Survey, Instrumental Support Survey, Fear-Somatic Arousal Survey, Self-Efficacy Survey, Loneliness Survey, Meaning and Purpose Survey, Perceived Stress Survey, Emotional Support Survey, Perceived Rejection, Positive Affect Survey, Perceived Hostility, ER Test: Number of Correct Anger Identifications, Number of Correct Sad Identifications, Number of Correct Neutral Identifications, Number of Correct Fearful Identifications and Number of Correct Responses

#### Tobacco, Alcohol and Substance Use

Tested positive for Tetrahydrocannabinol (THC), Average weekday any tobacco over past 7 days, Number days used any tobacco over past 7 days, Total any tobacco over past 7 days, Average weekday cigarettes over past 7 days, Average weekend any tobacco over past 7 days, Total cigarettes over past 7 days, Average weekend cigarettes over past 7 days, Times used any tobacco today, Number days drank over past 7 days, Total beer wine/cooler over past 7 days, Average weekday wine over past 7 days, Average weekend beer/wine/cooler over past 7 days, Total drinks over past 7 days, Total wine over past 7 days, Average weekday drinks over past 7 days, Average weekday beer/wine/cooler over past 7 days, Average weekday drinks over past 7 days, Total hard liquor over past 7 days, Average weekend wine over past 7 days, Average weekend hard liquor over past 7 days

#### Personality Five Factor Model (NEO-FFI)

Neuroticism, Agreeableness, Openness to Experience, Extraversion, Conscientiousness

# Achenbach Adult Self-Report (ASR) and Diagnostic and Statistical Manual of Mental Disorders (DSM)

Rule Breaking Behaviour, Thought Problems, Critical Items, Externalising Problems, Other Problems, Total Problems, Sum of Thought, Attention, and Other Problems, Somatic Complaints, Withdrawn, Aggressive Behaviour, Internalising Problems, Anxious/Depressed, Attention Problems, Intrusive Thoughts, Antisocial Personality Problems, Anxiety Problems, Depressive Problems, Somatic Problems, Avoidant Personality Problems, Hyperactivity Problems, AD/H Problems, Inattention Problems

#### Semi-Structured Assessment for the Genetics of Alcoholism (SSAGA)

Times used stimulants, Times used illicit substances, Times used illicit opiates, Times used illicit sedatives, Times used illicit hallucinogens, Times used marijuana, Ever used marijuana?, Marijuana Dependence, Times used cocaine, Whether age last smoked is current age, Whether age last smoked is current age, Smoking history, Childhood conduct problems, Suffered from possible agoraphobia, Panic Disorder, Major Depressive Episode, Number Depressive Symptoms, Max drinks in a single day in past 12 months, Age at first alcohol use, Frequency drunk in past 12 months, Drinks per drinking day in past 12 months, Frequency drunk in heaviest 12-month period, Drinks per day in heaviest 12-month period, Lifetime max drinks in single day, Alcohol Abuse number of symptoms, Alcohol Abuse Criteria Met, Number of DSM4 Alcohol Dependence Criteria Endorsed, Frequency of drinking 5+ drinks, heaviest 12-month period, Frequency of any alcohol use, Frequency of drinking 5+ drinks in past 12 months, Frequency of any alcohol use in past 12 months

# 6.9.3 Behavioural Measures used in CCA - UK Biobank Analyses

### Physical

Frequency of walking for pleasure in last 4 weeks, Frequency of other exercises in last 4 weeks, Duration walking for pleasure, Forced expiratory volume in 1-second: Best measure, Body mass index (BMI), Comparative height size at age 10, Comparative body size at age 10, Impedance of whole body, Handedness (chirality/laterality), Duration of walks, Number of days/week of vigorous physical activity 10+ minutes, Pulse rate, Usual walking pace, Hand grip strength (right), Hand grip strength (left), Ankle spacing width (left), Snoring, Number of days/week of moderate physical activity 10+ minutes, Duration of moderate activity, Number of days/week walked 10+ minutes, Duration of other exercises, Duration of vigorous activity, Frequency of stair climbing in last 4 weeks, Diastolic blood pressure, Systolic blood pressure, Heel bone mineral density, Forced vital capacity

#### Family History

Father's age at death, Maternal smoking around birth, Illnesses of mother, Father still alive, Mother still alive, Mother's age at death

### Cognitive

Mean time to correctly identify matches, Duration screen displayed, Time first key touched, Time last key touched, Time elapsed, Maximum digits remembered correctly, Digits entered correctly, Total errors traversing numeric path (trail #1), Number of incorrect matches in round, Number of rounds of numeric memory test performed, Duration to entering value, Total errors traversing alphanumeric path (trail #2), Duration to complete alphanumeric path (trail #2), FI8: chained arithmetic, FI7: synonym, FI6: conditional arithmetic, FI5: family relationship calculation, FI4: positional arithmetic, FI3: word interpolation, Prospective memory result, Number of puzzles correctly solved, Number of puzzles correct, PM: initial answer, Number of fluid intelligence questions attempted within time limit Duration to complete numeric path (trail #1), Interval between previous point and current one in numeric path (trail #1), Fluid intelligence score, Duration spent answering each puzzle, Time to complete test (for each cognitive test included), Number of symbol digit matches attempted, Number of correct matches in round, Number of symbol digit matches made correctly

### Psychological Well-Being, Health and Scores

Repeated disturbing thoughts of stressful experience in past month, General happiness with own health, Recent restlessness, Sleeplessness/insomnia, Fed-up feelings, Ever seen an un-real vision, Ever thought that life not worth living, Been in a confiding relationship as an adult, Frequency of tiredness / lethargy in last 2 weeks, Been involved in combat or exposed to war-zone, Miserableness, Ever sought or received professional help for mental distress, Sexual interference by partner or ex-partner without consent as an adult, Tense/'highly strung', Ever had prolonged loss of interest in normal activities, Ever had period of mania/excitability, Work/job satisfaction, Neuroticism score, Recent feelings of inadequacy, Recent changes in speed/amount of moving or speaking, Ever believed in un-real communications or signs, Frequency of depressed mood in last 2 weeks, Felt loved as a child, Ever unenthusiastic/disinterested for a whole week, Suffer from 'nerves', Ever had prolonged feelings of sadness or depression, Recent worrying too much about different things, Frequency of tenseness / restlessness in last 2 weeks, Avoided activities or situations because of previous stressful experience in past month, Ever highly irritable/argumentative for 2 days, Sensitivity / hurt feelings, Health satisfaction, Recent trouble concentrating on things, Physical violence by partner or ex-partner as an adult, Mood swings, Family relationship satisfaction, Ever addicted to any substance or behaviour, Trouble falling or staying asleep, or sleeping too much, Ever heard an un-real voice, Recent feelings of depression, Felt hated by family member as a child, Ever contemplated self-harm, Recent easy annoyance or irritability, Recent thoughts of suicide or self-harm, Risk taking, Ever believed in an un-real conspiracy against self, Irritability, Frequency of friend/family visits, Worry too long after embarrassment, Physically abused by family as a child, Ever suffered mental distress preventing usual activities, Ever worried more than most people would in similar situation, Recent feelings or nervousness or anxiety, Felt very upset when reminded of stressful experience in past month, Guilty feelings, Friendships satisfaction, Recent trouble relaxing, Ever self-harmed, Loneliness, isolation, Belittlement by partner or ex-partner as an adult, Ever felt worried, tense, or anxious for most of a month or longer, Victim of sexual assault, Recent poor appetite or overeating, General happiness, Ever had period extreme irritability, Recent feelings of foreboding, Recent feelings of tiredness or low energy, Worrier / anxious feelings, Happiness, Ever depressed for a whole week, Frequency of unenthusiasm / disinterest in last 2 weeks, Sexually molested as a child, Recent inability to stop or control worrying, Victim of physically violent crime, Someone to take to doctor when needed as a child, Financial situation satisfaction, Recent lack of interest or pleasure in doing things, Able to confide, Diagnosed with life-threatening illness, Been in serious accident believed to be life-threatening, Witnessed sudden violent death, Seen doctor (GP) for nerves, anxiety, tension or depression

#### Tobacco, Alcohol and Substance Use

Smoked cigarette or pipe within last hour, Alcohol usually taken with meals, Average weekly beer plus cider intake, Average weekly champagne plus white wine intake, Tobacco smoking, Smoking status, Alcohol intake versus 10 years previously, Average weekly intake of other alcoholic drinks, Past tobacco smoking, Ever smoked, Exposure to tobacco smoke outside home, Amount of alcohol drunk on a typical drinking day, Average weekly fortified wine intake, Average weekly red wine intake, Frequency of consuming six or more units of alcohol, Smoking/smokers in household, Frequency of drinking alcohol, Ever had known person concerned about, or recommend reduction of, alcohol consumption, Ever taken cannabis, Ever been injured or injured someone else through drinking alcohol, Average weekly spirits intake, Alcohol intake frequency

### Qualifications, Income and Employment

Qualifications, Age completed full time education, Year ended full time education, Own or rent accommodation lived in, Number of vehicles in household, Able to pay rent/mortgage as an adult, Average total household income before tax, Length of working week for main job, Current employment status, Distance between home and job workplace, Job involves heavy manual or physical work, Noisy workplace, Job involves shift work, Frequency of travelling from home to job workplace, Job involves mainly walking or standing

### Other Lifestyle

Difference in mobile phone use compared to two years previously, Ever had same-sex intercourse, Weekly usage of mobile phone in last 3 months, Lifetime number of sexual partners, Time spent outdoors in winter, Frequency of light DIY in last 4 weeks, Duration of light DIY, Length of time at current address, Childhood sunburn occasions, Time spent using computer, Time spent watching television (TV), Nap during day, Hands-free device/speakerphone use with mobile phone in last 3 month, Drive faster than motorway speed limit, Usual side of head for mobile phone use, Age first had sexual intercourse, Morning/evening person (chronotype), Number of older siblings, Getting up in morning, Private healthcare, Place of birth in UK - east co-ordinate, Place of birth in UK - north coordinate, Number in household, Number of full brothers, Number of full sisters, Breastfed as a baby, Sleep duration, Length of mobile phone use, Time spend outdoors in summer, Daytime dozing / sleeping (narcolepsy), Loud music exposure frequency

# 6.9.4 UK Biobank CCA top Non-imaging Loadings

The top 15 positive and negative non-imaging loading variables for the ten canonical modes extracted from the UK Biobank CCA analysis with 20,000 subjects. Each mode is statistically significant (at p = 0.0001), as assessed through permutation testing.

	Mode 1 (CR = $0.4$ , p = $0.0001$ )			Mode 2 (CR = 0.39, p = 0.0001)
0.47	- Private healthcare		0.72	- Ever believed in un-real communications or signs
	Belittlement by partner or ex-partner as an adult			Frequency of walking for pleasure in last 4 weeks
	Length of mobile phone use			Handedness (chirality/laterality)
	Number of days/week of moderat physical activity 10+ minutes			Length of time at current address
	Frequency of walking for pleasure in last 4 weeks			Number of correct matches in round
	Fed-up feelings			Recent worrying too much about different things
	Recent thoughts of suicide or self-harm			Seen doctor (GP) for nerves, anxiety, tension or depression
	Ever believed in un-real communications or signs			Number of days/week of vigorou physical activity 10+ minutes
0.16	- Avoided activities or situatiossful experience in past month	0.12 ਵ	- FI5 : family relationship calculation*	
	Handedness (chirality/laterality)		Felt very upset when remindedssful experience in past month	
	Recent worrying too much about different things			Irritability
	Ever depressed for a whole week			Job involves shift work
	Recent trouble concentrating on things			Time spent watching television (TV)
	Own or rent accommodation lived in	elati		Frequency of drinking alcohol
0.11	- Ever believed in an un-real conspiracy against self	orre	a.09	- FI3 : word interpolation*
-0.16	- FI5 : family relationship calculation*	<u></u> -0.0	-0.06	- Loud music exposure frequency
	Frequency of unenthusiasm / disinterest in last 2 weeks	ding		Recent trouble concentrating on things
	Ever smoked	Loa		Duration walking for pleasure
	Time spent watching television (TV)			Recent thoughts of suicide or self-harm
	Fluid intelligence score			Tobacco smoking
	Ever thought that life not worth living			Ever believed in an un-real conspiracy against self
	Someone to take to doctor when needed as a child			Frequency of stair climbing in last 4 weeks
-0.18	- Current employment status		-0.08	- Sleeplessness / insomnia
	Irritability			Ever highly irritable/argumentative for 2 days
	Felt hated by family member as a child			Alcohol usually taken with meals
	Recent feelings of foreboding			Length of mobile phone use
	Seen doctor (GP) for nerves, anxiety, tension or depression			Number of days/week of moderat physical activity 10+ minutes
	Duration spent answering each puzzle			Someone to take to doctor when needed as a child
	Number of days/week of vigorou physical activity 10+ minutes			Private healthcare
-0.35	- Number of correct matches in round		-0.42	- lob involves heavy manual or physical work

		Mode 3 (CR = 0.33, p = 0.0001)		
elation)	0.37 -	Felt hated by family member as a child	0.36	
		Belittlement by partner or ex-partner as an adult		
		Length of mobile phone use		
	0.2 -	FI5 : family relationship calculation*		
		Hand grip strength (right)		
		Difference in mobile phone usempared to two years previously		
		Total errors traversing alphanumeric path (trail #2)		
		Sleeplessness / insomnia		
		Handedness (chirality/laterality)	0.2	
		Frequency of other exercises in last 4 weeks		
		Time spent outdoors in winter		
		Place of birth in UK - north co-ordinate		
		Nervous feelings	(uoi	
		Recent changes in speed/amount of moving or speaking	elat	
orre	0.16 - Job involves heavy manual or physical work		Ĕ 0.1€	
<u>о</u>	-0.16 - Duration screen displayed			
din		Seen doctor (GP) for nerves, anxiety, tension or depression	idin	
Loa		FI3 : word interpolation	Loa	
		Maximum digits remembered correctly		
		Year ended full time education		
		Frequency of stair climbing in last 4 weeks		
		Happiness		
	-0.21 ·	Usual side of head for mobile phone use	-0.25	
		Distance between home and job workplace		
		Mother's age at death		
		Job involves shift work		
		Ever sought or received profesional help for mental distress		
		Hand grip strength (left)		
		Recent poor appetite or overeating		
	-0.35 -	Number of days/week of moderat physical activity 10+ minutes	-0.4	

#### Mode 4 (CR = 0.33, p = 0.0001) 6 - Felt hated by family member as a child Avoided activities or situatio...ssful experience in past month Recent thoughts of suicide or self-harm Handedness (chirality/laterality) Ever sought or received profes...ional help for mental distress Time to answer Mood swings Mother still alive 2 - Body mass index (BMI) Duration of moderate activity Frequency of walking for pleasure in last 4 weeks FI5 : family relationship calculation Frequency of other exercises in last 4 weeks Number of days/week walked 10+ minutes 6 - Time spent driving 8 - Sexually molested as a child Comparative height size at age 10 Maximum digits remembered correctly Number of days/week of vigorou... physical activity 10+ minutes Forced expiratory volume in 1-second (FEV1), Best measure Smoked cigarette or pipe within last hour Job involves heavy manual or physical work 5 - Recent poor appetite or overeating Duration spent answering each puzzle Number of correct matches in round Job involves shift work Belittlement by partner or ex-partner as an adult Fed-up feelings Length of mobile phone use -0.4 - Number of days/week of moderat... physical activity 10+ minutes

### Mode 5 (CR = 0.29, p = 0.0001)

	0.35 - Felt very upset when remindedssful experience in past month	
	Smoked cigarette or pipe within last hour	
	Forced expiratory volume in 1-second (FEV1), Best measure	
	Frequency of drinking alcohol	
	Duration spent answering each puzzle	
	Year ended full time education	
	Someone to take to doctor when needed as a child	
	FI3 : word interpolation	
	0.27 - Recent poor appetite or overeating	
	Time spent watching television (TV)	
	Job involves shift work	
	Felt hated by family member as a child	
ion)	Current employment status	ion)
elat	Ever thought that life not worth living	elat
orre	0.15 - Ever smoked	orre
<u>о</u> б	-0.16 - Fluid intelligence score	<u>о</u>
din	Happiness	din
Loa	Father still alive	Loa
	Guilty feelings	
	Been in a confiding relationship as an adult	
	Sexually molested as a child	
	Ever sought or received profesional help for mental distress	
	-0.22 - Number in household	
	Time spent outdoors in winter	
	Difference in mobile phone usempared to two years previously	
	Frequency of other exercises in last 4 weeks	
	Recent easy annoyance or irritability	
	Place of birth in UK - north co-ordinate	
	FI5 : family relationship calculation*	
	-0.35 - Number of days/week of moderat physical activity 10+ minutes	

# Mode 6 (CP = 0.24 p = 0.0001)

	Mode 6 (CR = $0.24$ , p = $0.0001$ )
0.23 -	Sleeplessness / insomnia
	Private healthcare
	Job involves heavy manual or physical work
	Frequency of walking for pleasure in last 4 weeks
	Mood swings
	Felt hated by family member as a child
	Age first had sexual intercourse
	Someone to take to doctor when needed as a child
0.15 -	Duration spent answering each puzzle
	Daytime dozing / sleeping (narcolepsy)
	Prospective memory result
	Impedance of whole body
	Work/job satisfaction
	Recent poor appetite or overeating
0.12 -	FI4 : positional arithmetic*
-0.12 -	Ever believed in an un-real conspiracy against self
	Felt very upset when remindedssful experience in past month
	Ever depressed for a whole week
	Frequency of unenthusiasm / disinterest in last 2 weeks
	Smoking status
	Trouble falling or staying asleep, or sleeping too much
	Duration of other exercises
-0.13 -	Frequency of drinking alcohol
	Ankle spacing width (left)
	Miserableness
	Duration of moderate activity
	Mother still alive
	Happiness
	Length of mobile phone use
-0.45 -	Belittlement by partner or ex-partner as an adult
	Mode 7 (CR = $0.22$ , p = $0.0001$ )
----------------------------	---
0.25 - Av	erage weekly red wine intake
Pla	ce of birth in UK - north co-ordinate
Eve	er highly irritable/argumentative for 2 days
FI5	: family relationship calculation*
Se	xually molested as a child
Du	ration of other exercises
Nu	mber of attempts
Dif	ference in mobile phone usempared to two years previously
0.17 - Eve	er heard an un-real voice
Eve	er believed in un-real communications or signs
Ha	ndedness (chirality/laterality)
Foi	rced vital capacity (FVC), Best measure
Q Pri	vate healthcare
elat Evi	er believed in an un-real conspiracy against self
0.15 - Irri	tability
ອີ -0.18 - <mark>So</mark>	meone to take to doctor when needed as a child
up Nu	mber of correct matches in round
FI3	: word interpolation*
Du	ration spent answering each puzzle
Ha	nd grip strength (left)
Us	ual side of head for mobile phone use
FI3	: word interpolation
-0.24 - Pro	ospective memory result
Lei	ngth of mobile phone use
Ye	ar ended full time education
Be	littlement by partner or ex-partner as an adult
Im	pedance of whole body
Fre	equency of depressed mood in last 2 weeks
Se	en doctor (GP) for nerves, anxiety, tension or depression
-0.42 - Wo	ork/job satisfaction

#### Mode 9 (CR = 0.19, p = 0.0001)

	0.26 -	Sleeplessness / insomnia
		Number of days/week of moderat physical activity 10+ minutes
		Fed-up feelings
		Trouble falling or staying asleep, or sleeping too much
		Frequency of drinking alcohol
		Usual side of head for mobile phone use
		Felt very upset when remindedssful experience in past month
		Seen doctor (GP) for nerves, anxiety, tension or depression
g (correlation)	0.14 -	Frequency of stair climbing in last 4 weeks
		Number of days/week walked 10+ minutes
		Handedness (chirality/laterality)
		Suffer from 'nerves'
		FI8 : chained arithmetic
		Daytime dozing / sleeping (narcolepsy)
	0.12 -	Time to complete test
	-0.14 -	Alcohol intake versus 10 years previously
ulpe		Private healthcare
õ		Sexually molested as a child
		Age completed full time education
		Maximum digits remembered correctly
		Sensitivity / hurt feelings
		Frequency of consuming six or more units of alcohol
	-0.17 -	Snoring
		Ever highly irritable/argumentative for 2 days
		Recent poor appetite or overeating
		Job involves shift work
		Happiness
		Felt hated by family member as a child
		Frequency of travelling from home to job workplace
	-0.33 -	Mood swings

		Mode 8 (CR = $0.21$ , p = $0.0001$ )
	0.44	Forced expiratory volume in 1-second (FEV1), Best measure
		Smoked cigarette or pipe within last hour
		Work/job satisfaction
		Frequency of depressed mood in last 2 weeks
		Impedance of whole body
		Snoring
		Daytime dozing / sleeping (narcolepsy)
		Prospective memory result
	0.21	Place of birth in UK - north co-ordinate
		Frequency of consuming six or more units of alcohol
on)		Usual side of head for mobile phone use
		Irritability
		Ever smoked
elati		Father's age at death
orre	0.19	Recent thoughts of suicide or self-harm
00	-0.12	Number of attempts
din		Recent changes in speed/amount of moving or speaking
Loa		Average weekly red wine intake
		Victim of physically violent crime
		Number of days/week of vigorou physical activity 10+ minutes
		Duration of other exercises
		Felt hated by family member as a child
	-0.2	Number of correct matches in round
		Year ended full time education
		FI3 : word interpolation
		Duration spent answering each puzzle
		Someone to take to doctor when needed as a child
		Trouble falling or staying asleep, or sleeping too much
		Frequency of other exercises in last 4 weeks
	-0.33	Ankle spacing width (left)

#### Mode 10 (CR = 0.19, p = 0.0001)

	0.27	- Felt very upset when remindedssful experience in past month
		Frequency of drinking alcohol
		Childhood sunburn occasions
		Felt hated by family member as a child
		Able to confide
		Number of attempts
		Diagnosed with life-threatening illness
		Recent changes in speed/amount of moving or speaking
	0.18	- Time to complete round*
		Past tobacco smoking
		Duration of light DIY
(uo		Loud music exposure frequency
		Duration of walks
elati		Recent feelings of tiredness or low energy
orre	0.16	- Smoked cigarette or pipe within last hour
0 0	-0.12	- Belittlement by partner or ex-partner as an adult
ģ		Ankle spacing width (left)
Loa		Number of symbol digit matches attempted
		Victim of sexual assault
		Duration to first press of snap-button in each round
		Time to complete round
		FI5 : family relationship calculation
	-0.16	Duration of other exercises
		Number of days/week of moderat physical activity 10+ minutes
		Irritability
		Ever thought that life not worth living
		Ever smoked
		Current employment status
		Duration spent answering each puzzle
	-0.37	Someone to take to doctor when needed as a child

**Figure 6.16:** The top non-imaging contributors for each subsampled analysis for the top ten canonical modes.

# Chapter 7

# **Conclusions and Future Directions**

This thesis has been an exciting journey and offers new methodology for mapping structural connectivity of the brain, across individuals and primate species, and exploring its associations with functional features, and ultimately human behaviour. Diffusion tractography is a unique tool used to estimate the structural architecture of the brain by probing the diffusion of water molecules. One of the primary goals of diffusion tractography is to extract the major white matter (WM) fibre bundles, which connect functionally specialised brain regions. To achieve this in a robust and generalisable manner, it is vital that tractography be standardised. Further, to understand the associations between brain connectivity and behaviour, the statistical techniques we apply to probe said associations must be robust. To this end, this thesis presents frameworks for the extraction of homologous WM fibre bundles across species, explores ways for using structural connectivity patterns to probe functional similarities of diverse brains and performs investigations into the stability of a commonly used brain-behaviour association technique.

Chapter 4 presents a new tool for the automated and standardised estimation of a set of 42 major WM fibre bundles, equivalent in the human and non-human primate (macaque) brain. The tool (XTRACT) is shown to be robust and generalisable across species and data-quality. We demonstrated this by applying XTRACT to very-high-quality data from the HCP and closer to typical quality data, using the UK Biobank datasets. We built high-quality atlases of the WM fibre bundles for the human brain, using an order of magnitude more data than previous efforts, and a corresponding atlas for the macaque brain, using a small although high-quality ex vivo dataset. Tractography protocols were indirectly validated in a number of ways: a) by comparing measures of tract lateralisation to the literature, b) by demonstrating that, although the tractography is generalisable and automated, the resultant tracts reflect the underlying individual variability, as assessed through twinship and in respecting anatomical abnormalities, c) by using the tracts to devise common connectivity spaces for mapping and linking diverse brains (Chapter 5).

XTRACT was released publicly as a part of FMRIB's software library (FSL) (version 6.0.2 and later) and has since been widely applied. For example, to explore associations between tract features and epilepsy symptoms (Jungilligens et al., 2020); to measure the long-term neurological impact of COVID-19 (Qin et al., 2021); to uncover links between structure and function specific to the temporal association cortex (Blazquez Freches et al., 2020); to explore expansion and lateralisation in the connectivity of the inferior parietal lobule in the macaque, chimpanzee and human brain (Cheng et al., 2021); and to generate atlases of WM fibre bundles (Hansen et al., 2020). In addition to these applications, colleagues have also developed XTRACT compatible tractography protocols for the chimpanzee (Bryant et al., 2020), pig (Benn et al., 2020) and human neonate (Thompson, 2020) brains.

Chapter 5 furthers the work presented in Chapter 4 to show that we can use measures of structural connectivity to build common connectivity spaces. In these spaces, connectivity patterns may be compared in the absence of geometrical alignment, even across species, and used to project brains and map features from one brain to the other, relying on the assumption that regions with similar connection patterns define similar functional regions across brains. We example this by projecting scalar maps of cortical myelination from the human brain to the macaque brain based on comparisons/similarities of connectivity patterns to homologous landmarks (either in WM or in GM). Such comparisons may be performed in a number of ways, exampled using two previously introduced approaches, yet the basic principles hold, showing robustness of the above ideas. Further, these ideas may be extended to incorporate connectivity patterns derived through fully data-driven approaches (such as non-negative matrix factorisation of tractography data), opening new avenues for exploring them in ways that are less reliant on prior assumptions and manual definitions.

Finally, Chapter 6 focuses on linking brain connectivity features and their variability across individuals with variability in behavioural traits. We explore the application and stability of multivariate statistical techniques, specifically canonical correlation analysis (CCA), and demonstrate significant challenges and limitations in its application. We identify that the ratio between the number of subjects available compared to the number of features used can have a significant effect on the robustness of the identified latent associations and their out-of-sample generalisability. Through forward predictions by a generative model, we show that, in order to achieve robust solutions, much larger (an order of magnitude higher) subject-tofeature ratios are required than previously used in the literature. We confirm these predictions by using 20,000 subjects from the UK Biobank and demonstrate stable associations between whole-brain connectivity features and behavioural/lifestyle traits that include measures of socioeconomic status, cognition and mental, psychiatric and physical health. Chapter 6 ends with a set of recommendations on the application of such techniques, arguing for greater transparency in the reporting of results and the careful consideration of alternative techniques.

## 7.1 Future Directions

The methods and explorations presented in this thesis enable exciting new opportunities and research questions, as briefly discussed in the following sections.

### 7.1.1 The Standardisation of Tractography

Even though tractography approaches have been around for the last 20 years, there are still multiple sources of variability, which deem standardisation of tractography protocols a vital requirement for future advances. Key sources of variability relate to the definition and nomenclature of tractography protocols and WM anatomy, differences in data acquisition and processing, and a lack of detail in and sharing of tractography protocols. These result in significant variability in the definitions and subsequent estimations of fibre bundles (Schilling et al., 2020b). These issues are augmented when the ontogenic (across ages) and phylogenic (across species) dimensions of brain connectivity are also considered. This thesis contributes towards this by providing a flexible framework through which tractography protocols may be developed and shared in a standardised manner. Protocols may be developed in any space and may contain combinations of hand-drawn, atlas-based and segmentation-based regions of interest. An initial set of protocols have been devised here, which are already unique in their cross-species equivalence.

Efforts to fully standardise and reach consensus in tractography are ongoing. For example, Schilling et al. (2020b) (to which we are contributing) seeks to first quantify, and then suggest remediations for, variability in tractography. These efforts are likely to be required for sometime.

### 7.1.2 Using Connectivity to Map Brain Diversity

One of the most under-appreciated features of the brain is its diversity, in its size, complexity and organisation. To reveal and understand these relationships we must be able to map features which are diverse across brains in a robust and consistent manner. As shown in this thesis, we can map structural connectivity and use this to study and link diverse brains, allowing us to identify similarities and divergence in their functional organisation.

This can unlock new ways in studying the brain and answering fundamental questions. For instance, exploring the brain throughout the lifespan - from early development to ageing (Betzel et al., 2014; Zhao et al., 2015), across evolutionary pathways (Assaf et al., 2020), or between health and disease (Thiebaut de Schotten et al., 2020). Each of these research areas will contribute to the understanding of the brain, be it in terms of human brain development and how variation across individuals leads to differences in behaviour; the re-organisation of the brain across species and how this leads to language and social behaviours for example; or in disease mechanisms - how a given disease leads to behavioural/cognitive impairment.

Connections provide a different way of "aligning" the cortex across brains, which is not driven by geometry and morphology, as these alone are not necessarily predictive of differences in functional organisation (Glasser et al., 2016). They provide a direct link to the functional specialisation of the brain (Passingham et al., 2002; Jbabdi and Behrens, 2012), which can be probed using MRI, and compared across cortical regions, across individuals, across species and between health and disease.

#### 7.1.3 Brain-Behaviour Associations

With the advent of big data in neuroscience, both in terms of the number of samples and the number of features, the opportunities to explore the brain are endless. Projects such as the ENIGMA consortium (Thompson et al., 2014) and the UK Biobank capture measures from a broad array of modalities; from behavioural traits, to physical measures, to brain MRI and genetic data, at the population-level. This allows for explorations on the associations between the brain and behaviour, body, genetics and more. However, as in many of the big data sciences, methods to utilise such data are under-developed, leading to limitations in current applications and open challenges (Sejnowski et al., 2014; Fan et al., 2014; Gandomi and Haider, 2015; Smith and Nichols, 2018; Bzdok et al., 2019). This thesis highlights some of those challenges specific to brain-behaviour associations, identifies solutions for best use of such approaches and seeks to reinforce the Peter Parker principle<sup>1</sup>.

The opportunities to explore the relationships between measures of structural connectivity and behaviour, and relationships between structural and functional connectivity, on such a scale is an exciting prospect. Further, identifying links between brain organisation and behavioural/cognitive mechanisms in health allows for the development of robust image-derived phenotypes, which may generalise to other cohorts, for example disease-specific cohorts in mental health. Such population-level explorations and discoveries can also lead to novel normative modeling approaches for brain connectivity, characterising typical variability in the population and paving the way for the quantitative mapping of individuals in such normative distributions.

<sup>&</sup>lt;sup>1</sup>S. Lee, J. Kirby, S. Ditko. Amazing Fantasy Vol 1 #15 - "With great power there must also come - - Great responsibility!", 1962, Marvel Comics  $\circledast$ 

# Bibliography

- S. Agarwal, H. I. Sair, S. Gujar, and J. J. Pillai. "Language Mapping with fMRI: Current Standards and Reproducibility". In: *Topics in Magnetic Resonance Imaging* (2019).
- [2] H. Akram, V. Dayal, P. Mahlknecht, D. Georgiev, J. Hyam, T. Foltynie, P. Limousin, E. De Vita, M. Jahanshahi, J. Ashburner, T. Behrens, M. Hariz, and L. Zrinzo. "Connectivity derived thalamic segmentation in deep brain stimulation for tremor". In: *NeuroImage: Clinical* (2018).
- [3] D. C. Alexander, C. Pierpaoli, P. J. Basser, and J. C. Gee. "Spatial transformations of diffusion tensor magnetic resonance images." In: *IEEE transactions on medical imaging* 20.11 (2001), pp. 1131–1139.
- [4] D. C. Alexander, T. B. Dyrby, M. Nilsson, and H. Zhang. "Imaging brain microstructure with diffusion MRI: Practicality and applications". In: NMR in Biomedicine July (2017), pp. 1–26.
- [5] F. Alfaro-Almagro, M. Jenkinson, N. K. Bangerter, J. L. R. Andersson, L. Griffanti, G. Douaud, S. N. Sotiropoulos, S. Jbabdi, M. Hernandez-Fernandez, E. Vallee, D. Vidaurre, M. Webster, P. McCarthy, C. Rorden, A. Daducci, D. C. Alexander, H. Zhang, I. Dragonu, P. M. Matthews, K. L. Miller, and S. M. Smith. "Image processing and Quality Control for the first 10,000 brain imaging datasets from UK Biobank." In: *NeuroImage* 166 (2018), pp. 400–424.
- [6] C. Amiez, C. R. Wilson, and E. Procyk. "Variations of cingulate sulcal organization and link with cognitive performance". In: *Scientific Reports* 32.4 (2019).
- [7] J. L. R. Andersson, M. S. Graham, E. Zsoldos, and S. N. Sotiropoulos. "Incorporating outlier detection and replacement into a non-parametric framework for movement and distortion correction of diffusion MR images." In: *NeuroImage* 141 (2016), pp. 556–572.
- [8] J. L. R. Andersson, M. Jenkinson, and S. M. Smith. "Non-linear optimisation. FMRIB technical report TR07JA1". In: *In prac.* (2007).
- [9] J. L. R. Andersson, S. Skare, and J. Ashburner. "How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging". In: *NeuroImage* 20 (2003), pp. 870–888.
- J. L. R. 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.
- [11] J. L. Andersson, M. Jenkinson, and S. Smith. *High resolution nonlinear registration with simultaneous modelling of intensities.* 2019.
- [12] D. B. Archer, D. E. Vaillancourt, and S. A. Coombes. "A Template and Probabilistic Atlas of the Human Sensorimotor Tracts using Diffusion MRI". In: Cerebral cortex (New York, N.Y. : 1991) (2018).

- [13] J. Ashmore, H. G. Pemberton, W. D. Crum, J. Jarosz, and G. J. Barker. "Implementation of clinical tractography for pre-surgical planning of space occupying lesions: An investigation of common acquisition and post-processing methods compared to dissection studies". In: *PLoS ONE* (2020).
- [14] Y. Assaf, A. Bouznach, O. Zomet, A. Marom, and Y. Yovel. "Conservation of brain connectivity and wiring across the mammalian class". In: *Nature Neuroscience* (2020).
- [15] Y. Assaf, H. Johansen-Berg, and M. Thiebaut de Schotten. The role of diffusion MRI in neuroscience. 2017.
- [16] J. A. Autio, M. F. Glasser, T. Ose, C. J. Donahue, M. Bastiani, M. Ohno, Y. Kawabata, Y. Urushibata, K. Murata, K. Nishigori, M. Yamaguchi, Y. Hori, A. Yoshida, Y. Go, T. S. Coalson, S. Jbabdi, S. N. Sotiropoulos, S. Smith, D. C. Van Essen, and T. Hayashi. "Towards HCP-Style Macaque Connectomes: 24-Channel 3T Multi-Array Coil, MRI Sequences and Preprocessing". In: *bioRxiv* (2019).
- [17] D. M. Barch, G. C. Burgess, M. P. Harms, S. E. Petersen, B. L. Schlaggar, M. Corbetta, M. F. Glasser, S. Curtiss, S. Dixit, C. Feldt, D. Nolan, E. Bryant, T. Hartley, O. Footer, J. M. Bjork, R. Poldrack, S. Smith, H. Johansen-Berg, A. Z. Snyder, D. C. Van Essen, and W.-M. H. C. P. Consortium. "Function in the human connectome: task-fMRI and individual differences in behavior." In: *NeuroImage* 80 (2013), pp. 169–189.
- [18] R. S. Barcikowski and J. P. Stevens. "A monte carlo study of the stability of canonical correlations, canonical weights and canonical variate-variable correlations". In: *Multivariate Behavioral Research* (1975).
- [19] R. S. Barcikowski and J. P. Stevens. "Studying canonical analysis: A reply to thorndike's comments". In: *Multivariate Behavioral Research* 11.2 (1976), pp. 255–258.
- [20] P. J. Basser. "Fiber-Tractography via Diffusion Tensor MRI (DT-MRI)". In: Proceedings of the 6th Annual Meeting ISMRM. Vol. 820. 1997, pp. 123–138.
- [21] P. J. Basser. "New histological and physiological stains derived from diffusiontensor MR images." In: Annals of the New York Academy of Sciences 820 (1997), pp. 123–138.
- [22] P. J. Basser, J. Mattiello, and D. LeBihan. "MR diffusion tensor spectroscopy and imaging." In: *Biophysical journal* 66.1 (1994), pp. 259–267.
- [23] 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.
- [24] D. S. Bassett and E. T. Bullmore. "Human brain networks in health and disease." In: *Current opinion in neurology* 22.4 (2009), pp. 340–347.
- [25] T. M. Beasley, S. Erickson, and D. B. Allison. "Rank-based inverse normal transformations are increasingly used, but are they merited?" In: *Behavior Genetics* (2009).

- [26] D. Beaton, G. Saporta, and H. Abdi. "A generalization of partial least squares regression and correspondence analysis for categorical and mixed data: An application with the ADNI data". In: *bioRxiv* (2019).
- [27] R. Becker and A. Hervais-Adelman. "Resolving the connectome Spectrallyspecific functional connectivity networks and their distinct contributions to behaviour". In: *bioRxiv* (2019).
- [28] 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.
- [29] M. Beckmann, H. Johansen-Berg, and M. F. Rushworth. "Connectivity-based parcellation of human cingulate cortex and its relation to functional specialization". In: *Journal of Neuroscience* (2009).
- [30] E. J. van Beek, C. Kuhl, Y. Anzai, P. Desmond, R. L. Ehman, Q. Gong, G. Gold, V. Gulani, M. Hall-Craggs, T. Leiner, C. C. Lim, J. G. Pipe, S. Reeder, C. Reinhold, M. Smits, D. K. Sodickson, C. Tempany, H. A. Vargas, and M. Wang. "Value of MRI in medicine: More than just another test?" In: *Journal of Magnetic Resonance Imaging* (2019).
- [31] T. E. J. Behrens and H. Johansen-Berg. "Relating connectional architecture to grey matter function using diffusion imaging." In: *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* 360.1457 (2005), pp. 903–911.
- [32] T. E. J. Behrens, H. Johansen-Berg, S. Jbabdi, M. Rushworth, and M. Woolrich. "Probabilistic diffusion tractography with multiple fibre orientations: What can we gain?" In: *NeuroImage* 34.1 (2007), pp. 144–155.
- [33] T. E. J. Behrens, H. Johansen-Berg, M. W. Woolrich, S. M. Smith, C. A. M. Wheeler-Kingshott, P. A. Boulby, G. J. Barker, E. L. Sillery, K. Sheehan, O. Ciccarelli, A. J. Thompson, J. M. Brady, and P. M. Matthews. "Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging." In: *Nature neuroscience* 6.7 (2003), pp. 750–757.
- [34] 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.
- [35] M. Belkin and P. Niyogi. "Laplacian eigenmaps and spectral techniques for embedding and clustering". In: Advances in Neural Information Processing Systems. 2002.
- [36] R. A. Benn, R. B. Mars, T. Xu, L. Rodríguez-esparragoza, P. Montesinos, G. Lopez-martin, J. Gonzalez-sanchez, E. P. Duff, and B. Ibañez. "A White Matter Atlas and Common Connectivity Space Facilitate the Pig as a Translational Model in Neuroscience". In: *BioArxiv* (2020).
- [37] J. Berman. "Diffusion MR Tractography As a Tool for Surgical Planning". In: Magnetic Resonance Imaging Clinics of North America (2009).

- [38] R. F. Betzel, L. Byrge, Y. He, J. Goñi, X. N. Zuo, and O. Sporns. "Changes in structural and functional connectivity among resting-state networks across the human lifespan". In: *NeuroImage* (2014).
- [39] J. D. Bijsterbosch, M. W. Woolrich, M. F. Glasser, E. C. Robinson, C. F. Beckmann, D. C. Van Essen, S. J. Harrison, and S. M. Smith. "The relationship between spatial configuration and functional connectivity of brain regions." In: *eLife* 7 (2018).
- [40] 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.
- [41] G. Blazquez Freches, K. V. Haak, K. L. Bryant, M. Schurz, C. F. Beckmann, and R. B. Mars. "Principles of temporal association cortex organisation as revealed by connectivity gradients". In: *Brain Structure and Function* 225.4 (2020), pp. 1245–1260.
- [42] A. D. Boes, S. Prasad, H. Liu, Q. Liu, A. Pascual-Leone, V. S. Caviness, and M. D. Fox. "Network localization of neurological symptoms from focal brain lesions". In: *Brain* (2015).
- [43] M. M. Bohlken, R. M. Brouwer, R. C. W. Mandl, M. P. den Heuvel, A. M. Hedman, M. De Hert, W. Cahn, R. S. Kahn, and H. E. Hulshoff Pol. "Structural Brain Connectivity as a Genetic Marker for Schizophrenia." In: JAMA psychiatry 73.1 (2016), pp. 11–19.
- [44] M. M. Bohlken, R. C. Mandl, R. M. Brouwer, M. P. van den Heuvel, A. M. Hedman, R. S. Kahn, and H. E. Hulshoff Pol. "Heritability of structural brain network topology: A DTI study of 156 twins". In: *Human Brain Mapping* 35.10 (2014), pp. 5295–5305.
- [45] R. Brown. A Brief Account of Microscopical Observations made in the Months of June, July, and August, 1827, on the Particles contained in the Pollen of Plants; and on the General Existence of Active Molecules in Organic and Inorganic Bodies. 1827.
- [46] K. L. Bryant, L. Li, and R. B. Mars. "A comprehensive atlas of white matter tracts in the chimpanzee". In: *bioRxiv* (2020).
- [47] K. S. Button, J. P. Ioannidis, C. Mokrysz, B. A. Nosek, J. Flint, E. S. Robinson, and M. R. Munafò. "Power failure: Why small sample size undermines the reliability of neuroscience". In: *Nature Reviews Neuroscience* (2013).
- [48] R. B. Buxton, K. Uludağ, D. J. Dubowitz, and T. T. Liu. "Modeling the hemodynamic response to brain activation". In: *NeuroImage* 23 (2004), S220– S233.
- [49] D. Bzdok, T. E. Nichols, and S. M. Smith. "Towards algorithmic analytics for large-scale datasets". In: *Nature Machine Intelligence* 1.7 (2019), pp. 296– 306.

- [50] B. J. Casey et al. "The Adolescent Brain Cognitive Development (ABCD) study: Imaging acquisition across 21 sites". In: *Developmental Cognitive Neu*roscience (2018).
- [51] M. Catani, F. Dell'Acqua, and M. Thiebaut de Schotten. "A revised limbic system model for memory, emotion and behaviour". In: *Neuroscience and Biobehavioral Reviews* 37.8 (2013), pp. 1724–1737.
- [52] 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.
- [53] 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.
- [54] M. Catani and M. Thiebaut de Schotten. "A diffusion tensor imaging tractography atlas for virtual in vivo dissections". In: *Cortex* 44.8 (2008), pp. 1105– 1132.
- [55] L. Cerliani, R. M. Thomas, S. Jbabdi, J. C. Siero, L. Nanetti, A. Crippa, V. Gazzola, H. D'Arceuil, and C. Keysers. "Probabilistic tractography recovers a rostrocaudal trajectory of connectivity variability in the human insular cortex". In: *Human Brain Mapping* (2012).
- [56] M. Chen. Gram-Schmidt orthogonalization. 2020.
- [57] L. Cheng, Y. Zhang, G. Li, J. Wang, W. D. Hopkins, C. C. Sherwood, G. Gong, L. Fan, and T. Jiang. "Divergent Connectional Asymmetries of the Inferior Parietal Lobule Shape Hemispheric Specialization in Humans, Chimpanzees, and Macaque Monkeys". In: *bioRxiv* (2021), p. 2021.01.26.428189.
- [58] Q. Chenot, N. Tzourio-Mazoyer, F. Rheault, M. Descoteaux, F. Crivello, L. Zago, E. Mellet, G. Jobard, M. Joliot, B. Mazoyer, and L. Petit. "A population-based atlas of the human pyramidal tract in 410 healthy participants". In: *Brain Structure and Function* (2019).
- [59] O. Ciccarelli, M. Catani, H. Johansen-Berg, C. Clark, and A. Thompson. "Diffusion-based tractography in neurological disorders: concepts, applications, and future developments". In: *The Lancet Neurology* 7.8 (2008), pp. 715– 727.
- [60] M. Cobb. The Idea of the Brain: A History. Profile Books, 2020.
- [61] D. M. Cole, S. M. Smith, and C. F. Beckmann. "Advances and pitfalls in the analysis and interpretation of resting-state FMRI data." In: *Frontiers in* systems neuroscience 4 (2010), p. 8.
- [62] G. Collin, R. S. Kahn, M. A. de Reus, W. Cahn, and M. P. van den Heuvel. "Impaired rich club connectivity in unaffected siblings of schizophrenia patients." In: *Schizophrenia bulletin* 40.2 (2014), pp. 438–448.
- [63] P. Comon. "Independent component analysis, A new concept?" In: Signal Processing (1994).

- [64] 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 of the United States of America* 96.18 (1999), pp. 10422–10427.
- [65] M. Corbetta, L. Ramsey, A. Callejas, A. Baldassarre, C. D. Hacker, J. S. Siegel, S. V. Astafiev, J. Rengachary, K. Zinn, C. E. Lang, L. T. Connor, R. Fucetola, M. Strube, A. R. Carter, and G. L. Shulman. "Common behavioral clusters and subcortical anatomy in stroke". In: *Neuron* (2015).
- [66] A. Córdova-Palomera, T. Kaufmann, K. Persson, D. Alnæs, N. T. Doan, T. Moberget, M. J. Lund, M. L. Barca, A. Engvig, A. Brækhus, K. Engedal, O. A. Andreassen, G. Selbæk, and L. T. Westlye. "Disrupted global metastability and static and dynamic brain connectivity across individuals in the Alzheimer's disease continuum." In: *Scientific reports* 7 (2017), p. 40268.
- [67] A. M. Dale, B. Fischl, and M. I. Sereno. "Cortical surface-based analysis: I. Segmentation and surface reconstruction". In: *NeuroImage* (1999).
- [68] A. De Benedictis, H. Duffau, B. Paradiso, E. Grandi, S. Balbi, E. Granieri, E. Colarusso, F. Chioffi, C. E. Marras, and S. Sarubbo. "Anatomo-functional study of the temporo-parieto-occipital region: Dissection, tractographic and brain mapping evidence from a neurosurgical perspective". In: *Journal of Anatomy* (2014).
- [69] J. DeFelipe. "From the connectome to the synaptome: an epic love story". In: Science 330.6008 (2010), pp. 1198–1201.
- [70] M. Demirtaş, J. B. Burt, M. Helmer, J. L. Ji, B. D. Adkinson, M. F. Glasser, D. C. Van Essen, S. N. Sotiropoulos, A. Anticevic, and J. D. Murray. "Hierarchical Heterogeneity across Human Cortex Shapes Large-Scale Neural Dynamics". In: *Neuron* (2019).
- [71] R. S. Desikan, F. Ségonne, B. Fischl, B. T. Quinn, B. C. Dickerson, D. Blacker, R. L. Buckner, A. M. Dale, R. P. Maguire, B. T. Hyman, M. S. Albert, and R. J. Killiany. "An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest". In: *NeuroImage* (2006).
- [72] E. A. DeYoe and R. V. Raut. "Visual Mapping Using BOLD fMRI Edgar". In: Neuroimaging Clin N Am. (2014).
- [73] R. Dinga, L. Schmaal, B. W. Penninx, M. J. van Tol, D. J. Veltman, L. van Velzen, M. Mennes, N. J. van der Wee, and A. F. Marquand. "Evaluating the evidence for biotypes of depression: Methodological replication and extension of Drysdale et al. (2017)". In: *NeuroImage: Clinical* (2019).

- [74] 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: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 36.25 (2016), pp. 6758–6770.
- [75] B. Draganski, F. Kherif, S. Klöppel, P. A. Cook, D. C. Alexander, G. J. Parker, R. Deichmann, J. Ashburner, and R. S. Frackowiak. "Evidence for segregated and integrative connectivity patterns in the human basal ganglia". In: *Journal of Neuroscience* (2008).
- [76] A. T. Drysdale, L. Grosenick, J. Downar, K. Dunlop, F. Mansouri, Y. Meng, R. N. Fetcho, B. Zebley, D. J. Oathes, A. Etkin, A. F. Schatzberg, K. Sudheimer, J. Keller, H. S. Mayberg, F. M. Gunning, G. S. Alexopoulos, M. D. Fox, A. Pascual-Leone, H. U. Voss, B. J. Casey, M. J. Dubin, and C. Liston. "Resting-state connectivity biomarkers define neurophysiological subtypes of depression". In: *Nature Medicine* 23.1 (2017), pp. 28–38.
- [77] R. R. Edelman. "The history of MR imaging as seen through the pages of radiology". In: *Radiology* (2014).
- [78] B. L. Edlow, A. Mareyam, A. Horn, J. R. Polimeni, T. Witzel, M. D. Tisdall, J. C. Augustinack, J. P. Stockmann, B. R. Diamond, A. Stevens, L. S. Tirrell, R. D. Folkerth, L. L. Wald, B. Fischl, and A. van der Kouwe. "7 Tesla MRI of the ex vivo human brain at 100 micron resolution". In: *Scientific Data* (2019).
- [79] 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.
- [80] N. Eichert, E. C. Robinson, K. L. Bryant, S. Jbabdi, M. Jenkinson, L. Li, K. Krug, K. E. Watkins, and R. B. Mars. "Cross-species cortical alignment identifies different types of anatomical reorganization in the primate temporal lobe". In: *eLife* (2020).
- [81] S. Eickhoff, T. E. Nichols, J. D. Van Horn, and J. A. Turner. "Sharing the wealth: Neuroimaging data repositories". In: *NeuroImage* (2016).
- [82] A. Einstein. "On the motion of small particles suspended in liquids at rest required by the molecular-kinetic theory of heat". In: Annalen der physik 17 (1905), pp. 549–560.
- [83] J. Fan, F. Han, and H. Liu. "Challenges of Big Data analysis". In: National Science Review (2014).
- [84] L. Fekonja, Z. Wang, I. Bährend, T. Rosenstock, J. Rösler, L. Wallmeroth, P. Vajkoczy, and T. Picht. "Manual for clinical language tractography". In: *Acta Neurochirurgica* (2019).
- [85] S. Finger. "Chapter 10 The birth of localization theory". In: *Handbook of Clinical Neurology*. 2009.

- [86] B. Fischl. "FreeSurfer." In: NeuroImage 62.2 (2012), pp. 774–781.
- [87] M. F. Folstein, S. E. Folstein, and P. R. McHugh. ""Mini-mental state". A practical method for grading the cognitive state of patients for the clinician". In: Journal of Psychiatric Research (1975).
- [88] A. Fornito, A. Zalesky, C. Pantelis, and E. T. Bullmore. "Schizophrenia, neuroimaging and connectomics." In: *NeuroImage* 62.4 (2012), pp. 2296–2314.
- [89] O. Friman, J. Cedefamn, P. Lundberg, M. Borga, and H. Knutsson. "Detection of neural activity in functional MRI using canonical correlation analysis". In: *Magnetic Resonance in Medicine* (2001).
- [90] K. J. Friston, C. D. Frith, R. S. Frackowiak, and R. Turner. "Characterizing dynamic brain responses with fMRI: A multivariate approach". In: *NeuroIm*age (1995).
- [91] K. J. Friston, J. B. Poline, A. P. Holmes, C. D. Frith, and R. S. Frackowiak. "A multivariate analysis of PET activation studies". In: *Human Brain Mapping* (1996).
- [92] A. Fry, T. J. Littlejohns, C. Sudlow, N. Doherty, L. Adamska, T. Sprosen, R. Collins, and N. E. Allen. "Comparison of Sociodemographic and Health-Related Characteristics of UK Biobank Participants with Those of the General Population". In: American Journal of Epidemiology (2017).
- [93] A. Gandomi and M. Haider. "Beyond the hype: Big data concepts, methods, and analytics". In: *International Journal of Information Management* (2015).
- [94] R. Garcia-Dias et al. "Neuroharmony: A new tool for harmonizing volumetric MRI data from unseen scanners". In: *NeuroImage* (2020).
- [95] E. Garyfallidis, M. Brett, M. M. Correia, G. B. Williams, and I. Nimmo-Smith. "QuickBundles, a Method for Tractography Simplification". In: Frontiers in Neuroscience (2012).
- [96] R. C. Gershon, M. V. Wagster, H. C. Hendrie, N. A. Fox, K. F. Cook, and C. J. Nowinski. "NIH toolbox for assessment of neurological and behavioral function." In: *Neurology* (2013).
- [97] M. F. Glasser and D. C. Van Essen. "Mapping human cortical areas in vivo based on myelin content as revealed by T1- and T2-weighted MRI." In: *The Journal of Neuroscience* 31.32 (2011), pp. 11597–11616.
- [98] 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.
- [99] M. F. Glasser, M. S. Goyal, T. M. Preuss, M. E. Raichle, and D. C. Van Essen. "Trends and properties of human cerebral cortex: Correlations with cortical myelin content". In: *NeuroImage* (2014).

- [100] 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, M. Jenkinson, and WU-Minn HCP Consortium. "The minimal preprocessing pipelines for the Human Connectome Project." In: *NeuroImage* 80 (2013), pp. 105–124.
- [101] W. Gong, C. F. Beckmann, and S. M. Smith. "Phenotype Discovery from Population Brain Imaging". In: *bioRxiv* (2020), p. 2020.03.05.973172.
- [102] N. Goyal, D. Moraczewski, and A. G. Thomas. "Computationally replicating the Smith et al. (2015) positive-negative mode linking functional connectivity and subject measures". In: *bioRxiv* (2020).
- [103] G. Grabner, A. L. Janke, M. M. Budge, D. Smith, J. Pruessner, and D. L. Collins. "Symmetric atlasing and model based segmentation: An application to the hippocampus in older adults". In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). 2006.
- M. D. Greicius, B. Krasnow, A. L. Reiss, and V. Menon. "Functional connectivity in the resting brain: a network analysis of the default mode hypothesis." In: Proceedings of the National Academy of Sciences of the United States of America 100.1 (2003), pp. 253–258.
- [105] A. Griffa, P. S. Baumann, J. P. Thiran, and P. Hagmann. "Structural connectomics in brain diseases". In: *NeuroImage* 80 (2013), pp. 515–526.
- [106] L. Griffanti, G. Salimi-Khorshidi, C. F. Beckmann, E. J. Auerbach, G. Douaud, C. E. Sexton, E. Zsoldos, K. P. Ebmeier, N. Filippini, C. E. Mackay, S. Moeller, J. Xu, E. Yacoub, G. Baselli, K. Ugurbil, K. L. Miller, and S. M. Smith. "ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging." In: *NeuroImage* 95 (2014), pp. 232– 247.
- [107] K. Grill-Spector. "What has fMRI taught us about object recognition?" In: Object Categorization: Computer and Human Vision Perspectives. 2009.
- [108] M. de Groot, M. W. Vernooij, S. Klein, M. A. Ikram, F. M. Vos, S. M. Smith, W. J. Niessen, and J. L. Andersson. "Improving alignment in Tract-based spatial statistics: Evaluation and optimization of image registration". In: *NeuroImage* 76 (2013), pp. 400–411.
- [109] L. Grosenick, T. C. Shi, F. M. Gunning, M. J. Dubin, J. Downar, and C. Liston. "Functional and Optogenetic Approaches to Discovering Stable Subtype-Specific Circuit Mechanisms in Depression". In: *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging* (2019).
- [110] R. C. Gur, J. D. Ragland, P. J. Moberg, T. H. Turner, W. B. Bilker, C. Kohler, S. J. Siegel, and R. E. Gur. "Computerized neurocognitive scanning: I. Methodology and validation in healthy people." In: *Neuropsychopharmacology : official publication of the American College of Neuropsychopharmacology* 25.5 (2001), pp. 766–776.

- [111] R. C. Gur, J. Richard, P. Hughett, M. E. Calkins, L. Macy, W. B. Bilker, C. Brensinger, and R. E. Gur. "A cognitive neuroscience-based computerized battery for efficient measurement of individual differences: standardization and initial construct validation." In: *Journal of neuroscience methods* 187.2 (2010), pp. 254–262.
- [112] E. M. Haacke, R. W. Brown, M. R. Thompson, and R. Venkatesan. Magnetic resonance imaging: physical principles and sequence design. Vol. 82. Wiley-Liss New York: 1999.
- [113] P. Hagmann, M. Kurant, X. Gigandet, P. Thiran, J. Van Wedeen, R. Meuli, and J.-P. Thiran. "Mapping human whole-brain structural networks with diffusion MRI". In: *PloS One* 2.7 (2007), e597.
- [114] F. Han, Y. Gu, G. L. Brown, X. Zhang, and X. Liu. "Neuroimaging contrast across the cortical hierarchy is the feature maximally linked to behavior and demographics". In: *NeuroImage* (2020).
- [115] C. B. Hansen, Q. Yang, I. Lyu, F. Rheault, C. Kerley, B. Q. Chandio, S. Fadnavis, O. Williams, A. T. Shafer, S. M. Resnick, D. H. Zald, L. E. Cutting, W. D. Taylor, B. Boyd, E. Garyfallidis, A. W. Anderson, M. Descoteaux, B. A. Landman, and K. G. Schilling. "Pandora: 4-D White Matter Bundle Population-Based Atlases Derived from Diffusion MRI Fiber Tractography". In: *Neuroinformatics* (2020).
- [116] M. M. Hansen, T. Miron-Shatz, A. Y. S. Lau, and C. Paton. "Big Data in Science and Healthcare: A Review of Recent Literature and Perspectives". In: *Yearbook of Medical Informatics* (2014).
- [117] J. Hau, S. Sarubbo, G. Perchey, F. Crivello, L. Zago, E. Mellet, G. Jobard, M. Joliot, B. Mazoyer, N. Tzourio-Mazoyer, and L. Petit. "Cortical terminations of the inferior fronto-occipital and uncinate fasciculi: Anatomical stem-based virtual dissection". In: *Frontiers in Neuroanatomy* (2016).
- [118] Y. He, Z. Chen, G. Gong, and A. Evans. "Neuronal networks in Alzheimer's disease." In: *The Neuroscientist : a review journal bringing neurobiology*, *neurology and psychiatry* 15.4 (2009), pp. 333–350.
- [119] E. E. Hecht, D. A. Gutman, B. A. Bradley, T. M. Preuss, and D. Stout. "Virtual dissection and comparative connectivity of the superior longitudinal fasciculus in chimpanzees and humans". In: *NeuroImage* 108 (2015), pp. 124– 137.
- [120] E. E. Hecht, D. A. Gutman, T. M. Preuss, M. M. Sanchez, L. A. Parr, and J. K. Rilling. "Process versus product in social learning: Comparative diffusion tensor imaging of neural systems for action execution-observation matching in macaques, chimpanzees, and humans". In: *Cerebral Cortex* (2013).
- [121] 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.

- [122] M. Helmer, S. Warrington, A.-R. Mohammadi-Nejad, L. Ji, A. Howell, B. Rosand, A. Anticevic, S. Sotiropoulos, and J. Murray. "On stability of Canonical Correlation Analysis and Partial Least Squares with application to brainbehavior associations". In: *bioRxiv* (2020).
- [123] R. Henson. "A mini-review of fMRI studies of human medial temporal lobe activity associated with recognition memory". In: Quarterly Journal of Experimental Psychology Section B: Comparative and Physiological Psychology (2005).
- [124] 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.
- [125] M. P. van den Heuvel and O. Sporns. *Network hubs in the human brain.* 2013.
- [126] M. Hiscock and M. Kinsbourne. "Lateralization of language across the life span". In: Handbook of the Neuroscience of Language. Elsevier, 2008, pp. 247– 255.
- [127] M. Horbruegger, K. Loewe, J. Kaufmann, M. Wagner, S. Schippling, M. Pawlitzki, and M. A. Schoenfeld. "Anatomically constrained tractography facilitates biologically plausible fiber reconstruction of the optic radiation in multiple sclerosis". In: *NeuroImage: Clinical* (2019).
- [128] H. Hotelling. "Relations Between Two Sets of Variates". In: *Biometrika* (1936).
- [129] H. Howells, M. Thiebaut de Schotten, F. Dell'Acqua, A. Beyh, G. Zappalà, A. Leslie, A. Simmons, D. G. Murphy, and M. Catani. "Frontoparietal tracts linked to lateralized hand preference and manual specialization". In: *Cerebral Cortex* 28.7 (2018), pp. 2482–2494.
- [130] K. Hua, J. Zhang, S. Wakana, H. Jiang, X. Li, D. S. Reich, P. A. Calabresi, J. J. Pekar, P. C. van Zijl, and S. Mori. "Tract probability maps in stereotaxic spaces: Analyses of white matter anatomy and tract-specific quantification". In: *NeuroImage* (2008).
- P. S. Huppi and J. Dubois. "Diffusion tensor imaging of brain development". In: Semin Fetal Neonatal Med 11.6 (2006), pp. 489–497.
- [132] C. R. Jack et al. "The Alzheimer's Disease Neuroimaging Initiative (ADNI): MRI methods". In: *Journal of Magnetic Resonance Imaging* (2008).
- [133] A. G. Jansen, S. E. Mous, T. White, D. Posthuma, and T. J. Polderman. "What Twin Studies Tell Us About the Heritability of Brain Development, Morphology, and Function: A Review". In: *Neuropsychology Review* 25.1 (2015), pp. 27–46.
- [134] S. Jbabdi and T. E. J. Behrens. "Specialization: the connections have it." In: *Nature Neuroscience* 15.2 (2012), pp. 171–172.
- [135] S. Jbabdi and H. Johansen-Berg. "Tractography: where do we go from here?" In: *Brain connectivity* 1.3 (2011), pp. 169–183.

- [136] 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.
- [137] 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.
- [138] M. Jenkinson, C. F. Beckmann, T. E. J. Behrens, M. W. Woolrich, and S. M. Smith. "FSL Review". In: *NeuroImage* 62.2 (2012), pp. 782–790.
- [139] J. L. Ji, M. Helmer, C. Fonteneau, J. B. Burt, Z. Tamayo, J. Demšar, B. Adkinson, A. Savić, K. H. Preller, F. Moujaes, F. X. Vollenweider, W. J. Martin, G. Repovš, J. D. Murray, and A. Anticevic. "Mapping Brain-Behavior Space Relationships Along the Psychosis Spectrum". In: *bioRxiv* (2020), p. 2020.09.15.267310.
- [140] H. Johansen-Berg and T. E. J. Behrens. "Just pretty pictures? What diffusion tractography can add in clinical neuroscience." In: *Current opinion in neurology* 19.4 (2006), pp. 379–385.
- [141] H. Johansen-Berg, T. E. J. Behrens, M. D. Robson, I. Drobnjak, M. F. S. Rushworth, J. M. Brady, S. M. Smith, D. J. Higham, and P. M. Matthews. "Changes in connectivity profiles define functionally distinct regions in human medial frontal cortex." In: *Proceedings of the National Academy of Sciences of the United States of America* 101.36 (2004), pp. 13335–13340.
- [142] H. Johansen-Berg, T. E. J. Behrens, E. Sillery, O. Ciccarelli, A. J. Thompson, S. M. Smith, and P. M. Matthews. "Functional-anatomical validation and individual variation of diffusion tractography-based segmentation of the human thalamus." In: *Cerebral Cortex* 15.1 (2005), pp. 31–39.
- [143] D. K. Jones, A. Simmons, S. C. Williams, and M. A. Horsfield. "Non-invasive assessment of axonal fiber connectivity in the human brain via diffusion tensor MRI." In: *Magnetic resonance in medicine* 42.1 (1999), pp. 37–41.
- [144] D. K. Jones. "Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI". In: *Magnetic Resonance in Medicine* (2003).
- [145] D. K. Jones. "Studying connections in the living human brain with diffusion MRI". In: *Cortex* 44.8 (2008), pp. 936–952.
- [146] F. R. Jumah and R. H. Dossani. Neuroanatomy, Cingulate Cortex. 2019.
- [147] J. Jungilligens, J. Wellmer, A. Kowoll, U. Schlegel, N. Axmacher, and S. Popkirov. "Microstructural integrity of affective neurocircuitry in patients with dissociative seizures is associated with emotional task performance, illness severity and trauma history". In: *Seizure* (2020).
- [148] C. Jutten and J. Herault. "Blind separation of sources, part I: An adaptive algorithm based on neuromimetic architecture". In: Signal Processing (1991).

- [149] A. Kamali, D. M. Yousem, D. D. Lin, H. I. Sair, S. P. Jasti, Z. Keser, R. F. Riascos, and K. M. Hasan. "Mapping the trajectory of the stria terminalis of the human limbic system using high spatial resolution diffusion tensor tractography". In: *Neuroscience Letters* 608 (2015), pp. 45–50.
- [150] N. H. Kashou. "Current Trends of fMRI in Vision Science: A Review". In: Functional Magnetic Resonance Imaging - Advanced Neuroimaging Applications. 2012.
- [151] V. Kiviniemi, J.-H. Kantola, J. Jauhiainen, A. Hyvärinen, and O. Tervonen. "Independent component analysis of nondeterministic fMRI signal sources". In: *NeuroImage* 19.2 (2003), pp. 253–260.
- [152] A. Klein, J. Andersson, B. A. Ardekani, J. Ashburner, B. Avants, M. C. Chiang, G. E. Christensen, D. L. Collins, J. Gee, P. Hellier, J. H. Song, M. Jenkinson, C. Lepage, D. Rueckert, P. Thompson, T. Vercauteren, R. P. Woods, J. J. Mann, and R. V. Parsey. "Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration". In: *NeuroImage* (2009).
- [153] M. L. Kringelbach. The human orbitofrontal cortex: Linking reward to hedonic experience. 2005.
- [154] P. W. Kuchel, G. Pagès, K. Nagashima, S. Velan, V. Vijayaragavan, V. Nagarajan, and K. H. Chuang. "Stejskal-tanner equation derived in full". In: *Concepts in Magnetic Resonance Part A: Bridging Education and Research* (2012).
- [155] S. Kullback and R. A. Leibler. "On Information and Sufficiency". In: The Annals of Mathematical Statistics (1951).
- [156] F. Latini, J. Mårtensson, E. M. Larsson, M. Fredrikson, F. Åhs, M. Hjortberg, H. Aldskogius, and M. Ryttlefors. "Segmentation of the inferior longitudinal fasciculus in the human brain: A white matter dissection and diffusion tensor tractography study". In: *Brain Research* 1675 (2017), pp. 102–115.
- [157] D. Le Bihan. "Diffusion MRI: what water tells us about the brain." In: EMBO molecular medicine 6.5 (2014), pp. 569–573.
- [158] 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* (1986).
- [159] L. F. Leach and R. K. Henson. "Bias and precision of the squared canonical correlation coefficient under nonnormal data conditions". In: *Journal of Modern Applied Statistical Methods* (2014).
- [160] C. Lebel, M. Gee, R. Camicioli, M. Wieler, W. Martin, and C. Beaulieu. "Diffusion tensor imaging of white matter tract evolution over the lifespan". In: *NeuroImage* (2012).
- [161] H. S. Lee. "Canonical correlation analysis using small number of samples". In: Communications in Statistics: Simulation and Computation (2007).

- [162] M. H. Lee, C. D. Smyser, and J. S. Shimony. "Resting-state fMRI: a review of methods and clinical applications." In: *American journal of neuroradiology* 34.10 (2013), pp. 1866–1872.
- [163] S. Leonelli. "The challenges of big data biology". In: *eLife* (2019).
- [164] J. P. Lerch, A. J. W. van der Kouwe, A. Raznahan, T. Paus, H. Johansen-Berg, K. L. Miller, S. M. Smith, B. Fischl, and S. N. Sotiropoulos. "Studying neuroanatomy using MRI." In: *Nature Neuroscience* 20.3 (2017), pp. 314–326.
- [165] C. Leuze, M. Goubran, M. Barakovic, M. Aswendt, Q. Tian, B. Hsueh, A. Crow, E. M. Weber, G. K. Steinberg, M. Zeineh, E. D. Plowey, A. Daducci, G. Innocenti, J. P. Thiran, K. Deisseroth, and J. A. McNab. "Comparison of diffusion MRI and CLARITY fiber orientation estimates in both gray and white matter regions of human and primate brain". In: *NeuroImage* (2021).
- [166] J. Li, T. Bolt, D. Bzdok, J. S. Nomi, B. T. T. Yeo, R. N. Spreng, and L. Q. Uddin. "Topography and behavioral relevance of the global signal in the human brain". In: *Scientific Reports* 9.1 (2019), p. 14286.
- [167] I. Lipp, G. D. Parker, E. C. Tallantyre, A. Goodall, S. Grama, E. Patitucci, P. Heveron, V. Tomassini, and D. K. Jones. "Tractography in the presence of multiple sclerosis lesions". In: *NeuroImage* (2020).
- [168] Y. Liu, M. Liang, Y. Zhou, Y. He, Y. Hao, M. Song, C. Yu, H. Liu, Z. Liu, and T. Jiang. "Disrupted small-world networks in schizophrenia." In: *Brain* : a journal of neurology 131.Pt 4 (2008), pp. 945–961.
- [169] C.-Y. Lo, P.-N. Wang, K.-H. Chou, J. Wang, Y. He, and C.-P. Lin. "Diffusion tensor tractography reveals abnormal topological organization in structural cortical networks in Alzheimer's disease." In: *The Journal of neuroscience :* the official journal of the Society for Neuroscience 30.50 (2010), pp. 16876– 16885.
- [170] F. López-Muñoz, J. Boya, and C. Alamo. "Neuron theory, the cornerstone of neuroscience, on the centenary of the Nobel Prize award to Santiago Ramón y Cajal". In: *Brain Research Bulletin* (2006).
- [171] S. Lui, X. J. Zhou, J. A. Sweeney, and Q. Gong. "Psychoradiology: The frontier of neuroimaging in psychiatry". In: *Radiology*. 2016.
- [172] J. Lv, M. Di Biase, R. F. Cash, L. Cocchi, V. L. Cropley, P. Klauser, Y. Tian, J. Bayer, L. Schmaal, S. Cetin-Karayumak, Y. Rathi, O. Pasternak, C. Bousman, C. Pantelis, F. Calamante, and A. Zalesky. "Individual deviations from normative models of brain structure in a large cross-sectional schizophrenia cohort". In: *Molecular Psychiatry* (2020), p. 75.
- [173] D. J. Mackay. "Probable networks and plausible predictions a review of practical bayesian methods for supervised neural networks". In: Network: Computation in Neural Systems (1995).

- [174] D. J. Madden, I. J. Bennett, A. Burzynska, G. G. Potter, N. kuei Chen, and A. W. Song. "Diffusion tensor imaging of cerebral white matter integrity in cognitive aging". In: *Biochimica et Biophysica Acta - Molecular Basis of Disease* (2012).
- [175] C. Maffei, S. Sarubbo, and J. Jovicich. "Diffusion-based tractography atlas of the human acoustic radiation". In: *Scientific Reports* (2019).
- [176] 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.
- [177] 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.
- [178] 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 (2009), pp. 777–785.
- [179] P. Mansfield. "Echo-Planar Imaging". In: *eMagRes* (2007).
- [180] P. Mansfield. "Multi-planar image formation using NMR spin echoes". In: Journal of Physics C: Solid State Physics 10.3 (1977), p. L55.
- [181] D. S. Marcus, M. P. Harms, A. Z. Snyder, M. Jenkinson, J. A. Wilson, M. F. Glasser, D. M. Barch, K. A. Archie, G. C. Burgess, M. Ramaratnam, M. Hodge, W. Horton, R. Herrick, T. Olsen, M. McKay, M. House, M. Hileman, E. Reid, J. Harwell, T. Coalson, J. Schindler, J. S. Elam, S. W. Curtiss, D. C. Van Essen, and W.-M. H. C. P. Consortium. "Human Connectome Project informatics: quality control, database services, and data visualization." In: *NeuroImage* 80 (2013), pp. 202–219.
- [182] S. Marek et al. "Towards Reproducible Brain-Wide Association Studies". In: bioRxiv (2020).
- [183] A. F. Marquand, K. V. Haak, and C. F. Beckmann. "Functional corticostriatal connection topographies predict goal-directed behaviour in humans". In: *Nature Human Behaviour* (2017).
- [184] R. B. Mars, N. Eichert, S. Jbabdi, L. Verhagen, and M. F. S. Rushworth.
  "Connectivity and the search for specializations in the language-capable brain". In: *Current Opinion in Behavioral Sciences* 21 (2018), pp. 19–26.
- [185] R. B. Mars, S. Jbabdi, J. Sallet, J. X. O'Reilly, P. L. Croxson, E. Olivier, M. P. Noonan, C. Bergmann, A. S. Mitchell, M. G. Baxter, T. E. J. Behrens, H. Johansen-Berg, V. Tomassini, K. L. Miller, and M. F. S. Rushworth. "Diffusion-weighted imaging tractography-based parcellation of the human parietal cortex and comparison with human and macaque resting-state functional connectivity." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 31.11 (2011), pp. 4087–4100.

- [186] 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. Dunbar, A. A. Khrapitchev, N. R. Sibson, K. L. Miller, and M. F. 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.
- [187] R. B. Mars, R. E. Passingham, and S. Jbabdi. "Connectivity Fingerprints: From Areal Descriptions to Abstract Spaces". In: *Trends in Cognitive Sciences* 22.11 (2018), pp. 1026–1037.
- [188] R. B. Mars, S. N. Sotiropoulos, R. E. Passingham, J. Sallet, L. Verhagen, A. A. Krapitchev, N. Sibson, and S. Jbabdi. "Whole brain comparative anatomy using connectivity blueprints". In: *eLife* 7.e35237 (2018).
- [189] R. B. Mars, J. Sallet, U. Schüffelgen, S. Jbabdi, I. Toni, and M. F. Rushworth. "Connectivity-based subdivisions of the human right "temporoparietal junction area": Evidence for different areas participating in different cortical networks". In: Cerebral Cortex (2012).
- [190] R. B. Mars, L. Verhagen, T. E. Gladwin, F. X. Neubert, J. Sallet, and M. F. Rushworth. "Comparing brains by matching connectivity profiles". In: *Neuroscience and Biobehavioral Reviews* (2016).
- [191] 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.
- [192] D. McRobbie, E. Moore, J. Graves, and M. R. Prince. MRI from Picture to Proton, Ed. 2. 2005.
- [193] N. Menjot de Champfleur, I. L. Maldonado, S. Moritz-Gasser, P. Machi, E. Le Bars, A. Bonafé, and H. Duffau. "Middle longitudinal fasciculus delineation within language pathways: a diffusion tensor imaging study in human". In: *European journal of radiology* 82.1 (2013), pp. 151–157.
- [194] T. Meynert. Psychiatry: a clinical treatise on diseases of the fore-brain based upon a study of its structure, functions, and nutrition. Ed. by B. Sachs. New York : G.P. Putnam's sons, 1885.
- [195] A. Mihalik, R. A. Adams, and Q. Huys. "Canonical Correlation Analysis for Identifying Biotypes of Depression". In: *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging* (2020).
- [196] A. Mihalik et al. "Brain-behaviour modes of covariation in healthy and clinically depressed young people". In: *Scientific Reports* (2019).
- [197] M. P. Milham et al. "An Open Resource for Non-human Primate Imaging". In: Neuron 100.1 (2018), pp. 61–74.

- [198] 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.
- [199] J. N. Morelli, V. M. Runge, F. Ai, U. Attenberger, L. Vu, S. H. Schmeets, W. R. Nitz, and J. E. Kirsch. "An image-based approach to understanding the physics of MR artifacts". In: *Radiographics* 31.3 (2011), pp. 849–866.
- [200] S. Mori, B. J. Crain, V. P. Chacko, and P. C. 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.
- [201] S. Mori and M. Aggarwal. "In vivo magnetic resonance imaging of the human limbic white matter". In: *Frontiers in Aging Neuroscience* (2014).
- [202] S. Mori, Y. Kageyama, Z. Hou, M. Aggarwal, J. Patel, T. Brown, M. I. Miller, D. Wu, and J. C. Troncoso. "Elucidation of white matter tracts of the human amygdala by detailed comparison between high-resolution postmortem magnetic resonance imaging and histology". In: *Frontiers in Neuroanatomy* (2017).
- [203] S. Mori, S. Wakana, P. C. Van Zijl, and L. M. Nagae-Poetscher. MRI Atlas of Human White Matter. Elsevier, 2005.
- [204] F. X. Neubert, R. B. Mars, J. Sallet, and M. F. Rushworth. "Connectivity reveals relationship of brain areas for reward-guided learning and decision making in human and monkey frontal cortex". In: *Proceedings of the National Academy of Sciences of the United States of America* (2015).
- [205] F. X. Neubert, R. B. Mars, A. G. Thomas, J. Sallet, and M. F. Rushworth. "Comparison of Human Ventral Frontal Cortex Areas for Cognitive Control and Language with Areas in Monkey Frontal Cortex". In: *Neuron* (2014).
- [206] T. E. Nichols, S. Das, S. B. Eickhoff, A. C. Evans, T. Glatard, M. Hanke, N. Kriegeskorte, M. P. Milham, R. A. Poldrack, J. B. Poline, E. Proal, B. Thirion, D. C. Van Essen, T. White, and B. T. Yeo. "Best practices in data analysis and sharing in neuroimaging using MRI". In: *Nature Neuroscience* (2017).
- [207] M. Nowell, S. B. Vos, M. Sidhu, K. Wilcoxen, N. Sargsyan, S. Ourselin, and J. S. Duncan. "Meyer's loop asymmetry and language lateralisation in epilepsy". In: *Journal of Neurology, Neurosurgery and Psychiatry* 87 (2016), pp. 836–842.
- [208] P. G. Nucifora, X. Wu, E. R. Melhem, R. E. Gur, R. C. Gur, and R. Verma. "Automated Diffusion Tensor Tractography. Implementation and Comparison to User-driven Tractography". In: *Academic Radiology* 19.5 (2012), pp. 622– 629.

- [209] L. J. O'Donnell and C. F. Westin. "Automatic tractography segmentation using a high-dimensional white matter atlas". In: *IEEE Transactions on Medical Imaging* (2007).
- [210] L. J. O'Donnell, C. F. Westin, I. Norton, S. Whalen, L. Rigolo, R. Propper, and A. J. Golby. "The fiber laterality histogram: A new way to measure white matter asymmetry". In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). 2010.
- [211] S. Ogawa, D. W. Tank, R. Menon, J. M. Ellermann, S. G. Kim, H. Merkle, and K. Ugurbil. "Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging". In: *Proceedings* of the National Academy of Sciences 89.13 (1992), pp. 5951–5955.
- [212] M. Oldehinkel, C. F. Beckmann, B. Franke, C. A. Hartman, P. J. Hoekstra, J. Oosterlaan, D. Heslenfeld, J. K. Buitelaar, and M. Mennes. "Functional connectivity in cortico-subcortical brain networks underlying reward processing in attention-deficit/hyperactivity disorder." In: *NeuroImage. Clinical* 12 (2016), pp. 796–805.
- [213] D. E. Osher, R. R. Saxe, K. Koldewyn, J. D. Gabrieli, N. Kanwisher, and Z. M. Saygin. "Structural Connectivity Fingerprints Predict Cortical Selectivity for Multiple Visual Categories across Cortex". In: *Cerebral Cortex* (2016).
- [214] Y. Ou, H. Akbari, M. Bilello, X. Da, and C. Davatzikos. "Comparative evaluation of registration algorithms in different brain databases with varying difficulty: Results and insights". In: *IEEE Transactions on Medical Imaging* (2014).
- [215] N. P. Oxtoby, S. Garbarino, N. C. Firth, J. D. Warren, J. M. Schott, D. C. Alexander, and A. D. N. Initiative. "Data-Driven Sequence of Changes to Anatomical Brain Connectivity in Sporadic Alzheimer's Disease." In: *Frontiers in neurology* 8 (2017), p. 580.
- [216] S. Pal, S. Mondal, G. Das, S. Khatua, and Z. Ghosh. *Big data in biology: The hope and present-day challenges in it.* 2020.
- [217] G. E. Palade. "Electron microscope observations of interneuronal and neuromuscular synapses". In: Anat Rec 118 (1954), pp. 335–336.
- [218] S. S. Panesar, F. C. Yeh, T. Jacquesson, W. Hula, and J. C. Fernandez-Miranda. "A Quantitative Tractography Study into the Connectivity, Segmentation and Laterality of the Human Inferior Longitudinal Fasciculus". In: *Frontiers in Neuroanatomy* 12 (2018).
- [219] S. S. Panesar, K. Abhinav, F. C. Yeh, T. Jacquesson, M. Collins, and J. Fernandez-Miranda. "Tractography for Surgical Neuro-Oncology Planning: Towards a Gold Standard". In: *Neurotherapeutics* (2019).
- [220] R. Pascalau, R. Popa Stănilă, S. Sfrângeu, and B. Szabo. "Anatomy of the Limbic White Matter Tracts as Revealed by Fiber Dissection and Tractography". In: World Neurosurgery 113 (2018), e672–e689.

- [221] 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.
- [222] R. Patriat, S. E. Cooper, Y. Duchin, J. Niederer, C. Lenglet, J. Aman, M. C. Park, J. L. Vitek, and N. Harel. "Individualized tractography-based parcellation of the globus pallidus pars interna using 7T MRI in movement disorder patients prior to DBS surgery". In: *NeuroImage* (2018).
- [223] L. Pauling. "Magnetic properties and structure of oxyhemoglobin". In: Proceedings of the National Academy of Sciences 74.7 (1977), pp. 2612–2613.
- [224] T. Paus. Primate anterior cingulate cortex: Where motor control, drive and cognition interface. 2001.
- [225] 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.
- [226] M. S. Pinto, R. Paolella, T. Billiet, P. Van Dyck, P. J. Guns, B. Jeurissen, A. Ribbens, A. J. den Dekker, and J. Sijbers. "Harmonization of Brain Diffusion MRI: Concepts and Methods". In: *Frontiers in Neuroscience* (2020).
- [227] B. R. Plantinga, Y. Temel, Y. Duchin, K. Uludağ, R. Patriat, A. Roebroeck, M. Kuijf, A. Jahanshahi, B. ter Haar Romenij, J. Vitek, and N. Harel. "Individualized parcellation of the subthalamic nucleus in patients with Parkinson's disease with 7T MRI". In: *NeuroImage* (2018).
- [228] J. B. Poline, J. L. Breeze, S. Ghosh, K. F. Gorgolewski, Y. O. Halchenko, M. Hanke, K. G. Helmer, D. S. Marcus, R. A. Poldrack, Y. Schwartz, J. Ashburner, and D. N. Kennedy. "Data sharing in neuroimaging research". In: *Frontiers in Neuroinformatics* (2012).
- [229] R. Pomponio et al. "Harmonization of large MRI datasets for the analysis of brain imaging patterns throughout the lifespan". In: *NeuroImage* (2020).
- [230] R. E. Propper, L. J. O'Donnell, S. Whalen, Y. Tie, I. H. Norton, R. O. Suarez, L. Zollei, A. Radmanesh, and A. J. Golby. "A combined fMRI and DTI examination of functional language lateralization and arcuate fasciculus structure: Effects of degree versus direction of hand preference". In: *Brain and Cognition* (2010).
- [231] R. H. R. Pruim, M. Mennes, D. van Rooij, A. Llera, J. K. Buitelaar, and C. F. Beckmann. "ICA-AROMA: A robust ICA-based strategy for removing motion artifacts from fMRI data." In: *NeuroImage* 112 (2015), pp. 267–277.
- [232] Y. Qin, J. Wu, T. Chen, J. Li, G. Zhang, D. Wu, Y. Zhou, N. Zheng, A. Cai, Q. Ning, A. Manyande, F. Xu, J. Wang, and W. Zhu. "Long-term micro-structure and cerebral blood flow changes in patients recovered from COVID-19 without neurological manifestations". In: *The Journal of Clinical Investigation* (2021).

- [233] M. E. Raichle, A. M. MacLeod, A. Z. Snyder, W. J. Powers, D. A. Gusnard, and G. L. Shulman. "A default mode of brain function." In: *Proceedings of the National Academy of Sciences of the United States of America* 98.2 (2001), pp. 676–682.
- [234] M. N. Rajah and M. D'Esposito. "Region-specific changes in prefrontal function with age: A review of PET and fMRI studies on working and episodic memory". In: *Brain* (2005).
- [235] R. Rathee, V. S. Rallabandi, and P. K. Roy. "Age-related Differences in White Matter Integrity in Healthy Human Brain: Evidence from Structural Mri and Diffusion Tensor Imaging". In: *Magnetic Resonance Insights* (2016).
- [236] J. Raven. "The Raven's Progressive Matrices: Change and Stability over Culture and Time". In: *Cognitive Psychology* (2000).
- [237] W. Reginold, J. Itorralba, A. Tam, A. C. Luedke, J. Fernandez-Ruiz, J. Reginold, O. Islam, and A. Garcia. "Correlating quantitative tractography at 3T MRI and cognitive tests in healthy older adults". In: *Brain Imaging and Behavior* (2016).
- [238] V. Ressel, H. J. van Hedel, I. Scheer, and R. O'Gorman Tuura. "Comparison of DTI analysis methods for clinical research: influence of pre-processing and tract selection methods". In: *European Radiology Experimental* (2018).
- [239] 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* (2014).
- [240] R. Rodríguez-Cruces, B. C. Bernhardt, and L. Concha. "Multidimensional associations between cognition and connectome organization in temporal lobe epilepsy". In: *NeuroImage* 213 (2020).
- [241] L. Roumazeilles, N. Eichert, K. L. Bryant, D. Folloni, J. Sallet, S. Vijayakumar, S. Foxley, B. C. Tendler, S. Jbabdi, C. Reveley, L. Verhagen, L. B. Dershowitz, M. Guthrie, E. Flach, K. L. Miller, and R. B. Mars. "Longitudinal connections and the organization of the temporal cortex in macaques, great apes, and humans". In: *PLoS Biology* (2020).
- [242] M. Rubinov and O. Sporns. "Complex network measures of brain connectivity: uses and interpretations." In: *NeuroImage* 52.3 (2010), pp. 1059–1069.
- [243] S. Sagiroglu and D. Sinanc. "Big data: A review". In: Proceedings of the 2013 International Conference on Collaboration Technologies and Systems, CTS 2013. 2013.
- [244] 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.

- [245] J. Sallet, R. B. Mars, M. P. Noonan, F. X. Neubert, S. Jbabdi, J. X. O'Reilly, N. Filippini, A. G. Thomas, and M. F. Rushworth. "The organization of dorsal frontal cortex in humans and macaques". In: *Journal of Neuroscience* (2013).
- [246] S. Sarubbo and L. Petit. "Editorial: Organization of the White Matter Anatomy in the Human Brain". In: *Frontiers in Neuroanatomy* (2019).
- [247] 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* (2012).
- [248] Z. M. Saygin, D. E. Osher, E. S. Norton, D. A. Youssoufian, S. D. Beach, J. Feather, N. Gaab, J. D. Gabrieli, and N. Kanwisher. "Connectivity precedes function in the development of the visual word form area". In: *Nature Neuroscience* (2016).
- [249] L. K. Scheffer et al. "A connectome and analysis of the adult drosophila central brain". In: *eLife* (2020).
- [250] D. Scheinost, S. Noble, C. Horien, A. S. Greene, E. M. Lake, M. Salehi, S. Gao, X. Shen, D. O'Connor, D. S. Barron, S. W. Yip, M. D. Rosenberg, and R. T. Constable. *Ten simple rules for predictive modeling of individual differences in neuroimaging*. 2019.
- [251] K. G. Schilling, L. Petit, F. Rheault, S. Remedios, C. Pierpaoli, A. W. Anderson, B. A. Landman, and M. Descoteaux. "Brain connections derived from diffusion MRI tractography can be highly anatomically accurate—if we know where white matter pathways start, where they end, and where they do not go". In: *Brain Structure and Function* (2020).
- [252] K. G. Schilling et al. "Tractography dissection variability: what happens when 42 groups dissect 14 white matter bundles on the same dataset?" In: *bioRxiv* (2020), p. 2020.10.07.321083.
- [253] J. D. Schmahmann and D. N. Pandya. Fiber Pathways of the Brain. 2006.
- [254] E. Schmidt. "Zur Theorie der linearen und nichtlinearen Integralgleichungen -I. Teil: Entwicklung willkürlicher Funktionen nach Systemen vorgeschriebener". In: Mathematische Annalen (1907).
- [255] F. D. Schönbrodt and M. Perugini. "At what sample size do correlations stabilize?" In: Journal of Research in Personality (2013).
- [256] S. Seabold and J. Perktold. "Statsmodels: Econometric and Statistical Modeling with Python". In: Proceedings of the 9th Python in Science Conference. 2010.
- [257] T. J. Sejnowski, P. S. Churchland, and J. A. Movshon. Putting big data to good use in neuroscience. 2014.
- [258] K.-K. Shen, S. Rose, J. Fripp, K. L. McMahon, G. I. de Zubicaray, N. G. Martin, P. M. Thompson, M. J. Wright, and O. Salvado. "Investigating brain connectivity heritability in a twin study using diffusion imaging data". In: *NeuroImage* 100 (2014), pp. 628–641.

- [259] R. B. Silberstein, A. Pipingas, M. Farrow, F. Levy, C. K. Stough, and D. A. Camfield. "Brain functional connectivity abnormalities in attention-deficit hyperactivity disorder." In: *Brain and behavior* 6.12 (2016), e00583.
- [260] V. Siless, K. Chang, B. Fischl, and A. Yendiki. "AnatomiCuts: Hierarchical clustering of tractography streamlines based on anatomical similarity". In: *NeuroImage* (2018).
- [261] 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. J. Behrens. "Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data." In: *NeuroImage* 31.4 (2006), pp. 1487–1505.
- [262] S. M. Smith, D. Vidaurre, C. F. Beckmann, M. F. Glasser, M. Jenkinson, K. L. Miller, T. E. Nichols, E. C. Robinson, G. Salimi-Khorshidi, M. W. Woolrich, D. M. Barch, K. Uğurbil, and D. C. Van Essen. "Functional connectomics from resting-state fMRI." In: *Trends in cognitive sciences* 17.12 (2013), pp. 666–682.
- [263] 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.
- [264] S. M. Smith, T. E. Nichols, D. Vidaurre, A. M. Winkler, T. E. J. Behrens, M. F. Glasser, K. Ugurbil, D. M. Barch, D. C. Van Essen, and K. L. Miller. "A positive-negative mode of population covariation links brain connectivity, demographics and behavior". In: *Nature Neuroscience* 18.11 (2015), pp. 1565– 1567.
- [265] S. M. Smith et al. "Resting-state fMRI in the Human Connectome Project". In: Neuroimage 80 (2013), pp. 125–143.
- [266] S. M. Smith and T. E. Nichols. "Statistical challenges in "Big Data" human neuroimaging". In: Neuron 97.2 (2018), pp. 263–268.
- [267] Y. Song, P. J. Schreier, D. Ramírez, and T. Hasija. "Canonical correlation analysis of high-dimensional data with very small sample support". In: Signal Processing (2016).
- [268] S. N. Sotiropoulos and A. Zalesky. "Building connectomes using diffusion MRI: why, how and but." In: NMR in Biomedicine 1492 (2017).
- [269] 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. Ven Essen, K. Ugurbil, T. E. Behrens, and WU-Minn HCP Consortium. "Advances in diffusion MRI acquisition and processing in the Human Connectome Project". In: *NeuroImage* 80 (2013), pp. 125–143.
- [270] 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.

- [271] S. N. Sotiropoulos, S. Jbabdi, J. L. Andersson, M. W. Woolrich, K. Ugurbil, and T. E. J. Behrens. "RubiX: combining spatial resolutions for Bayesian inference of crossing fibers in diffusion MRI." In: *IEEE transactions on medical imaging* 32.6 (2013), pp. 969–982.
- [272] 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.
- [273] 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.
- [274] N. Steno. Discours sur l'anatomie du cerveau. 1669.
- [275] K. Strand and S. Kossman. "Further Inquiry into the Stabilities of Standardized and Structure Coefficients in Canonical and Discriminant Analyses". In: *American Educational Research Association Annual Meeting.* 2000.
- [276] C. Sudlow, J. Gallacher, N. Allen, V. Beral, P. Burton, J. Danesh, P. Downey, P. Elliott, J. Green, M. Landray, B. Liu, P. Matthews, G. Ong, J. Pell, A. Silman, A. Young, T. Sprosen, T. Peakman, and R. Collins. "UK Biobank: An Open Access Resource for Identifying the Causes of a Wide Range of Complex Diseases of Middle and Old Age". In: *PLoS Medicine* (2015).
- [277] 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.
- [278] P. Tewarie, J. Meier, B. Prasse, K. Mandke, S. Warrington, C. J. Stam, M. J. Brookes, P. Van Mieghem, S. N. Sotiropoulos, and A. Hillebrand. "Predicting time-resolved electrophysiological brain networks from structural eigenmodes". In: Under review (2020).
- [279] 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. 1245–1246.
- [280] M. Thiebaut de Schotten, D. H. Ffytche, A. Bizzi, F. Dell'Acqua, M. Allin, M. Walshe, R. Murray, S. C. Williams, D. G. 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.
- [281] M. Thiebaut de Schotten, P. L. Croxson, and R. B. Mars. "Large-scale comparative neuroimaging: Where are we and what do we need?" In: *Cortex* 118 (2019), pp. 188–202.
- [282] M. Thiebaut de Schotten, F. Dell'Acqua, S. Forkel, A. Simmons, F. Vergani, D. G. M. Murphy, and M. Catani. "A Lateralized Brain Network for Visuo-Spatial Attention". In: *Nature Precedings* (2011).

- [283] M. Thiebaut de Schotten, F. Dell'Acqua, R. Valabregue, and M. Catani. "Monkey to human comparative anatomy of the frontal lobe association tracts". In: *Cortex* (2012).
- [284] M. Thiebaut de Schotten, C. Foulon, and P. Nachev. "Brain disconnections link structural connectivity with function and behaviour". In: *Nature Communications* (2020).
- [285] M. Thiebaut de Schotten, F. Tomaiuolo, M. Aiello, S. Merola, M. Silvetti, F. Lecce, P. Bartolomeo, and F. Doricchi. "Damage to white matter pathways in subacute and chronic spatial neglect: A group study and 2 single-case studies with complete virtual "in vivo" tractography dissection". In: Cerebral Cortex (2014).
- [286] M. Thiebaut de Schotten, M. Urbanski, R. Valabregue, D. J. Bayle, and E. Volle. "Subdivision of the occipital lobes: An anatomical and functional MRI connectivity study". In: *Cortex* (2014).
- [287] B. Thompson. "Finding a Correction for the Sampling Error in Multivariate Measures of Relationship: A Monte Carlo Study". In: *Educational and Psychological Measurement* (1990).
- [288] E. Thompson. "Mapping Connections in the Neonatal Brain with Magnetic Resonance Imaging". PhD thesis. University of Nottingham, 2020.
- [289] E. Thompson, A. R. Mohammadi-Nejad, E. C. Robinson, J. L. Andersson, S. Jbabdi, M. F. Glasser, M. Bastiani, and S. N. Sotiropoulos. "Non-negative data-driven mapping of structural connections with application to the neonatal brain". In: *NeuroImage* (2020).
- [290] P. M. Thompson et al. "The ENIGMA Consortium: Large-scale collaborative analyses of neuroimaging and genetic data". In: *Brain Imaging and Behavior* (2014).
- [291] R. M. Thorndike. "Studying canonical analysis: Comments on Barcikowski and Stevens". In: *Multivariate Behavioral Research* 11.2 (1976), pp. 249–253.
- [292] R. M. Thorndike and D. J. Weiss. "A study of the stability of canonical correlations and canonical components". In: *Educational and Psychological Measurement* (1973).
- [293] K. Toma and T. Nakai. "Functional MRI in human motor control studies and clinical applications". In: *Magnetic Resonance in Medical Sciences* (2002).
- [294] C. Traynor, R. A. Heckemann, A. Hammers, J. O'Muircheartaigh, W. R. Crum, G. J. Barker, and M. P. Richardson. "Reproducibility of thalamic segmentation based on probabilistic tractography". In: *NeuroImage* (2010).
- [295] O. Tymofiyeva, C. P. Hess, D. Xu, and A. J. Barkovich. "Structural MRI connectome in development: Challenges of the changing brain". In: *British Journal of Radiology* (2014).

- [296] A. C. Tziortzi, S. N. Haber, G. E. Searle, C. Tsoumpas, C. J. Long, P. Shotbolt, G. Douaud, S. Jbabdi, T. E. Behrens, E. A. Rabiner, M. Jenkinson, and R. N. Gunn. "Connectivity-based functional analysis of dopamine release in the striatum using diffusion-weighted MRI and positron emission tomography". In: *Cerebral Cortex* (2014).
- [297] L. Q. Uddin, J. S. Nomi, B. Hébert-Seropian, J. Ghaziri, and O. Boucher. Structure and Function of the Human Insula. 2017.
- [298] K. Uğurbil et al. "Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project." In: *NeuroImage* 80 (2013), pp. 80–104.
- [299] A. R. Vaidya, M. S. Pujara, M. Petrides, E. A. Murray, and L. K. Fellows. "Lesion Studies in Contemporary Neuroscience". In: *Trends in Cognitive Sciences* (2019).
- [300] D. C. Van Essen et al. "The Human Connectome Project: a data acquisition perspective." In: *NeuroImage* 62.4 (2012), pp. 2222–2231.
- [301] D. C. Van Essen. "Windows on the brain: The emerging role of atlases and databases in neuroscience". In: *Current Opinion in Neurobiology* 12.5 (2002), pp. 574–579.
- [302] D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. 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.
- [303] G. Varoquaux, P. R. Raamana, D. A. Engemann, A. Hoyos-Idrobo, Y. Schwartz, and B. Thirion. "Assessing and tuning brain decoders: Cross-validation, caveats, and guidelines". In: *NeuroImage* (2017).
- [304] J. L. Vincent, G. H. Patel, M. D. Fox, A. Z. Snyder, J. T. Baker, D. C. Van Essen, J. M. Zempel, L. H. Snyder, M. Corbetta, and M. E. Raichle. "Intrinsic functional architecture in the anaesthetized monkey brain". In: *Nature* (2007).
- [305] M.-A. T. Vu, T. Adalı, D. Ba, G. Buzsáki, D. Carlson, K. Heller, C. Liston, C. Rudin, V. S. Sohal, A. S. Widge, H. S. Mayberg, G. Sapiro, and K. Dzirasa. "A Shared Vision for Machine Learning in Neuroscience". In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 38.7 (2018), pp. 1601–1607.
- [306] 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.
- [307] S. Wakana, H. Jiang, L. M. Nagae-Poetscher, P. C. van Zijl, and S. Mori. "Fiber Tract-based Atlas of Human White Matter Anatomy". In: *Radiology* 230.1 (2004).
- [308] F. Wang, Z. Dong, Q. Tian, C. Liao, Q. Fan, W. S. Hoge, B. Keil, J. R. Polimeni, L. L. Wald, S. Y. Huang, and K. Setsompop. In vivo human wholebrain Connectom diffusion MRI dataset at 760 µm isotropic resolution. 2020.

- [309] H. T. Wang, D. Bzdok, D. Margulies, C. Craddock, M. Milham, E. Jefferies, and J. Smallwood. "Patterns of thought: Population variation in the associations between large-scale network organisation and self-reported experiences at rest". In: *NeuroImage* 176 (2018), pp. 518–527.
- [310] H. T. Wang, N. S. P. Ho, D. Bzdok, B. C. Bernhardt, D. S. Margulies, E. Jefferies, and J. Smallwood. "Neurocognitive patterns dissociating semantic processing from executive control are linked to more detailed off-task mental time travel". In: *Scientific Reports* (2020).
- [311] J. Wang, R. Khosrowabadi, K. K. Ng, Z. Hong, J. S. X. Chong, Y. Wang, C. Y. Chen, S. Hilal, N. Venketasubramanian, T. Y. Wong, C. L. H. Chen, M. K. Ikram, and J. Zhou. "Alterations in Brain Network Topology and Structural-Functional Connectome Coupling Relate to Cognitive Impairment". In: Frontiers in Aging Neuroscience (2018).
- [312] Y. Wang, A. Metoki, K. H. Alm, and I. R. Olson. "White matter pathways and social cognition". In: *Neuroscience and Biobehavioral Reviews* (2018).
- [313] 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* (2020).
- [314] D. Wassermann, N. Makris, Y. Rathi, M. Shenton, R. Kikinis, M. Kubicki, and C. F. Westin. "The white matter query language: a novel approach for describing human white matter anatomy". In: *Brain Structure and Function* 221.9 (2016), pp. 4705–4721.
- [315] J. Wasserthal, P. Neher, and K. H. Maier-Hein. "TractSeg Fast and accurate white matter tract segmentation". In: *NeuroImage* (2018).
- [316] S. L. Weinberg and R. B. Darlington. "Canonical Analysis When Number of Variables Is Large Relative to Sample Size". In: *Journal of Educational Statistics* (1976).
- [317] L. T. Westlye, K. B. Walhovd, A. M. Dale, A. Bjørnerud, P. Due-Tønnessen, A. Engvig, H. Grydeland, C. K. Tamnes, Y. Østby, and A. M. Fjell. "Lifespan changes of the human brain white matter: Diffusion tensor imaging (DTI) and volumetry". In: *Cerebral Cortex* (2010).
- [318] A. M. Winkler, M. A. Webster, D. Vidaurre, T. E. Nichols, and S. M. Smith. "Multi-level block permutation." In: *NeuroImage* 123 (2015), pp. 253–268.
- [319] A. M. Winkler, O. Renaud, S. M. Smith, and T. E. Nichols. "Permutation inference for canonical correlation analysis". In: *NeuroImage* 220 (2020), p. 10046.
- [320] D. M. Witten, R. Tibshirani, and T. Hastie. "A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis". In: *Biostatistics* (2009).

- [321] C. H. Xia, Z. Ma, R. Ciric, S. Gu, R. F. Betzel, A. N. Kaczkurkin, M. E. Calkins, P. A. Cook, A. García de la Garza, S. N. Vandekar, Z. Cui, T. M. Moore, D. R. Roalf, K. Ruparel, D. H. Wolf, C. Davatzikos, R. C. Gur, R. E. Gur, R. T. Shinohara, D. S. Bassett, and T. D. Satterthwaite. "Linked dimensions of psychopathology and connectivity in functional brain networks". In: *Nature Communications* (2018).
- [322] T. Xu, K. H. Nenning, E. Schwartz, S. J. Hong, J. T. Vogelstein, A. Goulas, D. A. Fair, C. E. Schroeder, D. S. Margulies, J. Smallwood, M. P. Milham, and G. Langs. "Cross-species functional alignment reveals evolutionary hierarchy within the connectome". In: *NeuroImage* (2020).
- [323] K. Yamada, K. Sakai, K. Akazawa, S. Yuen, and T. Nishimura. "MR tractography: A review of its clinical applications". In: *Magnetic Resonance in Medical Sciences* (2009).
- [324] A. Yamashita et al. "Harmonization of resting-state functional MRI data across multiple imaging sites via the separation of site differences into sampling bias and measurement bias". In: *PLoS Biology* (2019).
- [325] Q. Yang, F. Liu, W. Qin, C. Yu, and M. Liang. "Can we trust the results of canonical correlation analysis of brain imaging data a stability test". In: *Organization for Human Brain Mapping.* 2019.
- [326] T. Yarkoni, R. A. Poldrack, T. E. Nichols, D. C. Van Essen, and T. D. Wager. "Large-scale automated synthesis of human functional neuroimaging data". In: *Nature Methods* (2011).
- [327] C. H. Yeh, D. K. Jones, X. Liang, M. Descoteaux, and A. Connelly. "Mapping Structural Connectivity Using Diffusion MRI: Challenges and Opportunities". In: *Journal of Magnetic Resonance Imaging* (2020).
- [328] 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.
- [329] W. Zhang, A. Olivi, S. J. Hertig, P. van Zijl, and S. Mori. "Automated fiber tracking of human brain white matter using diffusion tensor imaging". In: *NeuroImage* 42.2 (2008), pp. 771–777.
- [330] J. Zhao, M. Thiebaut de Schotten, I. Altarelli, J. Dubois, and F. Ramus. "Altered hemispheric lateralization of white matter pathways in developmental dyslexia: Evidence from spherical deconvolution tractography". In: *Cortex* 76 (2016), pp. 51–62.
- [331] T. Zhao, M. Cao, H. Niu, X. N. Zuo, A. Evans, Y. He, Q. Dong, and N. Shu. "Age-related changes in the topological organization of the white matter structural connectome across the human lifespan". In: *Human Brain Mapping* (2015).
- [332] X. Zhuang, Z. Yang, and D. Cordes. "A technical review of canonical correlation analysis for neuroscience applications". In: *Human Brain Mapping May* (2020), pp. 1–27.

## Acronyms

- **AC** anterior commissure. 84
- ADC apparent diffusion coefficient. 28, 29
- **AF** arcuate fasciculus. 80, 88, 101, 102, 105, 112, 142
- **AR** acoustic radiation. 87, 123
- ATR anterior thalamic radiation. 86, 105, 109, 142
- **BOLD** blood oxygenation level dependent. 44
- CBD cingulum bundle: dorsal. 85, 105
- CBP cingulum bundle: peri-genual. 85
- **CBT** cingulum bundle: temporal. 85
- CCA canonical correlation analysis. i, 8, 151, 152, 153, 154, 155, 156, 157, 158, 159, 160, 162, 163, 164, 166, 167, 169, 170, 171, 172, 173, 174, 176, 177, 183, 184, 185, 186, 187, 188, 189, 190, 192, 193, 200, 205
- **CSF** cerebrospinal fluid. 24, 25, 33, 34
- **CST** corticospinal tract. 86, 109
- dMRI diffusion magnetic resonance imaging. 4, 23, 24, 25, 35, 36, 37, 38, 40, 45, 48, 52, 55, 58, 64, 88, 89, 160, 163, 165, 166, 167, 168
- **DTI** diffusion tensor imaging. 29, 31
- **DWI** diffusion weighted imaging. 29
- **EPI** echo planar imaging. 21, 22, 23, 38, 44, 55, 62, 63
- **FA** fractional anisotropy. 31, 32, 38, 39, 63, 65, 66, 67, 78, 89
- **FA** frontal aslant tract. 82, 105
- **FMA** forceps major. 83, 84, 110
- **FMI** forceps minor. 83, 84
- **fMRI** functional magnetic resonance imaging. 4, 7, 23, 43, 44, 45, 46, 47, 48, 55, 59, 60, 127, 158, 160, 163, 165, 166, 167, 168, 174, 175
- FX fornix. 85
- **GE** gradient echo. 20, 21, 22

- **GM** grey matter. 24, 25, 32, 34, 41, 42, 57, 67, 68, 75, 111, 126, 127, 135, 138, 140, 141, 142, 145, 175, 204
- HCP "WU-Minn" Human Connectome Project. ii, vi, xi, 6, 7, 9, 50, 51, 52, 53, 54, 55, 56, 58, 60, 61, 62, 63, 64, 65, 66, 67, 68, 69, 70, 72, 75, 78, 88, 89, 91, 92, 93, 95, 96, 98, 99, 100, 101, 102, 107, 108, 110, 111, 112, 114, 123, 128, 133, 138, 141, 151, 153, 154, 158, 159, 160, 161, 163, 164, 166, 168, 174, 176, 183, 184, 186, 190, 192, 193, 203
- ICA independent component analysis. 46, 47, 64
- **IDP** image-derived phenotype. 63
- **IFO** inferior fronto-occipital fasciculus. 81, 101, 102, 105, 110, 112, 123
- **ILF** inferior longitudinal fasciculus. 81
- KL divergence Kullback-Leibler divergence. 94, 95, 102, 103, 104, 105, 112, 127, 131, 132, 134, 137, 138, 140, 141, 142, 144, 145
- MCP middle cerebellar peduncle. 83, 109, 140
- MD mean diffusivity. 31, 32, 65
- MdLF middle longitudinal fasciculus. 81, 88, 101, 102, 109, 110, 112, 142
- MRI magnetic resonance imaging. i, 3, 4, 6, 7, 12, 13, 14, 15, 16, 18, 21, 23, 25, 32, 33, 38, 43, 47, 48, 50, 51, 55, 56, 57, 58, 59, 61, 62, 63, 70, 88, 91, 133, 144, 152, 158, 163, 166, 183, 184, 207
- NMF non-negative matrix factorisation. 125, 127, 128, 135, 138, 139, 140, 141, 142, 145, 147, 204
- **NODDI** neurite orientation dispersion and density imaging. 63, 65, 66, 68
- **OR** optic radiation. 87, 109
- **PCA** principle component analysis. 64, 156, 158, 165, 166, 167, 174, 175, 187
- **PGSE** pulsed gradient spin echo. 25, 27, 28
- **PLS** partial least squares. 8, 154, 158
- **RF** radio frequency. 12, 14, 16, 17, 18, 20, 21, 22, 25, 26, 27, 55, 61
- **ROI** region of interest. 46, 63, 73, 75, 79, 114, 144
- **RS-fMRI** resting-state fMRI. 45, 55, 57, 61
- **SE** spin echo. 21, 22

SLF superior longitudinal fasciculus. 79, 80, 101, 102, 105, 106, 112, 123, 142

- **SNR** signal-to-noise ratio. 14, 22, 29, 38
- **STR** superior thalamic radiation. 86, 105
- **UF** uncinate fasciculus. 82, 105, 142
- VOF vertical occipital fasciculus. 83
- WGB white-grey matter boundary. 67, 68, 91, 92, 94, 98, 101, 103, 129, 133, 139
- WM white matter. vi, 25, 29, 32, 33, 41, 48, 57, 67, 70, 72, 73, 74, 75, 76, 80, 82, 83, 84, 90, 91, 92, 94, 95, 96, 100, 101, 102, 104, 107, 108, 111, 113, 123, 125, 126, 127, 128, 129, 133, 135, 138, 141, 142, 145, 163, 164, 175, 203, 204, 205