Connectivity Driven Registration of Magnetic Resonance Images of the Human Brain

Aleksandar Petrović

Oxford Centre for Functional MRI of the Brain

Department of Clinical Neurology

University of Oxford

Balliol College

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas 2010

Abstract

Image registration methods underpin many analysis techniques in neuroimaging. They are essential in group studies when images of different individuals or different modalities need to be brought into a common reference frame. This thesis explores the potential of brain connectivity-driven alignment and develops surface registration techniques for magnetic resonance imaging (MRI), which is a noninvasive neuroimaging tool for probing function and structure of the human brain.

The first part of this work develops a novel surface registration framework, based on free mesh deformations, which aligns cortical and subcortical surfaces by matching structural connectivity patterns derived using probabilistic tractography (diffusion-weighted MRI). Structural, i.e. white matter, connectivity is a good predictor of functional specialisation and structural connectivity-driven registration can therefore be expected to enhance the alignment of functionally homologous areas across subjects.

The second part validates developed methods for cortical surfaces. Resting State Networks are used in an innovative way to delineate several functionally distinct regions, which were then used to quantify connectivity-driven registration performance by measuring the inter-subject overlap before and after registration. Consequently, the proposed method is assessed using an independent imaging modality and the results are compared to results from state-of-the-art cortical geometry-driven surface registration methods.

A connectivity-driven registration pipeline is also developed for, and applied to, the surfaces of subcortical structures such as the thalamus. It is carefully validated on a set of artificial test examples and compared to another novel surface registration paradigm based on spherical wavelets. The proposed registration pipeline is then used to explore the differences in the alignment of two groups of subjects, healthy controls and Alzheimer's disease patients, to a common template.

Finally, we propose how functional connectivity can be used instead of structural connectivity for driving registrations, as well as how the surface-based framework can be extended to a volumetric one. Apart from providing the benefits such as the improved functional alignment, we hope that the research conducted in this thesis will also represent the basis for the development of templates of structural and functional brain connectivity.

Acknowledgements

When I first came to Oxford I had little idea about what good research was and how much effort, patience and focus it required. In these respects, everything I have learned is from my two supervisors, Mark Jenkinson and Steve Smith (therefore, I blame them for any future bad papers I might write). This thesis and my research at FMRIB would not have been possible without their steady guidance. Both Steve and Mark put an extraordinary effort to keep me on the right track (and there were times when I strayed off the path in search of something, maybe beauty, to find an absolutely beautiful uselessness). Their insight into my research problems and their positive and optimistic attitude has meant a lot to me. The whole thesis was meticulously read by Mark and then Steve, and its clarity and focus were so much improved by their comments and suggestions. Aside from research, I thank them for fostering an exceptional collaborative spirit within the lab. The FSL courses, the graduate course, and retreats in particular will be fondly remembered.

In my second year I was sent off to MGH Martinos Center, Boston, to work with Bruce Fischl and Lilla Zöllei. I am most grateful for their hospitality and the crucial insights into the 'cortical aspect' of my thesis. Many of the ideas from Chapter 2 came during the brainstorming sessions at MIT and MGH. The initial future research (Chapter 5) was conducted jointly with Lilla Zöllei. They both remain valuable support and very dear colleagues.

In FMRIB, I was never alone and did not have to look far for an answer to any question. To thank but a few: Ivana Drobnjak knew what I was thinking before I said it; Brian Patenaude taught me how to squash meshes and use VTK libraries; Saad Jbabdi generously helped out with tractography. I thank Morgan Hough for teaching me Linux and scripting (and getting me addicted to Family Guy); Tim Behrens for his ideas; Ricarda Menke for the collaboration on the analysis of an Alzheimer's dataset; Valentina Tomassini for providing validation data (even with wrong initial left-right orientations); the IT Guys for scratch and the cluster. I also thank Thomas Yeo (MIT) for valuable discussions and advice about surface registration cost functions.

My College, Balliol, has been a continuous source of support and made my academic stay in Oxford a joy. Piers Nye, my college advisor, was always there to help, to let me teach/tutor, and to make sure things go smoothly. I would also like to thank my previous professors, especially Mića Stanković (whose doctoral thesis was the first thesis defence I attended) and Srdjan Stanković, for encouraging me to tackle challenging problems.

I am delighted to thank all the Loons and Massive Lads for the gin-coloured nights and friendships set in stone. To mention just a few: Michael Cole, Paul Engeham, Pádraig Belton, Jeremias Prassl, Steffen Deutschenbauer, Chris Noon, Tim Robbins, Robin Fellerman, Byron Spring, Sebastian Petzolt and all the others - without them, I would have drifted from the centre and my time in Oxford would have been infinitely less diverse.

I would further like to thank: the OUFC - for letting me discover a new passion in life and for exquisitely complementing my research. Late night trainings and sword-wielding, to which I always eagerly looked forward, brought clarity into my day-to-day existence. The OUWC - for wonderful walks all around the UK. CHFHB - for showing me how to hunt without a shotgun and for keeping my passion for nature alive within me. The OU - for expanding my interests to the worlds of politics and debate.

It is impossible not to mention all the ROH buddies, the Balliol Skiers and my fellow Holywell Manor Residents; the Old Guard: Valentina Gosetti, Anna Lewis, Jenn Reuer, Jarrod Hepburn, Marija Vlajić, Baker Graham and many other Balliolites for helping out with the bop costumes, for selling me concert tickets, baking pancakes, making sushi, organising it all, keeping me company in the Hall, tolerating me...

Special thanks to Dora Steel for being nice, to Aimé Lopez Aguilar for stuff, and to Alex Jeličić for long chats about village life. To the West Wing Pain Office, especially Kathy Fairhurst, for bringing light and music into the dark JR nights and for *endlessly* baking cakes and brewing tea.

I have seen, learned and experienced in these three years more than in much of my previous life. None of it would have been possible without substantial financial support from The Clarendon Fund. I shall feel indebted to them forever. Others contributed as well: Balliol College, ComTrade IT Solutions and Services, Serbian Academy of Sciences and Arts, the Guarantors of Brain and the Organisation for Human Brain Mapping.

Finally, I would like to thank my parents, my brother and my family for their love and unwavering support.

Contents

1 Introduction **17**

1.1 The Human Brain 18

1.1.1 Basic Neuroanatomy 18

1.1.2 Brain Connectivity 22

1.2 Magnetic Resonance Imaging 23

1.2.1 T1-weighted MR Imaging 25

1.2.2 Functional MRI 26

1.2.3 Diffusion-weighted Imaging 27

1.3 Brain Image Registration 28

1.3.1 Concept and Aim of Surface Registration 30

1.3.2 Limitations 31

1.3.3 Image Registration Using Structural Connectivity 32

1.4 FMRIB Software Package and FreeSurfer 33

1.5 Main Thesis Contributions 35

1.6 Thesis Overview 36

2 Methods for Connectivity Driven Registration of Cortical Surfaces 39

2.1 Introduction 40

2.1.1 FreeSurfer Cortical Geometry Driven Registration 41

2.1.2 Other Cortical Alignment Paradigms 43

2.1.2.1 Anatomical Landmarks 43

2.1.2.2 Functional MRI-driven Alignment 44

2.1.2.3 Application-driven Cortical Alignment 45

2.1.2.4 Relation to Registration Using Connectivity Information 45

2.2 White Matter Tractography 46

2.2.1 Deterministic Streamline Tractography 46

2.2.2 Probabilistic Streamline Tractography 47

2.3 Probabilistic Tractography (FSL/probtrackX) . . . 48
2.4 White Matter Connectivity Driven Alignment . . . 51
2.4.1 Motivation . . . 51
2.4.1.1 Functional Specialisation vs. Cortical Morphology . . 52
2.4.1.2 Cortical Morphology and White Matter Connectivity 52
2.4.1.3 Connectivity-based Parcellation . . . 53
2.4.1.4 Practical Benefits . . . 55
2.4.1.5 Limitations . . . 56
2.4.2 Notation and Terminology . . . 58
2.4.3 Registration Pipeline . . . 58
2.4.3.1 Formation of SCFVs . . . 59
2.4.3.2 Dimensionality Reduction . . . 67
2.4.4 Registration Algorithm . . . 68
2.4.4.1 Deformable Model Framework . . . 69
2.4.4.2 The Cost Function . . . 71
2.4.4.3 Regularisation vs. Intensity Force . . . 74
2.4.4.4 Multi-scale Framework . . . 76
2.4.4.5 Optimisation . . . 77
2.4.4.6 The Deformation Field . . . 82
2.5 Discussion . . . 84
2.5.1 Formation of SCFVs . . . 84
2.5.1.1 High- vs. Low-dimensional SCFVs . . . 84
2.5.1.2 Current Choices . . . 85
2.5.1.3 Other Methods for SCFV construction . . . 87
2.5.2 Registration Algorithm . . . 88
2.5.2.1 Other Registration Methods . . . 89

3 Connectivity Driven Registration of Cortical Surfaces - Validation and Applications 92
3.1 Validation Using Resting-state fMRI . . . 93
3.1.1 Cortical Parcellation Using Resting-state fMRI . . . 94
3.2 SMA/pre-SMA Boundary Alignment . . . 95
3.2.1 Data Acquisition . . . 95
3.2.2 SMA/pre-SMA Parcellation Using RS FC . . . 96
3.2.2.1 Calculating Functional Connectivity Profiles for the MFC . . . 97

3.2.3 SMA/pre-SMA Parcellation Using RSNs 101
3.2.3.1 RSNs and ICA 101
3.2.3.2 ICA-based RSN Clustering 103
3.2.3.3 Comparison to FC-based Parcellation 106
3.2.4 Boundary Alignment Before and After FACS 108
3.2.4.1 Comparing FACS and FreeSurfer 108
3.2.4.2 RSN Connectivity-driven Alignment 111
3.2.4.3 Results 112
3.3 Alignment of Boundaries Between Functional Regions 114
3.3.1 Adapted Edge Detection Algorithm 116
3.3.2 Whole-brain Analysis 120
3.3.3 Analysis of the MFC Area 125
3.3.4 Analysis of the Cingulate Area 126
3.3.4.1 Cingulate Area Clustering Using RSNs 132
3.3.4.2 Comparison to MFC Clustering Using RSNs 141
3.4 Applications 141
3.4.1 Construction of The SCFV Template 142
3.4.2 Analysis of an Alzheimer's Disease Data Set 144
3.4.2.1 Data Acquisition 145
3.4.2.2 AD vs. Control Groups 146
3.5 Discussion 149

4 Connectivity Driven Registration of the Surfaces of Subcortical Structures 154
4.1 FIRST and the Problem of Vertex-to-vertex Correspondence 155
4.1.1 Segmenting Subcortical Structures Using FIRST 156
4.1.2 Motivating Connectivity-driven Surface Registration 158
4.2 Registration Algorithms 159
4.2.1 Volumetric Structural Connectivity Descriptors 160
4.2.2 Free Deformation-based Registration Method 163
4.2.2.1 Artificial Examples and Parameter Identification . . 164
4.2.2.2 Variations of the Proposed Registration Framework . 170
4.2.3 Registration Based on Spherical Wavelets 172
4.2.3.1 The ABCA Test 175
4.2.3.2 Comparison on an AD Data Set 176
4.3 Subcortical Registration Using SCFVs From the Cortical Framework 179

4.3.1 SCFV Subcortical Templates . . . 180
4.3.1.1 SCFV Template Clustering . . . 180
4.3.1.2 Registration Assessment Using Thalamic Clusters . . 182
4.3.2 AD vs. Controls - FASS Results . . . 184
4.3.3 HDPM Clustering of Thalamic SCFVs . . . 193
4.3.3.1 Using Inter-subject Thalamic Cluster Correspondence to Drive Registration . . . 194
4.3.3.2 AD vs. Controls - Analysis of Cluster Displacements 197
4.4 Discussion . . . 203

5 Conclusions and Future Work 207
5.1 Conclusions . . . 207
5.1.1 Summary of Major Contributions . . . 207
5.1.2 Registration Methods . . . 209
5.1.3 Cortical Framework . . . 212
5.1.4 Subcortical Framework . . . 216
5.2 Future Work . . . 218
5.2.1 Assessment of Potential of Volumetric Connectivity-driven Registration . . . 221
5.2.1.1 Methods . . . 222
5.2.1.2 Data Description . . . 224
5.2.1.3 Results . . . 225
5.2.1.4 Discussion . . . 227

A Linear Interpolation 229

B Regularisation Cost Function 230

C Locally Affine Transforms 232

D Imposing Labelling Consistency 234

List of Figures

1.1 The structure of the neuron 19
1.2 Corpus callosum fibre bundle 20
1.3 Main brain lobes 21
1.4 Grey/white matter surface 22
1.5 T1-weighted MR brain image - axial slices 25

2.1 Segmentation, surface reconstruction and inflation steps 43
2.2 Probabilistic tractography - axon geometry and streamlining procedures 50
2.3 Tractography to a set of targets 55
2.4 Registration pipeline - preprocessing 59
2.5 Assignment of a SCFV to a mesh vertex 60
2.6 Cortical alignment using folding patterns geometry 61
2.7 Formation of cortical targets 63
2.8 Construction of cortical targets for probabilistic tractography - challenges 66
2.9 Part of a Similarity Matrix (SM) 72
2.10 Registration pipeline - the registration algorithm 73
2.11 Intensity Force - demonstration 75
2.12 Artificial test example 80
2.13 Cost function convergence properties 83
2.14 Cingulate clustering using SCFVs 86

3.1 The MFC mask 98
3.2 MFC anatomic parcellation 100
3.3 FC-based MFC clustering 101
3.4 Different MFC clustering paradigms 102
3.5 Examples of RSNs used for SMA/pre-SMA parcellation 104
3.6 Construction of RSN Feature Vectors 107
3.7 Comparison of RS FC- and ICA-based RSN clustering 1 109
3.8 SCFV-based clustering 110

3.9 Comparison of RS FC- and ICA-based RSN clustering 2 111
3.10 FACS vs. FreeSurfer - SMA/pre-SMA alignment 113
3.11 Approximating the gradient on a mesh 116
3.12 Canny edge detection on the $ic5$ - artificial example 117
3.13 Canny edge detection on the $ic5$ - real data 1 119
3.14 Canny edge detection on the $ic5$ - real data 2 120
3.15 Mean edge probability maps ($ic5$ - real data) 122
3.16 Comparison of different registration methods 1 124
3.17 Comparison of different registration methods 2 125
3.18 Comparison of different registration methods 3 126
3.19 Comparison of different registration methods 4 127
3.20 Comparison of different registration methods - MFC 1 128
3.21 Comparison of different registration methods - MFC 2 129
3.22 Cingulate cortex mask 130
3.23 Cingulate cortex - mean gradient image 131
3.24 Comparison of different registration methods - cingulate cortex 1 . . . 132
3.25 Comparison of different registration methods - cingulate cortex 2 . . . 133
3.26 Comparison of different registration methods - cingulate cortex 3 . . . 134
3.27 ICA-based RSN cingulate clustering 135
3.28 Within-subject clustering consistency 136
3.29 Inter-subject clustering consistency 137
3.30 Registration performance measures - cingulate clustering 1 138
3.31 Registration performance measures - cingulate clustering 2 139
3.32 Registration performance measures - cingulate clustering 3 140
3.33 Registration performance measures - MFC clustering 142
3.34 ICA decomposition of the template CM 1 144
3.35 ICA decomposition of the template CM 2 145
3.36 AD vs. controls - differences in the deformation fields 147
3.37 AD vs. controls - differences in 'connectivity' 148

4.1 Subcortical segmentation 156
4.2 Connectivity-based thalamic clustering 159
4.3 Construction of VSCDs 161
4.4 Artificial examples - ABCA test data 165
4.5 Artificial examples - the construction 165
4.6 Artificial examples - ABCA residual displacements 1 167

4.7 Artificial examples - ABCA residual displacements 2 168
4.8 Artificial examples - the improvement 169
4.9 Spherical wavelets - the basis field 174
4.10 Spherical wavelets - ABCA residual displacements 176
4.11 Spherical wavelets vs. free deformation - AD data set 177
4.12 Free deformation using VSCDs - AD data set 178
4.13 Clustering template CMs of subcortical structures 181
4.14 Inter-subject alignment of thalamic clusters 184
4.15 AD vs. controls - the thalamus 1 185
4.16 AD vs. controls - the thalamus 2 186
4.17 AD vs. controls - the thalamus 3 187
4.18 AD vs. controls - left putamen, pallidum and amygdala 1 188
4.19 AD vs. controls - left putamen, pallidum and amygdala 2 189
4.20 AD vs. controls - right putamen, pallidum and amygdala 1 190
4.21 AD vs. controls - right putamen, pallidum and amygdala 2 191
4.22 AD vs. controls - left putamen - individual results 192
4.23 Driving FASS using cluster correspondences 195
4.24 HDPM thalamic templates 196
4.25 Inter-subject alignment of thalamic clusters 197
4.26 AD and controls - thalamic clusters 198
4.27 AD vs. controls - thalamic clusters, left 200
4.28 AD vs. controls - thalamic clusters, right 201

5.1 Compound Jaccard coefficient 1 226
5.2 Compound Jaccard coefficient 2 227

C.1 Locally affine projection task 233

List of Tables

1.1 FSL Tools 34

2.1 FreeSurfer segmentation of subcortical structures 64
2.2 Multi-scale framework optimisation stages 77

5.1 CVS vs. FNIRT 225
5.2 CVS vs. FNIRT 226

List of Abbreviations

$ic3$	Regular icosahedral mesh with 642 vertices
$ic5$	Regular icosahedral mesh with 10242 vertices
J_d	Regularisation cost function
J_p	Data/intensity cost function
pdf	Probability Density Function
3D	Three Dimensional
AD	Alzheimer's Disease
BA	Brodmann Area
BOLD	Blood Oxygen Level Dependent
CD or CP	Connectivity Descriptor or Connectivity Profile
CM	Connectivity Matrix
CNS	Central Nervous System
CVS	Combined Surface and Volumetric Registration
DMN	Default Mode Network

DR	Dual Regression
DW	Diffusion Weighted
EEG	Electroencephalography
EPM	Edge Probability Map
FA	Fractional Anisotropy
FACS	FMRIB's Alignment of Cortical Surfaces
FASS	FMRIB's Alignment of Subcortical Surfaces
FC	Functional Connectivity
FCM	Functional Connectivity Map
FDT	FMRIB's Diffusion Toolbox
FIRST	FMRIB's Integrated Registration and Segmentation Tool
FLIRT	FMRIB's Linear Image Registration Tool
fMRI	Functional Magnetic Resonance Imaging
FMRIB	Oxford University Centre for Functional MRI of the Brain
FNIRT	FMRIB's Noninear Image Registration Tool
FSL	FMRIB Software Library
FWHM	Full Width Half Maximum
HDPM	Hierarchical Dirichlet Process Mixture Models
HRF	Hæmodynamic Response Function
ICA	Independent Component Analysis

IF	Intensity Force
MD	Mean Diffusivity
MELODIC	Model-free FMRI analysis using Probabilistic Independent Component Analysis
MFC	Medial Frontal Cortex
MGI	Mean Gradient Image
MNI	Montreal Neurological Institute
MRI	Magnetic Resonance Imaging
MT	Middle Temporal area
NMR	Nuclear Magnetic Resonance
PCA	Principle Component Analysis
pre-SMA	Pre-supplementary Motor Area
probtrackX	Probabilistic Fiber Tracking Tool
RF	Radio Frequency
RF	Regularisation Force
ROC	Receiver Operating Characteristic
ROI	Region of Interest
RS	Resting State
RS FC	Resting State Functional Connectivity
RSN	Resting State Network

RSNF	Resting State Networks Feature
RSNFM	Resting State Networks Feature Matrix
RSNFV	Resting State Networks Feature Vector
SCFV	Structural Connectivity Feature Vector
SM	Similarity Matrix
SMA	Supplementary Motor Area
VBM	Volumetric Based Morphometry
VSCD	Volumetric Structural Connectivity Descriptor

Chapter 1

Introduction

Medical imaging is an important diagnostic and research tool, able to generate visual representations of living tissue. Although many assessments are performed on a single subject basis, such as pre-surgical planning, medical images can be compared across subjects and modalities to understand population-specific variations. In the case of the human brain, the aims of the inter-subject image comparison are typically to investigate variations in the shape, size and function of corresponding brain structures. The delineation of these structures in different subjects would historically require an expert operator to manually label them, which is a time consuming process prone to different human-induced biases. This provides the main motivation for the development of automated registration (as well as segmentation) techniques that are able to align images so that their comparison can be done in one common space, i.e. a common coordinate system.

Magnetic Resonance Imaging (MRI) is a noninvasive tool capable of providing a detailed insight into the structure and function of a biological sample, such as the human brain (section 1.2). MR images provide a volumetric (three-dimensional) reconstruction of the brain, and, depending on the acquisition settings, can depict as diverse qualities as tissue chemical composition, anatomy or the functional characteristics. Registration protocols are typically capable of aligning MR images both across modalities as well as subjects.

The alignment of brain images, in particular, crucially depends on the meaning of 'correspondence'. Most commonly, this refers to the correspondence in structure, i.e. the thalamus of one brain is aligned to the corresponding thalamus of the other, the frontal lobe of one brain is aligned to that of the other, etc. Functional correspondence means that functionally equivalent areas are aligned across subjects, for example, the language area in one subject is aligned to the language area in the other. So far, the anatomical alignment has been traditionally used as a surrogate for functional alignment as well, since (limited) correspondence exists between some anatomical features, such as brain folding patterns, and the localisation of different functional areas. This thesis proposes how to use *brain connectivity* (explained in the next section), which is closely related to brain function, to drive image registration and improve functional alignment.

The first part of the introductory chapter explains basic brain anatomy relevant to this thesis and introduces elementary principles of image formation using Magnetic Resonance Imaging (MRI). The aim of connectivity-driven alignment and this thesis in general is presented in the second part.

1.1 The Human Brain

The focus of this section is on providing a background in neuroanatomy that underlies the images that we are attempting to register. We will briefly review the basic structure of the Central Nervous System (CNS) and then follow up with an introduction to the concept of brain connectivity which underlies connectivity-driven registration.

1.1.1 Basic Neuroanatomy

Neuronal cells (neurons) are the main information processing units of the CNS and they use chemical and electrical signals to pass information from one neuron to

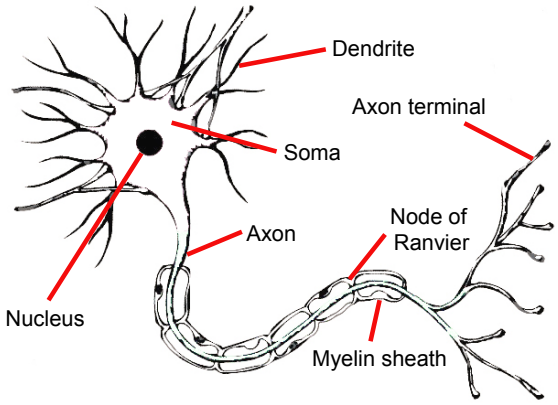

Figure 1.1: An illustration of a neuron.

the other. Neuronal and neuroglial cells make up the major cellular component of the nervous system. Unlike neurons, neuroglia do not directly participate in information processing, but can support neurons by providing myelin and other protective functions.

Figure 1.1 illustrates the main parts of a neuron. Each neural cell has a cell body (soma) with the organelles playing an important role in cell metabolism. The information in the form of a chemical signal is received by the dendrites before being processed in the cell body and transmitted electrically along the axon to the axon terminals. The terminals are connected to the dendrites (via synapses) of other neurons and the information/signal from an axon to a dendrite is (chemically) propagated onto another cell. Some axons, which can greatly differ in length, are surrounded by a myelin sheath which improves the electrical conductance of the signal through the axon (Bear et al., 2007). The axons are often referred to as *nerve fibres* or just fibres. Nerve cells tend to be grouped into areas dominated by cell bodies and dendrites, so called *grey matter*, and the areas dominated by nerve fibres, or the *white matter* (lipids within the myelin that covers many of the axons give white matter a paler appearance compared to the grey matter). Nerve fibres interconnect different parts of the brain and are essential to quick and efficient signal transmission. Many fibres following similar paths and connecting similar (functional) areas are organised

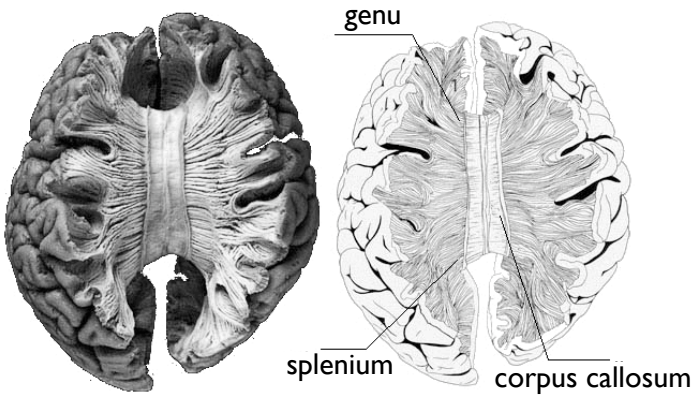

Figure 1.2: **(left)** A resection of a post mortem human brain (dorsal view) reveals the white matter fibre bundles of the corpus callosum, which connects the two hemispheres. **(right)** A stylised image from the left. The splenium is the anterior and the genu is the posterior aspect of the corpus callosum. Image reproduced by courtesy of Dr S. Jbabdi.

into *fibre bundles* or *tracts*. The corpus callosum is an example of a large, flat bundle of neural fibres connecting two hemispheres (Figure 1.2). Tracts are also known as *structural connections* and a term describing the pattern of structural connections and synapses in the brain is *structural connectivity* (the description of how the brain is structurally 'wired up').

At the organ level, the central nervous system consists of the brain and the spinal cord. The human brain can be divided into three main parts: the forebrain (consisting of the cerebral hemispheres and the diencephalon), the brain stem and the cerebellum. Also contained within the brain is a system of ventricles filled with cerebrospinal fluid. Of special interest here is the forebrain, which is dominated by two large hemispheres. Each hemisphere consists of the cortex, the inner core of white matter and the basal ganglia. The cortex is a highly convoluted layer of grey matter and forms the outer brain surface. It is primarily responsible for high-level functions such as learning and cognition, as well as much sensory processing and motor (muscle) control. The convolutions in the cortical sheath are called gyri and the folds between them sulci. We often refer to these as *folding patterns*. Different parts of a cortex (separated by large folds) perform specific functions and this is often referred to as *functional*

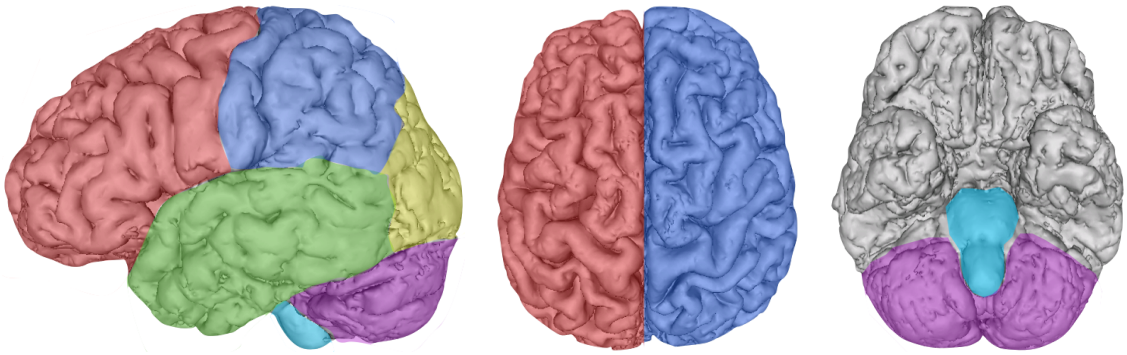

Figure 1.3: An illustration of the main divisions of the brain. **(left)** Red - frontal lobe, blue - parietal lobe, green - temporal lobe, yellow - occipital lobe, purple - cerebellum, turquoise blue - brain stem. The lateral side of the brain is shown. The left of the image corresponds to the anterior of the brain. **(middle)** The two brain hemispheres are coloured differently. The dorsal view is shown. **(right)** Purple - cerebellum, turquoise blue - brain stem. The ventral view is shown.

specialisation. Functional specialisation is, to some degree, related to the folding patterns. The four main lobes of a cortical hemisphere: frontal, parietal, temporal and occipital have different functional characteristics (Figure 1.3). For example, the frontal lobe contains motor areas, responsible for planning, control, and execution of movements. The parietal lobe processes somatosensory inputs whereas temporal and occipital lobes contain centres for processing auditory and visual information respectively. Functional specialisation also exists at smaller spatial scales (Bear et al., 2007). Functional alignment, which is one of the goals of connectivity-driven registration, aims to match functionally homologous areas across subjects. Beneath the cortex is the white matter and deep within white matter are subcortical structures. The hypothetical surface separating cortical grey matter from the white matter we shall call the grey/white matter surface. In Figure 1.4 (left), the grey/white matter surface of one hemisphere is delineated with a yellow line and the red line follows the outer border of the cortex (pial surface). The area between the two lines is the cortical grey matter. The grey/white matter surface can be reconstructed from e.g. T1-weighted MR images and an example of such a reconstruction is shown in Figure 1.4 (right).

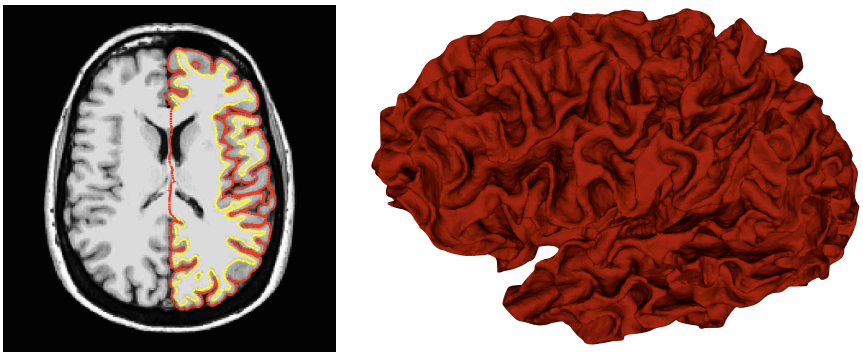

Figure 1.4: **(left)** An axial slice of a T1-weighted MR image with the delineated grey/white matter surface (yellow) and the outer cortical surface (red). **(right)** The grey/white matter surface has been reconstructed for the whole left hemisphere exemplifying its convoluted nature. Lateral view is shown. Anterior of the brain corresponds to the left of the image.

The terminology 'subcortical structures', as used in this document, refers to basal ganglia and the thalamus (which is part of the diencephalon), or in general, the groupings of grey matter within the cortical hemispheres that are not included as part of the cortex. Subcortical structures, unlike the cortex, have smooth surfaces and perform roles like relaying of cortical signals (the thalamus) or control of muscle tone, posture and movement (basal ganglia).

1.1.2 Brain Connectivity

In this work, Brain Connectivity refers to either a pattern of anatomical connections (structural connectivity, i.e. white matter fibres) or to the functional connectivity (explained in the next paragraph) between distinct elements of the brain (Sporns et al., 2005). Depending on the scale at which connectivity is assessed, these elements can refer to individual neurons, neuronal populations or anatomically defined brain regions. The resolution of brain MR images constrains the assessment of connectivity to brain regions that are at least several cubic millimetres in size. Structural connectivity is, in this thesis, assessed by measuring the spatial distribution of white matter fibre tracts using diffusion-weighted MR images (see section 1.2.3). An

illustrative atlas of white matter fibre bundles reconstructed using diffusion-weighted images was assembled by Wakana et al. (2004).

Functional connectivity represents statistical or causal relationships in neural activity between different brain regions. These relationships can be measured at varying temporal timescales using different information flow measures or can be modelled by supposing a correlative or a causal structure in the signal passing and processing mechanism. In this thesis, functional connectivity is assessed by processing signals correlated with blood oxygenation level (functional MRI) that are also related to neuronal activation (section 1.2.2). Both structural and functional connectivity patterns are related to functional specialisation (as an example, see Johansen-Berg et al. (2004); Kim et al. (2010)) and this relationship is extensively utilised in this work.

1.2 Magnetic Resonance Imaging

Magnetic Resonance Imaging (MRI) is a tomographic[1] imaging technique used to generate images of the internal properties of an object or a living tissue by externally measuring their interaction with magnetic fields, in particular, the nuclear magnetic resonance (NMR) signals. MRI has several distinct features that differentiate it from the other tomographic imaging techniques (Liang and Lauterbur, 1999):

- The source of the NMR signal comes directly from the imaged material (typically hydrogen nuclei in water and fat for biological samples). Therefore, no manipulation such as the injection of radioactive isotopes is needed to emit signals.

- MR operates in the radio-frequency (RF) range, which means no ionising radiation is used and therefore it can be considered a non-invasive imaging

[1]Tomography is concerned with creating images of the internal (anatomical or functional) organisation of an object without physically cutting it open (Liang and Lauterbur, 1999).

method.

- MRI can, without any mechanical adjustments to the machine, generate spatial distributions of the measured physical quantity in any orientation, i.e. any two-dimensional (2D) sectional 'cuts', three-dimensional (3D) volumetric images, etc.

- MR images are extremely rich in their information content and, depending on the image acquisition parameters, can represent many physical and chemical properties of the material such as the molecular motion (e.g. diffusion), tissue oxygenation, tissue density, etc.

From the systems perspective, the essential parts of the MRI machine are the main magnet, the gradient system and the RF coils. The main magnet provides a strong, static in time, external magnetic field (typically from 1.5 T to 7 T) that prepares the sample for imaging by changing the properties of the imaged object, i.e. the (nuclear) spin system. The imaging gradients produce magnetic fields of controlled spatial nonuniformity and are essential for spatial localisation (and therefore image formation) of the NMR signal. The RF system consists of the transmit and receive coils: the transmitter coil is used to excite the imaged object, i.e. the (nuclear) spin system, whereas the receiver coil detects the signal.

Depending on how the MRI system is configured, the obtained images can reveal specific physical/chemical properties of the imaged object. In this thesis, of particular interest is an image modality called Diffusion-weighted MRI (or DW MRI) which quantifies the diffusion of water molecules e.g. in the living brain tissue. Other modalities that we extensively use are T1-weighted images as well as functional MRI.

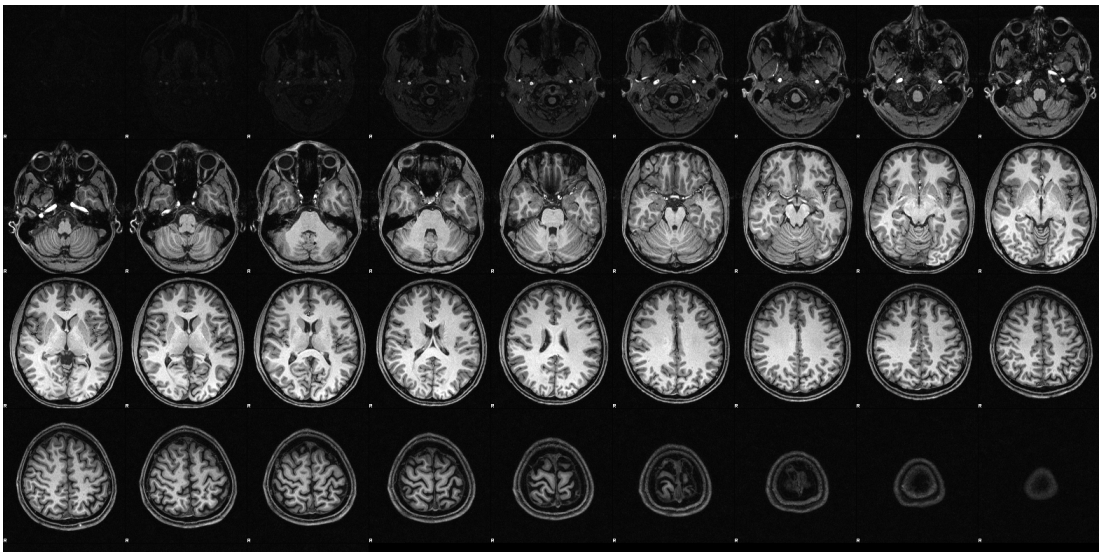

Figure 1.5: T1-weighted axial images (orthogonal to the long axis of the body) of the human brain. Two-dimensional images (slices) are stacked together to form a three-dimensional volumetric representation. The upper left corner slice corresponds to the bottom of the head, whereas the lower right slice represents the top of the head. The top of each slice corresponds to the anterior of the brain.

1.2.1 T1-weighted MR Imaging

MRI measures the magnetic properties of the (nuclear) spin systems which are, in most cases, protons in hydrogen nuclei within tissue. An external magnetic field produces a bulk alignment of magnetic moments (which are vector quantities related to nuclear spin) along the direction of the main static magnetic field. Net magnetic moment of the spin system precesses at the Larmor frequency which is proportional to the strength of the magnetic field. When an RF pulse is applied at Larmor frequency perpendicular to the static magnetic field, the net magnetisation tilts away from its default position (this is called the excitation). Due to the precession of the excited spin system, energy is emitted in the form of RF signals that are detected by the receiver coils. As the excited system relaxes, spins realign to the direction of the main magnetic field and the signal decays. The measured signal depends on the proton density, two relaxation constants, T1 and T2, and the relative timing of external RF pulses and field gradients. T1 is a measure of the time required for the net magnetisation to

return to equilibrium. T1 time (as well as T2 and proton density) is a tissue-specific property and the MRI pulse sequences can be designed so that the resulting image contrast exemplifies the differences in T1 time of different tissues such as the grey and white matter. An example of a T1-weighted image is shown in Figure 1.5. A three-dimensional image is commonly formed by stacking together a number of two-dimensional images (often called *slices*) of consecutive cross-sections of the brain. Figure 1.5 depicts a number of axial slices which, when put together, can be used to reconstruct a three-dimensional image. In this thesis, T1-weighted images are used for image segmentation and initial cortical geometry-based alignment.

1.2.2 Functional MRI

Functional Magnetic Resonance Imaging (fMRI) is an MRI modality used for mapping brain activity, by measuring haemodynamic changes that are indirectly related to neuronal activation (Ogawa et al., 1993; Jezzard et al., 2001). In this modality, magnetic properties of the blood related to oxygenation (and implicitly neuronal activity) are imaged using the Blood Oxygen Level Dependent (BOLD) contrast. BOLD signal arises from the changes in 'magnetic susceptibility' of the blood which depends on how the external magnetic field interacts with the material. The oxyhaemoglobin (HbO_2) and deoxyhaemoglobin (dHb) have different magnetic properties (HbO_2 influences magnetic flux in the same way as the other tissue, whereas dHb increases it - dHb therefore acts as an endogenous contrast agent) and blood oxygenation consequently leads to local distortions of the magnetic fields inside and around the blood vessels. Neuronal activity increases oxygen consumption (increase in dHb), but at the same time leads to a large increase in blood flow (up to 70% above baseline) and therefore in oxygenated blood. The increase in dHb is overcompensated by a large increase in HbO_2 and a couple of seconds after a burst in neuronal activity,

the relative concentration of dHb in the venules[2] drops. As dHb increases, the tissue T2* relaxation values drop, and since neuronal activation leads to less dHb, the T2*-weighted signal decays more slowly (compared to the surrounding tissue) which leads to a brighter T2*-weighted image. T2* is a time constant similar to T2, but which also includes extra dephasing effects (e.g., loss of phase coherence due to static magnetic field inhomogeneities and susceptibility effects). In this way, the neuronal activity is indirectly (i.e. through the blood-related or haemodynamic response) related to the MR image intensity. In this thesis we use BOLD contrast to assess resting state brain activity from which we infer functional brain connectivity.

1.2.3 Diffusion-weighted Imaging

Diffusion is a process of random motion/displacement of, e.g. water molecules (Mori, 2006). In the brain, diffusion is restricted by tissue boundaries, cell membranes, etc. Therefore, for example, water diffuses more freely *along* white matter fibres than across them. Consequently, diffusion in white matter fibre bundles has a preferred direction (i.e. is anisotropic). MR can be sensitised to measure the signal which is proportional to the amount of molecular motion in a specific spatial direction. MR images obtained in this way are called Diffusion-weighted (DW). The mechanism of diffusion-weighting crucially depends on the application of the *dephasing* and *rephasing* gradients. The dephasing gradient first systematically changes phases of signals coming from water molecules. The rephasing gradient, which is applied a short time (e.g. 30 ms) after the dephasing one, reverses the effect of the dephasing gradient and restores the signal[3]. In the case when no diffusion happens in the time between these two gradients, the net MR signal amplitude (which is the sum of signals coming from all water molecules) measured after the rephasing gradient is higher than in case

[2]A venule is a small blood vessel that allows deoxygenated blood to return from the capillary beds to the larger blood vessels.

[3]Typically, spin-echo sequences are used and the signal decays predominantly by T2 relaxation.

when water molecules diffuse. This happens because the diffusion changes the phase of the (precessing) MR magnetisation of individual water molecules, so after rephasing, phases fail to realign, which manifests as a relative signal loss. This signal loss can be associated with different spatial directions depending on which combination of gradients is applied. The drop in signal is only linked to motion (of water molecules) *in the direction of the gradient*. In other words, the DW signal can quantify diffusion in multiple spatial directions (all obtained as different images) allowing for the inference about e.g. the principal diffusion direction[4]. One can then use diffusion directionality information, such as the principal diffusion direction, to trace white matter fibre bundles in a process called tractography (Chapter 2, section 2.3). Throughout this thesis, tractography is used to measure white matter connectivity.

1.3 Brain Image Registration

In brain imaging, registration can be performed to align a group of subjects to a (common) template, to align images of a single subject taken at different timepoints, or to align different modalities imaging the same anatomy. In each of these cases, the registration algorithm deforms (i.e. warps) the *input* image to match the shape and appearance of the *reference*[5]. When aligned in this way, one can for example, compare activations in corresponding brain regions and perform group-based analysis in one common space.

A registration algorithm typically has three major parts (Hill et al., 2001; Zitova and Flusser, 2003):

- the measure quantifying the similarity between images;
- the warping model which specifies how the input image can be deformed to match the reference;

[4]This is the direction along which the greatest diffusion occurs.

[5]The input image is often called the 'source' whereas the reference image is often called the 'target'.

- the optimisation scheme that specifies how parameters of the warping model are estimated to maximise the input to reference image similarity measure.

Measures quantifying image similarity are numerous, but among the most commonly used ones are the correlation-based measures (Pratt, 2007), which measure similarity as the normalised cross-correlation between images, or the simple sum of squared differences measure. Comparison methods based on information theoretic similarity measures have been successful as well, in particular the mutual information similarity measure (Wells et al., 1996). However, whichever measure is used, ideally it should have a prominent global maximum for the 'perfect' image overlap.

The warping model essentially encodes the deformation and constrains the space of possible solutions. With respect to the warping models, two main classes are affine or non-affine[6]. Affine deformations have up to 12 degrees of freedom (free parameters) in 3D and allow for image translation, rotation, scaling and shears. Higher degrees of freedom are associated with non-affine registration methods. In general, non-affine registrations are more flexible in representing the deformation field compared to the affine ones, allowing highly dissimilar brains to be matched, such as the ones with normal and atrophic ventricles.

Finally, the optimisation algorithm finds the set of parameters encoding the deformation field such that the image similarity measure is maximised. Depending on the exact deformation model, parameters can be identified through numerical optimisation, e.g. using a Gauss-Newton minimisation scheme or Levenberg-Marquardt optimisation method. Whichever method is used, it should ideally converge to the solution quickly and be stable with respect to the initial conditions.

In brain imaging, affine registration with 6 degrees of freedom (a so called 'rigid' transform) is commonly used for within-subject alignment (e.g. using mutual information when aligning different modalities) whereas non-affine registration

[6]Note that some refer to these as linear and non-linear respectively.

methods are increasingly used for inter-subject (or subject to template) registration (e.g. see the differences between two state of the art affine and non-affine registration methods, FLIRT/FSL (Jenkinson et al., 2002) and FNIRT/FSL (Andersson et al., 2007a,b)).

1.3.1 Concept and Aim of Surface Registration

Brain image registration methods can be roughly divided into volumetric and surface-based. Volumetric methods deform the 3D representation of the brain whereas surface-based methods align surfaces such as the cortical grey/white matter surface. Volumetric representation of the brain consists of elements called voxels (an equivalent of a pixel in 3D) whereas surfaces are often represented with meshes defined by a set of points and edges between them. Although we discuss how methods developed in this thesis can be extended to the volumetric framework (Chapter 5), we are primarily concerned with the registration/matching of surfaces.

In brain imaging in particular, matching of cortical surfaces (or other surfaces) is usually performed in three stages: surface reconstruction, surface inflation and surface matching (Thompson and Toga, 1996; Van Essen et al., 1998; Dale et al., 1999; Fischl et al., 1999a; Yeo et al., 2010a) (for a review see (Audette et al., 2000)). In this way surfaces are matched in one common coordinate system, typically a parametrised sphere. Therefore, surface registration in this case reduces to shuffling the points of the mesh along the sphere. In other words, the deformation found by the registration algorithm reorders points on a pre-defined 2D manifold such as a sphere (details of the surface matching are explained using an example from the FreeSurfer software package in Chapter 2, section 2.1.1). Throughout this thesis we make extensive use of this methodology.

The aim of (sub)cortical surface registration is to bring functionally or structurally homologous areas into correspondence. Surfaces are used instead of 3D volumes for a

variety of purposes, but primarily because of easy visualisation, increased sensitivity in the inter-subject alignment of functional areas and for better accuracy in smoothing the fMRI results across the surface (Fischl et al., 1999a). To achieve registration, surface matching algorithms can use features such as cortical curvature to drive the alignment (Chapter 2, section 2.1.2). The optimisation is usually performed by either using a variant of a point matching algorithm that iteratively matches points (i.e. mesh vertices) which are similar in the input and the reference subjects (Besl and McKay, 1992; Zhang, 1994), or the meshes are iteratively deformed by a set of 'forces' that are derived through the maximisation of the image similarity function (Fischl et al., 1999a). The latter optimisation method is preferred when no anatomical landmarks (Chapter 2, section 2.1.2.1) are used and is exploited in this thesis.

1.3.2 Limitations

Matching of surfaces is typically driven by features derived from cortical geometry such as the surface curvature, average convexity, gyrification patterns, etc. Other algorithms use functional data to enhance the cortical geometry-driven surface registration (for a review of popular cortical surface registration methods see Chapter 2, section 2.1). In both cases the aim is to achieve anatomical or functional correspondence. However, cortical macroanatomy, i.e. cortical folding patterns, is not a perfect predictor of functional localisation (Van Essen and Dierker, 2007). Great variability of folding patterns makes it difficult to match functionally homologous features. For example, the question arises how to match functionally corresponding sulci in two subjects where, in one subject, the sulcus is split in two, and in the other remains a single, unbroken structure (Ono et al., 1990).

Although the aim of registration is ultimately application-specific, functional correspondence is most commonly sought after, especially in group fMRI studies. Consequently, several novel algorithms have appeared recently that try to enhance

functional alignment by using prior information from other imaging modalities such as functional MRI or *ex vivo* cytoarchitectonic maps (details of these approaches are discussed in Chapter 2, section 2.1.2). Although they both achieve some enhancement in functional alignment, mentioned registration methods either fail to achieve improvements across the whole cortical surface and/or depend on a highly specific training data set that needs to suit a particular application.

1.3.3 Image Registration Using Structural Connectivity

Recent improvements in registration accuracy and robustness (achieving better functional and structural correspondence) commonly involve the incorporation of other imaging modalities such as functional MRI or prior information about the cytoarchitectonic brain subdivisions. Recent developments of protocols imaging structural brain connectivity have opened up opportunities to use connectivity as a predictor of fine-scale functional segregation and improve joint functional-structural alignment. The in-depth motivation for using structural connectivity to drive registration is given in Chapter 2, section 2.4.1. Here we summarise the main points and explain two ways in which this thesis makes contributions:

- Firstly, by examining the potential for using structural connectivity-driven registration to enhance the alignment of functional homologues. In this sense we develop a comprehensive registration pipeline that takes in connectivity information measured through probabilistic tractography (Chapter 2, section 2.3). We validate this framework using fMRI data and explore potential applications (Chapter 3). In this sense, the major contribution lies in proposing how long-range structural connectivity can be used to drive registration. The proposed framework does not require extensive training data, it uses a commonly acquired modality (diffusion-weighted images) and has a potential to achieve enhancements in functional correspondence across the whole

cortical surface. Furthermore, we propose how surfaces of subcortical structures can be registered using structural connectivity-information (Chapter 4). The enhancement of vertex-to-vertex anatomical and/or functional correspondences in the subcortical framework is demonstrated on the example of a human thalamus and other structures.

- Secondly, the development of the connectivity-driven registration pipeline is the first and an essential step towards the construction of an atlas of brain connectivity for both cortical and subcortical surfaces (see the section on *Human Connectome* in Chapter 5, section 5.2). The human connectome is expected to provide a valuable mechanistic insight into dynamic brain data (Sporns et al., 2005) and accurate connectivity-based alignment should allow for the comparison of connectivity patterns among different individuals or groups of subjects. In this thesis we explore how connectivity-driven registration can provide a deeper insight into the differences between the healthy controls and the Alzheimer's disease patients. Finally, we propose how other types of connectivity, i.e. functional connectivity, can be readily used in the registration pipeline developed in this work.

1.4 FMRIB Software Package and FreeSurfer

FMRIB Software Package (FSL) and FreeSurfer are the two brain MR image processing software packages extensively used throughout this thesis. We prefer these two software solutions to other similar packages due to their reputation and accessibility. However, we stress that registration methods proposed in this thesis, although designed using FSL and FreeSurfer, are essentially independent of the exact implementation and could be part of any other state-of-the-art software package that has similar functionality. In other words, the cost function behind connectivity-driven registration (which is one of the major novel contributions of this work) can be

integrated into other registration algorithms working with surface and/or volumetric brain MR data.

Functional MRI	FEAT	Model-based FMRI analysis: data preprocessing (including MCFLIRT motion correction); first-level FILM GLM timeseries analysis; higher-level FLAME Bayesian mixed effects analysis.
	MELODIC	Model-free FMRI analysis using Probabilistic Independent Component Analysis (PICA). MELODIC automatically estimates the number of interesting noise and signal sources in the data and because of the associated "noise model", is able to assign significance ("p-values") to the output spatial maps. MELODIC can also analyse multiple subjects or sessions simultaneously using Tensor-ICA.
Structural MRI	BET	Brain Extraction Tool - segments brain from non-brain in structural and functional data, and models skull and scalp surfaces.
	FAST	FMRIB's Automated Segmentation Tool - brain segmentation (into different tissue types) and bias field correction.
	FLIRT	FMRIB's Linear Image Registration Tool - linear inter- and intra-modal registration.
	FNIRT	FMRIB's Nonlinear Image Registration Tool - nonlinear registration.
	FIRST	FMRIB's Integrated Registration and Segmentation Tool. FIRST uses mesh models trained with a large amount of rich hand-segmented training data to segment subcortical brain structures.
	FUGUE	Unwarps geometric distortion in EPI images using B0 field maps.
	SIENA	Structural brain change analysis, for estimating brain atrophy.
	FSL-VBM	VBM-style analysis using FSL tools, for voxelwise analysis of grey-matter density.
Diffusion MRI	FDT	FMRIB's Diffusion Toolbox - tools for low-level diffusion parameter reconstruction and probabilistic tractography, including crossing-fibre modelling.
	TBSS	Tract-Based Spatial Statistics (part of FMRIB's Diffusion Toolbox) - voxelwise analysis of multi-subject diffusion data.

Table 1.1: A list of FSL image processing tools used or mentioned in this thesis.

FSL (http://www.fmrib.ox.ac.uk/fsl/) is a set of tools for analysing structural, functional and diffusion-weighted MR brain images (see Table 1.1). Its flexible

design, easy integration with FreeSurfer and comprehensive functionality were the main motivating factors for developing connectivity-driven registration within this framework. Table 1.1 gives the list of tools describing their basic functionality, many of which are used in this thesis mostly in the preprocessing stage (Chapter 2, Figure 2.4).

FreeSurfer (http://surfer.nmr.mgh.harvard.edu/fswiki/FreeSurfer) is a set of tools primarily aimed at studying cortical and subcortical anatomy. Most of FreeSurfer pipeline is highly automated making it suitable for processing large data sets. The pipeline has two streams: the surface- and volume-based ones. The surface-based stream segments white/grey matter and pial surfaces, performs geometry-driven surface registration (see Chapter 2, section 2.1.1) and geometry-driven surface parcellation. The volume-based stream parcellates the volumetric image into functionally/structurally valid regions such as grey/white matter, subcortical structures, etc. FreeSurfer also contains graphical user interface for viewing surfaces (tksurfer) and volumes (tkmedit) as well as a set of tools for statistical analysis of surface data. Of particular interest for this thesis is the surface segmentation, registration and parcellation scheme.

1.5 Main Thesis Contributions

The main contribution of this thesis is in developing an image registration pipeline driven by brain connectivity measured using probabilistic tractography. We developed a rounded-up registration framework that can be further improved and validated, and demonstrated through validation using artificial and real data, that structural connectivity-driven registration improves the alignment of functional homologues. The novel registration cost function based on similarities in global connectivity profiles can be used in both surface and volumetric settings and is independent of the software package in which it is implemented. Moreover, the proposed connectivity profiles

similarity matrix (Chapter 2, section 2.4.4.2) can be constructed not only using structural, but also functional connectivity, for which we propose a specific solution. Therefore, we proposed how to measure connectivity relevant to registration and how to efficiently use it to construct a registration cost function. This thesis validates proposed registration framework for both cortical and subcortical structures using an independent imaging modality where possible. In its own right, connectivity-driven registration of subcortical surfaces is a novel idea and this thesis offers the proof of concept for it. Other contributions lie in the development of a Resting State Networks-based functional parcellation method for brain areas such as the Medial Frontal Cortex, as well as in demonstrating the possible benefits of connectivity-driven registration in the volumetric framework.

1.6 Thesis Overview

In the chapters to follow, we motivate, develop and validate a surface-based connectivity-driven registration pipeline and perform exploratory analysis involving Alzheimer's disease patients. The first part of Chapter 2 reviews existing surface registration paradigms as well as techniques for quantifying cortical connectivity (using Diffusion-weighted Images) relevant to this thesis. In the second part, we motivate connectivity-driven cortical alignment, propose a registration algorithm and mathematically formulate the registration task. In Chapter 3, the proposed cortical registration pipeline is validated with an independent, fMRI, data set. Potential applications are explored, i.e. through the analysis of connectivity-based co-registration of control and Alzheimer's patient groups. Connectivity-driven registration is extended to the surfaces of subcortical structures in Chapter 4. Detailed tests on artificial data are performed to ensure registration consistency. Finally, the exploratory analysis investigates the differences in warping patterns between two groups of subjects: Alzheimer's disease patients and healthy controls. In

Chapter 5 we draw conclusions, describe areas of future work and show some initial results investigating the potential of connectivity-driven registration in the volumetric setting.

"Wandering between two worlds, one dead,
The other powerless to be born,
With nowhere yet to rest my head,
Like these, on earth I wait forlorn."

– Matthew Arnold, Stanzas from the Grande Chartreuse, 1855

Chapter 2

Methods for Connectivity Driven Registration of Cortical Surfaces

In this chapter we explain how cortical white matter connectivity[1] (Chapter 1, section 1.3.3) can be used to drive registration of cortical surfaces. Section 2.1 introduces the concept of the cortical surface and explains how it is used in cortical geometry-driven registration methods, in particular with FreeSurfer. Our methods build upon these by adding additional information (derived from diffusion tensor images) into the registration pipeline.

Section 2.3 explains how cortical connectivity can be assessed by probabilistic tractography. The main features of the tractography algorithm are presented and the benefits and limitations of the probabilistic fibre tracking are discussed in the light of connectivity-driven registration.

Finally, section 2.4 motivates the use of structural connectivity for brain alignment, introduces our registration framework and presents the algorithmic pipeline. We discuss in detail how structural connectivity features are constructed and compared among different cortices as well as how cortical surfaces are deformed to achieve connectivity-driven matching.

[1]Further referred to as *connectivity* or *structural connectivity* if not stated differently.

2.1 Introduction

The human cerebral cortex has a topology of a two-dimensional sheet, a property which has been widely used for comparing cortical folding patterns in different individuals (Drury et al., 1996; Van Essen et al., 1998; Fischl et al., 1999b; Mangin et al., 2004; Lyttelton et al., 2007). These patterns have been shown to predict to a certain extent the functional parcellation/specialisation of the brain (Van Essen and Dierker, 2007; Fischl et al., 2008). As cortical folds (geometry of the cortical surface) appear with some consistency across the human population they have long been used to establish correspondence between functional homologues as well.

In the context of functional neuroimaging, the representation of the brain as the pial or white/grey matter surface has its advantages. Many of the observed brain activations (e.g. through an fMRI paradigm or EEG) can be projected to one of these surfaces and be compared across subjects through cortical geometry-based alignment. Moreover, the spherical topology of the cerebral hemispheres (surfaces thereof) allows for better visualisation of aforementioned activations. These surfaces can be 'inflated' revealing areas such as deep sulcal fundi.

From the brain registration point of view, comparing features on the surfaces changes the three-dimensional (3D) optimisation problem into a two-dimensional (2D) one. Finally, many of the image processing techniques, such as smoothing, work better on the surface than in the 3D volume due to a highly convoluted nature of the cerebral cortex. However, by reducing the representation of the brain to one of the cortical surfaces, other information (e.g. the shape of subcortical structures, image intensities within the brain volume) is lost. Therefore, the utility of the surface-based methods is limited. But nevertheless, surface processing and registration methods have extensively been combined with volume-based ones (such as in CVS (Postelnicu et al., 2009)) bringing together many of the advantages of these two worlds.

2.1.1 FreeSurfer Cortical Geometry Driven Registration

FreeSurfer is a software package with a set of automated procedures for accurate reconstruction of cortical surfaces as well as their inter-subject comparison and parcellation. It uses the two-dimensional structure of the cortical surface to facilitate the measurement and examination of important properties of the cerebral cortex such as the cortical thickness as well as retinotopic, tonotopic and somatotopic maps (Dale et al., 1999).

FreeSurfer cortical alignment is preceded by the following steps: image segmentation, surface reconstruction, surface inflation and parametrisation. During the image segmentation step, the white and grey matter boundaries are delineated and the pial and white/grey matter surfaces are modelled by a dense mesh (reconstruction). Original surface topology is preserved by correcting for topological defects (Segonne et al., 2005). The surface is then inflated to a significantly smoother one (such as the sphere) by minimising metric and topological distortions (Fischl et al., 1999a). Figure 2.1 illustrates the core processing steps.

Fischl et al. (1999b) introduced the procedure for the inter-subject surface-based alignment of cortical folding patterns. After introducing the surface-based coordinate system for the cortex they developed a means for generating an average folding pattern map through a nonrigid alignment procedure. This 'average' or 'template' space can be viewed as the target space for the individual subject registration.

The alignment is carried out (on the unit sphere) by minimising the cost function: the mean squared difference between the maps of average convexity of reconstructed cortical surfaces across a set of subjects and that of the reference. Convexity reflects the folding patterns of the cortex and is a scalar measure that can be associated with every point of the reconstructed/inflated surface. To account for the local inter-subject variability of convexity maps, the cost function is modulated by the variance of convexity across subjects. In this way, the cost function down-weights the

contribution of matching highly variable areas (in terms of cortical surface convexity) compared to matching more stable features. Our registration algorithms follow similar steps: the registration happens on the unit sphere, but convexity is replaced by connectivity information.

The registration problem is formulated as a nonlinear optimisation and can be summarised through the minimisation of the energy functional J. Energy functional can be written as: $J = J_p + \lambda_A J_A + \lambda_d J_d$. J_d and J_A perform the role of the mesh regulariser, whereas J_p quantifies similarity in folding patterns between the two brains/meshes. Typically, J_p measures similarity of surface convexity across corresponding vertices of the input and reference mesh in the common coordinate system (spherical registration space). J_d and J_A preserve local distances and areas respectively. The distance (Fischl et al., 1999b) term J_d quantifies the amount of vertex displacements due to registration and serves to give the surface some local stiffness, thus discouraging the introduction of excessive shear. The areal term J_A acts on triangular mesh elements and quantifies their area. By not allowing excessive change in triangle area it prevents folds and significant compression/expansion. λ_A and λ_d are the tuneable coefficients that balance contributions of regularising terms.

This type of surface-based alignment has been shown to perform better than traditional normalisation techniques (such as 3D Talairach transform (Talairach and Tornoux, 1988)) in aligning sulcal/gyral features as well as functional areas across subjects (Fischl et al., 1999b). In our work, we adopt this alignment paradigm by tuning the parts of the energy functional to the needs of connectivity-driven alignment.

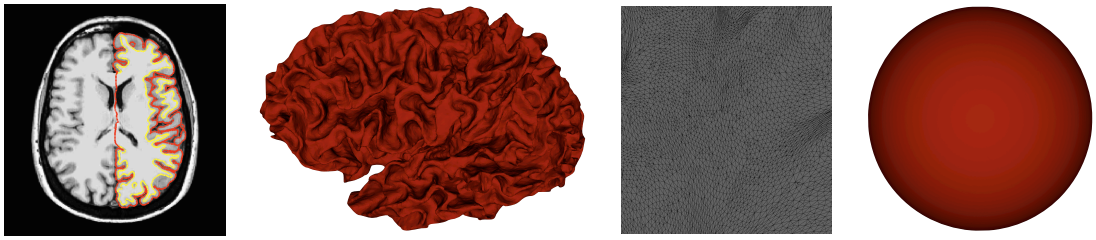

(a) **(b)** **(c)** **(d)**

Figure 2.1: **(a)** Red delineates the pial and yellow the white/grey matter surface of the left hemisphere. **(b)** The reconstructed white/grey matter surface of the left hemisphere. **(c)** Part of the high-density mesh representing the surface in (b). **(d)** The unit sphere onto which the reconstructed surfaces are mapped/projected. The unit sphere serves as the common coordinate system for surface registration. It is also important to note that all meshes encoding the same structure such as the white/grey matter surface and represented in different spaces (such as the original native space, the inflated space or the unit sphere) have exactly the same number of vertices and retain the original vertex labelling established in the mesh of the original reconstructed surface (after any topological corrections).

2.1.2 Other Cortical Alignment Paradigms

2.1.2.1 Anatomical Landmarks

FreeSurfer uses cortical geometry information to perform surface registration and this is done completely automatically. However, the ambiguities in sulcal/gyral correspondences between subjects reduce the consistency of the mappings and can result in ill-constrained deformations.

An alternative is to use a set of (automatically or manually defined) *anatomical landmarks* to constrain mappings to the average template. Anatomical landmarks are usually associated with main cortical sulci such as the central sulcus which are consistently found in the human population and can be either automatically or manually segmented in brain images (Mangin et al., 2004). This approach has been adopted by a number of software packages such as Caret (Van Essen et al., 2001) and BrainVisa (Cointepas et al., 2010). This is done in hope to better constrain deformations and impose some expert knowledge on the initial correspondence.

2.1.2.2 Functional MRI-driven Alignment

A common aim of cortical surface registration is to relate, across subjects, gross brain anatomy to cytoarchitecture and functional specialisation (Van Essen and Dierker, 2007). Naturally, other methods developed that make use of functional modalities to infer sought after mappings. For example, running specific functional localisers on a set of subjects can help establish better correspondence between those particular functional areas.

This trend culminates with the work of Sabuncu et al. (2010) where response patterns to a complex functional paradigm (viewing a motion picture film) are used to achieve high-resolution surface matching. This has been shown to improve the overlap of functional homologues within subjects, especially those involved with cognitive-visual processing. This work also proposes an algorithm capable of dealing with multidimensional functional data in a group-wise registration pipeline.

Challenges Although greatly important, this work leaves open crucial questions such as:

- How can the functional paradigm be optimised to give the greatest improvement in functional overlap over the largest cortical area?
- Is it possible to construct a functional template brain that could work across a range of functional paradigms across a range of scanning conditions?

Also, acquiring functional data is often tied to numerous acquisition challenges ranging from incorrect performance of the task to language barriers. However, recent work on utilising resting-state functional data and especially resting-state *functional connectivity* might successfully address some of these issues (Conroy et al., 2009).

Our work somewhat addresses the first question by trying to find a structural predictor of functional specialisation that is relatively consistent in the human

population and use that as the surrogate input in registration algorithms. In this way, the registration could be expected to have beneficial effects in aligning functionally homologous regions across the whole cortical surface.

2.1.2.3 Application-driven Cortical Alignment

It has been recently proposed by Yeo et al. (2010b, 2009) that the application-specific nature of the registration problem can be exploited to design registration algorithms tailored to the specific task. Yeo et al. (2010b) achieve the enhancement in the overlap of cytoarchitectonically homologous areas by combining the template derivation (including a manually labelled training set) and the registration steps into one. Therefore, the 'free' parameters in the cost function are optimised so as to achieve best match in functional regions by alignment of cortical folding.

The application-specific nature of this approach renders it highly suitable in instances when the corresponding training dataset is available and ensures the 'optimal' solution to the surface registration problem in such instances. However, the generic registration algorithms still have an advantage of most often being independent of training data and retain their appeal due to their wide applicability.

2.1.2.4 Relation to Registration Using Connectivity Information

Although classified as a 'generic approach', our method is different from the others presented in this section in that, in the long term, it tries to make use of the *Human Connectome* (Sporns et al., 2005) to drive registration. We believe that connectome could provide a valuable mechanistic insight into dynamic brain data. Whilst cortical structure-function relationship is unlikely to be a one-to-one mapping, the connectome could be a structural descriptor with high enough inter-subject consistency to allow successful alignment and be a better predictor of functional specialisation than some other brain features commonly used today (e.g. folding patterns). However, it yet

remains unclear to what extent (and at which spatial scale) the connectome can relate to brain function as well as how it can be best imaged for this purpose.

2.2 White Matter Tractography

Diffusion-Weighted MRI can help in discerning structural features of brain tissue. Water diffusion occurs predominantly along the white matter axons and is hindered across them due to the existence of the myelin sheaths (Beaulieu, 2002). Therefore, measuring diffusion in vivo (Catani et al., 2002) can tell us about the spatial position and direction of white matter axons and fibre bundles. As we have seen in Chapter 1, DW images quantify water diffusion in certain spatial directions.

Tractography algorithms are image processing methods that can help in resolving WM fibre bundles/tracts from DW images. Many variants of local tractography methods exist, but they can be roughly divided into two major categories (Sotiropoulos, 2010): deterministic and probabilistic. They all have in common that they try to infer on the underlying white matter tracts, i.e. their orientation, spatial extent, etc.

2.2.1 Deterministic Streamline Tractography

Each voxel of a DW image encodes information about the fibre orientations within it. The orientation estimates (whether there is one or more per voxel (Chao et al., 2008; Wedeen et al., 2008)) form a vector field. The 3D curve tangent to the estimated vector field is called a streamline and represents the estimate of the path that a certain WM tract takes. The basic idea of deterministic tractography is to form streamlines by following the trajectory encoded by the principle eigenvector of the diffusion tensor in each voxel (Basser et al., 2000). Propagation methods start from a seed voxel and interpolate the streamline across voxels until a certain stopping criterion is met, such as reaching an area with low FA, abrupt bending of the streamline, etc.

More advanced streamline methods try to utilise multiple fibre orientation estimates (within each voxel). This is beneficial for tracking in the areas of crossing fibres or areas where the uncertainty of the main fibre direction is high. Multiple fibre orientations can be estimated, for example as the local peaks of the orientation distribution function using DSI (Wedeen et al., 2008) or Q-Ball Imaging (Tuch, 2004). In this way, the streamline estimates are improved by being able to interpolate, upon entering the voxel, along the direction that produces the minimal curving of the tract (Parker and Alexander, 2003) which is closer to the true underlying anatomy.

Deterministic tractography gives binary answers to whether two brain regions are connected. To address this issue and to try to estimate the uncertainty associated with the connectivity of two regions, probabilistic streamlining was proposed.

2.2.2 Probabilistic Streamline Tractography

Probabilistic tractography aims to set confidence intervals on the existence of a fibre bundle. Typically, probabilistic tractography methods estimate the uncertainty associated with the existence of tracts stemming from a certain seed voxel (Behrens et al., 2003b; Parker et al., 2003). Effectively, many streamlines (spatial distributions thereof) originating from the seed can be generated by taking into account the (estimated) uncertainty associated with the underlying fibre orientations within voxels. The method of Behrens et al. (2003b), which is used in this thesis and implemented in FSL, uses Bayesian framework to estimate posterior distribution of model parameters such as the local fibre orientations.

Other variations (and implementations) of probabilistic tractography are often used. Uncertainty in fibre orientation measurements can be assessed using bootstrapping (Pajevic and Basser, 2003). Bootstrapping is a non-parametric method for estimating the distribution of a certain parameter by randomly sampling the set of repeated (indirect) measurements of that parameter. This technique has

been used by Jones and Pierpaoli (2005); Lazar and Alexander (2005) in *bootstrap tractography* where a streamline is generated per bootstrap sample. The set of suprameasurements in this case includes repeated acquisitions in all diffusion directions. Bootstrap tractography has advantages over probabilistic Bayesian framework in that it is sensitive to sources of variability and noise that cannot be modelled parametrically. This tractography method, as well as many other algorithms for fitting diffusion tensors are implemented in CAMINO software package (Cook et al., 2006). Apart from Bayesian and bootstrap approaches, procedures that empirically predict orientation uncertainty also exist (Parker and Alexander, 2003, 2005). However, in this thesis, we primarily use probabilistic tractography as implemented in FSL and described in detail in section 2.3. This is done to facilitate integration with the rest of the registration code and FreeSurfer. Nevertheless, we acknowledge that any other tractography method capable of quantifying the 'degree of connectivity' between two brain regions can be used instead.

2.3 Probabilistic Tractography (FSL/probtrackX)

Probabilistic tractography is a method for estimating cortical connectivity from Diffusion-Weighted MR images. First introduced by Behrens et al. (2003b); Parker et al. (2003) and further developed in (Behrens et al., 2007; Jbabdi et al., 2007) probabilistic tractography has been used to estimate global connectivity, i.e. to measure uncertainty associated with the existence of a connection between any two points.

In the first instance, the method models the diffusion signal on a voxel-by-voxel basis and represents the uncertainties in model parameters in the form of probability density functions (*pdfs*). These local *pdf* estimates are then used to derive a spatial *pdf* between any two points in the data field.

Modelling the Diffusion Signal Locally, diffusion signal from every voxel is related to the underlying fibre structure. When no crossing fibres are modelled, the signal is decomposed into a two-compartment partial volume model. One compartment models diffusion only in the fibre direction and the second models isotropic diffusion of free water. Figure 2.2 (left) illustrates an axon in green encapsulated in a symbolic blue voxel. The two-compartment model is represented by yellow and pink compartments, the yellow one elongated in the principle direction of the fibre bundle and the pink one modelling free water diffusion:

$$\mu_i = S_0 \left(\underbrace{(1-f)\exp(-b_i d)}_{isotropic} + \underbrace{f \exp(-b_i d \mathbf{r_i}^T \mathbf{R}\mathbf{A}\mathbf{R}^T \mathbf{r_i})}_{directed} \right), \tag{2.1}$$

where μ_i is the diffusion signal from the observed voxel, S_0 is the signal with no diffusion weighting and $d, b, \mathbf{r}, \mathbf{R}$ and $\mathbf{A}$ are diffusion parameters (encoding principle diffusion direction and diffusion weighting properties). f defines relative contributions of the two compartments in the voxel. The two compartments are labelled as *isotropic* and *directed*. The directed one can be further refined by including a distribution on the fibre direction parameters (whose parameters in turn become a part of the model). All model parameters can be estimated in a Bayesian setting given the data, e.g. through Markov Chain Monte Carlo sampling schemes. Therefore, the result of the local modelling is an estimate of a distribution *pdfs* on fibre direction parameters and the different compartment contributions, for each voxel.

The Model of Global Connectivity In the second step, *pdfs* of the local partial volume models are used to infer on a model of global connectivity. This is achieved through the construction of 'probabilistic streamlines' (in contrast to the 'deterministic' ones as in Figure 2.2 (middle)) by drawing samples from the posterior *pdfs* on fibre directions at each point in space. The probabilistic streamline originating at point A is formed by starting at A and following fibre direction encoded in sampled

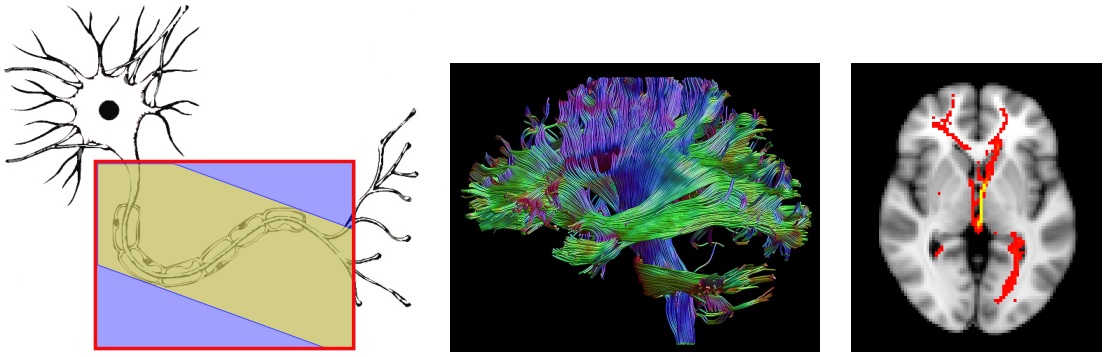

2-compartment model | Fibre Bundles 'Streamlines' | Probabilistic Tractography

Figure 2.2: **(left)** A scheme describing the two compartment model for predicting the diffusion-weighted signal. The axone (a schematic of a neurone is shown) is symbolic of the directed fibre bundle. Red quadrangle denotes the region of interest, i.e. a 'voxel', whereas the two compartments, *isotropic* and *directed*, are shown in blue and yellow respectively. **(middle)** Traditional fibre tracking - streamlining procedures cannot trace in the areas of high fibre direction uncertainty such as cortical grey matter, and are therefore constrained to tracing only the most prominent fibre bundles. They also do not explicitly quantify the uncertainty associated with the tracking procedure (*Fibre Bundles* reproduced by courtesy of Dr A. Leemans http://www.exploredti.com/gallery). **(right)** Probabilistic tracking output overlaid on an axial slice of the standard brain. The seed voxel is on the medial wall of the left thalamus. Colours ranging from yellow to red encode for the number of 'probabilistic streamlines' (yellow - high, red - low).

pdf until some stopping criterion is met. By forming many such streamlines it is possible to form a *spatial pdf* of the connection between A and any other predefined point.

Implementation Probabilistic tractography is implemented in the FDT package ProbtrackX of the FSL. Main inputs into the probabilistic tracking algorithm are the 'seed region' where the tracking commences as well as the target regions (if any) to which connectivity is being quantified. Other important free parameters are the number of samples drawn to characterise the *spatial pdf* of global connectivity (usually set to 5000 in our applications), the stopping criterion (usually set as the angle of the sharpest bend allowed or, less commonly, based on the local fractional anisotropy values) as well as the number and length of steps used for construction of 'probabilistic

streamlines'. In our work we use some additional properties such as the constrain on the tract length.

Results of Probabilistic Fibre Tracking The output of the probabilistic fibre tracking is a volumetric map encoding for each voxel the number of 'probabilistic streamlines' passing through it (this number depends on the total number of samples drawn). An example of such image can be seen in Figure 2.2 (right). The benefit of probabilistic over other frameworks is its ability to trace in the areas of high uncertainty in fibre direction and quantify that uncertainty by producing a spatial map of the 'tract spread'. It is predominantly for these reasons we opted for using this method to quantify white matter connectivity for registration purposes. Probabilistic framework for fibre tracking is set in continuous space which allows for sampling of *pdfs* on fibre directions at subvoxel resolution. This property, along with the ability to trace even within cortical grey matter have proven instrumental in using measures of connectivity to parcellate cortex and subcortical structures into areas of distinct connectivity patterns. We also make extensive use of these features.

2.4 White Matter Connectivity Driven Alignment

2.4.1 Motivation

In this section we give a brief overview of the current *state of the art* cortical surface registration methods and the challenges they encounter. We then explore how white matter connectivity could enhance the identification of corresponding functional locations in different brains. Finally, the benefits and limitations of using structural connectivity for image registration are put forward. The remainder of the Chapter introduces the connectivity-driven registration pipeline.

2.4.1.1 Functional Specialisation vs. Cortical Morphology

Ever since Brodmann published his seminal work on brain cytoarchitecture (distribution of neuronal cell bodies) (Brodmann, 1909), it has become common practice to relate microstructural organisation of the brain to its macroscopic features (such as folding patterns) and functional specialisation. Often, functional imaging results are described in terms of 'Brodmann Area' (BA) regions. Likewise, Brodmann areas are determined by looking for 'anatomical landmarks' as surrogate markers consistently identifiable across subjects such as the anterior and posterior central gyri, the central sulcus, etc.

In this way, the inter-subject comparison of brain imaging results is often performed with a high degree of subjectivity due to poor quantification of uncertainty of how well subject-specific (e.g. functional) maps correspond across individuals. Functionally distinct regions are often related to (or even equated with) template-based BA maps predicted by gross brain anatomy without examining the inter-subject consistency of these descriptors. In other words, one-to-one mapping is assumed between functional and BA maps, and cortical morphology. However, this relation is far from a one-to-one mapping due to different aspects of variability (Van Essen and Dierker, 2007) such as folding patterns, BA size variability, etc.

2.4.1.2 Cortical Morphology and White Matter Connectivity

Much work has been done to relate functional specialisation to cortical morphology and elucidate this relationship. This is especially important when disseminating fMRI results as the registration of structural scans is also hoped to align functionally corresponding areas. Recent studies such as the one by Fischl et al. (2008) have shown that cortical folding patterns can be good predictors of BAs closely corresponding to primary cortical areas. However, higher order cognitive areas such as Broca's area (BAs 44 and 45) are highly variable with respect to the folds which suggests a specific

hierarchy in relations between functional parcellation and gross brain morphology. This finding is important as the functionally specific regions that are arrayed across the cortical sheet are believed to be strongly linked to the underlying anatomy (Felleman and Van Essen, 1991). And it is this *mosaic of functionally defined regions, not necessarily BAs as such*, which we commonly try to compare across subjects. Thus, it is natural to wonder whether there could be a set of brain features, complementary to gross brain morphology, that could help boost predictability of functional specialisation.

In this sense, our aim is to devise a registration method that uses such feature sets to enhance inter-subject alignment of cortical surfaces. The pattern of cortical connectivity might be a more causal descriptor of the location of (functional) cortical areas than the folding patterns themselves (Van Essen, 1997). We could therefore speculate that a richer set of information, such as white matter connectivity, could help boost specificity of delineation and thus of the inter subject comparison of regions not well predicted by folding patterns (e.g. higher cognitive areas Fischl et al. (2008); Amunts et al. (1999)). However, for white matter connectivity-driven alignment to achieve this goal, it is crucial that connectivity patterns can be consistently identified across subjects. Equally importantly, structural connectivity needs to contain sufficient discriminatory power to discern the boundaries between functionally different regions. Quite suitably, diffusion-weighted images and different fibre tracking algorithms have opened up possibilities for achieving exactly that.

2.4.1.3 Connectivity-based Parcellation

Several studies investigate the reproducibility, generalizability and validity of connectivity-based parcellation of the human cortex (Klein et al., 2007; Tomassini et al., 2007a; Beckmann et al., 2009). The common conclusion is that diffusion tractography can be used as an objective tool for parcellation of the human cortex into functional regions. Probabilistic tractography in particular proved early on to be

of use in this respect (Behrens et al., 2003a; Johansen-Berg et al., 2005; Rushworth et al., 2006).

The common way to proceed with connectivity-based parcellation is to quantify connectivity of each subpart of the region of interest (ROI) to a set of predefined targets and then cluster connectivity descriptors (or connectivity feature vectors) using simple methods like k-means or spectral clustering into a number of clusters supported by the anatomy - see Figure 2.3. For example, Johansen-Berg et al. (2004); Klein et al. (2007) cluster medial frontal cortex (MFC) into the SMA and pre-SMA areas. SMA (Supplementary Motor Area) and pre-SMA have distinct functional properties: the former is important for temporal organisation of movements whereas the latter plays more abstract roles such as planning and preparation of movements. Connectivity-based clustering was successfully used to separate MFC into these two regions and the results were validated with respect to the underlying cytoarchitectonic maps as well as functional localisers. This is particularly important as there is no local landmark to help define this boundary (Geyer, 2004) in humans. We can therefore assume that in such areas connectivity-based registration would align functional boundaries better than cortical morphology-based one.

Similarly, Beckmann et al. (2009) parcellate the human cingulate cortex and compare results to a meta-analysis of functional studies reporting cingulate activations. Connectivity-based parcellation identified nine cingulate subregions which were to a certain extent related to the functionally distinct areas pooled from the meta-analysis.

Quite encouragingly, these two studies are supported by others using different fibre tracing techniques as well as different diffusion-sensitised sequences (Perrin et al., 2008; Anwander et al., 2007). Moreover, apart from cortical regions, many subcortical structures have successfully been parcellated using these methods. We shall discuss this in greater detail when presenting similar registration methods for the subcortical

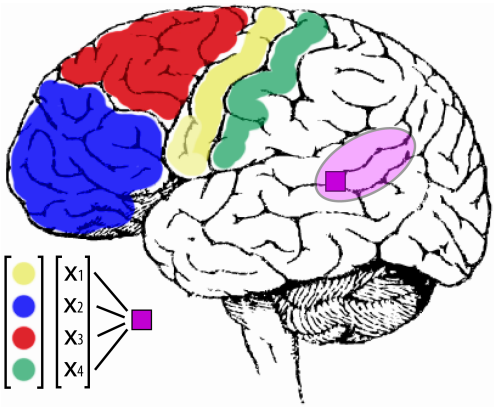

Figure 2.3: Connectivity-based parcellation method. Probabilistic tractography is performed from every voxel (pink square) within the region of interest (pink oval) to a set of cortical targets (shown in yellow, blue, red and green). Just those 'probabilistic streamlines' that reach the targets are considered. Once the streamline reaches the target, tracking terminates. Quantitative measurements are assembled into connectivity feature vectors e.g. $[x_1\, x_2\, x_3\, x_4]^T$ which are then compared through the clustering scheme.

framework (Chapter 4).

In summary, structural connectivity has been found a good predictor of functional specialisation in a number of cortical regions, some of which lack morphological landmarks for accurate delineation. Consequently, we hope that using such information in addition to cortical geometry could enhance alignment of functional homologues across subjects.

2.4.1.4 Practical Benefits

All the benefits of structural connectivity-driven registration come from achieving its ultimate goal: better inter-subject alignment of functional homologies. Depending on the spatial scale at which these functional regions are captured by diffusion-weighted data, the improvement in alignment should have positive impact on a range of functional studies. In relation to this, we see potential value particularly in the following:

- *Increased sensitivity of group fMRI studies.* Drawing conclusions from

group fMRI studies heavily relies on successful registration to the common template/prototype subject. Improving on this registration step means increasing statistical power of drawing group-wise conclusions.

- *Matching of atypical 'folding variants'.* Major sulci or gyri can have a very different configuration in some individuals compared to the most common pattern (Van Essen and Dierker, 2007; Ono et al., 1990). These folding idiosyncrasies might prevent one-to-one mapping between cortical surfaces based on folding patterns. Connectivity-driven alignment has the potential to provide a clearer way of aligning such areas in a functionally meaningful way.

- *Development of imaging biomarkers.* Changes in brain connectivity (e.g. due to a pathology such as Alzheimer's disease) in a group of subjects can be reflected in the deformation field matching that group to the common template. Since our registration is driven by the (dis)similarities in connectivity profiles, the resulting deformation fields should reflect these relations as well. Therefore, discrimination between groups (or even between individuals) on this basis could potentially help in identification and classification of some brain pathologies. Moreover, direct benefits could also come from spatially localising areas of connectivity different to that of the control group/template subject.

- *Atlas of brain connectivity.* A rigorous whole brain connectivity-driven registration framework can be directly employed for the construction and refinement of a connectivity brain atlas. For example, each point of such an atlas brain could encode the population-based knowledge of the connectivity of that point to the rest of the brain.

2.4.1.5 Limitations

As well as the benefits, the limitations and the unknowns of this registration framework are also numerous. Major concerns relate to the *unknown spatial scales* at

which connectivity information can discriminate/identify functionally distinct brain regions. Furthermore, even if such a scale is found or imposed, the question of consistency still remains. Are these regions robustly identifiable across the human population? Even further, is this true for the whole cortex or just some particular parts? In relation to the latter, the cortex-wide applicability of connectivity-based parcellation (see section 2.4.1.3) can be questioned. To the best of our knowledge, there have been no studies to date that achieve robust and replicable across-subjects parcellation of the whole cortical surface based solely on the connectivity information. The challenges with achieving this goal are numerous and rooted into the complexity of connectivity data as well as their inter-subject variability (Roca et al., 2009). Therefore, although we hope that connectivity-driven registration would yield alignment improvements across the whole cortex, its applicability might be restricted to a limited set of brain areas. On top of that, due to the mentioned uncertainty in spatial scales, we cannot *a priori* know the range at which alignment improvements happen.

Other challenges come from the imperfect nature of diffusion data. Although models of diffusion signal have become more complex and more reflective of the underlying anatomy (CHARMED, Q-BALL and Diffusion Spectrum Imaging (Assaf and Basser, 2005; Tuch, 2004; Wedeen et al., 2008)), many issues remain to be solved such as those of the multiple crossing fibres. In relation to probabilistic tractography, modelling crossing fibres helps resolve 'weak tracts' but many known pathways still remain beyond the reach of magnetic resonance imaging (Behrens et al., 2007). Moreover, the applicability of our technique also depends on the answers to some basic neuroscientific questions such as what is the relation between structural connectivity and functional specialisation. Finally, the functional organisation of the brain has pronounced plasticity (Lledo et al., 2006) which can modify the functional specialisation of certain areas (or even their spatial extent). However, this

adaptation and functional reordering are not necessarily reflected in similar changes in connectivity (although some indications of this exist - Scholz et al. (2009)).

2.4.2 Notation and Terminology

Throughout the rest of the chapter we refer to *a mesh* M_X. A mesh M_X is defined by a set (E_{M_X}, V_{M_X}), where E_{M_X} is a set of edges and V_{M_X} is a set of vertex labels in M_X. A vertex label $i \in V_{M_X}$ is a reference/name for an individual vertex of the M_X mesh. Similarly, $x_i,\ i \in V_{M_X}$, refers to a *point* x_i defined by its coordinates and associated with vertex i from mesh M_X. We often use meshes which are regular subdivisions of an icosahedron. A subdivision of an icosahedral mesh is a new mesh derived from the icosahedral one by dividing each of its sides into 4 equilateral triangles. These subdivisions are referred to as icx where x denotes the number of subsequent divisions of the icosahedron. For example, $ic5$ results in a mesh of 10242 vertices.

We further refer to the connectivity-driven registration of cortical surfaces as FACS (FMRIB's Alignment of Cortical Surfaces) and abbreviate Structural Connectivity Feature Vectors to SCFV. The Connectivity Matrix combines all SCFVs of a single mesh/cortical hemisphere and is abbreviated as CM.

2.4.3 Registration Pipeline

The prerequisite for successful connectivity-driven registration is the construction of connectivity descriptors which can be associated with every image element such as a voxel/vertex. These connectivity descriptors or Structural Connectivity Feature Vectors (SCFVs) must be directly comparable across subjects and therefore formed in a consistent fashion. Setting up of SCFVs belongs to the preprocessing step whereas their comparison and the actual mesh deformation are the core of what we call *the registration algorithm* (see Figure 2.4). Formation of SCFVs is described in the next section 2.4.3.1. The registration algorithm, i.e. the mesh deformation model, the construction of the cost function, regularisation, etc., is detailed in section 2.4.4.

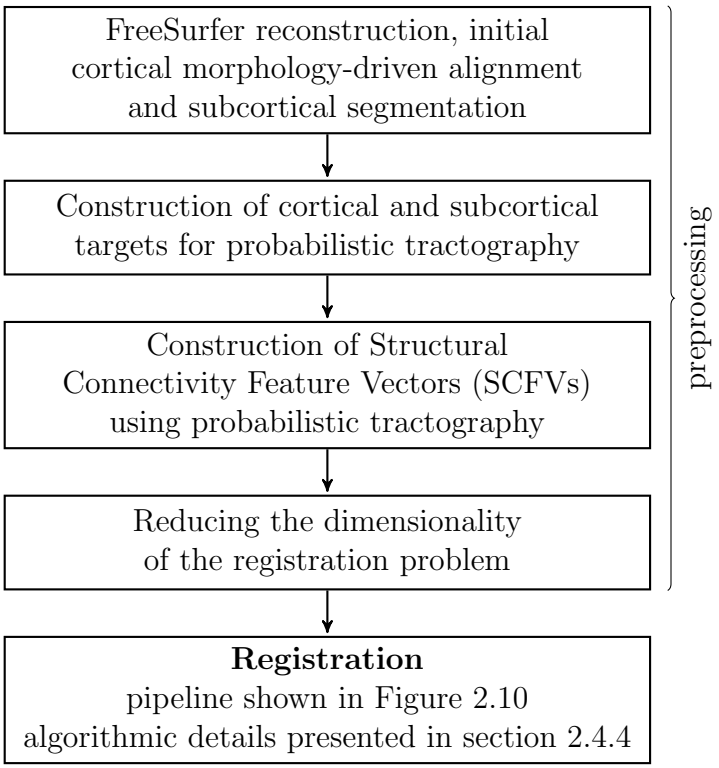

Figure 2.4: Flowchart of the connectivity-driven registration pipeline. Formation of the SCFVs belongs to the preprocessing step whereas their utilisation for alignment purposes is the core of the registration method.

Relevant to the alignment algorithm presented in this Chapter are the registration consistency tests described in Chapter 4, section 4.2.2.1. The artificial test data were devised to help assess registration consistency, robustness to noise and to help the identification of free parameters. Section 2.5.1 of this Chapter discusses the SCFV construction method whereas section 2.5.2 discusses our choice of the surface registration algorithm. Much of the rest of the thesis is concerned with the validation of the registration methods introduced here.

2.4.3.1 Formation of SCFVs

SCFVs encode connectivity properties that drive the registration. Eventually, SCFVs are compared across subjects and the result of that comparison becomes an essential part of the cost function optimised by the registration algorithm. Therefore, the aim

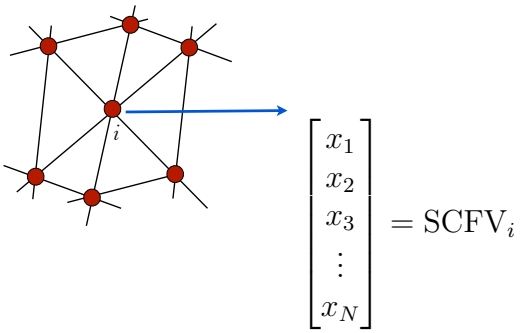

Figure 2.5: Every mesh vertex (i) is associated with a N-dimensional SCFV_i where $x_j,\ j \in [1, N]$ is the measure of connectivity of vertex i to target j.

is to have SCFVs whose information content is sufficient to discriminate[2] between functional regions such as the parcels of the cingulate, thalamic nuclei, etc. which can hardly be resolved with other common structural modalities/contrasts. At the same time, these descriptors should be as simple (or as low-dimensional) as possible to allow for efficient computation of the registration cost function.

As we have already seen in section 2.4.1.3, probabilistic tractography can be used to construct SCFVs with the desired properties. The basic method of fibre tracking to a set of targets was illustrated in Figure 2.3. Each element of the SCFV is a number quantifying the connectivity of the voxel/vertex (with which the vector is associated) to a set of targets defined in volumetric space (Figure 2.5). Therefore, the number of tractography targets determines the dimensionality of SCFVs as well as their discriminatory power. Depending on the application, this dimension can vary from an order of ten (e.g. see the parcellation of the human thalamus by Behrens et al. (2003c)) to thousands (such as when parcellating the cingulate gyrus (Beckmann et al., 2009) - every brain voxel effectively becomes a separate target).

We would like to use target-based probabilistic tractography to form our SCFVs.

[2]Successful functional parcellation/clustering using a given set of SCFVs is not a prerequisite for registration. Alignment is essentially a *local* procedure, not explicitly aware of (neither dependent upon) the boundaries defined by SCFVs. However, registration could be expected to implicitly improve the alignment of such boundaries.

Subject 1
before registration

Subject 2

Subject 1
registered to subject 2

Figure 2.6: Subjects 1 and 2 (left cortical hemispheres) have been registered to each other with FreeSurfer. Coloured unit spheres (registration coordinate system) indicate how folding patterns match after the alignment. Colour encodes for curvature.

The targets should span the whole cortex (as we expect our registration to have cortex-wide effects) as well as include relevant subcortical structures. Critically, all targets need to be consistently identifiable across the group of subjects. All these targets are constructed with the help of FreeSurfer (Figure 2.7).

The first step at the preprocessing stage is to run FreeSurfer segmentation and morphology-driven alignment for all subjects. This 1) provides necessary cortical meshes to work with (in our case these are encoding for white/grey matter boundary surfaces) and 2) initially registers cortical surfaces (Figure 2.6). Initial alignment is necessary for two reasons: firstly, as we shall see, it is essential for ensuring that SCFVs are 'consistently' defined across subjects and secondly, serves as the starting point for our registration algorithm (FMRIB's Alignment of Cortical Surfaces - FACS). FreeSurfer segmentation and registration is applied as part of the automatic pipeline and requires little or no manual intervention. One of the other useful products of the FreeSurfer processing pipeline is the segmentation of subcortical structures (Fischl et al., 2002). We use these segmentations too in constructing SCFVs.

Cortical Targets The first step in the formation of 80 cortical targets is FreeSurfer morphology-driven alignment of all subjects (cortical surfaces thereof) to the average space. Parcellation of each subject's cortical hemisphere is performed in the

average/template space (defined on the unit sphere such as in Figure 2.1d) to impose inter-subject parcellation consistency. Results are then projected back to the native space i.e. original white/grey matter surface segmentation. Parcels are chosen in the form of triangular divisions of the unit sphere (Figure 2.7a) and are then mapped back to the original surface (Figure 2.7a). Please note that these parcels do not follow any anatomical borders as their shape has been artificially imposed in the average space. This is done to minimise the correlation between the parcel and the parcellation error. In cases where parcels follow anatomical borders, it can often happen that some of them are labelled with greater inter-subject accuracy than the others (Desikan et al., 2006; Fischl et al., 2004). In our case, that could reflect itself as an imbalance in contribution of different parts of the SCFVs to the cost function driving the registration.

In the next step, parcellation labels are propagated into the volumetric image in the native space (Figure 2.7c) and d). That is done in the following way: the labelled white/grey matter surface is matched to its corresponding volumetric image and the labels are projected to a one-voxel-thick layer stretching from the white/grey matter surface towards the interior of white matter, down the local surface normals. We tried other projection strategies as well, such as forming a ribbon that was several millimetres thick, which starts off several millimetres away from the white/grey matter surface and projects into the surface interior. However, many variations of these projection strategies produced similar patterns in the matrices grouping SCFVs from all the vertices in a cortical hemisphere. In particular, the variation of the layer thickness and/or the layer depth by $(l \times \mathit{largest\ voxel\ dimension})$mm, for $l = 1, 2$ had negligible effect on sparseness of the aforementioned connectivity matrices (typically, 95% of all the entries in these matrices are zero values). In our experience, varying the layer thickness or depth by more than $l = 1, 2$ had an adverse effect on labelling targeted cortical white matter: it can often happen that

(a) Icosahedral parcellation of the unit sphere

(b) Parcellated white/grey matter surface

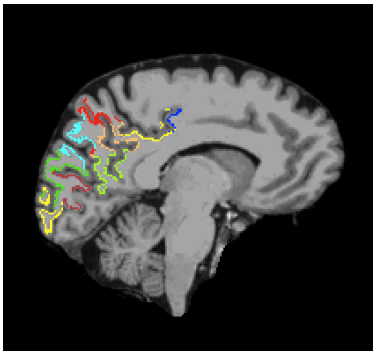

(c) Parcellation projection into the volumetric space

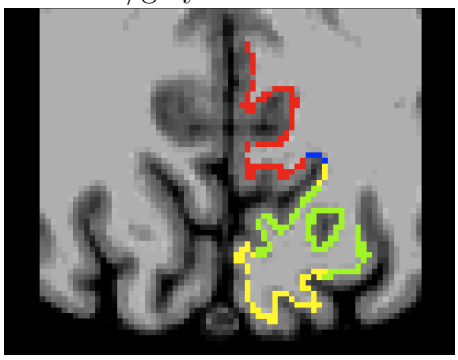

(d) Detail in (c) axial slice

Figure 2.7: **(a)** Triangular targets in the spherical space. Each triangular region (depicted by a different colour) is propagated into the volumetric image becoming a target for probabilistic tractography. **(b)** Parcellation from (a) shown in the native space of a template subject (left cortical hemisphere). **(c)** Several of the triangular parcels from (a) and (b) propagated into the volume of the corresponding volumetric structural image. Different colours correspond to different parcels. Parcels/targets become a one-voxel-thin layer projecting into the white matter. **(d)** A zoom into a detail from (c) - axial slice.

due to complex folds (e.g. narrow sulci) the 'eroded' mesh self-intersects producing unpredictable (and therefore inconsistent across subjects) volumetric labelings.

Subcortical Targets In the ideal case we would have as many subcortical structures for targets as possible, but the main constraint on this is that they need to be automatically segmented with good inter-subject consistency. Inclusion of subcortical targets such as the thalamus or the striatum is important as many major fibre pathways such as the thalamocortical tract exist between deep grey matter

Structure
thalamus
caudate
putamen
pallidum
hippocampus
amygdala
accumbens area
brain stem
white matter of the cerebellum

Table 2.1: The list of subcortical brain structures (plus the brain stem and the cerebellar white matter) segmented by FreeSurfer and used in SCFV construction. All structures (except the brain stem) are segmented in both left and right cortical hemispheres.

structures and the neocortex (Schmahmann and Pandya, 2006).

The choice of possible subcortical targets was limited by the subcortical segmentation properties of FreeSurfer (Fischl et al., 2002). However, other subcortical parcellation algorithms could be equally successfully used for this purpose as well, such as FIRST (Patenaude, 2007). The main advantage of using FreeSurfer is a simple and natural integration into the registration pipeline: all FACS processing stages already rely on FreeSurfer reconstruction and averaging steps.

The list of subcortical brain structures (plus the brain stem and the cerebellar white matter) included in the formation of SCFVs is shown in Table 2.1. Each of these structures has a direct or an indirect connection to the cortex (Schmahmann and Pandya, 2006) and can therefore help in better constraining our registration task by enriching the connectivity information content in the SCFVs.

Probabilistic Tractography to the List of Targets Cortical and subcortical targets are compiled into a list that can be directly fed into the FSL/ ProbtrackX probabilistic tracking algorithm. Tracking is performed from every vertex of the white/grey matter surface/mesh to each of the targets using default tracking parameters. We customised ProbtrackX to take as inputs FreeSurfer cortical meshes

in ASCII format and produce as the output a matrix of SCFVs associated with each of the vertices. We shall call this $V \times N$ matrix a *Connectivity Matrix* (CM), where V is the number of vertices in a mesh (typically 10^5) and N is the number of targets (dimensionality of a SCFV - here set to 89). Values in the connectivity matrix range from 0 to the number of probabilistic streamlines generated from each seed point (this is usually referred to as the number of samples and is by default set to 5000).

However, connectivity matrices/SCFVs generated in this way have an inherent problem that biases any registration for which they might be used. Namely, each of the seed points/vertices already belongs to one of the targets (by the very nature of cortical target construction). Therefore, the connectivity of that point to its nearest (or belonging) target would be 'artificially' high. In other words, the seed and the target are at the same place and hence the tract of length 0 is assumed between the two. As a result of this, all points within one of the triangular targets would have very well correlated SCFVs (compared to the correlations with SCFVs of the points belonging to the other targets). If our registration algorithm was to be driven by the similarities in these connectivity features, then the registration results would be biased towards restoring the original triangular icosahedral parcellation of the cortical meshes.

To visualise this problem we have clustered into 80 clusters the rows of the self-correlated CM generated in the described manner. Results of such clustering can be seen in Figure 2.8 (left). Coloured triangles (i.e. different clusters resulting from the K-means clustering of the self-correlated CM) correspond exactly to the original icosahedral cortical parcellation used for target generation (Figure 2.7a).

Several different solutions to this problem were considered. The CM clustering procedure was used to assess whether triangular patterns can be seen for different variations of CM construction. Initially, we removed (set to 0) all highest entries in each of the rows of a CM hoping that most of those would have been responsible for

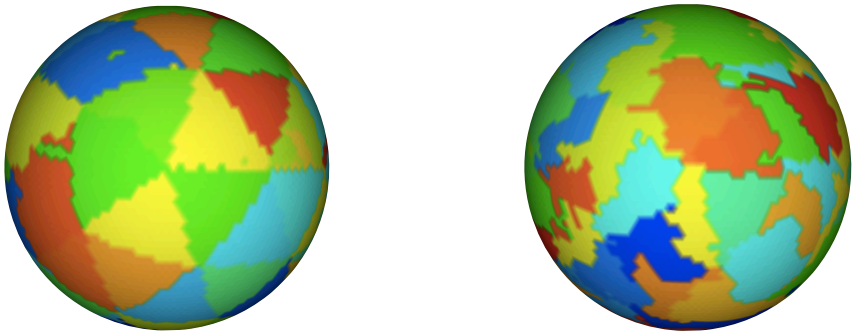

Before correction After correction

Figure 2.8: **(left)** Clustering of the self-correlated connectivity matrix for a prototype subject (left cortical hemisphere) before the tract length-based thresholding. The self-correlated connectivity matrix was clustered into 80 clusters using the K-means algorithm and different clusters are depicted in different colours. The triangular clustering pattern is an artefact associated with the way seeds and targets (see Figure 2.7a) for probabilistic tractography are constructed. It indicates that SCFVs of vertices belonging to the original triangular targets have 'artificially' high correlations with the other points from the same target. **(right)** The connectivity matrix (i.e. SCFVs) was constructed using only those tracts whose length (at the point of reaching the target) is greater than a certain threshold (e.g. 30 mm). The matrix was then clustered into 80 clusters (represented by different colours) just as in the image to the left. The tract thresholding helps to overcome the 'triangular cluster' problem; no artifactual triangular clusters can be observed. This threshold (length) value is determined empirically.

the within-target SCFV correlations. This procedure did indeed have beneficial effects on the clustering test, but failed to completely eradicate triangular cluster patterns. The reason for failure is that each seed point is usually immediately surrounded by more than one target voxel, sometimes even from a neighbouring target/parcel (possible due to a highly convolved nature of the cortex - see Figure 2.7d). Moreover, removing entries in SCFVs which correspond to all targets neighbouring an associated seed point can significantly reduce effective SCFV dimensionality.

Another approach that was tried was to threshold tracts *according to their length* and consider just those longer than a pre-determined threshold. This was promising as the short tracts (some of them of length close to 0) are mostly associated with the selfsame or immediate neighbouring targets (from the perspective of the seed

point). Furthermore, this method keeps those tracts that leave the current target, travel into the cortex and curve back to reach the same target again (if the curvature constraints allow for it). If the length threshold is low enough such tracts would rightly be considered a part of the CM formation, thus minimising the thresholding effects on the descriptive power of SCFVs.

Implementation We implemented tract length-based thresholding in ProbtrackX and performed the clustering test for different thresholds (from 10 mm to 50 mm in steps of 5 mm). For the threshold of 20 mm we already could not discern any remnants of the triangular cluster patterns, and for the threshold of 30 mm clustering remains unchanged relative to the clusterings with thresholds of 30 ± 5 mm[3]. The clustering test for the tract length threshold of 30 mm is shown in 2.8 (right). This thresholding procedure was adopted for all further analysis.

2.4.3.2 Dimensionality Reduction

The connectivity matrix CM can be efficiently calculated even for the high-density cortical mesh (e.g. ~10^5 vertices) with a processing time of approximately 10% of the time required for running all stages of FreeSurfer (not including the time necessary for the estimation of diffusion parameters at each voxel). However, as we shall see in the next section on the registration algorithm 2.4.4, the input to the registration pipeline is the 'correlation' between CMs of the input and the reference subjects. Its dimensionality (without any optimisation) would be approximately $10^5 \times 10^5$ - prohibitively large for storage and computation. Although this size can effectively be reduced (as not every pair of points in the input and the reference mesh need SCFV comparison) we decided to use a simpler solution and reduce the size of the CM matrix by subsampling spherical meshes in the registration space (Figure 2.1d). Effectively,

[3]Just as a comparison, the visual pathways (the optic nerve, the optic tract and optic radiations) run from the front to the back of the brain and can be successfully imaged using DW images (Hofer et al., 2010).

spherical meshes in the average registration space (after FreeSurfer cortical surface registration) are subsampled to a 10242 vertex, large icosahedral mesh (which is one of the upsampled tessellations of the icosahedron).

Original inter-vertex distance in FreeSurfer cortical meshes is below a millimetre. Typical vertex spacing for subsampled meshes in the native space is around 4 mm (similarly done by Dale and Sereno (1993)) which is still much smaller than the inter-subject variability in the location of some functional regions (after the registration). For example, after affine normalisation to the standard space, BAs 18 and 19 and some other areas can be misaligned on the order of centimetres (Amunts et al., 2000; Malikovic et al., 2007; Wilms et al., 2005; Van Essen and Dierker, 2007).

2.4.4 Registration Algorithm

The core registration algorithm relies on the *deformable model framework* which has been widely used in brain imaging in applications varying from brain extraction (Smith, 2002) to cortical surface registration (Fischl et al., 1999a). The defining characteristic of this framework is that it works with deformable meshes encoding surfaces of interest and the deformations are driven by a set of 'forces' derived from the cost function. In other words, the geometrical optimisation problem (deforming a mesh embedded in a 3D space) is solved by the iterative application of a set of forces/displacements to individual vertices.

As we shall later see, forces are derived directly from the cost function and play two roles: 1) to optimise the 'data-driven' part of the cost function and 2) to minimise deformations required for successful alignment. They are commonly known as the *intensity* and the *regularisation* forces respectively. In our case, the 'data-driven' part of the cost function is responsible for achieving the required connectivity-driven matching and is derived from the comparison of *connectivity matrices* of the input and the reference subjects.

2.4.4.1 Deformable Model Framework

Every registration requires two corresponding cortical surfaces (such as the white/grey matter surface of the left cortical hemisphere) with the associated connectivity matrices: one of the input and one of the reference subject/mesh. The input mesh is matched to the reference one to maximise similarities in SCFVs between corresponding vertices (groups of vertices).

Initially, both the input and the reference mesh are represented in the *registration space* (i.e. the parametrised unit sphere) and are registered to the average template using cortical geometry-driven FreeSurfer registration (see section 2.1.1 and Figure 2.6). In that way, an initial correspondence between the input and the reference mesh is established. It is this correspondence that we wish to further *refine* using connectivity information. Connectivity-driven registration is initialised by FreeSurfer-aligned meshes to simplify the optimisation task and prevent some of the local minima traps.

Re-parametrisation Before proceeding further, we re-parametrise and re-tessellate the input and the reference mesh[4] to enforce the same number of vertices (i.e. the same resolution) and to also establish the inter-subject initial 'anatomical' correspondence in vertex labels. Moreover, due to the nature of our regulariser (see section 2.4.4.3), both the input and the reference mesh have to be represented as regularly subdivided icosahedrons (in our implementation we use $ic5$ or icosahedron after 5 consecutive subdivisions, having 10242 vertices).

$$M_O \xrightarrow[inflation]{registration} M_R \xrightarrow[subsampling]{re-parametrisation} ic5$$

M_O is a mesh in the native space (e.g. the original reconstructed white/grey matter surface - Figure 2.1b) and M_R is the mesh in the registration space

[4]When referring to the input and the reference mesh, if not stated differently, we assume that these are both in the registration space of the unit sphere.

(after cortical geometry-driven alignment - Figure 2.6). This step also effectively achieves dimensionality reduction as described in section 2.4.3.2. Therefore, after re-parametrisation, the input and the reference mesh both have the same number of vertices and the geometry of a unit $ic5$. Furthermore, due to the initial FreeSurfer-imposed correspondence, vertex with label x in the reference mesh now 'anatomically' corresponds to the vertex with *the same label* in the input mesh.

To see this better, let us introduce the concept of *overlapping spheres/meshes*. Imagine the input and the reference mesh both with the same radii and the same origin. In such a case, re-parametrised input and reference meshes would perfectly overlap and appear as one (i.e. corresponding vertices and edges in both meshes overlap). It is useful to visualise the registration problem in this way as the input mesh can now be seen as the one *being deformed* to maximise similarities in SCFVs with the *underlying* reference mesh. We shall resort to this parallel again when describing the effect of registration forces.

During the subsampling step the information is kept on correspondences between vertices of the subsampled and the full-resolution meshes. This is facilitated through using the nearest neighbour interpolation method for subsampling. Tracking these correspondences (which are saved in a simple look-up table) allows for an efficient formation of a subsampled CM (the necessity of this was discussed in section 2.4.3.2).

Smoothing In addition to subsampling, when reducing the dimension of CMs, we also perform SCFV smoothing (effective CM smoothing, reduces noise in data) by making use of having a full-resolution CM at hand. This is done in the following way. Let us imagine $ic5$ and the mesh being subsampled (M_R) as the overlapping spheres.

$$\forall\ \mathrm{sSCFV}_i,\ i \in V_{ic5} :\ \mathrm{sSCFV}_i = \frac{1}{n} \sum_{j \in \mathcal{P}_n(j)} \mathrm{SCFV}_j, \tag{2.2}$$

where $\mathcal{P}_n$ is a set of n closest neighbours of vertex $i \in V_{ic5}$ in V_{M_R} and sSCFV stands for the smoothed SCFV. n is set to 7 in our implementation effectively reducing the

dimension of the high-density CM 7-fold (note that the dimension of the high-density mesh and therefore of the high-density CM, is already reduced ~10-fold[5], which is more than the dimension reduction induced by smoothing). Due to the high density of M_R and its locally regular structure, this smoothing procedure is in practice equivalent to smoothing using a distance-weighted kernel.

Smoothed/subsampled connectivity matrices (which we further still refer to as 'connectivity matrices' for simplicity) have the same number of rows for both the reference and the input subjects and are directly used in the construction of the 'data-driven' part of the cost function. Therefore, $ic5$ now becomes our new 'registration space'.

2.4.4.2 The Cost Function

Smoothed/subsampled CMs are formed for both the input (CM_I) and the reference (CM_R) subjects. The next step in the registration pipeline, as shown in Figure 2.10, is the derivation of the 'data-driven' part of the cost function which would measure (dis)similarities in CM_I and CM_R. We decided to quantify these similarities in smoothed SCFVs (further referred to as SCFVs; smoothing is assumed as a preprocessing step) using a 'correlation'[6] measure:

$$Corr(\mathrm{X},\mathrm{Y}) = 1 - \frac{(\mathrm{X}-\mathrm{Y})^T(\mathrm{X}-\mathrm{Y})}{\mathrm{X}^T\mathrm{X}+\mathrm{Y}^T\mathrm{Y}} = \frac{\mathrm{X}^T\mathrm{Y}}{(\mathrm{X}^T\mathrm{X}+\mathrm{Y}^T\mathrm{Y})/2}, \tag{2.3}$$

where X and Y are vectors, e.g. SCFVs. This measure is suitable as it is bound between 0 and 1 (as X and Y are always CM entries with no demeaning and hence always positive) and encodes the linear relationship between vectors (see the numerator). Importantly, it is sensitive to scaling - unlike traditional correlation, making it particularly suitable for comparing SCFVs where difference in scaling can indicate difference in 'connectivity strength'. Similarly motivated measures, such as

[5]Reduction in number of mesh vertices from approximately 10^5 (high definition FreeSurfer reconstruction) to approximately 10^4, i.e. $ic5$ mesh.

[6]Please note this is not a traditionally defined correlation, i.e. the vectors are not demeaned.

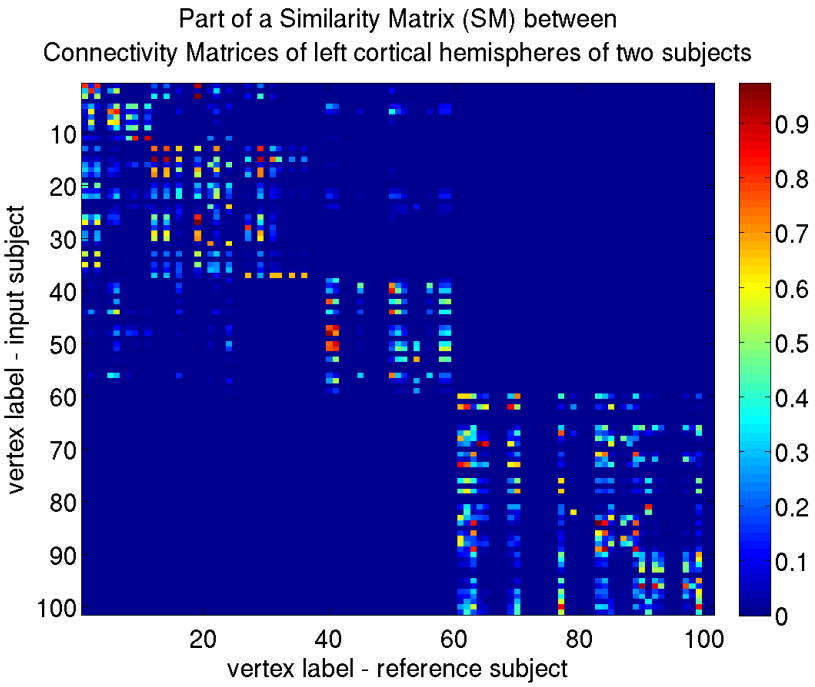

Figure 2.9: Part of a Similarity Matrix (Equation (2.4)) between smoothed Connectivity Matrices of left cortical hemispheres of the input and the reference subject. Colourbar represents similarity coefficient values. High values reflect high similarity in SCFVs, i.e. in white matter connectivity. High values along the main diagonal indicate good initial vertex-to-vertex alignment. The SM is precalculated before running registration.

ETA^2 have been widely used for comparison of functional connectivity descriptors (Cohen et al., 2008). Therefore, the *Similarity Matrix* between the input and the reference mesh (SM_{IR}) is given by

$$\mathrm{SM}_{IR}(i,j) = Corr\left(\mathrm{SCFV}_I(i), \mathrm{SCFV}_R(j)\right),\ i,j \in [1,N], \tag{2.4}$$

where $\mathrm{SCFV}_I(i)$ is the (smoothed) SCFV of the ith vertex from the input mesh, $\mathrm{SCFV}_R(j)$ is the (smoothed) SCFV of the jth vertex from the reference mesh, and N is the number of vertices in these meshes (number of vertices in $ic5$). SM_{IR} can be precomputed and depending on the dimensionality of the SCFVs this computation can be highly parallelised if necessary (e.g. in the case when SCFVs are 3D volumes - discussed in section 2.5.1 and further used in Chapter 4, section 4.2.1).

We can now write the 'data-driven' or the 'intensity-part' of the cost function

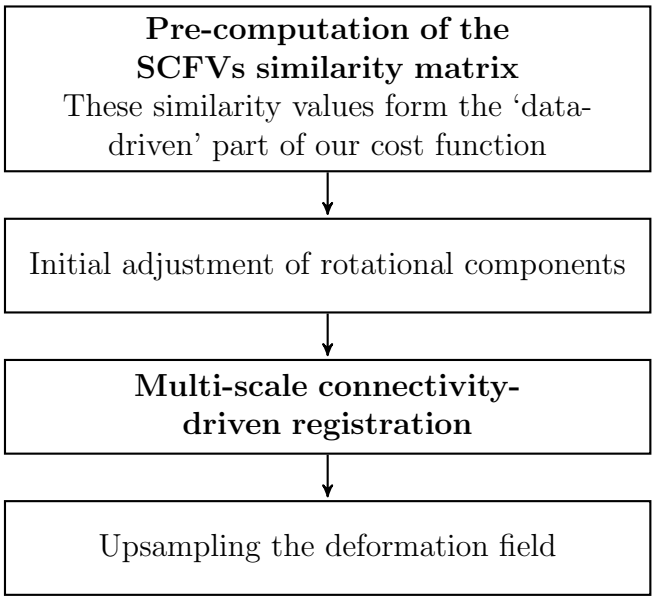

Figure 2.10: Main registration steps.

associated with vertex $i \in V_{M_I}$ where M_I is the input and M_R is the reference mesh (both equivalent in topology and number of vertices to $ic5$):

$$J_p(x_i, \mathrm{SM}_{IR}, \sigma) = const - \frac{\sum_{j \in V_{M_R}} d_\sigma(x_i, y_j) \mathrm{SM}^*_{IR}(x_i, y_j)}{\sum_{j \in V_{M_R}} d_\sigma(x_i, y_j)},\ i \in V_{M_I},\ j \in V_{M_R}. \quad (2.5)$$

$\mathrm{SM}^*_{IR}(x_i, y_j)$ is the interpolated value of the SM_{IR}. Interpolation of the similarity matrix is performed using local linear interpolation (see Appendix A). $d_\sigma(x_i, y_j)$ is a function of the geodesic distance $D(x_i, y_j)$. Assuming that M_I and M_R are overlapping spheres (note that M_R always remains $ic5$, i.e. regularly spaced and is considered a constant), $D(x_i, y_j)$ is the geodesic distance between points x_i and y_j, $i \in V_{M_I}$ and $j \in V_{M_R}$. In other words, $D(x_i, y_j)$ is the length of the great circle of the sphere (passing through x_i and y_j) between x_i and y_j. In our implementation $d_\sigma(x_i, y_j)$ is calculated as a bell function with a spread parameter σ:

$$d_\sigma(x_i, y_j) = e^{-\left(\frac{D(x_i, y_j)}{\sigma}\right)^2},\ i \in V_{M_I},\ j \in V_{M_R}. \quad (2.6)$$

Parameter σ is instrumental in implementing the multi-scale registration framework as further explained in section 2.4.4.4. *const* is a sufficiently large constant to ensure the positivity of $J_p(x_i, \mathrm{SM}_{IR}, \sigma)$.

Full Cost Function The 'data-driven' part of the cost function (Equation (2.5)) is balanced with the regularisation cost function J_d which tries to preserve original (regular) mesh spacing. Using proper regularisation is essential for escaping overfitting and degenerate solutions. Regularisation also helps reduce the effects of noise and lowers the likelihood of achieving deformations that would compromise the invertibility of the final deformation field. The full cost function associated with vertex $i \in V_{M_I}$ can therefore be written as

$$J(x_i, \mathrm{SM}_{IR}, \sigma) = \lambda J_p(x_i, \mathrm{SM}_{IR}, \sigma) + \mu J_d(x_i), \tag{2.7}$$

where λ and μ are scaling coefficients balancing contributions of different parts of the cost function. The redundancy in having two separate scaling parameters (instead of one) is helpful when realising the multi-scale optimisation in which the overall scaling needs to be re-adjusted for each step (see paragraph on Combined Force 2.4.4.3).

2.4.4.3 Regularisation vs. Intensity Force

Forces in the deformable model framework deform the input mesh so that the full cost function from Equation (2.7) is *minimised* for every vertex of the input mesh w.r.t. its position. In other words, forces induce deformations and are calculated as spatial derivatives of respective cost functions. All our deformations are constrained to lie on the sphere S^2 embedded in the 3D space[7].

Intensity Force If $w : S^2 \to S^2$ is the deformation mapping point $x_i,\ i \in V_{M_I}$ to its deformed point $x'_i = x_i + w(x_i)$, the intensity force (IF), which is a vector, is given by:

$$\mathrm{IF}(x_i, \mathrm{SM}_{IR}, \sigma) = \frac{\partial J_p(x'_i, \mathrm{SM}_{IR}, \sigma)}{\partial w(x_i)}. \tag{2.8}$$

This partial derivative is calculated numerically as described in section 2.4.4.5.

[7]In current implementation we use Cartesian coordinates for derivation of forces and constrain/re-project deformations to the unit sphere. This is possible as every deformation is kept small in relation to the unit radius.

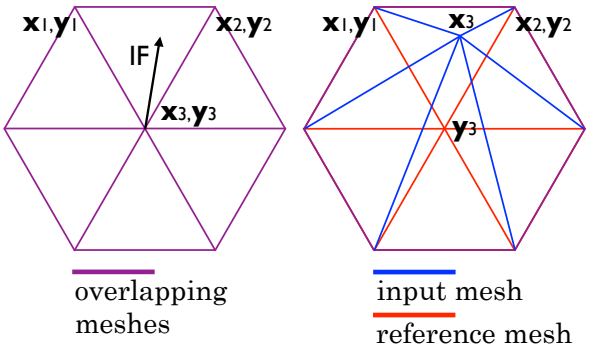

Figure 2.11: The effect of the Intensity Force (IF) on an input mesh vertex (x_i – in the input mesh (blue); y_i – in the reference mesh (red)). The deformation of x_3 due to the IF is shown. **(left)** Reference and input meshes initially overlap (depicted in purple). **(right)** IF displaces x_3 away from y_3 and towards y_2 as e.g. $\mathrm{SM}^*(x_3, y_2) > \mathrm{SM}^*(x_3, y_3)$. The regularisation force would have the opposite effect, bringing x_3 closer to y_3.

Regularisation Force Similarly, the Regularisation Force (RF) can be calculated as:

$$\mathrm{RF}(x_i) = \frac{\partial J_d(x_i')}{\partial w(x_i)}. \tag{2.9}$$

The form we have chosen to use for $J_d(x_i)$ is similar to the Thin Plate Spline energy functional introduced by Bookstein (1989) (appendix B elaborates on this similarity), but the force can also be directly written as:

$$\mathrm{RF}(x_i) = \frac{1}{n}\left(\sum_{\substack{1\le k\le n\\ k\in\mathcal{R}_n(i)}} w(x_k)\right) - w(x_i) + \underbrace{\frac{1}{n}\left(\sum_{\substack{1\le k\le n\\ k\in\mathcal{R}_n(i)}} x_k\right) - x_i}_{\approx 0}, \tag{2.10}$$

where $\mathcal{R}_n(i)$ is a set of n closest neighbours (vertices) of vertex $i \in V_{M_I}$ in V_{M_I}. The second term is close to zero for a regularly spaced mesh such as $ic5$. This regulariser has a direct geometrical interpretation: it brings the vertex closer to the position given by the mean of its neighbours (see Figure 2.11). Such regularisers have already been used in mesh regularisation problems such as in automatic brain segmentation (Smith, 2002).

Combined Force The intensity and the regularisation force are scaled, combined and applied to the associated vertex x'_i,

$$x'_i = x_i + \lambda \mathrm{IF}(x_i, \mathrm{SM}_{IR}, \sigma) + \mu \mathrm{RF}(x_i) \tag{2.11}$$

which is re-projected to the unit sphere. Scaling coefficients λ and μ are empirically determined through a series of tests on artificial data (details given in Chapter 4, section 4.2.2.1).

2.4.4.4 Multi-scale Framework

The aim of the multi-scale (or multi-resolution) framework is to achieve registration for different levels of data *smoothing*, and in that way progressively registering first prominent global features at the highest smoothness level and then detailed features (no smoothing) at the lowest smoothness level. Such multi-stage alignment was introduced to help escape local minima and achieve quicker convergence (Woods et al., 1993). Coarse-to-fine registration is here effectively achieved by re-running the registration algorithm for different levels of 'SM smoothness'.

The multi-resolution nature of the registration problem can already be seen in Equation (2.5). Parameter σ specifies the extent of the area in which the similarities in SCFVs are calculated. Therefore, for large σ (i.e. large, relative to the circumference of the great circle of the registration sphere), the vertex under deformation (in the input mesh) 'sees' all vertices of the reference mesh. Conversely, for small σ, the deformed vertex 'sees' just its immediate neighbourhood in the reference mesh. In this way, by varying σ from large to small we hope to achieve alignment of progressively smaller (more focal) connectivity features.

The first step in the multi-scale registration is to remove any rotational components (if any such exist) of the deformation between the input and the reference

step	σ	λ	μ
1	$0.6R_g$	240.0	0.3
2	$0.4R_g$	120.0	0.3
3	$0.2R_g$	72.0	0.4
4	$0.1R_g$	24.0	0.4
5	$0.0R_g$	2.0	0.5

Table 2.2: Multi-scale optimisation is conducted in 5 stages each of which is run until convergence with a specific set of σ, λ and μ parameter values. These parameters are optimised empirically through a set of tests on artificially generated data (section 2.4.4.5 as well as Chapter 4, section 4.2.2.1). R_g is the circumference of the great circle of the unit sphere and is equal to 2π.

mesh. This is done by maximising

$$J_{p-rot}(\mathrm{SM}_{IR}, \alpha, \beta, \gamma) = \sum_{i \in V_{M_I}} \sum_{j \in V_{M_R}} \mathrm{SM}^*_{IR}(\mathrm{R}(\alpha, \beta, \gamma)x_i, y_j),\ i \in V_{M_I},\ j \in V_{M_R}, \quad (2.12)$$

w.r.t. the set of Euler angles (α, β, γ). R is the rotation of the input mesh parametrised by (α, β, γ). Optimisation is performed using a simple gradient descent algorithm and its result initialises further registration.

The next step is a multi-stage optimisation with varying σ. In five distinct optimisation steps, each performed until convergence, σ, λ and μ are varied as shown in Table 2.2. For $\sigma = 0.0$ the distance function d (see Equation (2.6)) is set to 0 everywhere except for the neighbouring vertices[8] of the vertex under deformation, where it is set to 1. Therefore, the neighbouring vertices participating in the interpolation of the SM are weighted equally (the interpolation of SM effectively assumes the role of distance weighting). The way parameters σ, λ and μ are determined is introduced in the next section 2.4.4.5 and further detailed in Chapter 4, section 4.2.2.1.

2.4.4.5 Optimisation

The optimisation algorithm tries to find the minimum of the joint cost function (Equation (2.7)) by iteratively applying the intensity and the regularisation forces

[8]Here, the neighbouring vertices are the vertices of the triangular mesh element (of the reference sphere) to which the deformed vertex of the input mesh belongs.

together. The optimisation problem is decoupled since every vertex is moved independently from the others, i.e. forces are calculated for and applied to every vertex separately. In the current implementation, the combined force is calculated by separately calculating the gradient of J_p and J_d (as given in Equation (2.5) and appendix B respectively) and then applying the combined force (Equation (2.11)).

Spatial derivatives of J_p and J_d are found for each vertex in the local tangent plane (of that vertex) to the sphere. In this way, the 3D optimisation problem is reduced to the 2D one. J_p and J_d are sampled along the tangent plane axes and the derivatives are computed using these measurements. A similar procedure has been used for the analysis of cortical thickness and cortical convexity-driven registration (Fischl et al., 1999a). However, this iterative minimisation does not guarantee convergence to the global minimum, but has been shown to perform well (in terms of convergence properties) on cortical surface registration tasks and has been widely used throughout FreeSurfer (Fischl et al., 1999a). Forces are first calculated in the tangent plane of every vertex and scaled into displacements. Once the displacements for all vertices have been calculated, they are applied to the vertices of the mesh in random order After being mapped to a new position, each vertex is re-projected to the unit sphere.

During every deformation step we check for local mesh self-intersections whose existence would change mesh topology and compromise invertibility of the final deformation. Therefore, if a local mesh self-intersection is detected after the application of a displacement vector, such a movement is not performed, or is performed with a scaled displacement vector (which does not produce changes in mesh topology, but still minimises the cost function). By checking for mesh self-intersections iteratively with the application of deformations to every vertex, we eliminate the need for using computationally expensive postprocessing topology correction procedures.

Many of the mesh processing techniques, such as calculation of the displacement

forces, require tracking of vertex-neighbour relationships between the input and the reference meshes. In other words, when the input and the reference mesh overlap (as in this case), we need to track for every vertex of the input mesh (which is being deformed) what are its n closest neighbours *in the reference mesh*. Depending on the application, n can vary from 1 to the number of vertices in the mesh. These relationships are stored in a lookup table to speed up computations and are updated every 10 iterations (which is typically much fewer than required for changing neighbour relationships). In our implementation, the lists of neighbours are sorted using *Quick-sort* algorithm (Hoare, 1962).

Finally, we should mention that before the multi-scale registration commences (section 2.4.4.4), the input and the reference meshes need to be *relaxed*. This entails the application of the regularisation force only until both $ic5$ meshes are brought into a new stable state. Through this step we address the problem of $ic5$ spacing inhomogeneity: most of the mesh vertices have 6 neighbours, whereas some of them have 5. If no mesh relaxation is performed, those vertices with 5 neighbours would experience a slightly different regularisation force compared to those with 6 due to slightly different mesh properties. The relaxation aims to produce a constant (i.e. zero) regularisation force over the undeformed meshes.

Free Parameters The free parameters of our model are σ, λ and μ (see previous section), the number of scale steps and the number of iterations. All of these parameters were first determined in a subcortical registration framework through the analysis of artificial data/tests (Chapter 4, section 4.2.2.1). These tests quantified registration consistency and robustness and values of σ, λ and μ were initially determined to maximise these measures. Since free parameters are decoupled from mesh spacing/resolution, they were readily applied to the geometry of $ic5$. Fine tuning (required due to different properties of the Similarity Matrix - e.g. it encodes

A
input

B = Ψ(A)
reference

$\mathcal{W}(A)$
A warped to B using FACS

Figure 2.12: **(left)** The white/grey matter surface of the input subject (A) in registration space subsampled to *ic*5 and parcellated into 35 parcels (depicted in different colours) using FreeSurfer cortical parcellation algorithm. **(middle)** The reference mesh (B) generated by applying a warp field Ψ (taken from the cortical geometry-driven registration of two different control subjects) to A. The shape of parcels changes due to warping. **(right)** A after being registered to B using FACS algorithm ($\mathcal{W}$(A)). The correspondence in parcel labels between B and $\mathcal{W}$(A) is higher than between B and A.

cortical connectivity, not that of subcortical structures as in Chapter 4) was done on a test constructed using real data.

The test aimed to measure how well we can restore the (known) mapping between two surfaces using the *full connectivity information* as proposed in this Chapter. The data of one subject's (A) left cortical hemisphere was processed as indicated in Figure 2.4 (preprocessing) and the Connectivity Matrix and the corresponding white/grey matter surface were obtained. Figure 2.12 (left) shows the white/grey matter surface in registration space subsampled to *ic*5 and parcellated into 35 parcels using FreeSurfer cortical parcellation algorithm (Fischl et al., 2004). This subject/mesh plays the role of the *input subject* for this test. Subject B (*the reference subject*) was generated from A (high-density registration sphere) by applying a warp field Ψ (taken from the cortical geometry-driven registration of two different control subjects) to A and subsampling the resulting mesh and the corresponding Connectivity Matrix to *ic*5 as described in section 2.4.4.1. The input mesh after

warping is shown in Figure 2.12 (middle). Therefore, the CM associated with B (CM_B) was generated from the CM associated with A (CM_A), and the implicit correspondence between the two is given by warp Ψ. A was then registered to B using FACS following the methodology previously described (i.e. re-parametrisation, SCFV smoothing, multi-scale registration, etc).

After generating mapping $\mathcal{W}$: A $\rightarrow$ B using FACS (Figure 2.12 (right)) we assessed the similarity between Ψ and $\mathcal{W}$ by measuring:

1. Vertex-to-vertex similarities in corresponding SCFVs between B and $\mathcal{W}$(A) using Equation 2.3;

2. Percentage of overlap in corresponding cortical parcels (Figure 2.12) and $\mathcal{W}$(A).

The highest similarity/overlap was achieved for the set of parameters given in Table 2.2:

1. Mean vertex-to-vertex similarity in corresponding SCFVs between A and B before FACS was 0.38. After removing the rotational component as explained in section 2.4.4.4, mean similarity grew to 0.65 and after application of the multi-scale part of the non-linear registration settled at 0.78. In the ideal case this value should be close to 1.0. However, CM_A and CM_B do not have perfect one-to-one correspondence as both have been obtained by subsampling the high-density CM before (A) and after (B) applying warp Ψ. Therefore, even in the case of an ideal matching, the mean similarity would be lower than 1.0. Moreover, due to the nature of the test, FACS was not initialised by FreeSurfer alignment, significantly increasing the complexity of the warp Ψ. This, in turn, decreases the likelihood of finding the global optimum;

2. Percentage of overlap in cortical parcels between A and B before FACS was 0.72 and grew to 0.93 after application of all FACS steps. Note that the registration

is driven using a different set of data compared to the one used for cortical parcellation.

The number of iterations to achieve convergence at each step was 20 except at the last stage ($\sigma = 0.0$) when 30 iterations were needed. Figure 2.13 shows mean J_p, J_d and and $\lambda J_p + \mu J_d$ cost functions (averaged over the whole sphere) for this registration problem. The mean value of the regularisation cost function grows as the mesh undergoes greater deformation, whereas the intensity cost function always decreases as expected. Convergence is achieved when the mesh stops to change configuration (measured as the sum of all vertex displacements) due to further application of forces, i.e. enters a stable state. The full (compound) cost function is minimised at each of 5 stages.

Challenges Although conceptually simple, our optimisation approach suffers from the need to empirically determine free parameters. Moreover, systematic assessment of its robustness to model parameters is challenging. Chapter 4, section 4.2.2.1 presents further tests and examines the effect of noise in the Similarity Matrix on registration results. Given the similar nature of the two registration problems (subcortical and cortical), these tests were not all repeated for the cortical framework. Additionally, the time consuming parameter optimisation and artificial example construction is done more efficiently on smaller scale problems (smaller meshes and smaller Connectivity Matrices) in the subcortical framework.

2.4.4.6 The Deformation Field

After performing registration, the deformation of the input mesh is encoded as a deformation (or a vector) field. Each vertex of the original input mesh (relaxed $ic5$) is associated with a vector encoding its displacement. The deformation field therefore maps $S^2 \rightarrow S^2$. However, this deformation might need to be applied to

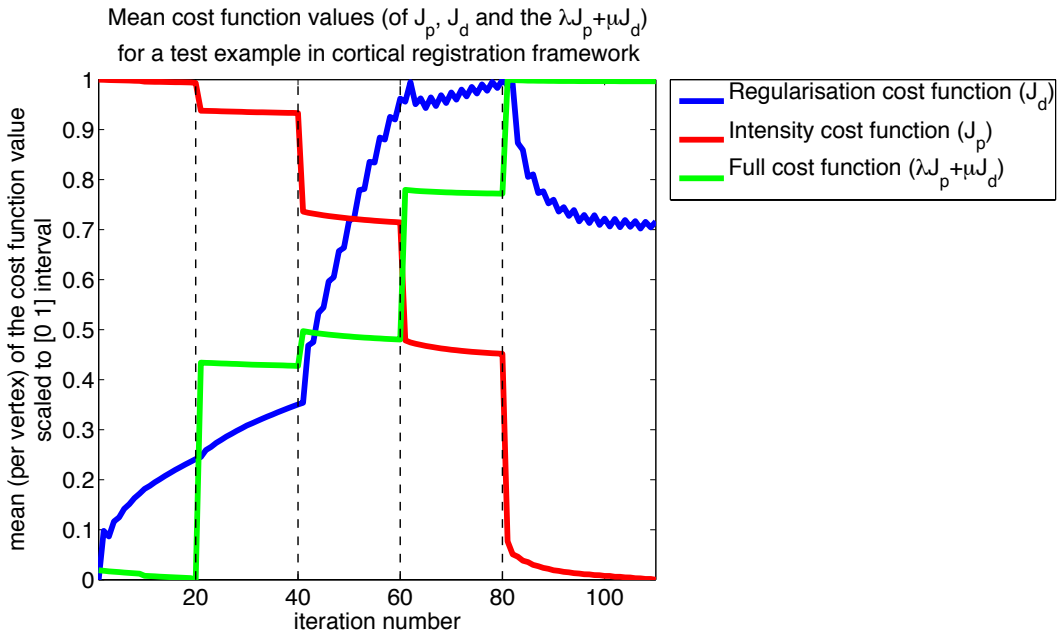

Figure 2.13: Convergence properties of J_p, J_d and $\lambda J_p + \mu J_d$ cost functions for the registration problem presented in Figure 2.12. The mean value (across space) of the regularisation cost function grows as the mesh undergoes greater deformation (divergence from the regularly spaced mesh), whereas the intensity cost function always decreases as expected. The full (compound) cost function is minimised at each stage. Stable, or constant, periods in the cost function plots indicate that the regularisation and the intensity force cancel out and the stable deformation is reached. Vertical dashed lines indicate transitions between different stages in the multi-scale registration.

another subject, or we might want to upsample it or even apply it to the original white/grey matter surface in the native space.

All these tasks can be solved by upsampling and projecting the deformation field and both of these rely on efficient deformation field interpolation. To achieve this we used two methods capable of locally interpolating displacement vectors:

1. Radial Basis Functions (in our case Thin Plate Splines (Bookstein, 1989)) can be used to locally interpolate every component (i.e. x, y and z coordinates) of the displacement field;

2. Locally affine transforms can be calculated between corresponding *groups* of vertices. Details of the procedure are given in appendix C.

Both of these methods were used to the same effect, but the latter one produced fewer mesh self-intersections as the result of interpolation and was adopted throughout FACS.

2.5 Discussion

2.5.1 Formation of SCFVs

2.5.1.1 High- vs. Low-dimensional SCFVs

The construction of targets for probabilistic tractography can vary between two extremes: one, where just main brain lobes are labelled as targets, and the other, where every brain voxel is considered 'a target'. The latter is referred to as whole-brain tractography and effectively no target definition is needed - connectivity between the seed voxel and every other voxel in the brain is quantified. Each of these two extremes has its positive and negative sides when it comes to potential usage for SCFVs construction.

In the first case, brain lobes can be identified with high consistency across subjects and the SCFVs would therefore be meaningfully and accurately comparable. On the other hand, low dimensionality of such connectivity descriptors might limit their descriptive power and thus the registration specificity. In the second case, the descriptive power of high-dimensional SCFVs could be greater (due to richer connectivity information), but the success of inter-subject comparison would be directly affected by the initial volumetric voxel-to-voxel registration quality to a standard space. In other words, in order to compare these 3D SCFVs effectively across subjects, one would need to transform and compare them in a common space. However, the volumetric registration step can often result in misalignments of cortical folds (if simple affine registration methods are used). Furthermore, having volumetric SCFVs imposes serious challenges on storage and memory. For example, for a mesh with ~10^4 vertices we would have ~10^4 volumetric images to store and eventually

compare. Finally, having in mind the ultimate purpose of SCFVs, extra care has to be taken to ensure that calculation of similarities between SCFVs is computationally feasible since the number of comparisons squares with the number of SCFVs/vertices (details on problem dimensionality were discussed in section 2.4.3.2). Ideally, we would like to find the balance between SCFVs' descriptive power, the resolution of their (volumetric) representation and the computational speed of comparing them.

The final utility of SCFVs is measured in terms of performance of the connectivity-driven registration algorithm. In this sense, optimising for SCFV dimensionality can be seen as an additional optimisation step to be performed once the registration-validation pipeline has been established. Because this thesis is primarily concerned with introducing and validating connectivity-driven registration framework, this supplementary optimisation step is left for future research. The registration-validation pipeline, however, works in the same way even with high-dimensional features.

2.5.1.2 Current Choices

Still, choices of SCFVs we made fall in between the two extremes mentioned earlier. The number of parcels into which a cortical hemisphere is commonly parcellated is between 34-85 (Desikan et al., 2006; Fischl et al., 2004) (although the number of functionally-specific regions into which the cortex can be subdivided is estimated to be between 100 and 200 (Werner and Chalupa, 2004))[9]. We use a similar number of parcels (though not following anatomical borders) which eventually become tractography targets. Moreover, these parcels are expected to be determined with greater inter-subject consistency than, for example, when using 3D connectivity descriptors. This comes from the very nature of surface-based registration and parcellation algorithm implemented in FreeSurfer (Fischl et al.,

[9]This number is rather application specific: for the purposes of identifying the structural core of the human cerebral cortex, some authors, like Hagmann et al. (2008), map structural connectivity among about 1000 cortical ROIs.

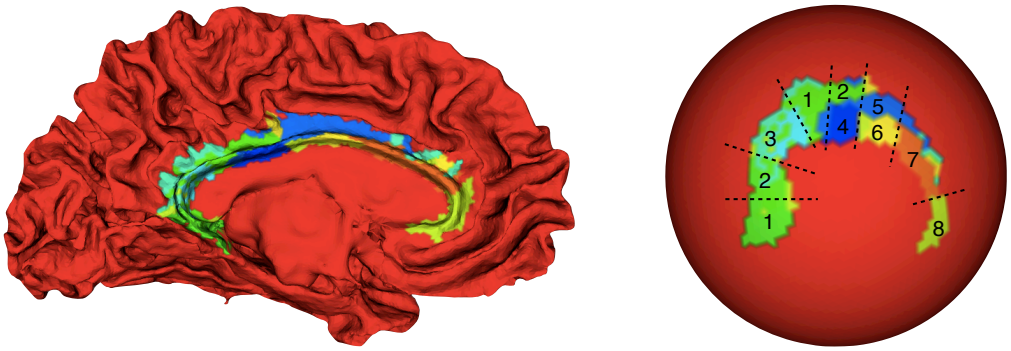

Clustering in the native space
left-posterior, right-anterior

Equivalent clustering
in the registration space (sphere)

Figure 2.14: **(left)** Clustering of the cingulate area using SCFVs (9 clusters) shown in the native space (white/grey matter surface) of the template subject. Different colours correspond to different clusters. Parcellation was performed using K-means algorithm across 11 control subjects and the results were combined to produce the 'prototype clustering' presented here (for details on the clustering procedure see Chapter 3, section 3.2.4.1). **(right)** Image from left shown in the registration space. Dashed black lines roughly indicate some of the boundaries between clusters (8 different clusters can be discerned in the image - associated with numbers). Most of the boundaries run in the dorsal-ventral direction which is consistent with the functional and cytoarchitectonic division of the cingulate area (Vogt, 2009).

1999b). We also include a number of subcortical targets which are segmented for every subject separately increasing the labelling consistency over simple standard-space-based segmentations (Fischl et al., 2002). In this way every entry in the 89-dimensional SCFV is associated with a well defined brain region which augments the interpretability of registration results. For example, one could determine which structures'/parcels' connectivity is driving registration differences between two groups (Chapter 3, section 3.4.2). Furthermore, 89-dimensional features can be efficiently compared and stored for a large number of vertices.

A side-result of the validation framework was to show that current choices of SCFVs relate to some of the connectivity-based parcellations previously performed just with high-dimensional connectivity features. Figure 2.14 presents the cingulate clustering using our SCFVs. The cingulate was clustered in 9 clusters as proposed by (Beckmann et al., 2009) and this clustering somewhat resembles histological findings

(Vogt, 2009).

Finally, in Chapter 3, in the section on Connectivity Profiles Template (section 3.4.1) we show that the SCFVs describing the whole cortex (all SCFVs of a single hemisphere grouped into a matrix) can be decomposed using Independent Component Analysis (ICA) into maps that resemble those of well-defined brain structures/functional areas such as the hippocampus, the cingulate gyrus, the frontal and occipital poles, etc. This finding indicates that SCFVs indeed have certain potential to discriminate between brain regions on the basis of their connectivity and also gives an insight into the spatial scales at which these distinct functional/structural regions exist.

Still, many open questions remain:

1. At which spatial scales can we *a priori* expect improvements in alignment of functional homologues across subjects?

2. What effects does triangular parcellation have on final registration conclusions?

3. What effects does tract thresholding (by length) have on connectivity-driven registration applicability?

4. What are the relative merits of cortical and subcortical parts of SCFVs in relation to the registration outcome?

While Chapter 3, which is devoted to the validation of our registration pipeline, sheds some light on the first of these questions, the others are left for the future work to investigate.

2.5.1.3 Other Methods for SCFV construction

It has been recently noted (Oguz et al., 2009) that the construction of the SCFVs can be facilitated through white matter surface smoothing procedures. Oguz et al. (2009) encountered a problem similar to the 'triangular SCFV clustering' explained above.

Their solution was to smooth and erode the white matter surface and use connectivity information related to that surface to drive registration. This also helps reduce the effects of differing tractography performance in highly convoluted cortical regions. Although an elegant solution, the smoothing and erosion steps affect the one-to-one mapping between the original white matter surface (which is being registered) and the eroded surface introducing possible registration biases. Other important questions relating to the construction, usage and within-registration-framework update of the SCFVs are discussed in the next section 2.5.2.

2.5.2 Registration Algorithm

The registration algorithm presented here falls in the category of many similar iterative mesh deformation procedures. It is well-suited for the formulation of connectivity-driven registration problem in that the similarity/correlation-based cost function was readily implemented. The registration run-time (around 30 minutes on an Intel Core 2 Duo 2.4GHz processor) is negligible compared to the time required for preprocessing such as running FreeSurfer and probabilistic tractography (around 30 hours on the same architecture).

Nevertheless, we rely on soft regularisation constraints to ensure deformation invertibility. Some other registration algorithms such as Diffeomorphic Spherical Demons (Yeo et al., 2010a) impose diffeomorphic mapping by construction and could perhaps help lower the number of free parameters as well. We repeated the test presented in section 2.4.4.5 using adapted Spherical Demons with results similar to the ones obtained using our method (i.e. the percentage of overlap in corresponding parcels after registration increased from 72% to 83% - in case of FACS, the increase was from 72% to 93%). Note that the cost function used was the squared difference between SCFVs (not SCFV similarity as in case of FACS) and no fine tuning of registration parameters was performed. In this sense, adapting Spherical Demons

to work with the similarity-based cost functions might yield an improvement in registration performance (and reduce the registration about time 10-fold).

One of the greatest challenges of the current framework is the dependence between construction of the SCFVs and the registration itself. Inter-subject correspondence of Cortical Targets for probabilistic tractography is imposed using the initial FreeSurfer alignment of cortical surfaces. However, it is the improvement over that registration that we aim to achieve. Therefore, ideally, with every step of our registration algorithm we should re-compute cortical targets and Connectivity Matrices. Unfortunately, processing time required for such evaluation is prohibitive. On the other hand, we could see the current framework as the *first iteration* where after every registration step the set of cortical targets would be recomputed using updated surface correspondence. Also, recent interest in registration driven using functional connectivity has shed some new light on this problem and there are promising solutions that propose efficient update techniques that could be applied to Connectivity Matrices in our case (Conroy et al., 2009; Cathier and Mangin, 2006).

2.5.2.1 Other Registration Methods

We have also tried other registration/optimisation paradigms to achieve connectivity-driven alignment. These always used the same SCFV similarity measure to form the intensity cost function, but differed in the way the deformation field is encoded and the multi-scale optimisation performed. In particular, apart from the algorithm explained in this Chapter (and small variations thereof) we used a method which encodes the deformation field using spherical wavelets as basis functions (details given in (Petrovic et al., 2009) and Chapter 4 section 4.2.3). It also achieved multi-scale registration and regularisation in a very different way compared to the one explained here. However, all these variations in algorithm design, although differing in execution time and convergence properties, produce similar results in tests we devised. This supports the proof of concept of connectivity-driven registration by indicating that

the registration results are largely unaffected by the actual algorithms used. Still, the method presented here proved in these tests to be somewhat superior to the other ones examined, especially in its convergence properties and the execution time.

"*Seize the moments of happiness, love and be loved! That is the only reality in the world, all else is folly. It is the one thing we are interested in here.*"

– Leo Tolstoy, War and Peace, 1865

Chapter 3

Connectivity Driven Registration of Cortical Surfaces - Validation and Applications

The aim of connectivity-driven registration is to align cortical surfaces of different individuals based on either structural or functional connectivity. If functional or structural connectivity are good predictors of functional segregation, then connectivity-driven registration could be expected to improve *functional correspondence*. The term of functional correspondence is explained by Van Essen and Dierker (2007): "Cortical locations in different individuals are in functional correspondence if they are part of the same cortical area and are matched in terms of whatever is mapped within that area." We move away from using gross anatomical landmarks, or cortical surface geometry, as predictors of functional specialisation and employ brain connectivity instead. If structural connectivity is a better predictor of functional segregation than e.g. the convexity of the cortical surface, one would expect an improvement in achieving *functional correspondence* after (structural) connectivity-driven registration. Consequently, the goal of validation should be to test exactly that. In this sense, we use functional MRI to assess the registration driven by structural/diffusion-weighted (DW) data. Section 3.1 introduces the validation concept used throughout this work. Sections 3.2 and 3.3 test functional

correspondence after FACS and FreeSurfer for a number of cortical areas using resting-state fMRI.

In the rest of the Chapter we briefly explore possible applications of FACS. First (section 3.4.1), the template of structural connectivity feature vectors is introduced (SCFV Template) and then used (section 3.4.2) to analyse a dataset comprising early-stage Alzheimer's disease patients and healthy controls. We demonstrate how differences in the FACS deformation fields of the two groups (registered to the SCFV Template) can help elucidate differences in cortical connectivity between them. Finally, the Chapter ends with a discussion (section 3.5) about validation methods.

3.1 Validation Using Resting-state fMRI

An ideal way to perform validation would be to compare the overlap of a set of corresponding functional areas (e.g. the primary visual area, the supplementary motor area, etc) in different individuals before and after FACS. These areas can be identified using *functional localisers* which are fMRI paradigms designed to elicit response exclusively in one functional area. For example, a visual motion experiment can be used to identify the middle temporal (MT) area in different subjects. The measure of the inter-subject overlap of this area after the alignment through different registration paradigms, can be used to assess the registration quality[1]. However, this type of validation depends on the set of functional localisers, i.e., is limited to the restricted areas of grey matter covered by any given suite of task paradigms. Alternatively, recent advances in using resting-state fMRI for functional parcellation have made it possible to perform similar validation while using the entire grey matter (as all grey matter areas relate to one or more 'resting state networks'). Resting-state fMRI (RS fMRI) is a functional imaging method where the subject is imaged while 'resting', i.e. is not instructed to perform any task (e.g. such as finger tapping or backwards

[1]Here we assume that the registration quality can be assessed by measuring functional correspondence.

counting) in contrast to paradigm-driven fMRI. We show how RS fMRI data can be used to parcellate the medial frontal cortex (MFC) into the supplementary motor area (SMA) and the pre-supplementary motor area (pre-SMA). Then, these parcellations are used to test whether FACS yields an improvement over FreeSurfer in aligning the SMA/pre-SMA boundary across subjects.

3.1.1 Cortical Parcellation Using Resting-state fMRI

In Chapter 2, section 2.4.1.3 we described how functional cortical parcellation can be performed using structural connectivity. Resting-state fMRI has the same potential. In the study of Kim et al. (2010), *functional connectivity* (FC) is used to define the SMA and pre-SMA subregions in human MFC. Resting-state functional connectivity (RS FC) is a method for examining regional interactions that occur when a subject is not performing an explicit functional task (Cohen et al., 2008; Margulies et al., 2007; Beckmann et al., 2005; Fox et al., 2005). A common way of quantifying RS FC is by correlating preprocessed RS fMRI timeseries in the voxels of interest and producing spatial maps of these correlations. However, before applying any of these procedures, RS fMRI has to be preprocessed to reduce the noise level. The preprocessing step typically consists of motion correction, temporal filtering, smoothing and regressing signals of no interest (such as the fMRI signal in the ventricles) (Kim et al., 2010).

In order to achieve RS FC-based parcellation, one can compare how RS FC of one voxel/vertex (to the rest of the brain or to a certain brain region) differs from that of the other. Therefore, if we label the vector describing the RS FC of one voxel to every other voxel in a cortical target as the *RS FC Profile* (or just a FC Profile) of that voxel, then the discrimination of voxels in a certain region on the basis of similarities/differences in RS FC Profiles is possible (section 3.2.2).

It has also been noted that the boundaries between functionally distinct regions can be revealed by examining the *gradients* of the images formed by correlating FC

Profiles. This approach for whole-brain functional parcellation was pioneered by (Cohen et al., 2008) and is also used in this work (section 3.3).

3.2 SMA/pre-SMA Boundary Alignment

The initial method, based on the SMA/pre-SMA parcellation was performed in two ways: that of Kim et al. (2010), resulted in poor inter-subject clustering consistency and is described in section 3.2.2. The other method we developed uses a set of Resting State Networks (RSNs) (Beckmann et al., 2005) to parcellate the MFC and is described in detail in section 3.2.3. The latter approach yielded an improvement in the inter-subject clustering consistency and produced fewer artifactual clusterings (Figure 3.7) than that of (Kim et al., 2010). Parcellations obtained by this approach were used to compare performance of FACS to FreeSurfer cortical alignment (section 3.2.4).

3.2.1 Data Acquisition

The data were kindly provided by Dr Valentina Tomassini, FMRIB Centre, Oxford and were originally collected as part of the Multiple Sclerosis-related study.

Diffusion-weighted Data Diffusion-weighted and T1-weighted images were acquired in 11 healthy subjects (9 women and 2 men; $mean \pm SD$ age, 43.1 ± 2.6) on a 1.5 T Siemens (Erlangen, Germany) Sonata MR scanner. All subjects gave informed written consent in accordance with ethical approval from the Oxford Research Ethics Committee. Diffusion-weighted data were acquired using echo planar imaging (72 2 mm-thick axial slices; matrix size, 128×104; field of view, 256×208 mm^2; giving a voxel size of $2 \times 2 \times 2$ mm). Diffusion weighting was isotropically distributed along 60 directions using a b value of 1000 s/mm^2. For each set of diffusion-weighted data, five volumes with no diffusion weighting were acquired at points throughout the acquisition. Two sets of diffusion-weighted data were acquired for subsequent

averaging to improve the signal-to-noise ratio. The total scan time for the diffusion-weighted imaging protocol was 45 min.

RS fMRI and T1-weighted Data During the 11 min scan, participants lay supine in a 1.5 T Siemens Sonata MR scanner. They were instructed to close their eyes and lie still. Cushions were used to reduce head motion. Whole-brain blood oxygen level-dependent (BOLD) fMRI data sets were collected, using the following parameters: 45 axial slices, in-plane resolution 3×3 mm, slice thickness 3 mm, no gap, repetition time = 3400 ms, echo time = 41 ms, 200 volumes. A structural scan was acquired for each participant, in the same session, using a T1-weighted 3D FLASH sequence (repetition time = 12 ms, echo time = 5.65 ms, and flip angle = 19°, with elliptical sampling of k-space, giving a voxel size of $1.3 \times 1.3 \times 1.0$ mm in 5 min and 5 s). Field maps were acquired for both the DW and the RS fMRI data (TE1 = 5.19 , TE2 = 9.95 ms, voxel size $2.3 \times 2.3 \times 2.5$ mm).

3.2.2 SMA/pre-SMA Parcellation Using RS FC

We utilised the idea of Kim et al. (2010) to parcellate the MFC into SMA and pre-SMA using RS FC. The clustering procedure was applied to 11 healthy control subjects (data acquisition described in the previous section). However, in this analysis, we performed slightly different preprocessing steps to those of Kim et al. (2010) and also customised the method to work with cortical surfaces (in contrast to 3D volumes).

Data Preprocessing Preprocessing used in this work was similar to that described by Beckmann et al. (2005). In other words, the following steps were applied: motion correction (Jenkinson et al., 2002), removal of nonbrain structures from the echo planar imaging volumes (Smith, 2002), B0-field distortion correction using FUGUE/FSL, spatial smoothing by a Gaussian kernel of 8 mm FWHM, mean-based intensity normalisation of all volumes by the same factor (i.e., 4D grand-mean),

high-pass temporal filtering (Gaussian-weighted least-squares straight line fitting) (FWHM = 148 s), and Gaussian low-pass temporal filtering (FWHM = 5.6 s). Preprocessed and B0-distortion corrected functional images were co-registered to their structural counterparts using FLIRT/FSL (Jenkinson et al., 2002). Each subject was fed into the FACS preprocessing pipeline (see Figure 2.4), obtaining the FreeSurfer reconstructions of the grey/white matter and the pial surface boundaries, as well. FreeSurfer registration was performed to bring all cortical surfaces into cortical geometry-based alignment.

3.2.2.1 Calculating Functional Connectivity Profiles for the MFC

We would like to perform MFC parcellation for the *grey/white matter surface* of each subject. Figure 3.1b) shows the MFC mask that was drawn in the MNI space (average T1-weighted brain image constructed from 152 healthy subjects at the Montreal Neurological Institute, Montreal, QC, Canada), mapped to the grey/white matter surface of one subject (affine registration). Therefore, each vertex of the grey/white matter surface/mesh belonging to this mask should be associated with a FC Profile. FC Profiles could then be clustered into putative SMA and pre-SMA. Figure 3.2 - right shows the SMA/pre-SMA clusters drawn using anatomical landmarks. One would expect that the clustered FC Profiles would result in similar parcellations.

To start forming FC Profiles of every vertex in the MFC mask, we first need to associate a RS timeseries with all vertices in the grey/white matter mesh. This assumes that the FC will be calculated between each vertex in the MFC mask and every other vertex of the mesh. Sampling of the preprocessed RS fMRI is done in the middle of the cortical grey matter sheet: for each vertex of the grey/white matter mesh, sampling is performed half-way along the line connecting two corresponding points of the pial and the grey/white matter surfaces. These surfaces, as well as the point-to-point correspondence between them, is established during FreeSurfer reconstruction. In this way, each vertex i of the grey/white matter mesh is associated

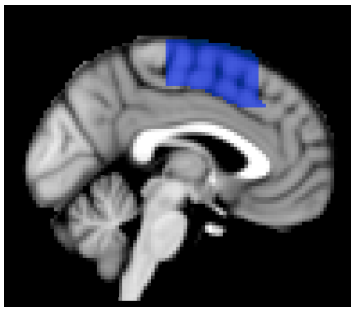

(a) The MFC mask in the MNI space

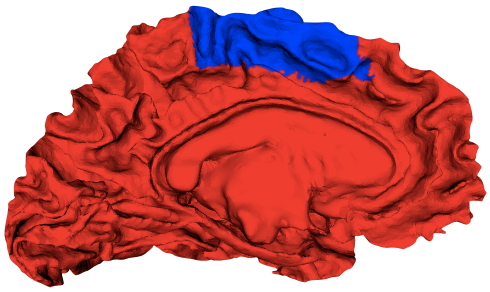

(b) The MFC mask from (a) projected to the grey/white matter surface of one subject

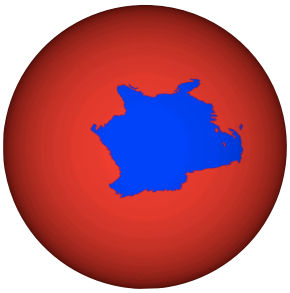

(c) The MFC mask from (b) on a high-resolution registration sphere

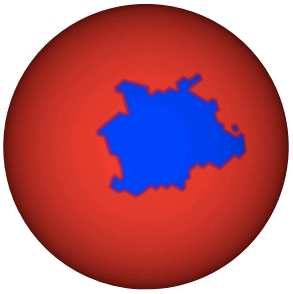

(d) The MFC mask from (c) on an $ic5$ registration sphere

Figure 3.1: **(a)** The MFC mask (blue) in the MNI space. The mid-sagittal slice of the left cortical hemisphere is shown. The same manually-drawn mask was used as in Johansen-Berg et al. (2004). **(b)** The MFC mask from (a) mapped to one subject's grey/white matter surface using affine registration. The sagittal view of the left cortical hemisphere is shown. The right side of the image corresponds to the anterior of the brain. **(c)** The surface/mesh from (b) in the FreeSurfer registration space after cortical geometry-driven registration to the template. **(d)** The surface/mesh from (c) after subsampling to the $ic5$. The MFC area is typically about 300 vertices in size.

with an n-dimensional vector $L(i,t)$ encoding RS fMRI signal at n timepoints. Thus, for each subject's cortical hemisphere (here we consider just left hemispheres), we have an $N \times n$ matrix of fMRI samples, where N is the number of mesh vertices. As high-density meshes have $N \approx 10^5$, the dimensionality of the problem is reduced by subsampling. Figure 3.1c) and d) show the effects of subsampling the high-density mesh to the $ic5$. Note the similarity of this step to subsampling the connectivity matrix in Chapter 2 section 2.4.3.2. After subsampling, a matrix of fMRI samples has the dimension $10242 \times n$.

We can now proceed to form FC Profiles for all vertices of the MFC mask. As suggested by Kim et al. (2010), FC between vertices i and j can be quantified with the Pearson correlation coefficient $r(i,j)$:

$$r(i,j) = \frac{\frac{1}{n}\sum_{t=1}^{n}\left(L(i,t) - \bar{L}(i)\right)\left(L(j,t) - \bar{L}(j)\right)}{S_L(i)S_L(j)}, \tag{3.1}$$

where $\bar{L}(.)$ and $S_L(.)$ are the mean and standard deviation of $L(.,t)$ respectively. n is the total number of time points. Finally, Fisher's Z-transform is applied to each correlation value (Kim et al., 2010). In this way a 10242-long FC Profile vector is associated with every vertex of the MFC mask. All such vectors can be assembled into a matrix $\mathrm{Z}_{N_{MFC}\times 10242}$ where N_{MFC} is the number of vertices in the MFC mask (see Figure 3.1d). The degree of similarity between FC Profiles of the MFC mask vertices is quantified through the *FC Profile Similarity Matrix* (S) which is calculated as $\mathrm{S} = \mathrm{ZZ}'$. Each entry of matrix $\mathrm{S}(i,j)$ encodes the correlation of FC Profiles of vertices i and j. This matrix becomes the input into the clustering algorithm (K-means, 2 clusters)[2].

Clustering Results The clustering procedure described so far was applied to all 11 subjects. Clustering results are summarised in Figure 3.3. 5 out of 11 subjects exhibit artifactual clustering separating the clusters in the superior-inferior instead of anterior-posterior direction (e.g. bottom right). We were unable to attribute the superior-inferior clustering pattern to any plausible functional separation and, therefore, attributed it to either imaging or processing artefacts. A recent study by Jo et al. (2010) suggests that such artefacts could be RF coil-related and that they can induce spurious correlations especially affecting assessment of the seed-based FC. They showed that the RF amplifiers or receive coil arrays can induce local image intensity shifts that can markedly bias FC calculations. We repeated the

[2]Note that the clustering algorithm does not use any vertex neighbourhood information to perform clustering.

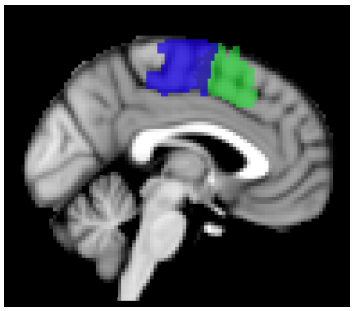

SMA/pre-SMA parcellation in the MNI space

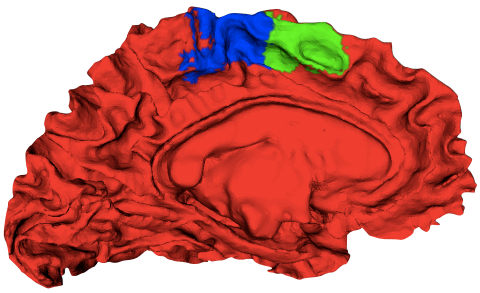

SMA/pre-SMA parcellation projected to the grey/white matter surface of one subject

Figure 3.2: The MFC (see Figure 3.1b) has been divided into the SMA (blue) and pre-SMA (green) using anatomical landmarks (Johansen-Berg et al., 2004). **(left)** The SMA/pre-SMA landmark-based parcellation in the MNI space - the mid-sagittal slice of the left cortical hemisphere is shown. **(right)** The parcellation from the left mapped to one subject's grey/white matter surface using affine registration. The sagittal view of the left cortical hemisphere is shown. The right side of the images corresponds to the anterior of the brain.

whole analysis with slightly different data preprocessing (FWHM = 5 mm Gaussian smoothing instead of 8 mm) and regressed out of RS fMRI data signals of no interest (global signal over the whole brain, averaged signals from the ventricular and white matter masks as well as 6 motion parameters and their first derivatives). We hoped that this would reduce the impact of, among others, potential cardiac and respiratory artefacts. However, the repeated analysis failed to achieve any improvements in this respect. An example of the effect of additional preprocessing on one subject exhibiting artifactual MFC clustering is shown in Figure 3.4.

Before proceeding further with using the SMA/pre-SMA parcellation for FACS validation, we wanted to try to improve the clustering consistency and, if possible, eliminate the artifactual clustering in the superior-inferior direction. An attempt to do so using the Independent Component Analysis of Resting State Networks (RSNs) is presented in the next section.

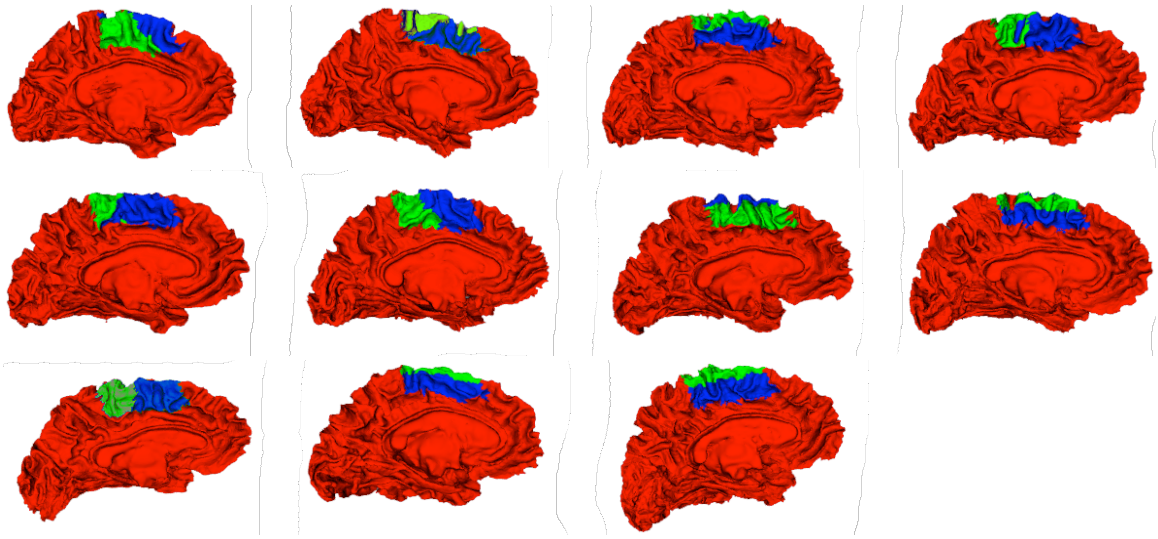

Figure 3.3: FreeSurfer reconstructions of the grey/white matter surfaces of 11 healthy control subjects are shown in red. The MFC was clustered into two clusters (green and blue), putative SMA/pre-SMA, using FC-based clustering (Kim et al., 2010). The right side of the images corresponds to the anterior of the brain. 5 out of 11 subjects exhibit artifactual clustering separating the clusters in the superior-inferior instead of anterior-posterior direction (e.g. bottom right). The others show clusterings similar to the one obtained by manual labelling in the native space (e.g. compare top-right to the Figure 3.2-right).

3.2.3 SMA/pre-SMA Parcellation Using RSNs

3.2.3.1 RSNs and ICA

Low frequency oscillations in BOLD signal ($\approx 0.01 - 0.1$ Hz) have been shown to have similar spatial structure as some task-related activations (Lowe et al., 2000). Coherent spatiotemporal patterns of these low frequency oscillations are called *resting state maps (or networks)* when consistently identifiable across individuals not performing an explicit task. One of the first such networks to be discovered, the Default Mode Network (DMN) (Raichle et al., 2001), is active during rest and is suppressed when goal-directed behaviour is induced. After this discovery, others followed (Beckmann et al., 2005; Damoiseaux et al., 2006; Smith et al., 2009) expanding the set of what are now believed to be RSNs.

One of the popular methods for the detection and analysis of RSNs is the model-

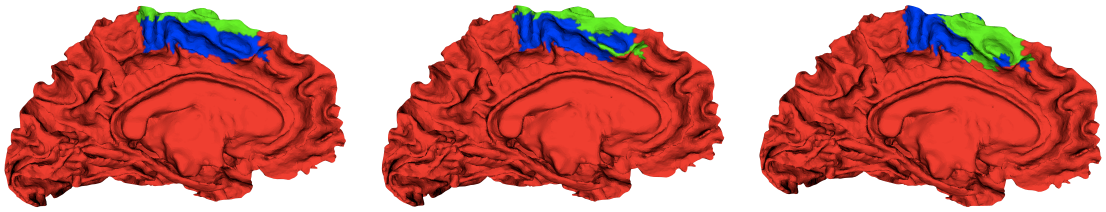

FC-based clustering *before* regressing out nuisance signals

FC-based clustering *after* regressing out nuisance signals

ICA-based RSN clustering of the MFC

Figure 3.4: **(left)** Results of the FC-based clustering before regressing out signals of no interest. **(middle)** Results for the same subject as in the image on the left, except that the data smoothing was lowered and nuisance signals regressed out. **(right)** The same subject as in the images left and in the middle, but clustered using ICA-based RSN parcellation method described in section 3.2.3. The medial side of the left cortical hemisphere (grey/white matter surface) is shown in all images. The right side of the images corresponds to the anterior of the brain.

free, independent component analysis (ICA) (Damoiseaux et al., 2006; De Luca et al., 2006; Beckmann et al., 2005). ICA is able to decompose the RS fMRI data into a set of independent (in space and time) maps and simultaneously extract a variety of different resting state networks. Moreover, this technique allows for straight-forward separation of other nuisance signals (when identified) such as the head motion or physiological confounds. This property of the ICA could prove useful in the attempt to achieve systematic noise removal in the RS fMRI data, resulting in better SMA/pre-SMA clustering. We also want to determine *the same set* of RSNs in each individual and consistently use just those networks for the SMA/pre-SMA clustering. However, the ICA neither guarantees that the same set of RSNs would be identified in each subject, nor that there would be one-to-one correspondence in the RSN maps between subjects. Nevertheless, a procedure that achieves exactly this, the *Dual Regression*, has recently been proposed (Filippini et al., 2009).

Dual Regression (DR) is a procedure capable of identifying *subject-specific* resting state maps so that they correspond across subjects[3]. For example, a map with label

[3]In contrast to the DR, the traditional group-ICA can identify a single set of RSNs specific for a given group, but is not able to obtain subject-specific maps corresponding to the group results.

X - a part of the DMN in one subject, is identified as a map with the same label and corresponds to the same part of the DMN in the other subject. The first step in the DR is the *spatial regression* of the (group)ICA results (e.g. the 'template' set of RSNs) against the individual fMRI data sets. This step results in a set of matrices describing the temporal dynamics of each component in each subject. In the second step, the timecourses are *temporally regressed* against the associated fMRI data sets resulting in a set of subject-specific spatial maps (i.e. spatial Z-statistic maps). As a result, for each of the RSNs from the 'template' set, the corresponding RSN is identified in every subject.

3.2.3.2 ICA-based RSN Clustering

The ability to identify RSNs *as well as* the true artifactual signals through ICA could be used for explicit data filtering. One way to achieve this is to separate artifactual ICA components in each subject and regress them out of the fMRI data. The explicitly filtered data could subsequently be used for FC-based clustering. However, it might be the case that instead of the filtered timeseries, the spatial RSN maps alone bear enough information to distinguish between SMA and pre-SMA (Petrovic et al., 2010). Using spatial maps for clustering can be viewed as a case of dimensionality reduction where the spatiotemporal information is reduced to a set of purely spatial data. To clarify, the full spatiotemporal dataset is used by the DR to identify subject specific RSNs consisting of both spatial maps *and* the associated timecourses. However, we decided to use only spatial components of a set of non-artifactual RSNs to calculate *RSN-based functional connectivity*[4]. The general idea is to form (for each vertex of the grey/white matter surface mesh) an equivalent of the FC Profile, but using spatial information of the non-artifactual RSNs. These features, called *RSN Feature Vectors*, can then be used in the same way and to the same effect as the FC Profiles

[4]Note that the RS fMRI timecourses play an essential role in the DR and are therefore instrumental in obtaining spatial RSN maps too.

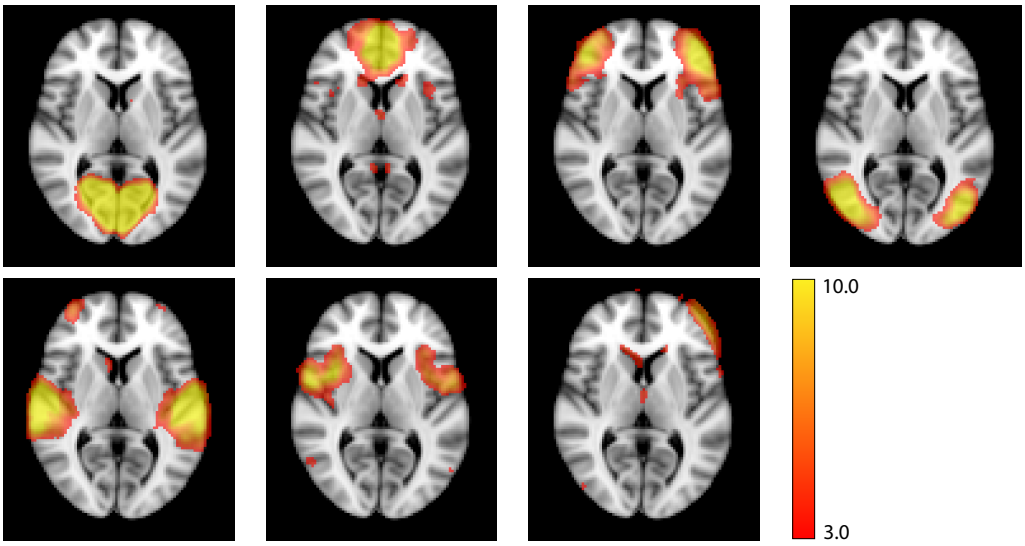

Figure 3.5: A sample of RSNs (all non-artifactual except the one in the lower right corner) used in the formation of RSN-based Feature Vectors. RSN spatial maps (red-yellow) are overlaid over the MNI template brain. Axial slices are shown. The colourbar represents Z-statistic values. The displayed RSNs are a sample from the 70 RSNs obtained through the group-ICA analysis of 36 healthy control subjects. The RSN corresponding to each of these maps was then identified (through the Dual Regression) in each of the 11 subjects in which SMA/pre-SMA clustering is performed. The lower right corner image shows an artifactual RSN component (most likely a motion artefact).

(e.g. for clustering). The adopted approach, as we shall later see (section 3.2.4), yields RSN Feature Vectors that can be used for RSN-based surface registration as well, in replacement of the SCFVs.

RSN Feature Vectors A set of 70 RSNs (both artifactual as well as non-artifactual, so called 'valid') was identified through the group-ICA analysis of 36 healthy control subjects (the same subjects and the same procedure was used by Smith et al. (2009))[5]. FSL/MELODIC software was used to perform the analysis (Beckmann and Smith, 2004). The spatial maps of some of the RSNs are shown in Figure 3.5.

[5]We could have used the same 11 subjects (as for the SMA/pre-SMA clustering) to obtain the group-ICA spatial maps. However, in the larger cohort it is possible to identify more RSNs with greater consistency. For our analysis this in turn means greater descriptive power of RSN Feature Vectors.

The set of 70 RSNs (also called 'the template RSN set') identified in the described manner was fed into the first stage of the DR along with the preprocessed (as in section 3.2.2) data of 11 subjects (called 'the target subjects') used for the SMA/pre-SMA clustering. The output of the DR was a set of 70 subject-specific maps, each of which has a counterpart in the template RSN set.

In the next step, we identified which of the 70 RSN maps in the template set are non-artifactual *and* exist in the left cortical hemisphere (as only the analysis of the left cortical hemisphere is presented). The same identification procedure was used in work of Smith et al. (2009)[6]. In addition to that, we eliminated all maps with the signal found exclusively (after thresholding Z-statistic maps at Z = 3.0) in the cerebellum (not included in cortical surface modelling) or the right cortical hemisphere. Through this procedure 33 RSN maps were retained and automatically identified in each of the target subjects using Dual Regression. Consequently, the original temporal fMRI data were reduced to 33 volumes (for each subject) encoding the Z-statistic of 33 spatial RSN components.

Each of these 4D data sets (comprising the 33 volumes) consisting of non-thresholded Z-statistic volumes (see Figure 3.6) was sampled as described in section 3.2.2.1 to yield a 33-dimensional vector for each of the vertices of the grey/white matter surface meshes. These vectors serve a purpose equivalent to that of the timeseries in the FC-based clustering. To construct RSN Feature Vectors (equivalents of FC Profiles), we proceed as follows: for every subject, each of the 33-dimensional vectors associated with every vertex of the MFC mask is 'correlated' (using ETA^2 measure) with vectors associated with each of the other vertices of the whole $ic5$ subsampled mesh. Let us denote with $P(i,k)$ the Z-statistic value of the kth RSN map associated with vertex i. The *RSN-based functional connectivity* between vertices

[6] 45 non-artifactual RSNs in total were identified.

i and j is then calculated as:

$$eta^2(i,j) = 1 - \frac{\sum_{l=1}^{n}\left((P(i,l)-m(l))^2 + (P(j,l)-m(l))^2\right)}{\sum_{l=1}^{n}\left((P(i,l)-M)^2 + (P(j,l)-M)^2\right)}, \tag{3.2}$$

where $m(l) = \frac{P(i,l)+P(j,l)}{2}$, $M = \frac{1}{n}\sum_{l=1}^{n} m(l)$ and n is the number of retained RSN maps, i.e. 33. The RSN Feature Vector associated with vertex i is given by $eta^2(i,.)$. ETA^2 measure was chosen over traditional correlation as it is sensitive to scaling and offset[7] and has already been proved to perform well in comparing functional connectivity descriptors (Cohen et al., 2008). The RSN Feature Vectors of the MFC were directly clustered with K-means into two clusters. All clusterings (both RSN- and FC-based ones) were found to be highly stable with respect to K-means parameters (MATLAB 2007a, The MathWorks) such as the number of replications, initialisation and the distance measure (e.g. correlation vs. squared Euclidean distance).

The results show improvement over RS FC-based clustering. Before, 5 subjects exhibited spurious superior-inferior clustering, whereas in case of the RSN-based approach just two subjects showed such a pattern. Figure 3.4 displays one subject in which artifactual clustering existed before and after data filtering (in the temporal domain), but improved after RSN-based method was adopted (Figure 3.4 - right). The overall comparison of the results in the two clustering approaches is presented in the next section.

3.2.3.3 Comparison to FC-based Parcellation

In order to compare the results of the two parcellation methods, all individual clusterings were represented in the common registration space (on the $ic5$ mesh - after FreeSurfer cortical geometry-driven alignment). If we suppose that the

[7] If the value of each voxel in one image is exactly double the value of another, they will have a correlation coefficient of 1, but are still different from one another at every point and will have ETA^2 values that may be much less than unity. Similarly, if the value of each voxel in one map is X units greater than another, they will again have a correlation coefficient of 1, but the ETA^2 measure should reflect these differences.

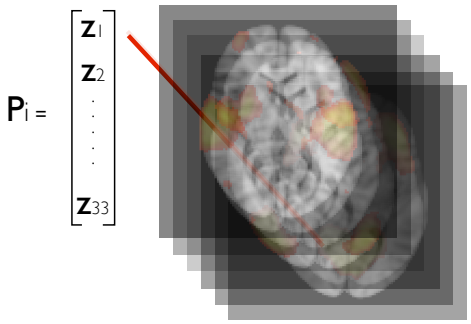

Figure 3.6: For each subject, the non-artifactual RSNs (i.e. their non-thresholded Z-statistic spatial maps) were combined into a 4D volume (like raw fMRI data). The sampling of the 4D volume was performed in the middle of cortical grey matter as described in section 3.2.2.1. These samples were associated with each vertex of the grey/white matter surface mesh. P_i (also later referred to as an RSN Feature) represents the vector associated with vertex i consisting of 33 samples (z_i, $i \in [1, 33]$) of non-artifactual RSNs consistently determined for every subject through Dual Regression. The red line marks the sampling point.

SMA/pre-SMA region is aligned across subjects, it would be possible to infer (from the inter-subject overlap of corresponding clusters) about inter-subject clustering consistency. For each of the 11 subjects, we manually labelled the identified clusters as either SMA or pre-SMA using (for the labelling criterion) the similarity with the SMA/pre-SMA clustering shown in Figure 3.3. After labelling, all parcellations are in correspondence (e.g. SMA overlaid on SMA) and in the common space too, facilitating the quantification of cluster overlap across subjects. Let us label every vertex belonging to the SMA cluster with 1 and all other vertices with 0. Since each subject's grey/white matter surface mesh is registered to the template and represented with the $ic5$, if the labels of all corresponding vertices across 11 subjects are summed up, the result is a spatial map quantifying the inter-subject overlap of the SMA cluster. High values in a certain vertex indicate that that vertex is *consistently* labelled as SMA across subjects and vice versa. The same procedure is repeated for the pre-SMA region. The overlap statistics can be brought back (by nearest neighbour interpolation) to the native space (i.e. the grey/white matter surface) of one subject, to be taken as the template. Figure 3.7 shows such a statistic. It is

clear that the ICA-based RSN parcellation results in a more consistent inter-subject clustering than the RS FC-based parcellation. These differences in the two methods can also be numerically quantified. Figure 3.9 represents the measure of the incorrect inter-subject cluster overlaps across all subjects. The ICA-based RSN clustering again yields smaller error rates.

For comparison purposes only, we also clustered SCFVs associated with vertices in the MFC in the same way as the RSN Feature Vector clustering was performed. The results are presented in the same fashion as those in Figure 3.7. Hence, Figure 3.8 shows the overlap statistic of the SMA and pre-SMA clusters found by clustering SCFVs (similar to clustering the SCFV Similarity Matrix). Although clusters appear not as well defined as those in Figure 3.7 c) and d), the boundary between regions roughly follows the superior-inferior direction.

With improved ICA-based RSN SMA/pre-SMA clustering we can now proceed to use the results in comparing registration methods. It is of particular interest to see whether FACS improves the inter-subject alignment of the border between SMA and pre-SMA compared to FreeSurfer. Related findings are presented in the next section.

3.2.4 Boundary Alignment Before and After FACS

3.2.4.1 Comparing FACS and FreeSurfer

The first step in comparing FACS and FreeSurfer is running FACS preprocessing (see Figure 2.4) for all 11 subjects and then registering them to a chosen reference (Figure 2.10). During preprocessing, FreeSurfer registers all subjects to its common template. However, for FACS, we randomly chose one subject to act as the reference and align all others (10 in total) to it. After the FreeSurfer and FACS registrations have been performed, a measure of incorrect overlaps (or mixing) of SMA and pre-SMA clusters across subjects can be calculated. In particular, the graph shown in Figure 3.9 can be replicated for both FACS and FreeSurfer (for the ICA-based RSN

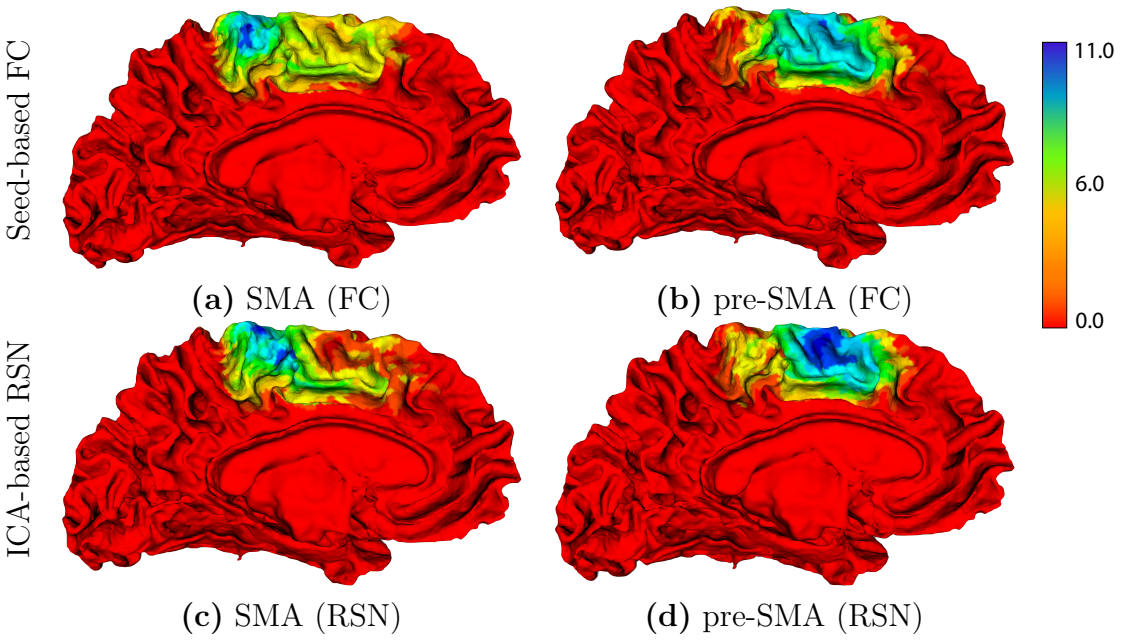

(a) SMA (FC) **(b)** pre-SMA (FC)

(c) SMA (RSN) **(d)** pre-SMA (RSN)

Figure 3.7: A comparative performance of two clustering approaches (RS FC- and ICA-based) is presented by displaying overlap statistics on a prototype subject. **(a)** and **(b)** show seed-based FC parcellation; **(c)** and **(d)** show ICA-based RSN parcellation; **(a)** and **(c)** show posterior (putative SMA) and **(b)** and **(d)** show anterior (pre-SMA) clusters. Surfaces were co-registered using FreeSurfer cortical surface registration. Colourbar represents the number of subjects for which the corresponding cluster overlaps, i.e. blue when all subjects (11 in total) have cluster 'X' at that point and red when there is no overlap of corresponding clusters.

parcellation only).

Practically, both FACS and FreeSurfer alignment result in two sets of high-density spherical meshes in two registration spaces (the FreeSurfer average space and the FACS reference subject space). All of these meshes have RSN Feature Vectors (RSNFVs) associated with each of the vertices. We subsample these meshes to $ic5$ and cluster associated RSNFVs into putative SMA and pre-SMA. The inter-subject overlap of corresponding clusters is measured as explained in section 3.2.3.3. The clustering procedure is performed a number of times (with different initialisations) to test for the consistency of the actual clustering algorithm. Once the clusters have been identified, they need to be assigned either SMA or pre-SMA labels. The labelling can either be done manually or be automated. So far, we manually assigned

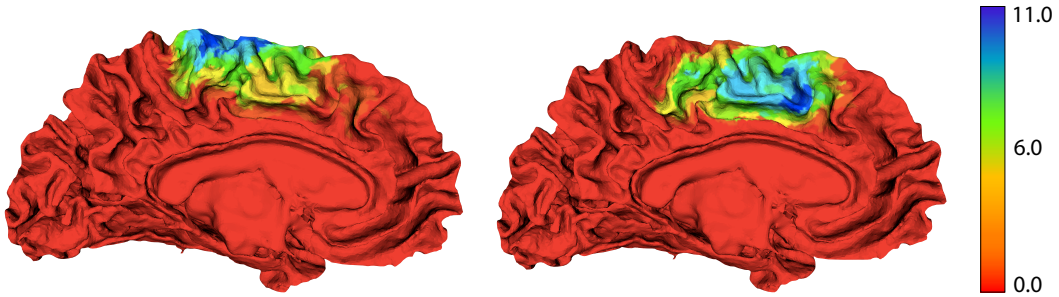

SMA (SCFVs) preSMA (SCFVs)

Figure 3.8: Clustering of the SMA/pre-SMA using SCFVs. The clustering was performed in the same was as that in Figure 3.7. **(left)** mean posterior cluster (putative SMA); **(right)** mean anterior cluster (putative pre-SMA). Colourbar represents the number of subjects for which the corresponding cluster overlaps, i.e. blue when all subjects (11 in total) have cluster 'X' at that point and red when there is no overlap of corresponding clusters.

correct labels - SMA or pre-SMA - to the identified clusters. This means that after K-means identified two clusters e.g. labelled as 1 and 2 (further explained in the next paragraph) we decided for each subject whether 1 or 2 should be denoted as SMA (and the other as pre-SMA). However, this becomes unfeasible for large number of independent clusterings and is done automatically for each registration method separately.

Automatic Cluster Labelling The K-means algorithm generates two cluster labels (e.g. 1 and 2) for each subject, but these do not necessarily correspond across subjects. For example, SMA in one subject may be labelled as 1 and in the other as 2. However, we would like the SMA in all subjects to be labelled consistently as e.g., 1. The Algorithm D.0.1 in Appendix D describes how this can be achieved. In short, the *prototype labelling* e.g. SMA = 1, pre-SMA = 2 from a randomly chosen subject is imposed on all others and this is embedded in an iterative procedure in which re-labelling and the derivation of the prototype labelling interchange. Other labelling schemes are possible too, such as the one which assigns cluster labels on the basis of the similarity in prototype (e.g. mean) RSNFVs, or even hierarchical clustering

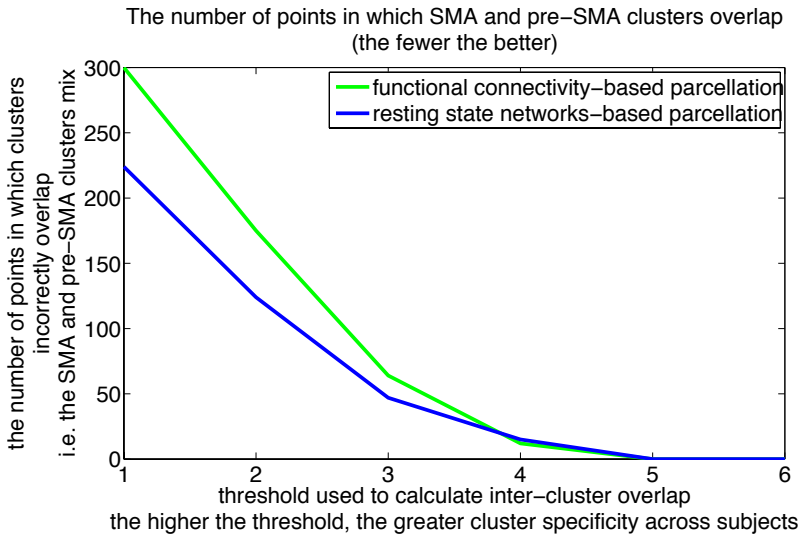

Figure 3.9: The measure of incorrect overlaps of SMA and pre-SMA for the whole dataset, across two methods. The data were subsampled to $ic5$ and the same reference was used in both methods. At a given vertex, given a threshold T, an incorrect overlap occurs when T or more subjects are labeled as SMA *and* T or more subjects are labeled as pre-SMA. The X-axis represents the threshold T and the Y-axis shows the number of incorrect overlaps. Note that the typical number of vertices on the boundary between clusters is around 50 (on the $ic5$).

(Jbabdi et al., 2009a). However, even the simple re-labelling procedure presented here yields satisfactory results. We have manually checked 20 (out of 200) randomly chosen re-labellings to find that all of them have labels assigned as would have been manually done. Nevertheless, we should note that the presented algorithm does not guarantee that all labels will appear in each subject after relabelling (e.g. both the SMA and the pre-SMA could be labelled with the same label within one subject). Still, we hope that this does not affect the comparison of FACS and FreeSurfer as the same re-labelling procedure is performed in both cases.

3.2.4.2 RSN Connectivity-driven Alignment

Apart from using SCFVs to drive registration in FACS, we can feed into the registration pipeline RSNFVs to the same effect. The similarities in RSN Feature Vectors is encoded in the Similarity Matrix (section 2.4.4.2, Equation 2.4) reflecting *RSN-based (functional) connectivity*. Substitution of SCFVs with RSNFVs is possible

because of the way RSNFVs were constructed: the Dual Regression ensures inter-subject correspondence and thus comparison of RSNFVs.

Along with white matter connectivity-driven alignment, we performed RSN-driven alignment too by just substituting the SCFVs with RSNFVs. However, we did not perform fine tuning of registration parameters, which is required due to the differently-structured SM. Our registration method is sensitive to the range of values in the SM. The free parameters (such as the forces scaling coefficients) were fine-tuned for the case of SM constructed using SCFVs. Therefore, substituting the SCFVs with RSNFVs and not fine tuning free parameters, might produce suboptimal results/registrations. Nevertheless, our aim is to demonstrate the possibilities of the FACS registration framework and show how DW data, i.e. structural connectivity, can be straightforwardly replaced by RS fMRI data, i.e. RSN-based connectivity, without modifying the general registration pipeline. The results of the inter-subjects, SMA/pre-SMA cluster alignment using FACS with DW data, FACS with fMRI data and FreeSurfer are shown in Figure 3.10.

3.2.4.3 Results

Figure 3.10 shows the quantification of the inter-subject incorrect cluster overlap across three registration methods. Note that we use cortical geometry (T1-weighted data) to drive FreeSurfer registration, structural connectivity (DW data) for FACS and fMRI data for RSN-driven registration. Therefore, RSN-driven registration is tested on non-independent data (unlike the other methods) and the results for RSN-driven registration are therefore shown just for reference. The whole analysis (i.e. clustering) was performed 200 times showing statistically lower incorrect cluster overlap rate for FACS compared to FreeSurfer. Clustering was run multiple times with different random initialisations, which affects the final cluster label assignments. This helps with assessing the effects of variability of clustering solutions to the quantification of overlaps. The RSN-driven alignment, as expected, performed better

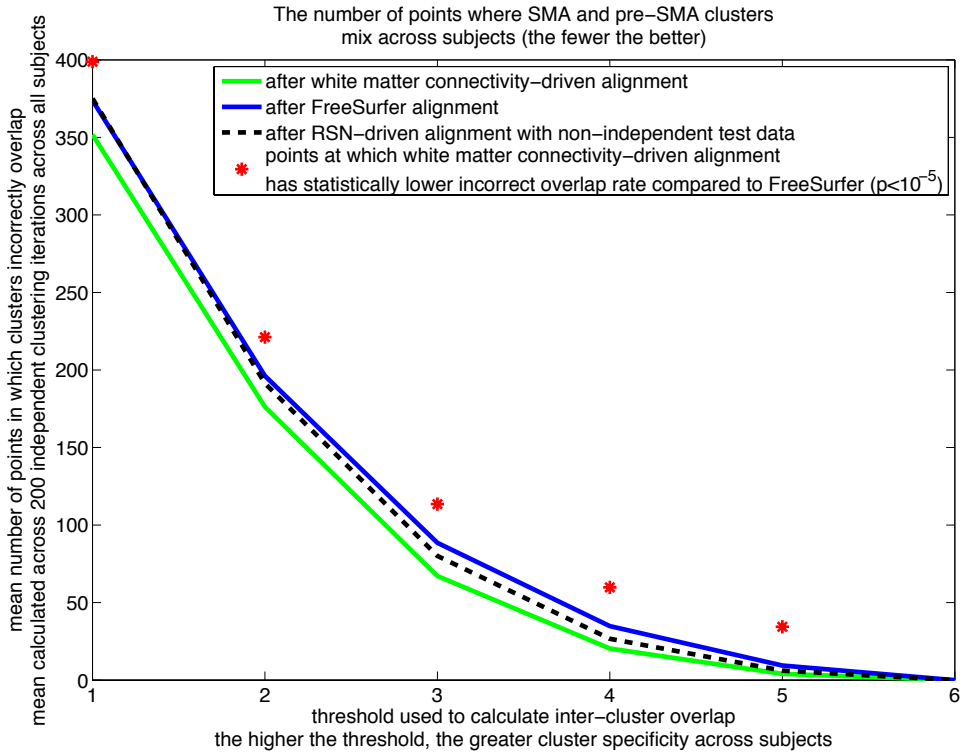

Figure 3.10: The figure shows the measure of incorrect overlaps of SMA and pre-SMA clusters for the whole dataset (11 subjects), across three registration methods (FACS, FreeSurfer and FACS using RSN Feature Vectors instead of SCFVs). The data were subsampled to $ic5$ and the same reference were used across all methods. For a given vertex, with a threshold T, an incorrect overlap occurs when T or more subjects are labeled as SMA and T or more subjects are labeled as pre-SMA. The X-axis represents the threshold, T, while the Y-axis represents the number of incorrect overlaps. Clustering was performed 200 times (with different random initialisations) for each subject and the statistical analysis over all parcellations shows significantly lower incorrect overlap rate with FACS compared to FreeSurfer. Note that the typical number of vertices on the boundary between clusters is around 50 (on the $ic5$).

than FreeSurfer, but poorer than FACS, possibly due to the lack of fine tuning. It is also worth pointing out that although clustering and registration use the same fMRI data, they utilise them in different ways having the alignment results only indirectly reflected in the clustering.

We would hope that these results hold not only for the SMA/pre-SMA region but can be extended to the others, e.g. such as the cingulate. Ideally, the alignment of *borders* between functionally different regions should be improved across the whole

cortical hemisphere. In the rest of the Chapter we explore this hypothesis. However, we remain constrained by the ability of our 'test data', i.e. RS fMRI to discern between different functional regions.

3.3 Alignment of Boundaries Between Functional Regions

Resting-state functional connectivity can be used not only to examine the SMA/pre-SMA boundary, but to reveal the edges (Cohen et al., 2008) separating functionally different regions across the whole cortex. In the context of validating FACS, we are especially interested how well these boundaries are aligned across subjects. This is important as it might be the case that the registration performs better in improving the overlaps of boundaries and does not do as well in the areas *within* those boundaries. One of the reasons for such behaviour can be that the SCFVs within distinct functional areas are too similar to each other and the alignment heavily relies on successful regularisation. On the other hand, the SCFVs on functional boundaries can have greater discriminatory power. The analysis done so far (SMA/pre-SMA clustering) does not explicitly look into these boundaries. To examine this further, we first detect functional boundaries and then measure the inter-subject alignment across the whole brain as well as in some specific areas such as the cingulate and the MFC.

Cohen et al. (2008) argue that the gradient of the spatial connectivity maps can provide useful insights in this respect. Their work on whole-brain (surface) parcellation using RS FC can be summarised in several steps:

1. Find the Functional Connectivity Map (FCM) for each vertex/voxel (referred to as 'seed') of the cortical surface (e.g. one cortical hemisphere). FCM effectively associates with each seed a spatial map (represented as a vector) encoding ETA^2 similarity of the seed's RS FC Profile to that of all the other vertices/voxels. RS

FC Profiles can, for this purpose, be high-dimensional, e.g. encoding correlations between RS timeseries of the seed voxel and *every other* voxel in the brain image;

2. Perform binary edge detection on the maps obtained in the previous step using gradient-based edge detectors such as the Canny edge detector (Canny, 1986);

3. Repeat the first two steps for *all* seed vertices/voxels;

4. Construct the mean 'edge map' across all seeds;

5. Use 'region growing' segmentation techniques such as the Watershed algorithm (Vincent and Soille, 1991) on the map constructed in previous step to define putative functional areas.

We would like to use a similar pipeline to try and discern edges between functional regions. The alignment of these edges (or regions between them) could subsequently be used for examining the effects of registration. However, we would like to build on the previous finding (section 3.2.4.3) that ICA-based RSN clustering performed better than RS FC-based parcellation, at least in the case of the SMA/pre-SMA segmentation. Therefore, the (Cohen et al., 2008) boundary detection framework could be adapted so that in step 3) of their pipeline, RSN Feature Vectors are directly used, which already have the structure of the spatial maps encoding similarities in *large-scale RSN spatial patterns* between the seeds. Note that in the case of Cohen et al. (2008), the spatial maps encoded large-scale similarities in seed-based functional connectivity profiles. In our case, the large-scale RSN spatial pattern associated with a seed vertex represents the membership (or 'strength of association') of that vertex to a predefined set of RSNs. In both cases 'large-scale' refers to the phenomena spanning a large portion of the brain surface. Therefore, if the local changes in similarities of these RSN spatial patterns carry the information about the boundaries between functional regions, it should be possible to devise a method for detecting

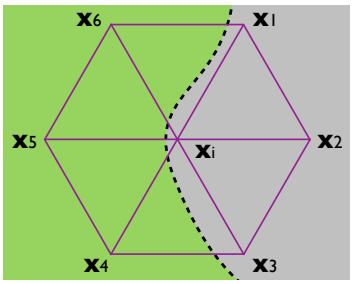

Figure 3.11: The basic mesh element (in purple) with a central vertex i. The black dashed line marks the boundary between two regions. The boundary direction is approximated by the triplet (x_1, x_i, x_3), where x_i encodes the position of a vertex i.

these boundaries. The adapted gradient-based edge detection algorithm which tries to achieve this in the cortical surface framework is introduced in the next section.

3.3.1 Adapted Edge Detection Algorithm

The edge detection on the sphere (e.g. *ic*5 mesh) presented here follows the same steps proposed by Canny (1986) for a planar image:

1. smoothing,
2. non-maximum suppression,
3. thresholding with hysteresis.

The main difference in the two methods is that, in this case, vertices assume the role of voxels (or pixels in a 2D implementation). This underlies the differences in the calculation of the gradient image as well as in smoothing. In a common 2D implementation of smoothing, the image is convoluted with a 3×3 Gaussian kernel (Cohen et al., 2008). In our implementation, a support of a Gaussian kernel is defined by a *basic mesh element* consisting of a central vertex and its immediate neighbours. Such a mesh element to which smoothing is applied can be seen in Figure 3.11. The new smoothed value for a vertex i (s'_i) is calculated as a weighted sum of its original

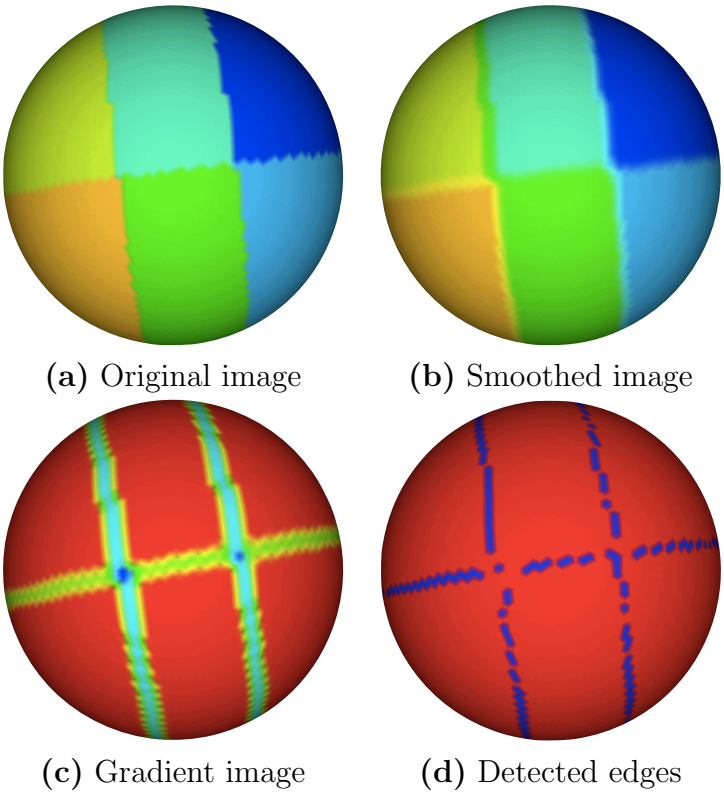

(a) Original image **(b)** Smoothed image

(c) Gradient image **(d)** Detected edges

Figure 3.12: **(a)** Different colours correspond to different image values. **(b)** An image from (a) after smoothing. **(c)** The gradient magnitude of the image in (b) - red = 0, blue = 1. High gradient values indicate the presence of an edge. **(d)** The extracted binary one-vertex-thick edges (in blue).

value (s_i) and values associated with its immediate neighbours:

$$s'_i = \frac{\sum_{j \in \mathcal{R}(i) \cup \{i\}} \frac{1}{\sqrt{2\pi}\sigma} s_j e^{-\frac{D^2\left(v_i, v_j\right)}{2\sigma^2}}}{\sum_{j \in \mathcal{R}(i) \cup \{i\}} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{D^2\left(v_i, v_j\right)}{2\sigma^2}}}, \tag{3.3}$$

where v_i encodes position of vertex i, D is the geodesic distance and σ is the smoothing parameter set to $0.1R$ in our implementation (R is the registration sphere radius, note that the smoothing takes place on a unit $ic5$, i.e. $R = 1$). $\mathcal{R}(i)$ is, as before, the set of immediate neighbours of vertex i.

The gradient is approximated using regular local mesh topology as well. Without the loss of generality, let us consider a mesh element from Figure 3.11 with a central vertex i and its six regularly spaced neighbours labelled from 1 to 6. The image

gradient magnitude in vertex i (g_i) is approximated by

$$\begin{aligned} g_i \approx (&|s_1 - s_3| + |s_1 - s_4| + |s_1 - s_5| + \\ &+ |s_2 - s_4| + |s_2 - s_5| + |s_2 - s_6| + \\ &+ |s_3 - s_5| + |s_3 - s_6| + |s_4 - s_6|)/9, \end{aligned} \tag{3.4}$$

where s_i is a scalar value associated with vertex i. In the case when the basic mesh element consists of 5 vertices, as happens on the $ic5$, the Equation (3.4) is modified as $g_i \approx (|s_1 - s_3| + |s_1 - s_4| + |s_2 - s_4| + |s_2 - s_5| + |s_3 - s_5|)/5$. The edge (gradient) direction is determined by the brute-force search over triplets of vertices (s_j, s_i, s_k)[8], $j, k \in [1, 6],\ 1 < |j - k| < 5$ by finding which pair (s_j, s_k) has the lowest contribution to the gradient: $\min_{j,k} |s_j - s_k|,\ j, k \in [1, 6],\ 1 < |j - k| < 5$. Having fully determined the gradient, the non-maximum suppression (eliminating vertices which are not local maxima) and the thresholding with hysteresis are performed in the usual manner. The thresholding uses lower and the upper thresholds set to 0.1 and 0.3 respectively (note that for $0 \leq s_i \leq 1$, as the case of RSN Feature Vectors, and so the gradient is bound to the $[0, 1]$ interval). These values were determined by analysing an artificial example presented in Figure 3.12. Figure 3.12 a) shows an artificial image overlaid on an $ic5$ mesh where differently coloured regions correspond to different image values (and vice versa). The smoothed image is shown in b). Figure 3.12 c) shows the absolute gradient of the original image (red = 0, blue = 1) and d) shows extracted binary edges in blue.

In practice, the scalar values are the entries of RSN Feature Vectors associated with every mesh vertex. An example of such an image (an RSN Feature Vector) for the left cortical hemisphere and a randomly chosen vertex is shown in Figure 3.13 a). Vertices with high values (in blue) depict areas with high ETA^2 similarity in the RSN descriptors with the seed vertex. Such maps are generated for each of the mesh

[8]The triplet defines an edge (if any) passing through the basic mesh element.

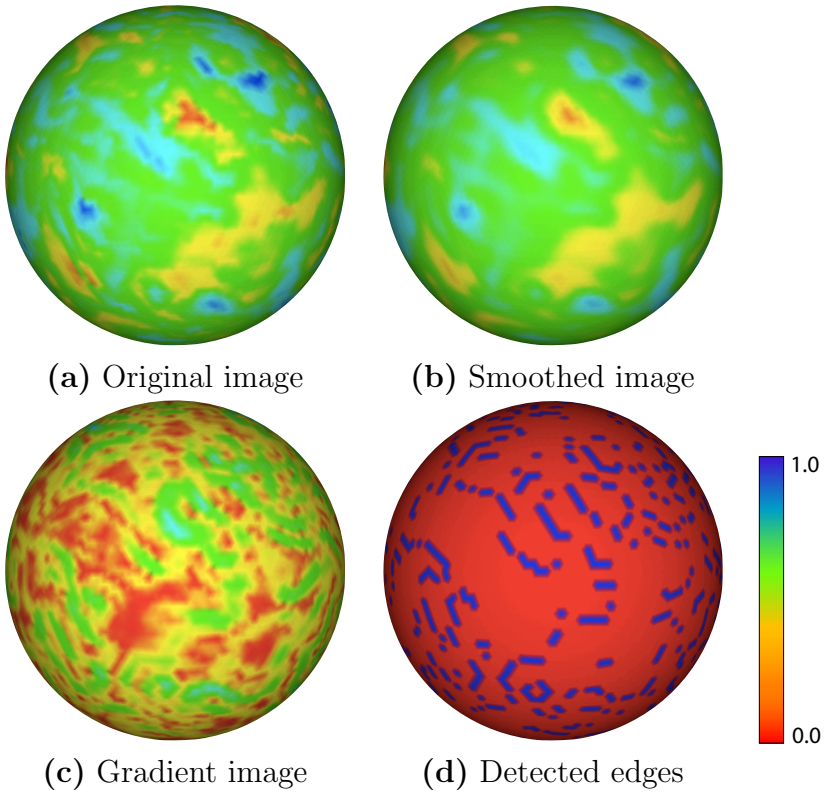

(a) Original image **(b)** Smoothed image

(c) Gradient image **(d)** Detected edges

Figure 3.13: The main steps in edge detection (as in Figure 3.12) shown for real data. **(a)** An image showing ETA^2 similarities between RSN Feature Vectors associated with a randomly chosen vertex (left cortical hemisphere). **(b)** An image from (a) after smoothing. **(c)** The gradient magnitude of the image in (b). **(d)** The final binary edge map (in blue).

vertices. Figure 3.13 b) shows a) after smoothing. The gradient magnitude image is shown in c) and the detected binary edges are shown in d) (in blue). Binary edge images can be formed for each vertex and averaged across edge images from all seed vertices to form an edge 'probability' map[9] (Figure 3.14).

Through the rest of the Chapter we make use of this edge detection procedure. However, we found that simple thresholding of gradient images (such as the one in Figure 3.13c) or 3.12c)) produces very similar results, and we use these gradient images to the same effect (e.g. Figure 3.23) with no differences in the final conclusions.

[9]'Probability' is here used as a descriptive quantifier of the spread of identified edges. It is not meant to represent the actual probability of the edge existence.

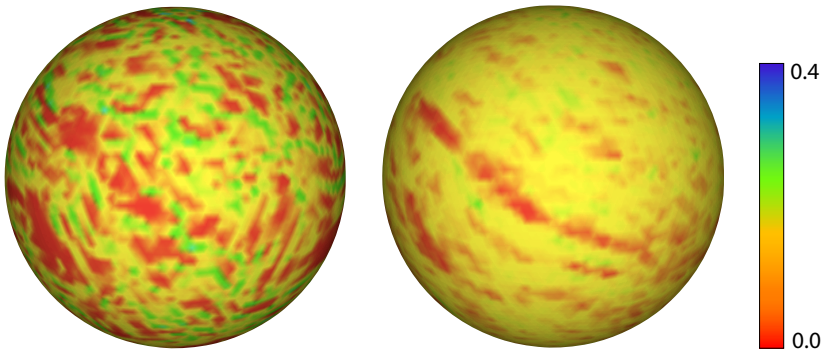

Figure 3.14: **(left)** A single subject 'edge probability map' (EPM). After images like the one in Figure 3.13 d) have been generated for all vertices, a mean image across all seed vertices represents an EPM. **(right)**. The mean of EPMs across 11 subjects. The colourbar indicates 'edge probability'.

3.3.2 Whole-brain Analysis

The binary edge maps are derived for each vertex of the $ic5$ (of a subject's left cortical hemisphere) and can be subsequently combined to produce the 'edge probability map'. The edge probability map *value* in vertex i, EPM_i is calculated as $\mathrm{EPM}_i = \frac{1}{n}\sum_{j=1}^{n}\mathrm{EM}_j(i)$ where $\mathrm{EM}_j(i)$ is the binary *Edge Map* value in vertex i for a map associated with vertex j. The EPMs can be calculated and in the same way averaged across all subjects (Figure 3.15). This assumes that the subjects have been co-registered and that the vertex-to-vertex correspondence exists between the $ic5$ representations of each subject's grey/white matter surface mesh (e.g., see Figure 3.1d). The registration can be performed using different registration methods. So far we utilised FreeSurfer and FACS. As before, FACS is driven using diffusion data (white matter connectivity) whereas FreeSurfer using T1 images. However, for this analysis two more methods were included: the affine alignment with 12 degrees of freedom (FLIRT/FSL, (Jenkinson et al., 2002)) and the nonlinear alignment (FNIRT/FSL, (Andersson et al., 2007a,b)) with default parameters. FNIRT is a nonlinear registration tool that uses B-splines representation of the registration warp and optimises the sum of squared differences of T1 image intensities as its objective function. For all methods, FreeSurfer was used to reconstruct the grey/white matter

surface. The correspondence between surfaces of different subjects is determined by the registration algorithm. In the case of FLIRT and FNIRT, the images were registered to a 1mm MNI template brain and the generated transformations were then applied to the reconstructed grey/white matter surfaces. These surfaces, being in the same space, were then inflated to a unit sphere minimising metric distortions (the same process is applied in the FreeSurfer pipeline - see Chapter 2 section 2.1.1). The high-density meshes were subsampled to $ic5$ as previously described. The mean EPMs across all 11 subjects for four different registration methods are shown in Figure 3.15. The regions depicted in red (low values) are consistently labelled as 'valleys' if the edges were to be seen as 'hills'. These areas in between the edges are the putative *functional regions* segmented through the edge detection procedure and it is the inter-subject correspondence of these functional regions that we are trying to achieve. By comparing Figure 3.15a) and d) one can see that d) has a conspicuous 'red stripe' stretching from the upper left to the lower right corner. This region corresponds to the central sulcus and is more prominent in d) (FACS) than in a) (FLIRT) where it can hardly be detected. Linear (affine) registration methods are known to perform poorly for aligning sulcal/gyral features compared to the non-linear methods such as FreeSurfer or CVS (Postelnicu et al., 2009) and this might be reflected here too, although through aligning functional areas.

A quantitative analysis of the similarities between each subject's EPMs can also be performed. We compared the thresholded EPMs between each pair of subjects across four registration methods. One way to measure similarity between thresholded (and therefore binary) co-registered images is using a Jaccard coefficient (Jaccard, 1912). It describes the overlap between the labels of images V and T as:

$$J(V,T) = \frac{|V \cap T|}{|V \cup T|}. \tag{3.5}$$

Figure 3.16 demonstrates how this measure can be used to assess the performance of registration algorithms in aligning (boundaries between) putative functional regions.

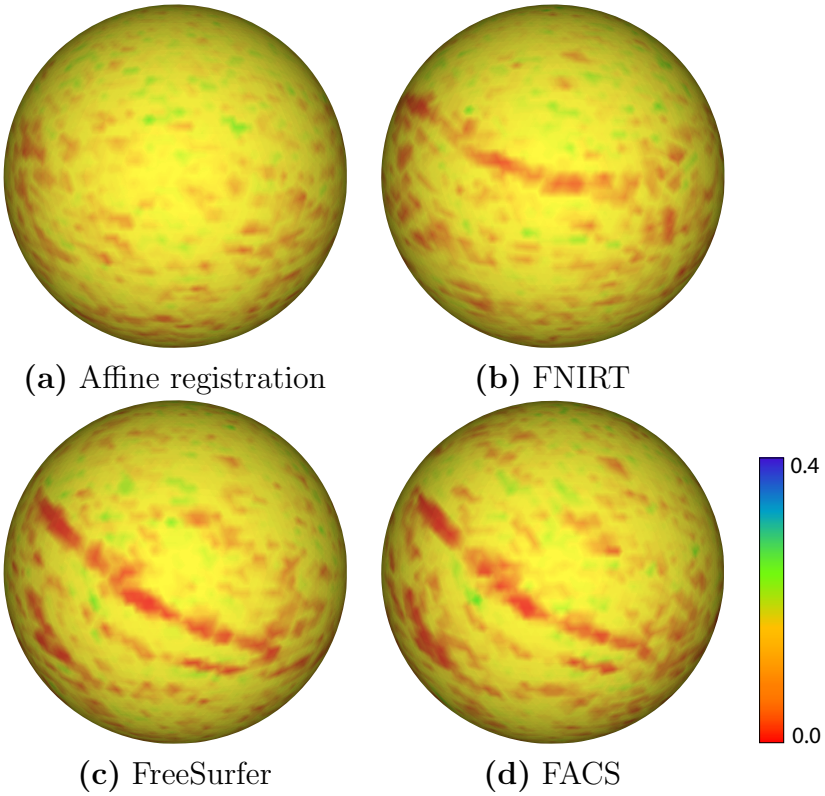

(a) Affine registration **(b)** FNIRT

(c) FreeSurfer **(d)** FACS

Figure 3.15: Images showing the mean Edge Probability Maps across all subjects for four different registration methods: **(a)** affine registration, **(b)** nonlinear FNIRT registration, **(c)** FreeSurfer registration and **(d)** FACS alignment. The 'red' regions, i.e. low values, indicate high inter-subject consistency in labelling that area as a distinct putative functional region. Conversely, one can view 'red' regions as the ones having consistently determined edges. Affine registration (a) features very few such regions (supposedly due to poor inter-subject alignment of sulcal/gyral features) whereas both FreeSurfer (c) and FACS (d) have a prominent one stretching from the upper left to the lower right corner. This 'red stripe' anatomically corresponds to the central sulcus.

The individual EPMs were first inverted, i.e. subtracted from 1 (the functional regions now have high values, and edges between them low values) and then thresholded with a variable threshold (steps of 0.01 were used) so that the values above the threshold are labelled as 1 and the others as 0. In this way, ideally, all separate functional regions would be labelled with different labels and prepared for comparison across subjects using the Jaccard measure (see section 3.3.4.1 for MFC clustering). However, in this case, all functional regions (separated by the identified edges) within the cingulate cortex are labelled with the same label (1) and therefore it is *not a priori*

ensured that corresponding regions will be compared across individuals. Nonetheless, the comparison of non-corresponding functional regions is likely to be accompanied by the increased number of 'valleys'-edges overlaps lowering the Jaccard overlap measure. In the next step, the Jaccard overlap is calculated between each pair of thresholded images. The mean of the Jaccard overlap coefficient across all image pairs is recorded and plotted for different thresholds and across four registration methods (Figure 3.16). For higher thresholds, e.g. > 0.9 (where just the functional regions defined with high inter-subject consistency are considered), the differences between registration methods become detectable. The affine alignment has the lowest Jaccard overlap values, followed by FNIRT. FreeSurfer and FACS, although similar in their performance, both score better than the other methods. These results suggest that FACS does not necessarily outperform FreeSurfer *across the whole cortical hemisphere* at all thresholds in aligning corresponding functional boundaries. The reasons why this might be the case are discussed in section 3.5.

Another way to quantify the overlap of binarised images is using specificity and sensitivity measures. For each pair of images, the true positive, true negative, false positive and false negative overlaps of one image to the other (the template) were calculated. These values were then used to calculate the sensitivity and specificity for each pair of images for a given threshold. In the next step, the mean was found for each measure across all pairs of subjects (for each threshold; the same steps as before were used). Mean sensitivity and specificity were combined to perform an ROC (receiver operating characteristic) analysis shown in Figure 3.17. Again, the conclusions reached using Jaccard measures were confirmed: FACS and FreeSurfer perform similarly and both do better than FNIRT and the affine registration.

Similar inferences can be drawn from analysing thresholded *mean gradient images* (MGIs) instead of thresholded EPMs[10]. Figures 3.18 and 3.19 show such analysis.

[10]On MGI construction see section 3.3.4.

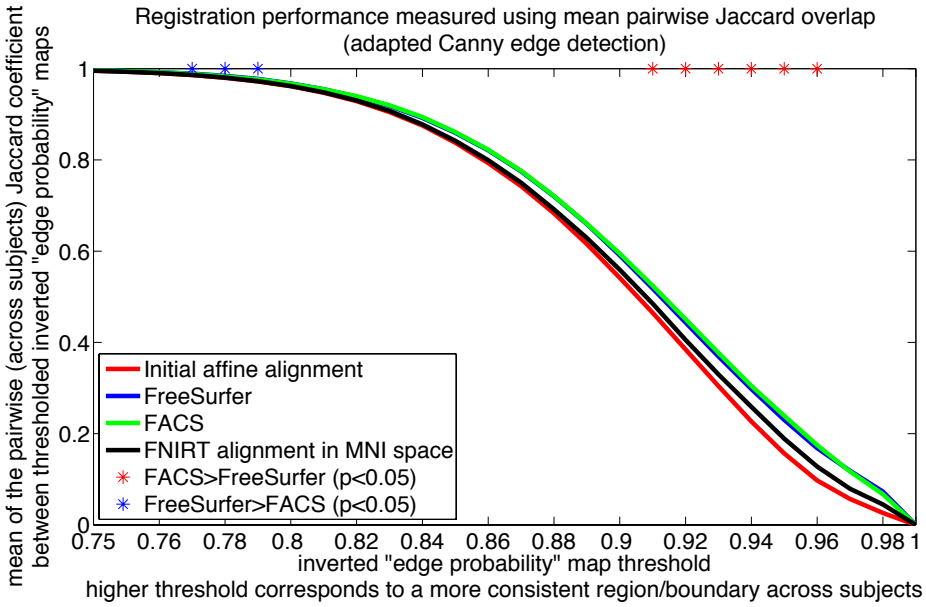

Figure 3.16: Comparison of the inverted Edge Probability Maps (EPMs) across all subjects for different thresholds. Analysis shown for the whole-brain case. Pairwise Jaccard overlap measure was used to quantify similarity between binarised EPMs. High Jaccard coefficient indicates good overlap. T-tests were used to assess the significance of the results for each threshold value. FNIRT, FACS and FreeSurfer always performed better than affine alignment and FACS and FreeSurfer always performed better than FNIRT. FACS and FreeSurfer performed similarly without convincing superiority of either method (although FACS did better for higher threshold values than FreeSurfer).

In this case, FACS outperforms FreeSurfer for all threshold levels with better discrimination between other registration methods as well (compared to the analysis using EPMs). This could suggest that MGIs have better discriminatory power than EPMs, which is also indicated in the subsequent analysis, for example, in the analysis of the cingulate area in section 3.3.4.

The results obtained so far led us to think that the improvements in registration, which can be measured with RSN-derived features, might not necessarily span the whole cortical surface, but can be rather localised to specific areas, perhaps those where both diffusion and functional data have been found to successfully discern different functional regions. To examine this further, we analysed the MFC and the cingulate area where it has been confirmed that probabilistic tractography can be

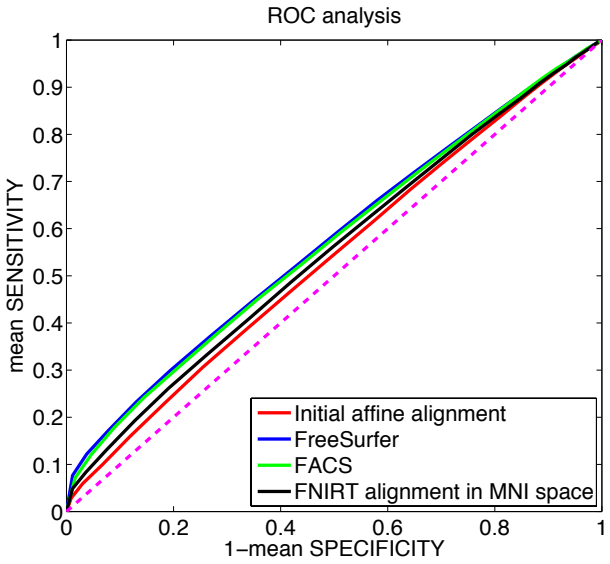

Figure 3.17: The thresholded EPMs can be compared using sensitivity and specificity measures. The ROC analysis confirms findings from Figure 3.16. FreeSurfer and FACS perform similarly whereas FNIRT and the affine alignment perform worse than the others (but better than chance - pink dashed line).

used for functional parcellation (Beckmann et al., 2009). Comparable findings are presented in sections 3.3.3 and 3.3.4.

3.3.3 Analysis of the MFC Area

The similarity of MGIs across subjects was now investigated just in the MFC area. We would like to examine the overlap of the SMA/pre-SMA boundary (and any other functional boundaries that might exist within the MFC) and compare the results to the whole-brain analysis (section 3.3.2) as well as to the analysis of the cingulate area (section 3.3.4). To achieve this we made use of the MFC mask shown in Figure 3.1 a). The Jaccard overlap was calculated for the inverted MGI values within the MFC mask (not for the whole cortical hemisphere as in Figure 3.16). These results are shown in Figure 3.20. As for the whole-brain, the analysis was performed for different thresholds and the T-test assessed the performance of FACS compared to FreeSurfer. The results indicate that FACS outperforms FreeSurfer for several threshold levels and FreeSurfer never does statistically better than FACS. This is in line with the results

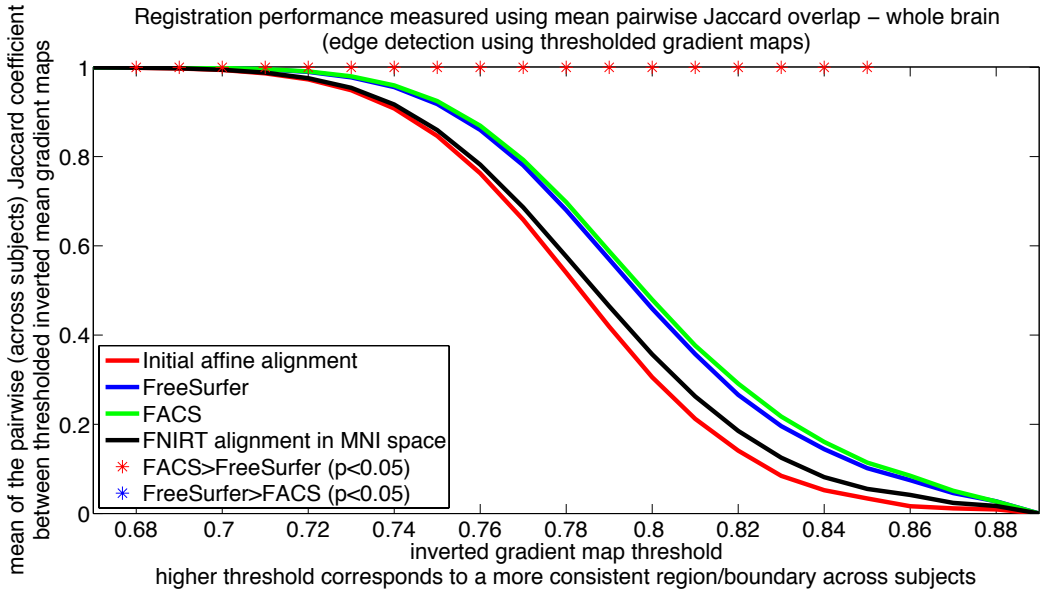

Figure 3.18: Comparison of the inverted Mean Gradient Images (MGIs) across all subjects for different thresholds. Analysis shown for the whole-brain case (left cortical hemisphere). Pairwise Jaccard overlap measure was used to quantify similarity between binarised MGIs. High Jaccard coefficient indicates good overlap. T-tests were used to assess the significance of the results for each threshold value. FNIRT, FACS and FreeSurfer always performed better than affine alignment and FACS and FreeSurfer always performed better than FNIRT. FACS outperforms FreeSurfer for almost all threshold values.

on SMA/pre-SMA clustering presented in section 3.2.3.2. However, the ROC analysis for the MFC area (Figure 3.21) shows no clear improvement of FACS over FreeSurfer. The overall sensitivity/specificity values are low and similar across methods rendering low discriminatory power of this assessment method.

3.3.4 Analysis of the Cingulate Area

The cingulate area can be analysed in a manner comparable to that of the MFC (section 3.3.3). The cingulate area, just like the SMA/pre-SMA, consists of several functional clusters (separated by edges we try to detect) that can be identified using both fMRI and DW data (details discussed further in this section). Therefore, we can repeat the analysis of the inter-subject overlap of the functional clusters (and borders) for the cingulate area as well.

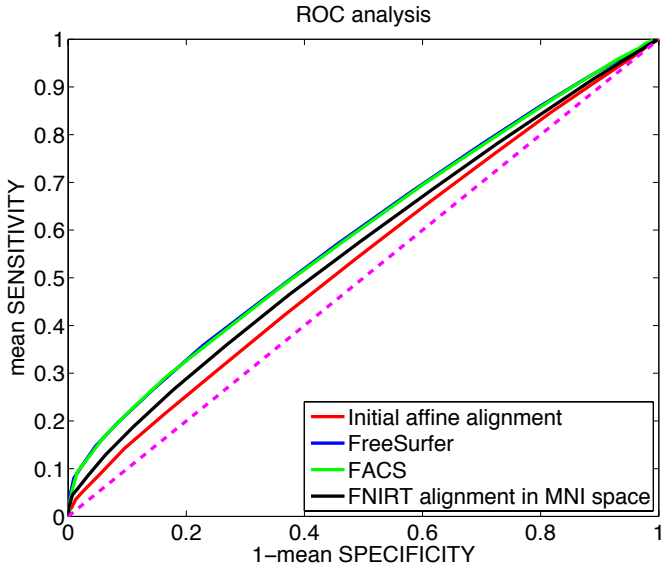

Figure 3.19: The thresholded MGIs can be compared using sensitivity and specificity measures. FreeSurfer and FACS perform similarly whereas FNIRT and the affine alignment perform worse than the others (but better than chance - pink dashed line).

It has been shown by Beckmann et al. (2009) that the cingulate cortex can be parcellated using structural connectivity information into a number of parcels closely corresponding to functionally distinct regions. In particular, they have been able to identify 9 clusters within the areas encompassing the cingulate gyrus and the cingulate and paracingulate sulci using structural 'connectivity fingerprints' (Tomassini et al., 2007b). These clusters were related to the functional areas identified through the meta-analysis of functional studies reporting activations in the cingulate. The boundaries between these clusters run in the dorsal-ventral direction which is consistent with the cytoarchitectonic findings (Vogt, 2009). We have used SCFVs to a similar effect (see Figure 2.14). Moreover, although not validated, the clustering of the cingulate cortex using RSNs (reported in section 3.3.4.1) into putative functional areas yields plausible results (an informative study on the relations between functional and structural connectivity can be found in (Honey et al., 2009). The cingulate cortex in particular exhibits similar connectivity patterns measured either using RS fMRI data or DW MRI data). In this sense, the MFC and the cingulate cortex appear

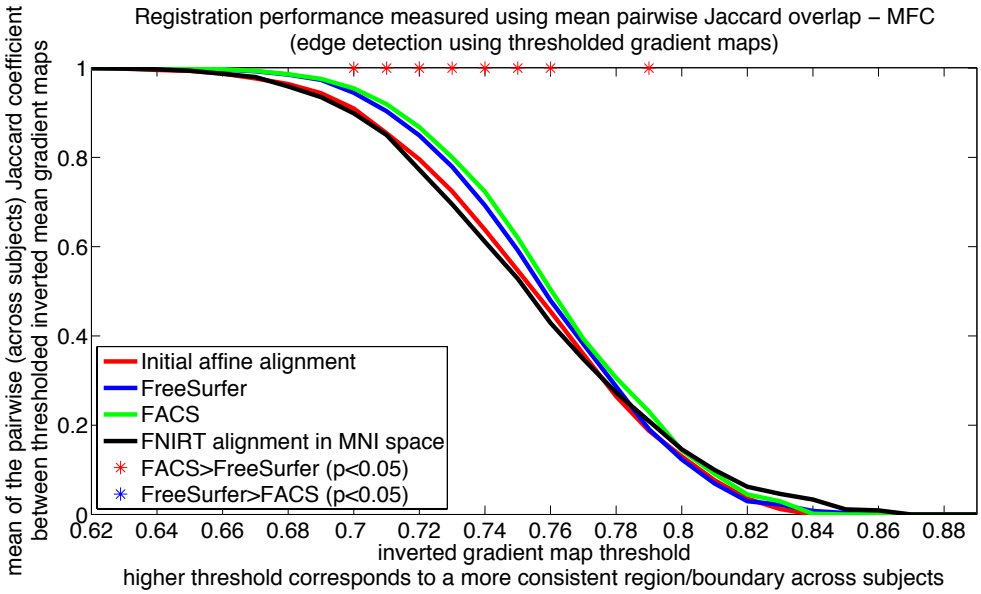

Figure 3.20: The results analogous to those in Figure 3.18 are presented (for the MFC area). Instead of analysing the whole surface of the cortical hemisphere, here we concentrated exclusively on vertices within the MFC mask (Figure 3.1 a). MGIs are used instead of EPMs (EPM-based analysis resulted in poor discrimination between methods; results not shown). Comparison of the inverted MGIs across all subjects for different thresholds. Pairwise Jaccard overlap measure was used to quantify similarity between binarised MGIs. High Jaccard coefficient indicates good overlap. The significance of the results for each threshold value was assessed by a T-test. FACS outperforms FreeSurfer for several threshold levels.

particularly suitable for the RSN-based validation of structural connectivity-driven registration. In the MFC and the cingulate cortex, the diffusion-weighted as well as the RS fMRI data reveal similar functional boundaries justifying their use for cross-validation purposes.

Results To perform the analysis similar to that of section 3.3.2, but for the cingulate cortex only, a liberal cingulate cortex mask was first constructed (we tried to observe the guidelines given in (Beckmann et al., 2009)). The mask was drawn in the native space of the grey/white matter surface of the template subject (see Figure 3.22) and mapped to each of the (other 10) subjects using a given registration procedure (four different algorithms were analysed as before).

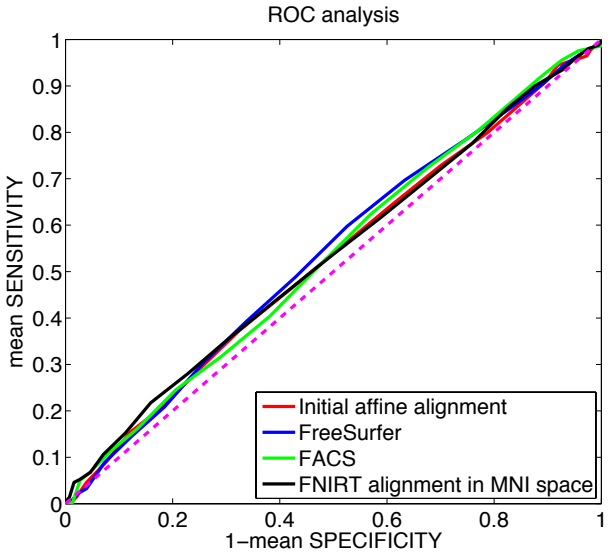

Figure 3.21: The results analogous to those in Figure 3.19 are presented (for the MFC area). The thresholded MGIs can be compared using sensitivity and specificity measures. However, the performance across all methods looks similar (with quite low sensitivity/specificity values) and makes it difficult to assess which algorithm performs best.

Similar to the whole-brain analysis, we now use the mean gradient images (MGIs) in addition to EPMs. The MGIs are (mean) raw gradient images, i.e. without performing Canny edge detection with hysteresis. Effectively, simple thresholding of the raw gradient image has similar effects to applying Canny edge detection. Therefore, the following analysis demonstrates that MGIs can be utilised to draw parallel, and perhaps more robust, conclusions (for this particular application) to those of the EPM analysis. The (mean) gradient maps carry information about the sharp transitions (boundaries) in the image and are also scaled to the $[0, 1]$ interval due to nature of the input data (ETA^2-correlated RSN Feature Vectors). Therefore, high gradient values indicate the presence of an edge or a 'hill' whereas low values are associated with the 'valleys'. As with the EPMs, the inverted MGI assigns high values to the 'valleys' (now becoming 'islands') which we hope to indicate segmented functional regions.

The mean gradient image for a given subject and a given registration algorithm

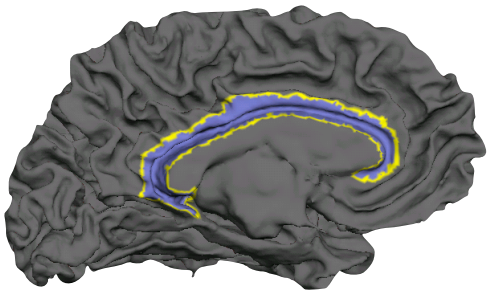

Figure 3.22: The cingulate cortex mask drawn in the native space of the grey/white matter surface of the template subject. The yellow line delineates the mask boundary. The analysis of four registration methods was constrained to the vertices belonging to the drawn mask. The right side of the image corresponds to the anterior of the brain.

was constructed by finding the mean of the gradient maps associated with each of the vertices of the $ic5$ representation of the co-registered grey/white matter surface (see sections 3.3.2 and 3.3.1). An example of the mean (across subjects) of the MGIs showing the cingulate cortex is presented in Figure 3.23 (for FACS and FreeSurfer registration methods). The red colour corresponds to low gradient values which prominently occur in the cingulate sulcus. We speculate that the better the registration, the better the alignment of the functional regions (seen as the red 'valleys' between blue/green edges).

The quantitative comparison between the affine alignment, FNIRT, FreeSurfer and FACS is performed as in section 3.3.2. The results, which summarise the quality of the inter-subject overlap of thresholded MGIs, are shown in Figures 3.24 and 3.26. FACS outperforms FreeSurfer in a more convincing fashion than in the case of the whole-brain analysis: the Jaccard coefficient is consistently higher for FACS than for FreeSurfer for every threshold. Similar conclusions can be drawn from the ROC analysis. FNIRT and the affine registration both perform worse than FACS and FreeSurfer.

These results, interpreted in conjunction with the previous conclusions, suggest that FACS yields improvements over FreeSurfer in the inter-subject overlap of certain

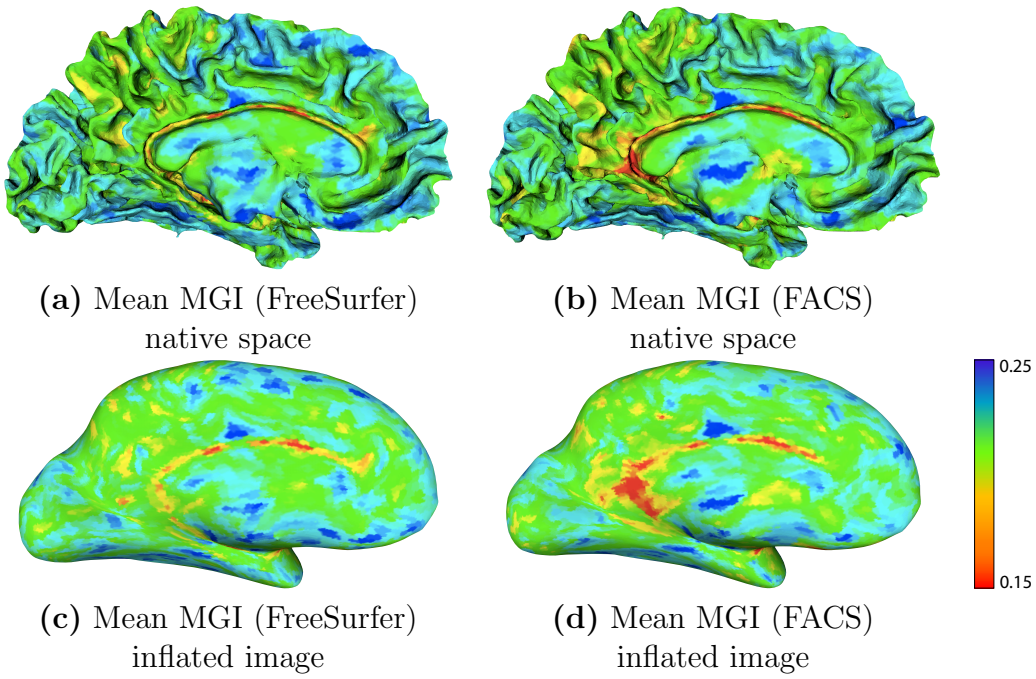

(a) Mean MGI (FreeSurfer) native space

(b) Mean MGI (FACS) native space

(c) Mean MGI (FreeSurfer) inflated image

(d) Mean MGI (FACS) inflated image

Figure 3.23: The mean gradient images (MGIs) carry the information about image boundaries. High gradient values indicate the existence of a boundary/edge. Thresholded MGIs can be used to assess the inter-subject overlap of the putative functional regions defined by the image boundaries (see Figures 3.24 and 3.26). **(a)** The mean (across subjects) of the MGIs for FreeSurfer registration. Results are mapped to the grey/white matter surface of the template subject. **(b)** Same as in (a) for FACS. **(c)** The inflated surface from (a). **(d)** The inflated surface from (b). The prominent red curved 'stripe' seen in the inflated images anatomically resembles and corresponds to the cingulate sulcus. The right side of each image corresponds to the anterior of the brain.

putative functional regions such as the SMA, pre-SMA and the parcellations of the cingulate. At the same time, it is not clear whether these improvements span the whole cortical surface or are just constrained to a number of areas. We are limited by the power of the resting-state functional data (RSNs in particular) to discriminate between functional regions and therefore to perform inter-subject comparison. At the same time, these constraints open up exciting questions about how RSN data and the structural connectivity can be jointly used for registration purposes so that they complement each other. The next section 3.3.4.1 investigates the inter-subject overlap of the *clusters* into which the cingulate cortex can be parcellated using RSNs.

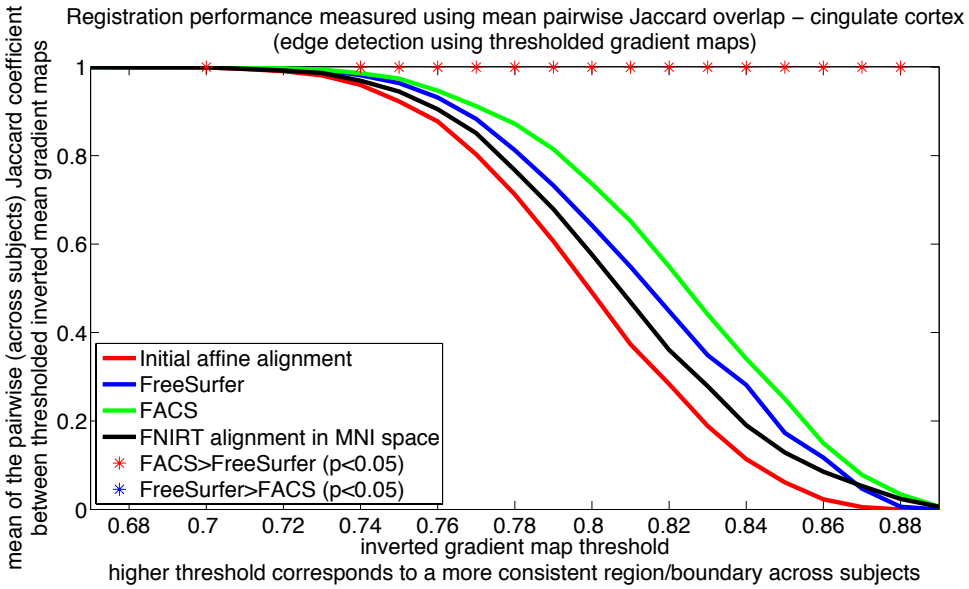

Figure 3.24: The results for mean pairwise Jaccard overlap in the cingulate cortex, analogous to those in Figures 3.18 and 3.20. Instead of analysing the whole surface of the cortical hemisphere, here we concentrated exclusively on vertices within the cingulate mask (Figure 3.22). MGIs are used instead of EPMs (both analysis were performed - see Figure 3.25). Comparison of the inverted MGIs across all subjects for different thresholds. Pairwise Jaccard overlap measure was used to quantify similarity between binarised MGIs. High Jaccard coefficient indicates good overlap. The significance of the results for each threshold value was assessed by a T-test. FACS outperforms FreeSurfer for almost every threshold level.

3.3.4.1 Cingulate Area Clustering Using RSNs

Each RSN Feature (RSNF) consists of the samples of Z-statistic values of a number of RSNs (see Figure 3.6)[11]. So far, we used 33 RSNs found to be predominantly in the left cortical hemisphere and assumed to be of non-artifactual origin. These RSNFs are referred to as 'valid' (or valid RSN features/components). However, the original RS fMRI data are decomposed into 70 RSNs and the RSNFs can be formed using all of them. Such RSNFs are called 'full' (or full RSN features/components) to stress that no data exclusion was performed. The MFC clustering was performed using valid features only as this 'filtering' procedure proved beneficial to obtaining consistent clustering results. In the analysis of functional boundaries and their inter-

[11]Note the distinction from RSN Feature Vectors that are ETA^2-correlated RSNFs.

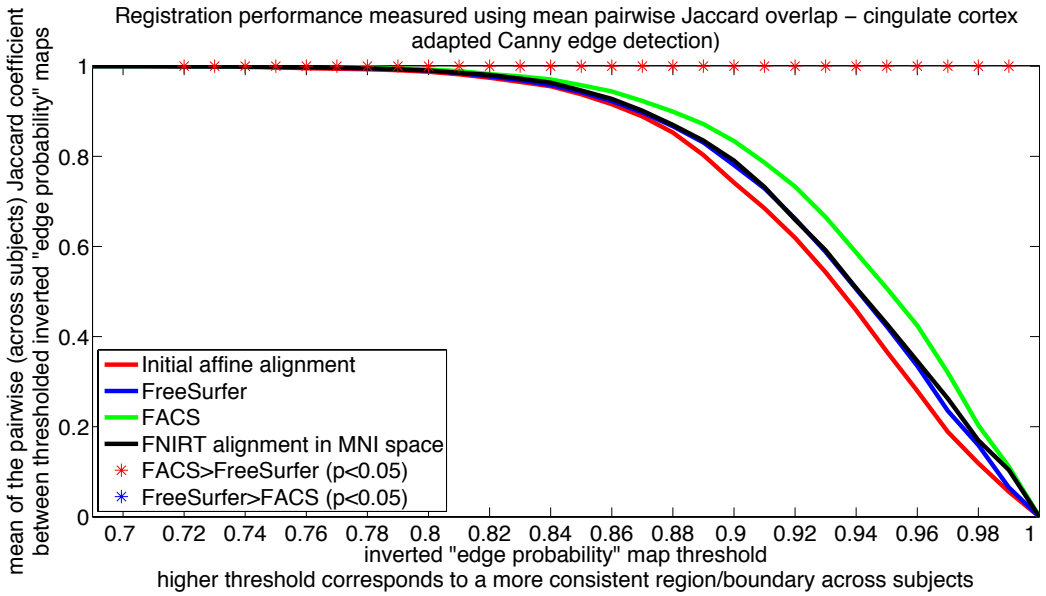

Figure 3.25: Continuation of Figure 3.24. Thresholded 'edge probability' maps are used to the same effect as the mean gradient images to demonstrate differences in performance between FACS and FreeSurfer. However, this edge detection method does not discriminate well between FreeSurfer and FNIRT (unlike using MGIs).

subject overlap (both in the whole-brain setting and for the cingulate cortex only) conducted to this point, we repeated the analysis with full RSN features too, to find no differences to the final conclusions and presented results. However, for the analysis in this section, the full and valid RSN features yield different results and are thus shown in parallel.

RSNFs (both full and valid) can be used to perform cingulate clustering. Each vertex i belonging to the cingulate mask has an n-dimensional RSN Feature associated with it (RSNF_i). If all such vectors associated with every cingulate mask vertex are stacked as columns of a RSN Feature Matrix ($\text{RSNFM}_{n\times N}$, N is the number of vertices in the mask), the square matrix $\text{ETA}^2(\text{RSNFM}, \text{RSNFM})$ can be clustered just like the matrix consisting of RSN Feature *Vectors* was clustered in the SMA/pre-SMA example (section 3.2.3.2). ETA^2 measures similarities between rows of the RSNFM.

Beckmann et al. (2009) report 9 prominent clusters in the cingulate cortex found through the joint analysis of diffusion and functional data. Following their example,

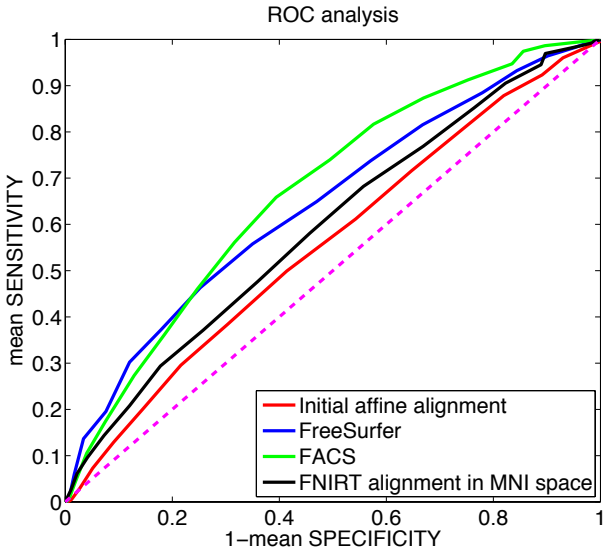

Figure 3.26: The results for the ROC analysis of the corresponding overlaps in the cingulate cortex analogous to those in Figures 3.17 and 3.21. The thresholded MGIs can be compared using sensitivity and specificity measures. The ROC analysis confirms findings from Figure 3.24. If judged by the *area under the curve* measure, FACS performs better than FreeSurfer and both outperform FNIRT and the affine alignment.

we performed K-means clustering (9 clusters) of the cingulate mask of each of the 11 subjects (for both full and valid features) and co-registered the results using four different registration methods. Imposing the inter-subject cluster correspondence was done through a procedure detailed in the appendix D. Figure 3.27 shows (left and middle) the single subject ICA-based RSN clustering as well as the 'prototype clustering' (right) found by the analysis of corresponding clusters across all subjects. As in the case of diffusion based clustering (Figure 2.14), the boundaries run in the dorsal-ventral direction which is consistent with the functional and cytoarchitectonic division of the cingulate area (Vogt, 2009).

The ICA-based RSN clustering can be performed using both valid and full features. To assess how the within- and the inter-subject clustering consistency differs, two types of analysis were performed (Figures 3.28 and 3.29). The within-subject clustering consistency was measured by assessing (in feature space) for each point how similar it is to the other points from its own cluster, compared to the points

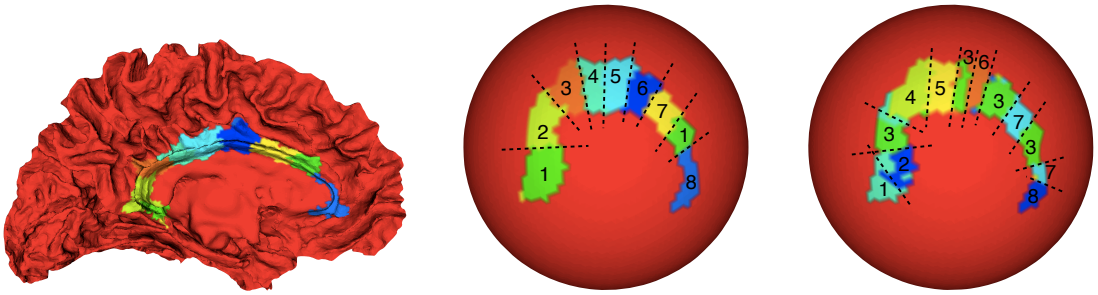

Grey/white matter surface
a single subject example
left-posterior, right-anterior

Registration space
a single subject example
(as on the left)

Registration space
'prototype' clustering
across subjects

Figure 3.27: Different colours/labels correspond to different clusters. **(left)** An example of the ICA-based RSN clustering of the cingulate into 9 clusters. Results shown in the native space of the grey/white matter surface. **(middle)** The clustering from (left) shown in the registration space of the $ic5$. 8 different clusters can be discerned. Note that the clustering algorithm does not use any spatial relations (e.g. vertex neighbourhood information). **(right)** The 'prototype clustering' showing the common clustering pattern across all subjects. 8 different cluster can be discerned (note that same cluster label can appear at different spatial locations - this is not prevented by the clustering algorithm or by the procedure imposing the inter-subject clustering consistency. However, possible 'split-clustering' should not affect comparison of registration methods described further in this section).

from the other clusters (commonly used silhouette value, scaled to $[-1, 1]$ interval). The silhouette value for point i is calculated as:

$$S(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}, \tag{3.6}$$

where $a(i)$ is the average distance of point i to all other points of the same clusters and $b(i)$ is the minimum (across clusters) of the average distance of point i to all other points in all other clusters (for details on computation see (Kaufman and Rousseeuw, 1990)). Such a measurement is performed for every point within the clustered area for a number of clustering iterations (Figure 3.28). In each clustering iteration the clustering algorithm is randomly re-initialised. The results contradict the expectations: the full consistently outperform valid features. Such a finding is robust with respect to the clustering cost function, the number of replications as well as the point-to-cluster similarity measure (correlation coefficient tried). The detailed

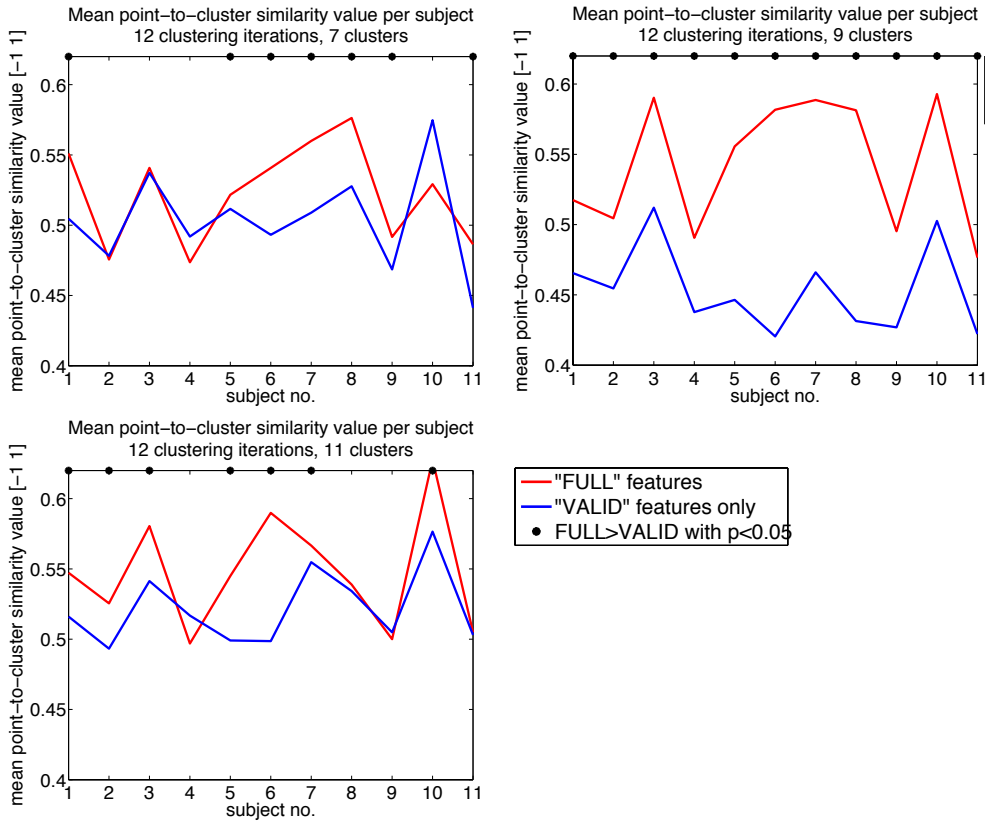

Figure 3.28: For each subject (11 in total) the within-subject clustering consistency was assessed by calculating the silhouette value for multiple independent clustering initialisations/iterations. The silhouette value of a given point measures how similar it is to the other points from its own cluster, compared to the points from the other clusters. The higher the score, the greater clustering consistency. Mean silhouette values across subjects suggest the best separation between full and valid features for the 9-cluster case. In all cases, full features (red) yield better within-subject clustering consistency than the valid features (blue).

investigation into the causes for such results is left for future work. However, one of the possible explanations for the described behaviour might be that the excluded RSN components still carry information important for cingulate clustering. The delineation of components as 'artifactual' was performed manually, and there is no guarantee that these components do not contain some true, and possibly valuable, signal. Both the full and valid features give visually very similar parcellations making it difficult to conclude which one of them more closely follows the expected divisions.

The silhouette value was calculated for varying number of clusters to examine

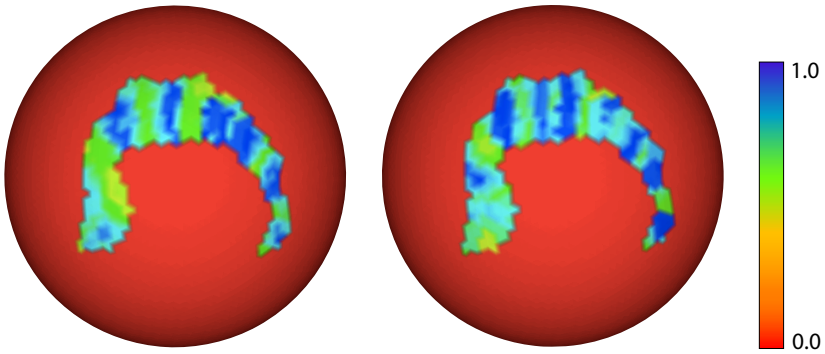

'Valid' features 'Full' features

Figure 3.29: After imposing inter-subject clustering consistency, the overlap of corresponding clusters across subjects was measured for multiple independent clustering initialisations. Each vertex in the cingulate mask reflects (see the colourbar) the normalised number of *correct* cluster overlaps at that point. The blue areas can be associated with consistent clusters, whereas the green and yellow ones correspond to the regions of inconsistent clustering. **(left)** Clustering using valid features. **(right)** Clustering using full features.

how the results differ w.r.t. full and valid features. The conclusions described above remain true for 7-, 9- and 11-cluster cases (note that 9 is the number of clusters estimated using structural connectivity analysis (Beckmann et al., 2009)). Further discussion on the effect of the number of clusters on the conclusions about FACS performance are given at the end of this section.

In addition to the within-subject clustering consistency, Figure 3.29 presents the results assessing the inter-subject consistency. Each vertex in the cingulate mask reflects (see the colourbar) the normalised number of *correct* cluster overlaps at that point. For example, should a certain vertex belong to the cluster X in every subject and for every clustering iteration, the number associated with it would be 1 (shown in blue)[12]. The blue areas can be associated with consistent clusters, whereas the green and yellow ones correspond to regions of inconsistent (or less consistent) clustering, i.e. putative cluster borders. The full feature-based clustering has overall higher clustering consistency measures, although the difference between the two methods is

[12]We assume that the inter-subject correspondence has already been established and the clusters relabelled.

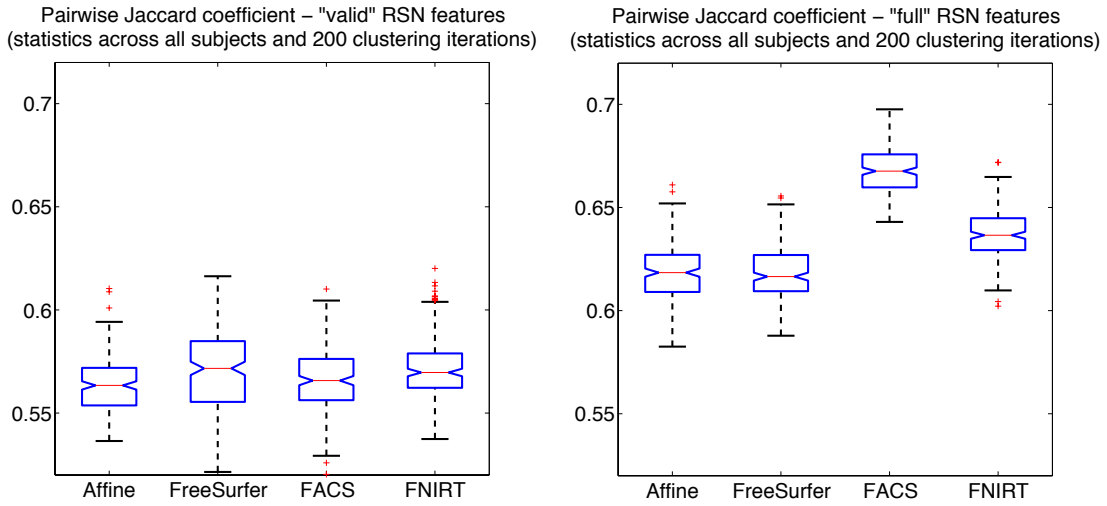

Figure 3.30: Extended pairwise Jaccard measure was used to assess the inter-subject overlap of corresponding clusters within the cingulate area for four different registration methods. **(left)** Clustering (9 clusters) using valid features. FreeSurfer performs significantly better than FACS. **(right)** Clustering (9 clusters) using full features. FACS performs significantly better than FreeSurfer, FNIRT and affine registration. The red line shows the median of the sample. Two medians are significantly different at the 5% significance level if their intervals do not overlap. The interval endpoints are the extremes of the 'notches'.

not striking.

Finally, we investigated how different registration methods align clusters extracted with both full and valid features. To measure the similarity between the two cingulate parcellations, we used an extended pairwise Jaccard measure (Postelnicu et al., 2009) capable of handling multiple mask labels (the traditional Jaccard measure compares binary images). This is necessary since the number of clusters in each parcellation was 9 (therefore 9 different labels were used). For each pair of subjects (fixing one as the template for the purposes of comparison only) the extended pairwise Jaccard coefficient was calculated. This calculation was performed for 200 independent random clustering initialisations (for four registration methods) and the results are summarised as *box plots* in Figure 3.30.

Results The results suggest that for full features, FACS performs significantly better than FreeSurfer. At the same time, the affine alignment and the FreeSurfer

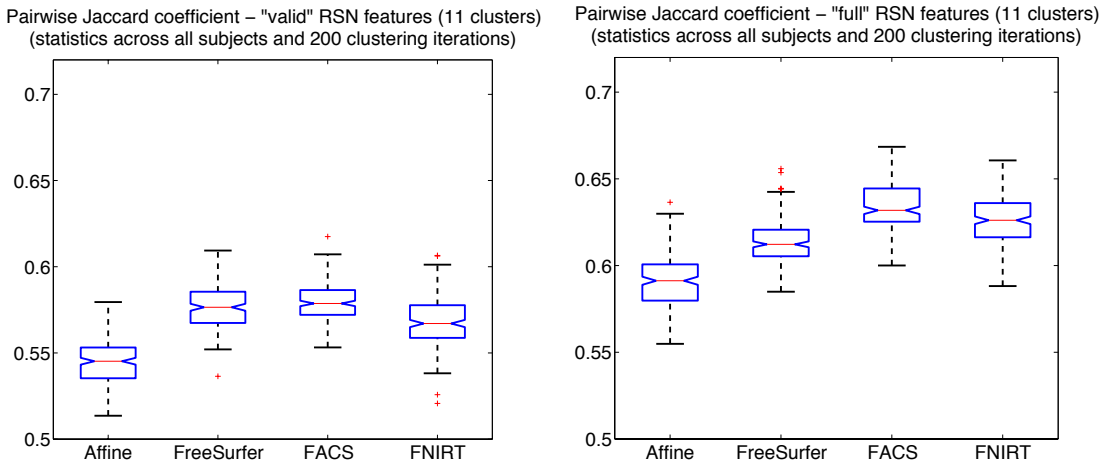

Figure 3.31: The inter-subject overlap of corresponding cingulate parcels for different registration methods (measured using the extended pairwise Jaccard coefficient - see Figure 3.30). The analyses for full **(right)** and valid **(left)** features are shown for 11 clusters. FACS performs the best for both full and valid features.

perform similarly, whereas FNIRT falls in between FACS and FreeSurfer. In the case of valid features however, FreeSurfer performs significantly better than FACS and affine alignment and FACS perform similarly. The overall values of the extended pairwise coefficient for full features analysis are higher than for valid features only. This is consistent with the findings presented in Figures 3.29 and 3.28 indicating better clustering with full features. Somewhat surprisingly, in the case of full features, FNIRT performs better than FreeSurfer. Nevertheless, note that the main part of the cingulate mask is the cingulate sulcus, which within itself is not 'rich' in features that are useful for FreeSurfer registration. There are indications that in such cases FNIRT might outperform FreeSurfer-based registration (Chapter 5, section 5.2.1) especially in the deep grey matter structures.

Furthermore, the results presented in this section are not robust with respect to the number of clusters into which the cingulate cortex is parcellated. The analysis was repeated with 7 as well as 11 clusters, in both cases showing better within- and inter-subject clustering parameters for full compared to valid features. However, the registration performance differed - Figures 3.31 and 3.32. There could be multiple

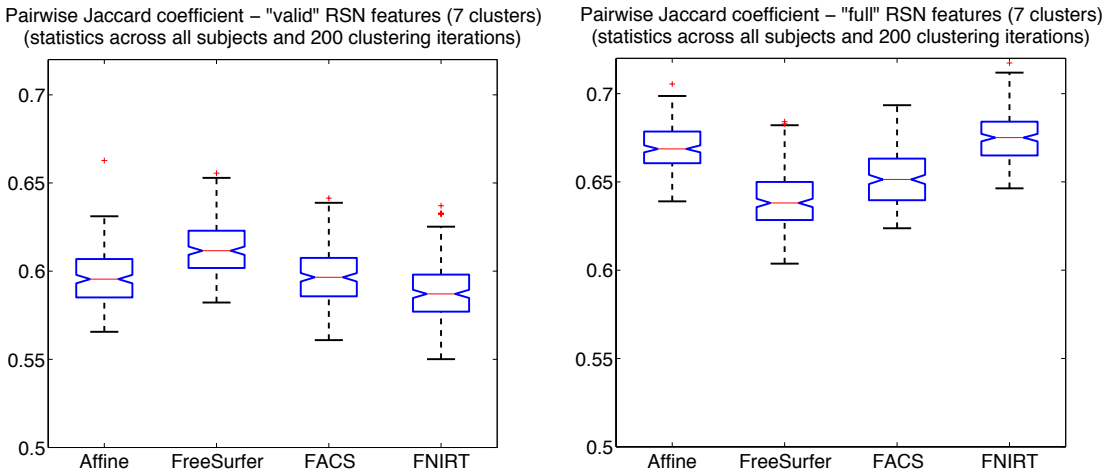

Figure 3.32: The inter-subject overlap of corresponding cingulate parcels for different registration methods (measured using the extended pairwise Jaccard coefficient - see Figure 3.30). The analyses for full **(right)** and valid **(left)** features are shown for 7 clusters.

reasons for these inconsistencies. Some of them certainly relate to the ability of the RSN parcellation method to discern 'real' functional regions, the others might be tied to the nature of the processing/clustering algorithms (e.g. not including vertex neighbourhood information, potential pitfalls in imposing clustering consistency as performed here). So far, neither the actual number of clusters identifiable through the RSN analysis nor the functional plausibility of the clusters themselves have been confirmed and validated. To investigate this further, it would be beneficial to have the 'ground truth' about the functional divisions (as well as their number) within the cingulate cortex and make comparisons with the ICA-based RSN parcellations for both the full and valid features. This thesis is concerned with validating the proposed *registration method* and the actual parcellation algorithms are developed as side tools to help achieve the primary goal. Although we proposed a novel method for the ICA-based RSN parcellation of the cingulate, the detailed assessment of its validity is yet to be performed and falls beyond the scope of the current research. Still, the analysis of the 9- and 11-cluster cases points out that FACS aligns putative cingulate subdivisions better than FreeSurfer. Figures 3.30 and 3.31 suggest that

FACS performs better than FreeSurfer for the full features case, whereas Figure 3.28 shows that full features result in better clustering compared to the valid ones. The 7-cluster case (Figure 3.32) is not consistent with these findings, possibly due to poor clustering performance. Moreover, we should also note that the cingulate cortex is an area of high structural variability (the number and shape of sulci/gyri) and the parcellation into a single canonical arrangement might not be feasible (Paus et al., 1996). This could additionally complicate the registration assessment using the parcellation-based approach, and speak in favour of the 'functional edges' and the comparison of 'edge probability' maps (section 3.3.4).

3.3.4.2 Comparison to MFC Clustering Using RSNs

In section 3.2.3.2 we performed the SMA/pre-SMA clustering using (valid) RSN Features. However, to do so, the manual cluster label assignment was performed. Here we repeat the clustering procedure but with automatic SMA/pre-SMA cluster label assignment. Therefore, the results that follow are directly comparable to those presented in e.g. Figure 3.31 examining the ICA-based RSN clustering of the cingulate. Figure 3.33 demonstrates that the automatic cluster labelling conclusions are analogous to those presented in Figure 3.7. FACS outperforms all other registration methods and more convincingly so for the 'valid' compared to 'full' features.

3.4 Applications

The potential applications of cortical connectivity-driven registration are numerous. Some of them were already mentioned in Chapter 2 section 2.4.1.4. Here we concentrate on two:

1. The construction of the cortical structural connectivity template (section 3.4.1);

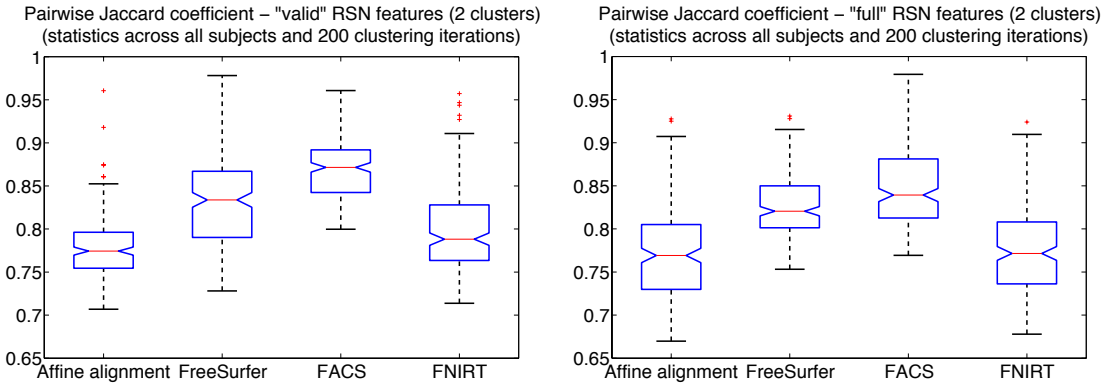

Figure 3.33: ICA-based RSN clustering of the MFC into 2 clusters (putative SMA and pre-SMA, see Figure 3.7). The inter-subject overlap of corresponding parcels for different registration methods was measured using the extended pairwise Jaccard coefficient. The analyses for full **(right)** and valid **(left)** features are shown. FACS performs the best for both full and valid features. This is consistent to the analogous findings presented in Figure 3.10.

2. The application of such a template in the analysis of differences in brain connectivity between Alzheimer's disease (AD) patients and healthy controls (section 3.4.2).

3.4.1 Construction of The SCFV Template

So far in this Chapter we used one randomly chosen subject to act as a template. This is tied with some disadvantages such as potential alignment biases towards that subject's cortical anatomy (structural connectivity). Here we propose a simple procedure for construction of a 'template of structural connectivity' in the form of a *SCFV Template* or a *Connectivity Matrix (CM) Template* (for SCFV and CM construction see Chapter 2 section 2.4.3.1). Although not completely free of potential one-subject-biases, this iterative template construction procedure forms an average of SCFVs of a group of subjects, reducing the influences of one subject's anatomy over the template. Subjects forming the template can, for example, be healthy controls age-matched to a group of subjects with pathologies. The template can therefore be entirely study-specific.

The template construction procedure is similar to that of Sabuncu et al. (2010) used for deriving the template for fMRI-driven cortical registration. The algorithm proceeds as follows: the CM, which consists of SCFVs, of an individual subject is selected as the initial group template. The alignment then proceeds by sequentially registering the remaining subjects to the group template using the pairwise FACS algorithm introduced in Chapter 2 section 2.4.4. After each registration, the template is refined by averaging corresponding SCFVs of the subjects that have been co-registered. Thus, subject n, $2 \leq n \leq N$, is registered to the average of subjects 1 through $(n-1)$, where N is the total number of subjects[13]. The connectivity matrix obtained after final averaging is called the *template connectivity matrix*

We used this algorithm to construct a template $\text{CMT}_{10242\times 89}$ for the group of 11 control subjects (section 3.2.1). To examine the structure of this matrix the ICA decomposition was performed on CMT into 80 'spatial' components, each of which is a vector of length 10242. The dimension, 80, was chosen to test once more for possible residuals of the 'triangular parcellation' used for SCFV construction. If these triangular spatial patterns were to be seen in the CM (see Figure 2.8), that would indicate the presence of biases originating from the nature of SCFV construction (see section 2.4.3.1). However, such triangular patterns were not observed. Quite the contrary, some of the identified components had anatomically meaningful structure (unlike the original cortical parcellation which does not follow anatomical borders). This suggests that the connectivity information encoded in the SCFVs has the ability to discriminate between different cortical structures, such as for example, the insular cortex and the hippocampal sulcus (Figures 3.34 and 3.35). Interestingly, some of the isolated components also resemble the RSNs derived from fMRI data, and many are relatively focal, contiguous cortical 'parcels', while others are more distributed across

[13]Note that the template obtained in this way can further be refined and the remaining potential biases related to the choice of the initial subject/template can be minimised as in (Sabuncu et al., 2010).

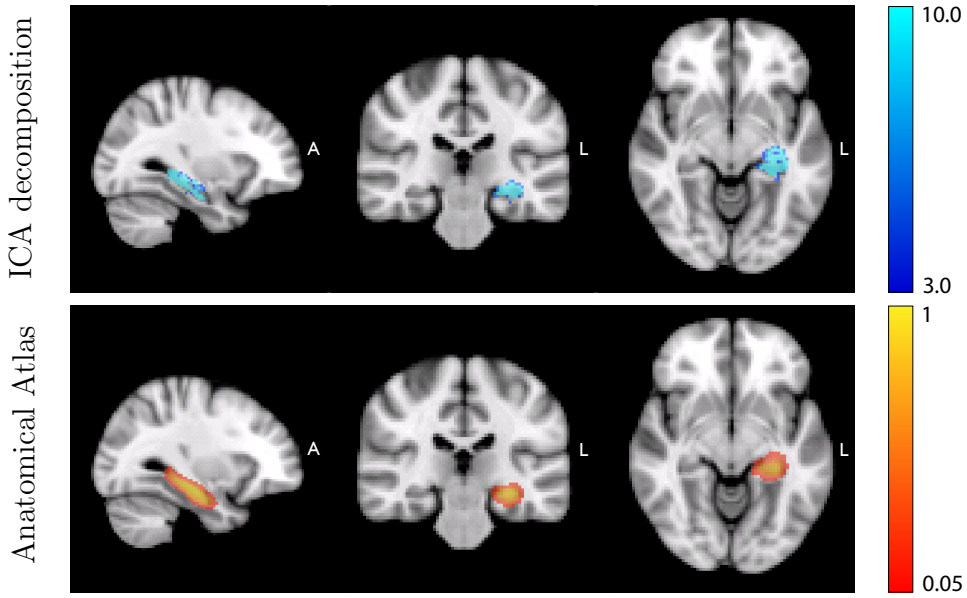

Figure 3.34: ICA estimation of the independent spatial patterns in the Template Connectivity Matrix. The surface-based results were projected into the cortical grey matter of the volumetric image (MNI152 template brain) using nearest neighbour interpolation (between vertices and voxels; *mri_surf2vol* command was used, part of FreeSurfer). **(top row)** Sagittal, coronal and axial views of the component resembling the left hippocampal sulcus overlaid on the MNI152 template brain. The colourbar represents the Z-statistic values. **(bottom row)** Coronal, sagittal and axial views of the left hippocampus probability map, for comparison purposes, overlaid on the MNI152 template brain. The hippocampus probability map was generated using Harvard-Oxford Subcortical Structural Atlas, part of the FSL. The colourbar represents probability values. A - anterior, L - left.

the cortex (though still generally appearing to have plausibly spare, as opposed to random, spatial structure).

3.4.2 Analysis of an Alzheimer's Disease Data Set

One of the other possible applications of the cortical connectivity-driven registration is the comparison of deformation fields between the two groups. Here we present results of the comparison of the healthy control and Alzheimer's Disease (AD) cohorts and show the differences in connectivity related to the differences in the deformation fields.

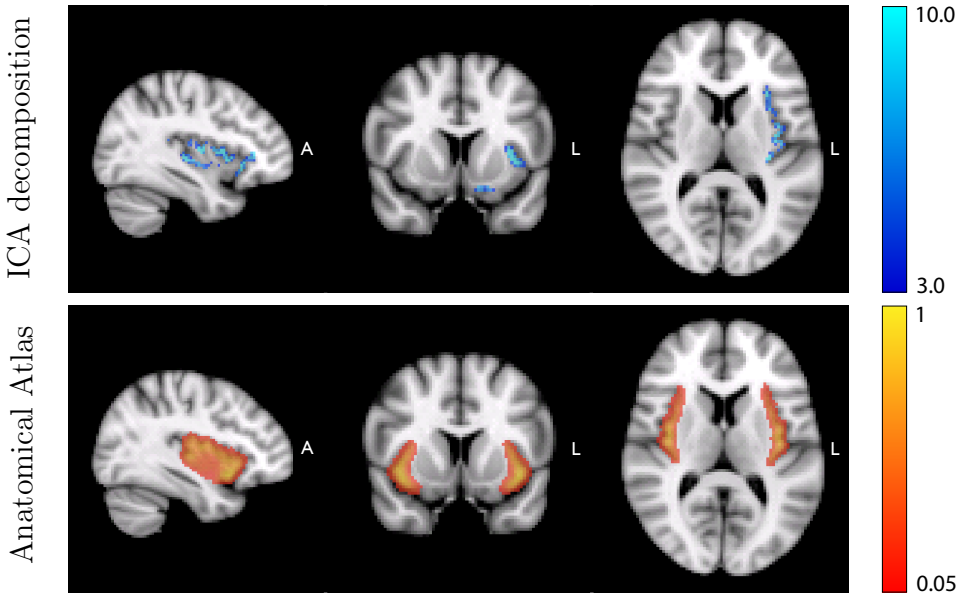

Figure 3.35: ICA estimation of the independent spatial patterns in the Template Connectivity Matrix. The surface-based results were projected into the cortical grey matter of the volumetric image (MNI152 template brain) using nearest neighbour interpolation (between vertices and voxels; *mri_surf2vol* command was used, part of FreeSurfer). **(top row)** Sagittal, coronal and axial views of the component resembling the left insular cortex overlaid on the MNI152 template brain. The colourbar represents the Z-statistic values. **(bottom row)** Coronal, sagittal and axial views of the insular cortex probability map, for comparison purposes, overlaid on the MNI152 template brain. The insular cortex probability map was generated using Harvard-Oxford Cortical Structural Atlas, part of the FSL. The colourbar represents probability values. A - anterior, L - left.

3.4.2.1 Data Acquisition

The data were generously provided by Dr Achim Gass, Dr Andreas Monsch (Basel), Prof. Paul Matthews, Dr Brandon Whitcher and Dr Anil Rao (GSK). This imaging study was part of the EAGLE (Early Alzheimers Disease Genetics - A Longitudinal Evaluation) conducted in Basel, Switzerland. In total, 135 subjects were studied (53 probable AD patients without vascular component, 56 MCI and 26 controls subjects). Here we analysed only 24 subjects (12 controls and 12 mild probable AD with MMSE $\geq$ 21).

For all 24 subjects, whole-brain diffusion-weighted data were acquired with a 3 T Siemens Allegra MR tomograph (Erlangen, Germany) with a standard quadrature

head coil and maximum 40 mT/m gradient capability. DTI images were acquired using echo-planar imaging (SE-EPI: TE/TR = 89 ms/7000 ms, 54 axial slices, bandwidth = 2056 Hz/vx, voxel size $2.5 \times 2.5 \times 2.5$ mm^3) with 30 isotropically distributed diffusion gradient orientations at a b-value of 900 s/mm^2 and six $b = 0$ images. Scanning was performed twice to improve the signal-to-noise ratio.

3.4.2.2 AD vs. Control Groups

We analysed 24 subjects $(12 + 12)$ of healthy controls and early Alzheimer's patients matched for age, gender, handedness, and years of education (for diffusion data acquisition parameters and other study details see section 3.4.2.1). First, the study-specific template was constructed for both left and right cortical hemispheres of control subjects as described in section 3.4.1. Each of the 24 subjects was registered to the Template CM. The registration output of an input subject to the template is a deformed $ic5$ mesh. Comparing the deformation fields between the two groups therefore practically means comparison of the deformed $ic5$ meshes (called the *output* meshes). Let us imagine all output meshes as the overlapping spheres (see section 2.4.4.1). If no registration was performed (or all subjects were identical to the template), all vertices with the same labels would have the same position (would 'overlap'). However, due to registration, set(s) of vertices are shifted around. Therefore, if a set of vertices[14] in in one group of subjects is shifted to a *consistently* different position compared to where the same set of vertices is mapped in the other group of subjects, then the deformation fields differ for that particular set of vertices. The assessment of the statistical significance of vertex shifts can be done for each vertex separately. Let $\vec{v}_i^k$ $k \in [1, 12]$ be a vector encoding the position of vertex i in an output mesh of subject k from the AD group. Similarly, let $\vec{p}_i^k$ $k \in [1, 12]$ be a vector encoding the position of vertex i in an output mesh of subject k from the control group. For each vertex i, we ask whether sets of *vectors* $\{\vec{v}_i^1, \ldots, \vec{v}_i^{12}\}$ and

[14]Vertices are defined by their associated labels.

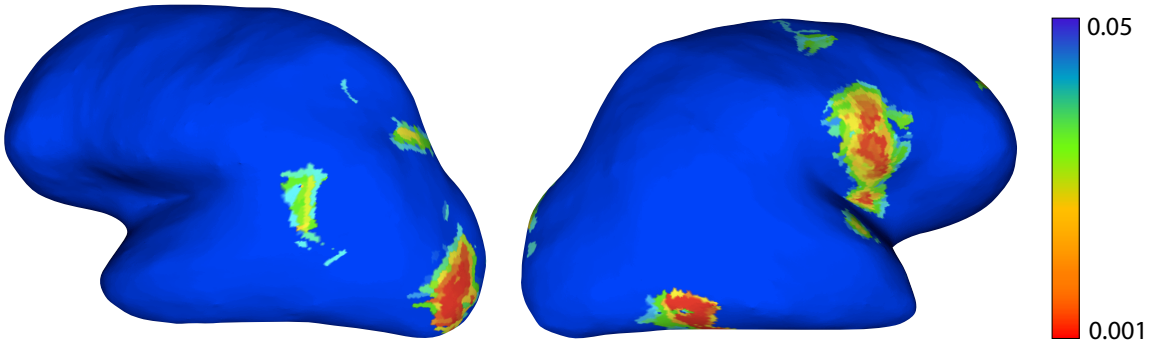

Left hemisphere - lateral side
inflated surface, left - anterior

Right hemisphere - lateral side
inflated surface, right - anterior

Figure 3.36: The differences in the deformation fields between the AD and the control groups were examined using the multivariate Hotelling's T^2 test on vectors encoding vertex positions after co-registration. The F-statistic was formed for each vertex and the p-value on differences in mean positions of that vertex in the two groups was computed. **(left)** The lateral side of the inflated left hemisphere of the prototype subjects. **(right)** The lateral side of the inflated right hemisphere of the prototype subjects. Non-blue areas have different displacement fields in the two groups. The colourbar represents the uncorrected p-values.

$\{\vec{p}_i^1, \ldots, \vec{p}_i^{12}\}$ are statistically different. A commonly used statistical test capable of answering this question is the multivariate Hotelling's T^2 test (for the construction of the Hotelling's T^2 statistic see (Lu et al., 2005)). Therefore, the statistic significance between the *means* of the two sets of vectors can be expressed as a p-value and associated with a corresponding vertex.

Such a comparison was performed between the AD and control groups and the results (in terms of spatial p-value maps) were upsampled (nearest neighbour interpolation) and registered to the inflated space of one randomly chosen subject (for visualisation purposes) - Figure 3.36. The non-blue areas have a p-value lower than 0.05. Note that no correction for multiple comparison was performed. When FDR (false discovery rate) correction is applied (Benjamini and Hochberg, 1995), no vertices pass the $q < 0.05$ thresholding (see discussion at the end of the Chapter). The differences in the deformation fields in these areas are driven by the differences in connectivity of those areas to the rest of the brain.

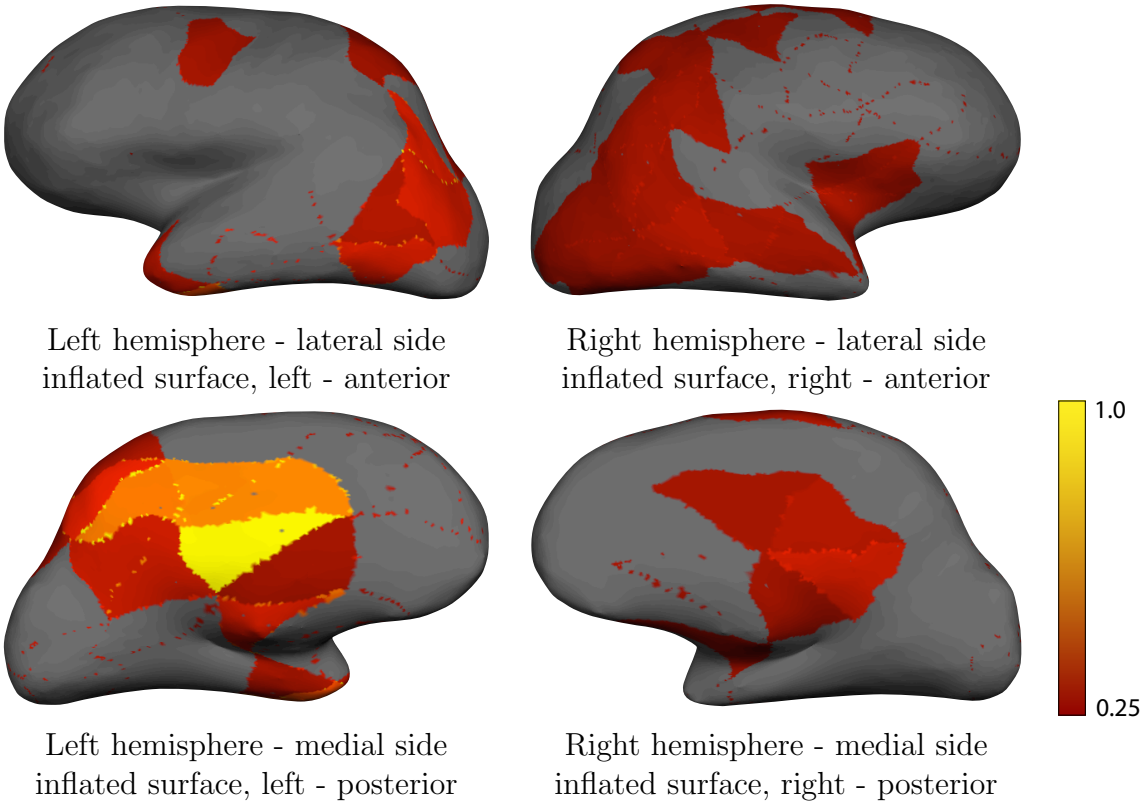

Figure 3.37: The connectivity of the areas with different deformation fields (for $p < 0.05$ - Figure 3.36) to the set of triangular cortical targets was computed. The yellow-red regions show triangular cortical targets which are 'most strongly' connected with the areas identified in Figure 3.36. The measures of connectivity to triangular targets are taken from the Template CM and were scaled to the $[0, 1]$ interval (see the colourbar). The lower threshold was arbitrarily set to 0.25 for display purposes.

To investigate the connectivity of those areas we visualised the outputs of probabilistic tractography from areas/vertices with $p < 0.05$. The mean of the SCFVs (from the Template CM) for vertices in those areas was calculated and the values across the mean SCFV were normalised to the $[0, 1]$ interval. The values of the mean SCFV can be related to the original triangular parcels (Figure 2.7). Figure 3.37 shows coloured triangular parcels for the left and right hemispheres (inflated surfaces). The colour of a parcel corresponds to a normalised mean connectivity measure (obtained through probabilistic tractography) of the connectivity between that parcel and the identified areas of 'significant' difference in the deformation fields between the groups.

The areas identified in Figure 3.37, primarily the temporal and parietal lobes as

well as the cingulate cortex, have been already reported to have compromised white matter connectivity in AD. Huang et al. (2007) report decreased fractional anisotropy (AD) in the temporal and parietal lobes of the early AD patients. Similarly, a voxel-based morphometry study (VBM) of Rose et al. (2008) report elevated mean diffusivity measures (MD) in AD patients in the medial temporal lobe as well as the parietal lobe and the posterior cingulate gyrus. Furthermore, Rose et al. (2000) and Xie et al. (2006) report significant reduction in the integrity of the association white matter fibre tracts (in AD patients), such as the splenium of the corpus callosum, superior longitudinal fasciculus, and cingulum and confirm the involvement of parietal and temporal lobes. Salat et al. (2010) find that white matter connectivity is altered in AD especially in the parahippocampal region and the regions with primary and secondary connections with the middle temporal lobe. The findings from Figure 3.37 suggest similar areas, although not with strong bilaterality. Still, the interpretability of our results is open and many questions remain. In particular, it is left to future investigations to answer why the differences in the deformation fields (Figure 3.36) exhibit even lower left-right hemisphere symmetry than the results from Figure 3.37. Likewise, the findings have limited statistical interpretability, as no multiple comparison correction was applied although there exist several vertices with p-values well beneath 0.05 (e.g. < 0.001). It would also be of interest to identify which white matter fibre bundles correspond to the areas highlighted in Figure 3.36. Finally, the question of the suitability of the template and the robustness of the results w.r.t. the template need to be further examined. Some of these questions are addressed again in Chapter 4 section 4.3.2.

3.5 Discussion

SMA/pre-SMA Analysis The results presented in section 3.2.4.3 indicate that FACS outperforms FreeSurfer in aligning two neighbouring but different functional

areas. The interpretability of these results depends on an important issue of the interpretability of the ICA-based RSN clustering. However, the putative SMA/pre-SMA clusters were found using a novel clustering procedure which yields results better than those of Kim et al. (2010) and therefore strengthens our conclusions about the examined registration methods. Still, our conclusions depend on how well can the ICA-based RSN clustering reveal true functional boundaries. Therefore, we hope that the development and validation of more accurate functional parcellation techniques would also result in better assessment of registration methods.

One should also acknowledge the existence of potential registration biases introduced through the Dual Regression procedure. The DR relies on the affine registration of functional data to the standard space. The quality of this registration affects the extraction of spatiotemporal maps which are then used for ICA-based RSN clustering and RSN-based boundary detection. Nevertheless, our goal was to demonstrate how RSN Feature Vectors (which are consistent across subjects) can be utilised for cortical parcellation, and equally importantly, to show that these Feature Vectors can be used instead of the SCFVs to achieve RSN-based connectivity-driven registration. The existence and the effects of potential biases, as well as the influences of data quality (e.g. the number of diffusion directions, voxel size) on results and their reproducibility should be carefully examined in future work.

Regarding the SMA/pre-SMA clustering of this particular dataset, further research is needed to investigate the nature of the clustering artefact (Figure 3.3). Additional investigations should also shed light on how RSN- and structural connectivity-driven alignment procedures can be combined. Although both of these approaches have the same goal of aligning functional boundaries across individuals, it might be the case that the RSN-driven (or FC-driven for that matter) registration performs better in some areas than the structural connectivity-driven one.

Analysis of Boundary Detection and Alignment The challenges of performing the detection of edges between functional regions are tied to the already mentioned interpretability of the resting-state data as well as the suitability of the particular edge detection method employed here. The high variability of the edges detected through procedures proposed in section 3.3.1 prevents straight-forward comparison of corresponding edges across subjects. The 'edge probability maps' reveal this variability (Figures 3.13 and 3.15) through high degree of uncertainty in edge position. The results of comparing such maps across different registration methods indicate that FACS and FreeSurfer perform similarly when the whole cortex is analysed, whereas FACS outperforms FreeSurfer when analysing just the cingulate area or the MFC. Therefore, the question remains of what are the other areas (if any) where the registration improvements can be seen, as well as why these improvements are not seen in the other areas. These questions relate to the already discussed issues of the applicability of connectivity-driven registration (Chapter 2). However, when 'mean gradient maps' (which were found to result in better discrimination between registration methods) are used instead of 'edge probability maps', FACS outperforms FreeSurfer even in the whole-brain case (in addition to the MFC and cingulate areas).

Applications With respect to the applications of the connectivity driven registration to the inter-group comparison (section 3.4.2), we have briefly presented one way of approaching the task of comparing the deformation fields. The in-depth analysis of obtained results hinges on the interpretability of the maps showing the displacement differences (Figure 3.36). However, in this case, possibly due to the small group size, the differences between groups do not survive an FDR multiple-comparison correction. Still, FDR does not take into account any measure of spatial location, and since the results that are shown are highly clustered, it is likely that a cluster-based multiple-comparison technique, which we did not have an implementation of, would

be more sensitive. Additional tests, especially addressing the issues of repeatability and robustness w.r.t. the template, should complement these findings and open doors to the construction of novel imaging biomarkers.

The template CM The clustering of the template CM yielded parcels that highly resemble functional/structural cortical subdivisions suggesting that the identified areas have distinct connectivity patterns common to a group of subjects. Since the templates can be derived for different cohorts separately, their comparison has a potential to yield an additional insight into group connectivity patterns. As we have shown, these patterns can either be assessed using deformation fields, or can potentially be examined by comparing corresponding parcellations. However, further research is needed into how these parcellations can be used for group comparison.

"Remember now thy Creator in the days of thy youth, while the evil days come not, nor the years draw nigh, when thou shalt say, I have no pleasure in them;"

– Ecclesiastes 12:1

Chapter 4

Connectivity Driven Registration of the Surfaces of Subcortical Structures

In this Chapter we apply connectivity-driven registration to the surfaces of subcortical structures, in particular the human thalamus. So far, the connectivity of cortical surfaces has been used to drive registration. The basic premise was that the white matter structural connectivity reflects, to some extent, the cortical functional parcellation. The same can be said of subcortical structures, such as the thalamus. Connectivity patterns can define areas within subcortical structures, otherwise unidentifiable on T1-weighted images, that correspond to functionally distinct modules such as, for example, the thalamic nuclei. This property can be utilised to enhance the inter-subject registration of subcortical structures and motivation for this is given in section 4.1. Adaptations of the registration algorithms to the subcortical framework, in particular potential changes to the formation of the SCFVs, are described in section 4.2. Here we also examine in greater detail the artificial examples and the identification of model parameters (section 4.2.2). In addition, a registration paradigm different from FACS is introduced and compared to the already described methods (4.2.3).

The connectivity-driven registration of subcortical structures is applied to the human thalamus and the effects of co-registration of Alzheimer's disease and control

groups are explored in sections 4.2.3.2 and 4.3 and discussed at the end of the Chapter.

4.1 FIRST and the Problem of Vertex-to-vertex Correspondence

Deep grey matter brain structures such as the basal ganglia can be identified in T1-weighted MR images and automatically segmented out for further processing (morphometric measurements, etc). Some segmentation algorithms, e.g. the one employed by FreeSurfer (Fischl et al., 2002), assign binary labels to voxels of a volumetric MR image (e.g. 1 - belonging or 0 - not belonging to the structure of interest). They may use voxel neighbourhood information and spatial priors on the structure's shape, but they do not explicitly model the shape of structures themselves. Essentially, the classification of voxels as belonging (or not) to a structure of interest is done mainly using the image intensity information (Figure 4.1 (left)). This type of segmentation is typically employed when analysing the volumes of structures across subjects. On the other hand, other segmentation algorithms, such as FSL/FIRST (FMRIB's Integrated Registration and Segmentation Tool), explicitly model the shape of structures allowing for systematic inter-subject shape comparison (Patenaude, 2007) (Figure 4.1 (middle and right)). However, from the registration point of view, it is often important to ensure that functionally corresponding parts of subcortical structures are compared across subjects. For example, in an fMRI experiment resulting in thalamic activations, we would expect that when comparing activations across subjects, the activations belonging to the limbic nuclei of one subject are compared to the activations within the limbic nuclei of the others. This is assured through successful registration, either volumetric- or surface-based (depending on which segmentation algorithm is used). In this Chapter we particularly investigate how to improve the correspondence between the surfaces of subcortical structures extracted by FIRST.

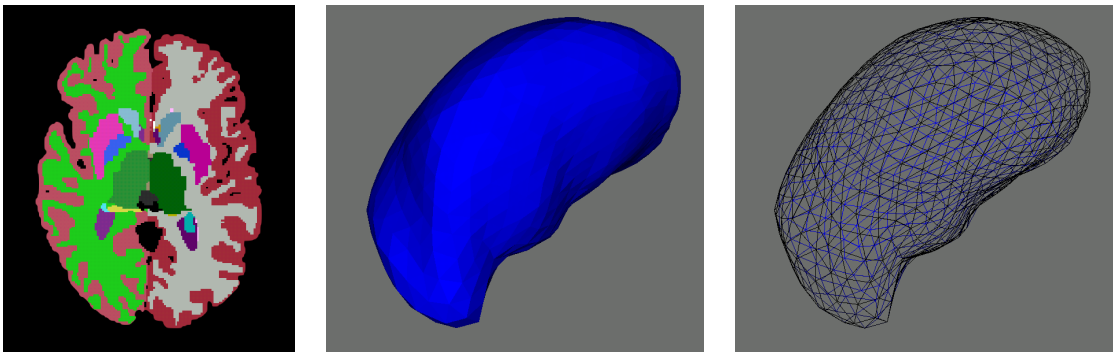

Figure 4.1: **(left)** FreeSurfer segmentation of the brain (axial slice). Different colours represent different brain structures - e.g. dark green shows the left thalamus. The shape of structures is not explicitly modelled. **(middle)** FIRST segmentation of the left thalamus shown as a solid surface. FIRST models the shape of 17 different deep grey matter structures and uses a fully probabilistic framework to capture the relation between the shape and image intensities. **(right)** The surface from (middle) shown as a mesh.

4.1.1 Segmenting Subcortical Structures Using FIRST

FIRST is a software package which segments 15 (subcortical) structures[1] from T1-weighted MR images. Moreover, FIRST explicitly models the shape of these structures using a surface-based point distribution model. Figure 4.1 (middle) shows an extracted thalamic surface whereas Figure 4.1 (right) shows the equivalent mesh representation of that surface. An important feature of FIRST is that it uses a training data set of manually labelled subcortical structures to learn a statistical relationship between the shape and appearance of structures. A Bayesian model relates the shape and the image intensities. The shape space is constrained to the linear combination of the mean shape and the eigenvectors of vertex coordinates (their variation across the training set). The conditional distribution of a shape given the image intensities (structure appearance) is used to infer on the shape of the structure. In that way, the contributions of shape constraints and the image intensities are automatically determined.

[1]These include the left and right thalamus, putamen, caudate, pallidum, hippocampus, amygdala, accumbens area as well as the brainstem.

FIRST-segmented subcortical structures (surfaces thereof) by construction have one-to-one vertex correspondence across subjects. This correspondence is primarily imposed by the way subcortical surfaces are parametrized using the point distribution model. In other words, the vertex labelled i in one subcortical mesh should, by construction, be in the same anatomical location as the vertex labelled i in the other corresponding subcortical mesh (i.e. thalamic surfaces). This correspondence is essential to successful segmentation since the image intensities are sampled at each vertex location and then compared to the training data to determine the final shape. Furthermore, when the segmentation is completed and shape representations formed, they can be jointly analysed assuming inter-subject vertex-to-vertex correspondence. In other words, the shape variation across a population is inferred for each vertex assuming that that vertex always corresponds to the same anatomical location in all subjects.

Challenges Although every precaution has been taken to ensure vertex-to-vertex correspondence throughout the FIRST pipeline (Patenaude, 2007), T1-weighted images provide limited information about functionally distinct areas within subcortical structures such as the thalamus (Behrens et al., 2003c). In the example of the thalamus, thalamic nuclei e.g. the anterior nuclear group, the dorsomedial nuclear group, etc. cannot be easily discerned in the T1-weighted images. Moreover, the image intensities within the thalamus (or the caudate or the globus pallidus) appear similar to each other offering little information about the anatomical/functional boundaries within structures. Therefore, although the whole of the structures can be successfully segmented, the inter-subject vertex-to-vertex correspondence is primarily driven by the T1-weighted image intensities, which does not guarantee correspondence in function/functional areas.

4.1.2 Motivating Connectivity-driven Surface Registration

The question arises whether the inter-subject vertex-to-vertex correspondence can be improved[2] using another imaging modality once the segmentations have been performed. This would help any subsequent statistical analysis involving the comparison of shape (or other image features) across subjects.

Chapters 2 and 3 have shown how DW images can be exploited to this effect in the cortical framework. White matter connectivity information can shed some light on the functional segregation of the brain not otherwise discernible in the T1-weighted images. Section 2.4.1.3 of Chapter 2 explains how structural connectivity measured using probabilistic tractography can help delineate functionally different cortical areas. In a similar fashion, the connectivity of the subcortical structures, in particularly the thalamus, can be used to find out more about fine-grained thalamic anatomy such as the thalamic nuclei. Behrens et al. (2003c) use parcellation methods similar to those explained in section 2.4.1.3 of Chapter 2 to cluster human thalami into anatomically plausible clusters (see Figure 2.3). Practically, each thalamic voxel is assigned a 'connectivity fingerprint' quantifying its structural connectivity to a set of targets (e.g. the motor and pre-motor cortices, etc). These connectivity fingerprints can be clustered just as in the case of the MFC or the SMA/pre-SMA areas (Figure 4.2).

Therefore, an analogous reasoning as in the cortical framework can be made here: connectivity information complements the T1-weighted images to provide additional insights into functional parcellations of subcortical structures. Therefore, just as we used SCFVs to drive registration of cortical grey/white matter surfaces, we might be able to use similar connectivity descriptors for the subcortical meshes provided by FIRST, and carry out connectivity-driven surface registration as before. The primary aim would be to ensure that the inter-subject comparison of subcortical

[2]This is equivalent to enhancing the inter-subject surface registration.

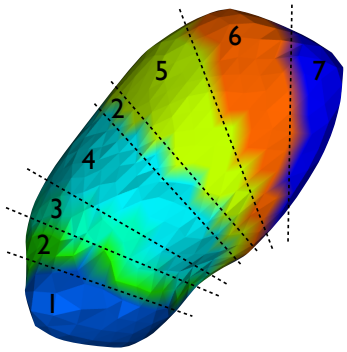

Figure 4.2: Connectivity-based thalamic parcellation was performed as described in section 2.4.1.3 of Chapter 2. SCFVs originating at each vertex of the thalamic mesh were clustered into 7 clusters (depicted in different colours and assigned numerals). Similar parcellations are reported by Behrens et al. (2003c) and found to have direct anatomic interpretation. The medioventral portion of the thalamus is shown. Top of the image corresponds to the anterior of the thalamus.

surfaces/meshes results in comparing features of *functionally homologous subcortical areas*. This is important as many fMRI experiments result in activations in subcortical structures and the group statistical analysis relies on the existence of such correspondence. Moreover, the between-group differences in vertex positions after the connectivity-driven registration to the common template could tell us about the areas of differing connectivity in the two (or more) groups.

The next section explores how connectivity-driven registration algorithms developed so far can be modified to suit the registration of subcortical surfaces. First, the formation of SCFVs is revisited and new, volumetric descriptors are introduced. Then the registration method is tested and refined on a set of artificial examples, most of which are relevant to the cortical framework as well. Finally, we introduce another alignment method, very different to the one so far used, and compare it to the results obtained earlier.

4.2 Registration Algorithms

The subcortical framework differs from the cortical primarily in the dimensionality of the registration problem. The cortical surfaces (e.g. the grey/white matter

surface) were represented as meshes with more than 10^5 vertices. The registration problem, which relates to the number of vertices undergoing free deformations, was reduced by an order of magnitude by scaling to the $ic5$ mesh. The reduction of the mesh size as well as the utilisation of low-dimensional SCFVs helped us achieve efficient registration. In the case of subcortical structures, the surfaces extracted by FIRST have, as a rule, 642 vertices (except the hippocampus) because they all come from a deformed $ic3$ (see a section on notation 2.4.2 in Chapter 2). This opens up new possibilities for the construction of Structural Connectivity Descriptors. Due to the (relatively) small number of vertices, these descriptors can be higher-dimensional, e.g. volumetric (3D). As discussed throughout Chapter 2 (especially section 2.5.1.1), the descriptive/discriminative power of connectivity descriptors is important for successful connectivity-driven registration. In section 4.2.1 we introduce Volumetric Structural Connectivity Descriptors which can be used in the subcortical framework instead of SCFVs. Later on, the connectivity-driven registration methods are described in relation to the subcortical framework (sections 4.2.2 and 4.2.3).

4.2.1 Volumetric Structural Connectivity Descriptors

To distinguish them from the SCFVs used before, we shall denote a volumetric equivalent of SCFVs, the Volumetric Structural Connectivity Descriptors, as VSCDs.

The first step in the processing pipeline of the FMRIB's Alignment of Subcortical Surfaces (FASS) is to segment matching structures from T1-weighted images of both (or more, if group registration is performed) subjects. This is done using the FIRST. Meshes delineating segmented structures have a fixed number of vertices (e.g., the thalamus has 642 vertices with mean inter-vertex separation of 2 mm) which, in our approach, become seeding points for whole-brain probabilistic tractography (Behrens et al., 2003b).

In the second step, DW-images are used to obtain estimates of tracts originating

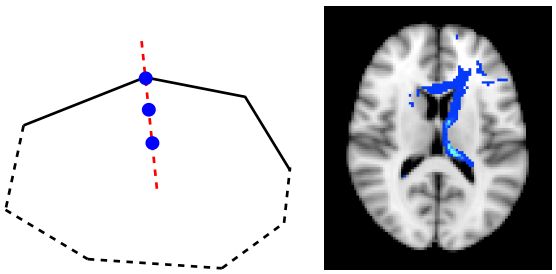

Figure 4.3: **(left)** The VSCD construction scheme. Black (solid and dashed) lines delineate the subcortical structure. The dashed red line represents the local normal and the blue dots are the 'seed' points from which VSCDs are computed. These points lie on the local normal of the corresponding mesh vertex and the neighbouring ones are 1.5 mm apart. Therefore, VSCDs are calculated for points on the normal *within* the structure and averaged to reduce sensitivity to misplaced structure boundaries. **(right)** An example of the VSCD for one vertex on the thalamic surface. In blue are shown thresholded tractography outputs overlaid on the MNI template brain (an axial slice is shown).

from each vertex. We label probabilistic tractography output from vertex i as the 'Volumetric Structural Connectivity Descriptor' of i: VSCD_i. The result of probabilistic tractography is an image where every voxel $V(x, y, z)$ has an integer value describing the number of tractography samples starting at the seeding point (in our case mesh vertex i) and passing through voxel $V(x, y, z)$. The maximum value is the number of random tractography samples at the seed, which is 5,000. Therefore, VSCD_i is an image where every voxel encodes the number of white-matter tracts originating at the coordinates of vertex i and passing through that voxel (Figure 4.3 - right).

In addition to forming VSCDs for vertices of the extracted meshes, we also construct VSCDs for points sampled along the local normals of the subcortical surface (Figure 4.3 - left). VSCDs are then combined to produce more robust estimates of structural connectivity. FIRST extraction can result in misplaced borders and if we were to sample exclusively on the extracted surface, it may happen that the VSCD is calculated for a point that lies e.g. within the ventricle (and would result in a close-to-zero VSCD). Therefore, the introduction of additional VSCDs coming from the

estimated *interior* of the subcortical structure should help in such cases by increasing the chances of sampling the structure of interest. The construction of additional VSCDs is performed in the following manner. For every vertex i of the subcortical mesh a local normal was calculated as given in Smith (2002). The VSCDs were then constructed for two more points along these normals, towards the structure's interior (labelled as VSCD_{i1} and VSCD_{i2}). The distance between neighbouring points was chosen to be 1.5 mm, as we found that such a sampling scheme almost always results in sampling within the structure's interior (tested on the thalami of 20 healthy control subjects)[3]. Finally, VSCD_i, VSCD_{i1} and VSCD_{i2} are combined to produce the mean VSCD (mVSCD - Figure 4.3):

$$\mathrm{mVSCD}_i = \frac{1}{3}\left(\mathrm{VSCD}_i + \mathrm{VSCD}_{i1} + \mathrm{VSCD}_{i2}\right). \tag{4.1}$$

For simplicity, through the rest of the text we shall call mVSCD simply VSCD. The sampling-averaging procedure introduced here is also used with SCFVs (instead of VSCDs) in the subcortical framework - section 4.3.

Under the assumption that similar white-matter tracts come from anatomically similar locations, comparing VSCDs should enable us to determine this anatomical correspondence. Therefore, in the next step, VSCDs of all vertices for both subjects' structures (input and reference) are transformed into standard space where their comparison can take place. First, registration between FA (fractional anisotropy, obtained using *dtifit* from FSL/FDT) and T1-weighted images is carried out by applying the affine registration tool (FSL/FLIRT). Subsequently, T1-weighted structural images of the input and reference subjects are nonlinearly registered to the MNI space using FSL/FNIRT. By combining these two registration steps, all VSCDs are registered to MNI space where they can be directly compared.

[3]Changes to this scheme are possibly necessary depending on the structure. E.g. the hippocampus has 'thin' regions that sometimes can be thinner than 3 mm.

In the next step, the similarity between VSCDs for the input and the reference structure is calculated using the 'correlation measure' - Equations 2.3 and 2.4. Just as in the cortical framework, the computation of the VSCD Similarity Matrix (SM) is performed 'off-line' prior to initialising the registration. Comparing VSCDs is more computationally demanding than manipulating SCFVs (around 3h on an Intel Core 2 Duo 2.4GHz processor for calculating the SM). In the current implementation of FASS, the calculation of the SM is parallelised (if VSCDs are used) to minimise the preprocessing time. The values of the SM become the inputs into the cost function calculation as described in Chapter 2, section 2.4.4.2.

4.2.2 Free Deformation-based Registration Method

The goal of FASS is similar to that of the cortical framework. Let every vertex of, for example a thalamic mesh, have a unique label (e.g. number from 1 to N, where N is the number of vertices in a mesh). Suppose then, that in one subject, a set of vertices $S_1 \subseteq \{1, \ldots, N\}$ corresponds to the surface of anterior dorsal thalamic nucleus. In another subject, the same nucleus is delineated by the set of vertices $S_2 \subseteq \{1, \ldots, N\}$. The main idea of the algorithm is to establish correspondence between S_1 and S_2 by moving (or sliding) vertices around the surface. In order to achieve this, we associated with every vertex a description of connectivity (labelled as a Connectivity Descriptor (CD) and can be either SCFV or a VSCD) of that point to the rest of the brain. The algorithm then tries to match these descriptors according to their similarity (similarity matrix is pre-calculated) and achieve anatomical correspondence.

Having in mind the aforementioned alignment goal, the FASS registration engine is the same as in FACS. The major preprocessing stage difference is in the formation of the Similarity Matrix. However, the cost function, the intensity and the regularisation forces as well as the multi-stage registration framework differ only in the number of vertices on the registration sphere. To initially estimate free parameters such as the

force scaling coefficients (Chapter 2, section 2.4.4.3) we devised an artificial data set further described in the following section. In particular, we optimise free parameters so that the registration consistency is maximised.

4.2.2.1 Artificial Examples and Parameter Identification

Artificial Data Set In order to test and tune the FACS/FASS algorithms, a set of artificial test data was devised. We would like to use the artificial data set to validate the effectiveness of the registration algorithms, estimate free parameters and evaluate the robustness of the alignment method to the input noise (existing in the SM). The artificial data set therefore aims to simulate real registration cases, such as the connectivity-driven alignment of the thalamus. To achieve this, we constructed several meshes with associated 'artificial connectivity data' and examined how our algorithms co-register them.

Without any loss of generality, we assume that all artificial meshes/data have the same shape, e.g. that of the $ic3$. Note that the $ic3$ is topologically equivalent to meshes encoding FIRST-extracted subcortical surfaces (therefore no upsampling of the registration results is needed). The main components of the artificial data set are the similarity matrices describing the 'connectivity' relations between vertices of different meshes. In other words, CDs have to be artificially devised.

Figure 4.4 introduces one of the artificial data sets used in this thesis [4]. Each mesh has six regions, each of which contains vertices with unique (compared to the other regions) simulated CDs. Colours represent areas with the same CDs and are therefore perfectly 'correlated' (see Equation 2.3) across artificial subjects. On the other hand, vertices with different CDs (i.e., different colours) are completely uncorrelated. In order to make examples more realistic, vertices near the boundaries between two regions with different CDs are allowed to be correlated with *both regions* to a certain

[4] A mesh and the associated artificial CDs are referred to as an 'artificial subject' throughout the text.

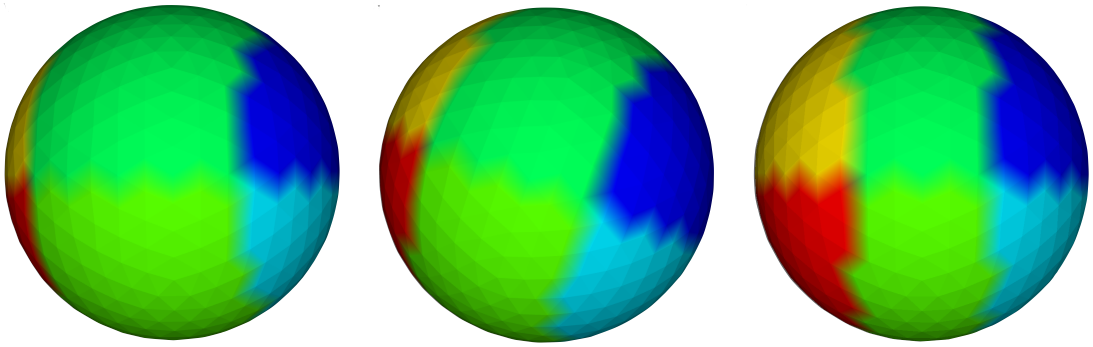

Figure 4.4: **(left)** A **(middle)** B **(right)** C. A, B and C represent three different artificial subjects. They all have 6 differently coloured regions that can be imagined as e.g. the different thalamic nuclei (each of them having distinct connectivity profiles). Areas of the same colour have artificial CDs perfectly 'correlated' across subjects. Sizes of coloured regions vary between subjects as well as the overall rotation of the spheres. The diameter of every sphere is 10 cm and the lengths of mesh edges are approximately 7 mm each. Every mesh has 642 vertices ($ic3$).

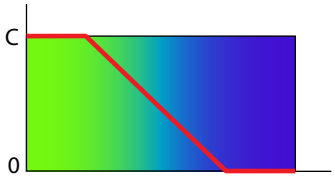

Figure 4.5: Representation of the connectivity patterns at boundaries of the regions in an artificial example (i.e. artificial CDs). The red line represents the 'number of connections' (for vertices along the horizontal axis) to the 'green region' (see Figure 4.4 and text). C is an arbitrary constant. The width of the shaded region is typically 10 mm. The diameter of every sphere is 10 cm and the lengths of mesh edges are approximately 7 mm each.

extent. This stops the boundaries being very sharp and instead represents a more realistic case of partial-volumed boundary, where connectivity changes in a piece-wise linear fashion.

This is achieved by describing every artificial CD with a six-dimensional vector $[p_1\ p_2 \dots p_6]^T$ where $p_i,\ i \in [1,6]$ represents 'the number of connections' from that vertex to region i, and this is set to be a linear function of the distance of that vertex to the region boundary (for illustration see Figure 4.5). These vectors replace SCFVs or VSCDs from real data. As in the real cases, vectors with similar entries at corresponding rows have high similarity and vice versa. If we look at p_i as the

number of tracts to region i, high similarity between vectors indicates that those vertices (seeding points) connect to specific regions in a highly similar way.

The ABCA Test We decided to tune free parameters by maximising *registration consistency* evaluated on an artificial data set presented here. For this purpose we used the A, B and C artificial subjects depicted in Figure 4.4. Artificial data were used to run pairwise A to B, B to C and C to A connectivity-driven registrations. If we denote the corresponding mappings by $\mathcal{W}_A$, $\mathcal{W}_B$ and $\mathcal{W}_C$: $A \xrightarrow{\mathcal{W}_A} B \xrightarrow{\mathcal{W}_B} C \xrightarrow{\mathcal{W}_C} A$, then in an ideal case, $\mathcal{W}_C \circ \mathcal{W}_B \circ \mathcal{W}_A = I$ where I is an identity mapping (Christensen and Johnson, 2003) (we shall call this an *ABCA test*). Therefore, the residual deformation field of the ABCA test is a good measure of the registration consistency although it cannot be used to assess absolute accuracy. The performance of the ABCA test was assessed by examining the histogram of residual displacements after combining three deformation fields $\mathcal{W}_C \circ \mathcal{W}_B \circ \mathcal{W}_A$ (Petrovic et al., 2009). Apart from the ABCA test we also looked at how well each registration step (A to B, B to C and C to A) was carried out by comparing the vertex-to-vertex correlations (Equation 2.3, the input and the reference meshes are the overlapping spheres) between the input and the reference subject before and after registration. In addition to that, the visual inspection of registration results was also performed.

The model parameters (Chapter 2, Table 2.2) were empirically estimated by minimising the ABCA test residual displacements while maximising the A to B, B to C and C to A individual registration performances (Figures 4.6 and 4.7). We first fixed the number of multi-scale steps to 5 (e.g. as in the state-of-the-art Spherical Demons (Yeo et al., 2010a)) and estimated λ and μ, scalings for the intensity and the regularisation forces respectively, for each step using a brute-force search method. Parameter λ was varied in increments of 2.0, whereas μ in increments of 0.1. Smaller increments were tested for both parameters but did not result in resolvable differences

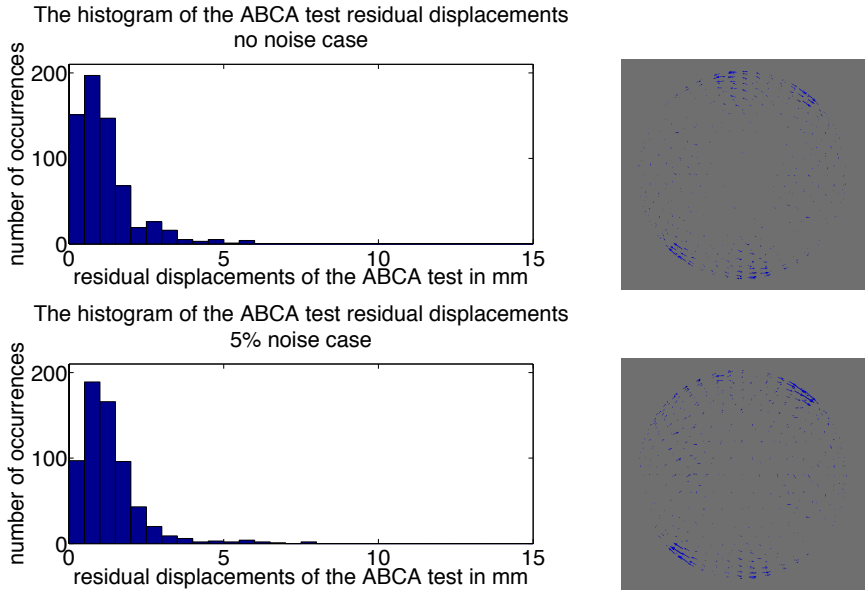

Figure 4.6: The analysis of residual displacements after performing the ABCA test. The displacements are calculated as the Euclidean distance between the corresponding vertices of the input mesh (A) and the input mesh after application of the $\mathcal{W}_C \circ \mathcal{W}_B \circ \mathcal{W}_A$ composition of the deformation fields. The analysis was performed for different levels of uniform noise added to the SMs. **(top left)** The histogram of residual displacements after the ABCA test for the no noise case. The mean residual displacement is ~1 mm and maximal recorded is ~6 mm, which is less than the average mesh edge length (~7 mm). **(top right)** The displacement field (residual displacement vectors shown in blue) corresponding to the histogram on the left. **(bottom left)** The histogram of residual displacements after the ABCA test for the 5% noise case. The mean residual displacement is ~1.5 mm and maximal recorded is ~6 mm. **(bottom right)** The displacement field corresponding to the histogram on the left.

in the final ABCA test results. Other optimisation schemes were also tried, such as optimising (using brute-force search) each multi-scale step separately: λ and μ were first estimated for step 1, then fixed to their estimated values for step 1 and estimated for step 2, and so on. However, such an optimisation scheme yielded similar results in terms of ABCA test performance but resulted in poorer individual A to B, B to C and C to A registrations.

To examine the sensitivity of the registration algorithm to the input noise we added noise to the Similarity Matrices encoding similarities in artificial CDs of the A to B, B to C and C to A mappings (for each SM separately). In the real

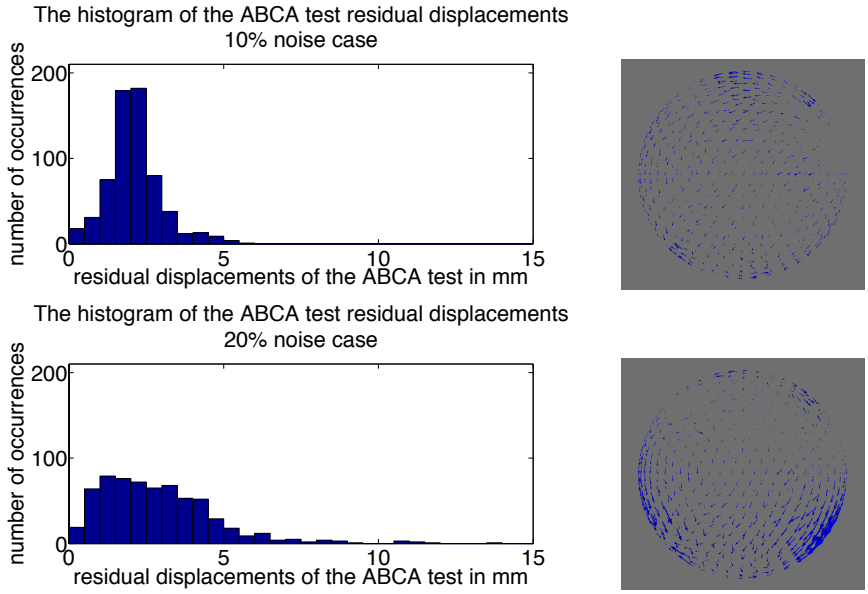

Figure 4.7: Continuation of Figure 4.6. **(top left)** The histogram of residual displacements after the ABCA test for the 10% noise case. The mean residual displacement is ~2 mm and maximal recorded is ~6 mm, which is less than the average mesh edge length (~7 mm). **(top right)** The displacement field (residual displacement vectors shown in blue) corresponding to the histogram on the left. **(bottom left)** The histogram of residual displacements after the ABCA test for the 20% noise case. The mean residual displacement is ~3 mm and maximal recorded is ~14 mm. **(bottom right)** The displacement field corresponding to the histogram on the left. Large residual displacement vectors can be observed.

data, the noise comes from several sources of which the most prominent are the noise associated with probabilistic tracking/diffusion data and the noise related to inaccurate segmentations. In the artificial subjects, the noise model has a uniform distribution $X \sim U(0,t)$ where t specifies the noise level, e.g. $t = 0.2$ corresponds to the 20% noise level (Figures 4.6 and 4.7). The SM entry with added noise of level t (SMn^t) was calculated as:

$$\mathrm{SMn}^t(i,j) = \frac{\mathrm{SM}(i,j) + X}{1+t}. \tag{4.2}$$

The SMn is therefore scaled to the $[0,1]$ interval as for the original SM. In addition to the no noise case, the noise levels of 5%, 10% and 20% were examined.

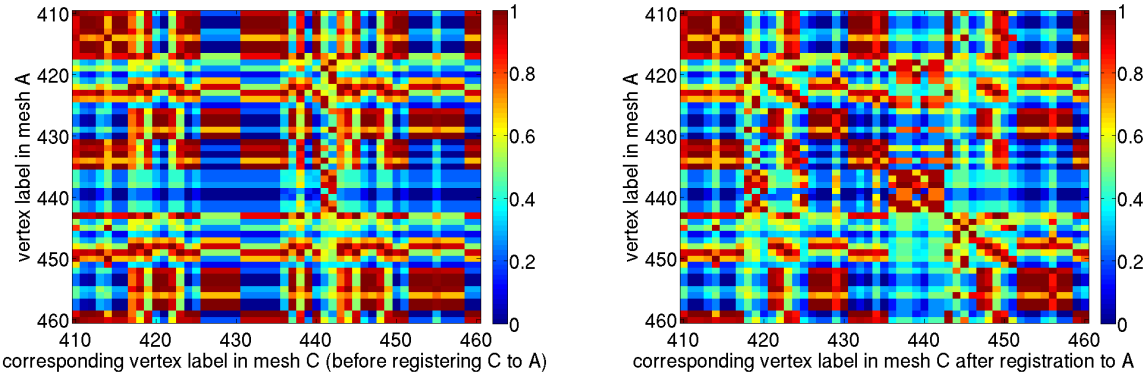

Figure 4.8: Demonstration of how similarities between artificial CDs of neighbouring vertices change after connectivity-driven registration. **(left)** A part of the similarity matrix between meshes A and C (artificial data from Figure 4.4). Values on the main diagonal show how correlated CDs of neighbouring vertices from both meshes are. The colourbars represent similarity coefficient values. **(right)** After registration of C to A, positions of vertices in C changed with respect to those in A. Similarities in CDs between neighbouring vertices from A and C should have increased due to minimising the cost function (Equation 2.5). This is demonstrated by having high similarity values on the main diagonal (similarities between vertices at 'anatomically' equivalent locations - vertex-to-vertex correspondence) compared to those in the image on the left.

No Noise Case For every registration step of the ABCA test, the final mean vertex-to-vertex correlations between the input and the reference meshes were approx. 0.98 (5% to 10% increase from the initial mean vertex-to-vertex correlation - as an example see Figure 4.8) indicating successful alignment/convergence. The mean residual displacement after performing the ABCA test with no noise added to the SMs is ~1 mm (see the histogram in Figure 4.6 - top) whereas the maximum residual displacement is ~6 mm. Therefore, all residual displacement values are smaller than the length of the mesh triangle side (~7 mm). As the ABCA test involves three registrations, the individual registration errors are likely to be less than the total ABCA residual error.

Mean (true) displacement values, by construction, for each of the $A \rightarrow B$, $B \rightarrow C$ and $C \rightarrow A$ registration steps are 5 mm, 6 mm and 10 mm respectively. For reference, the maximum registration displacement is recorded in the $C \rightarrow A$ registration and is

~22 mm (this is expected as the difference in sizes of coloured areas between artificial subjects A and C is the greatest - see Figure 4.4).

SM with Added Noise Figures 4.6 - bottom and 4.7 show results of the ABCA test for the increasing amount of uniform noise added to the SMs. As expected, the residual displacements increase with the amount of noise. For the 10% noise case, the mean residual displacement increases to ~2 mm and to ~3 mm for the 20% noise case. The maximum residual displacement for the 20% noise case is ~14 mm (which is almost twice the mean mesh edge length) compared to ~6 mm for no noise. Still, for 5% and 10% noise levels, the maximal residual displacement does not exceed mean mesh edge length. As in the no noise case, this indicates that the individual registration errors are likely to be less than the total ABCA residual error.

ABCA Test on Real Data from Healthy Controls The ABCA test was also performed on thalami from three healthy control subjects. The average vertex displacement from the registrations (A to B, B to C and C to A) was ~1.5 mm (5 mm max). The average residual displacements (A to B to C to A) were ~1 mm (5 mm max) and the average mesh edge length ~2 mm. As the ABCA test involves three registrations, the individual registration errors (or at least the inconsistencies) for the real data are still likely to be less than the total ABCA residual error.

The ABCA test was repeated for several configurations of the A, B and C artificial subjects. We varied the size of the regions of distinct connectivity as well as the global rotations. All these tests yielded results similar to those presented here with the same conclusions.

4.2.2.2 Variations of the Proposed Registration Framework

Apart from the registration method introduced in Chapter 2 and further validated in this Chapter, we also examined several variations of the proposed method. In

particular, we formulated the registration task differently, in that the deformation forces were consecutively applied to groups (or 'patches') of vertices instead of single vertices. A group moved together as the regularisation and the intensity forces were propagated through the patch. The multi-scale solution to the registration problem was imposed by varying the size of the patch, i.e. at the coarse scale, the patch had a large number of vertices (matching of large, global features) and at the finest scale, the patch consisted of one vertex only. A large number of patches were identified on the mesh at each scale and a complex update scheme was devised to deform separate patches independently. The whole procedure was embedded within an iterative scheme. We hoped that such a formulation would speed up convergence. However, the complexity of the update scheme (keeping track of the neighbourhood information for a large number of vertices) considerably slowed down computation. Moreover, a large number (at least 10) of scale steps had to be implemented to achieve satisfactory results for the individual artificial tests, e.g. A to B mapping. This additionally increased the number of free parameters further complicating the parameter optimisation scheme. To achieve the ABCA test results (artificial data) comparable to those presented here, one had to symmetrise the cost function to help avoid suboptimal solutions. The combined cost function was calculated for the inverse mapping and integrated with that of the forward mapping to produce the symmetrised cost function (Christensen, 1999). The main caveat of this modification was that the computational cost was that a single registration for the 642-vertex mesh (e.g. a thalamic mesh) took 4 to 6 hours on an Intel Core 2 Duo 2.4GHz processor (compared to 10 min for the current registration framework). The registration of larger meshes, such as the $ic5$ required for cortical registration, would be computationally even more intense. We have also experimented with using radial basis functions for the interpolation of the SCFV similarity matrix (instead of the linear interpolation currently used). However, in all registration paradigms that

we examined, utilising linear interpolation resulted in smaller residual errors in the ABCA test. The research into different deformation strategies, interpolation functions and the multi-scale optimisation implementation provided us with key insights into the nature of the registration task at hand, which was crucial to the development of the registration framework presented in this thesis.

In addition to investigating other registration strategies, we also briefly examined the effects of the SCFV correlation function (Equation (2.3)) to the ABCA test results. For example, the correlation function consisting solely of the denominator of the expression in Equation (2.3) worsens the ABCA test results primarily for two reasons. First, due to the lack of explicit normalisation, and second, because of the insensitivity to matching SCFVs with low-valued entries. These differences in performance prompted us to carefully motivate the currently used measure (Chapter 2, section 2.4.4.2).

4.2.3 Registration Based on Spherical Wavelets

Although the performance of our connectivity-driven registration on the ABCA test was satisfactory (in terms of the magnitudes of residual displacements), we wanted to test how the method compares to a very different class of registration paradigms - those modelling the deformation field using a set of basis functions. In volumetric image registration such a representation of a deformation field is common (Andersson et al., 2007b; Ashburner and Friston, 1997), but is seldom found in surface-based registration. The main idea of using basis fields to represent non-linear deformations is to represent the deformation field (displacements) as a linear combination of a set of spatial basis functions. The registration in such case reduces to an optimisation problem trying to find the coefficients scaling individual basis functions. The multi-resolution solution naturally arises from the possibility to represent the deformation field using basis functions of different 'coarseness'. Consequently, such a formulation

of the registration task is very different to what we have been employing so far. Therefore, it would be of interest to explore how a different field representation and optimisation scheme can cope with the connectivity-based registration problem.

To investigate this we devised a novel surface-based registration method relying on the representation of the deformation field using spherical wavelets (Schröder and Sweldens, 1995b,a). They are particularly suited for representing functions (e.g. deformations) on spherical meshes that come from the icosahedral subdivisions.

In the proposed registration method we represent the deformation field using a set of biorthogonal wavelet basis functions defined on a sphere (Schröder and Sweldens, 1995b,a). The basis set is constructed of scaling functions at the coarsest scale and wavelet functions at finer scales. Both scaling (φ) and wavelet functions (ψ) are defined as $\varphi_{j,k}, \psi_{j,k} : \mathbb{S} \rightarrow \mathbb{R}$, where $\mathbb{S}$ is a unit sphere, j defines the scale of the function and k refers to the spatial index which describes where on the surface the function is centered. At a particular scale j, wavelet functions are combinations of scaling functions at scales j and $j+1$. A given function $f : \mathbb{S} \rightarrow \mathbb{R}$ can be expressed as a linear combination of the basis functions

$$f(x) = \sum_k \lambda_{0,k} \varphi_{0,k}(x) + \sum_{j \geq 0} \sum_m \gamma_{j,m} \psi_{j,m}(x). \tag{4.3}$$

Scaling coefficients $\lambda_{0,k}$ represent the low-pass content of the signal f whereas coefficients $\gamma_{j,m}$ represent the localised band-pass content of the signal. The possibility to represent the function on the sphere with basis functions (wavelets) of different coarseness/spatial extent can be used to achieve the multi-scale registration. In other words, the idea of multi-scale registration using spherical wavelets relies on approximating the function f at different scales - coarse (low) scale encodes large deformations while higher scales determine fine features of the deformation field (for applications in shape analysis see, e.g., Yu et al. (2007); Nain et al. (2007)). If a spherical mesh has N vertices, a total of N basis functions are created, composed of

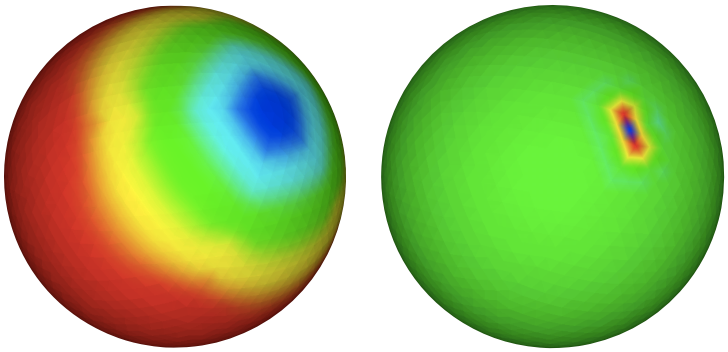

Figure 4.9: **(left)** An example of a scaling function associated with one of the mesh vertices. Red indicates low and blue high values. The scaling functions (note their large spatial extent) approximate the deformation field at coarse (low) scale encoding large scale deformations. The scaling and the wavelet functions form the basis field that models the displacements/deformations. The coefficients that multiply the basis functions to form the deformation field are found by optimising the CD similarity function SFun (Equation 4.4). **(right)** An example of a spherical wavelet basis function associated with one of the vertices. Wavelet basis functions model the deformation field at fine spatial scales.

N_0 scaling functions (where N_0 is the initial number of vertices of the base mesh - e.g. icosahedron) and N_r wavelet functions. If each of these basis functions is evaluated in each of the N vertices and these data stacked into a matrix $\Phi_{N\times N}$, every finite energy scalar function F evaluated at N vertices can be transformed into a vector of basis coefficients Γ using Forward Wavelet Transform $\Gamma = \Phi^{-1}F$ and recovered using *Inverse Wavelet Transform*: $F = \Phi\Gamma$.

If F encodes displacements in Euler angles for every vertex of the input mesh, coefficients in Γ can be numerically solved for by maximising the similarity function between CDs of the input (M_I) and reference (M_R) meshes. The similarity function (SFun) is calculated for all vertices of the reference mesh and can be written as:

$$\mathrm{SFun}_{IR} = \mu\frac{1}{N}\sum_{l\in Y}\mathrm{SM}^*(x_l, y_l) + \nu E_m(F), \tag{4.4}$$

where $i \in V_{M_I},\ j \in V_{M_R}$ and $\mathrm{SM}^*(x_l, y_l)$ is the interpolated value of the SM, where the position of x_l directly depends on the deformation field, F. The interpolation of SM values is achieved using linear interpolation as in Equation 2.5. E_m is the

membrane energy of the deformation field, taking on the role of the regulariser, while μ and ν are scaling parameters found empirically. In order to penalise local topology breaches, parameter ν is increased (by 10%) when the deformation field F results in a mesh-intersection.

We use variable metric non-linear optimisation to estimate coefficients λ and ψ (Γ) Press et al. (1995). This is performed in four steps (on four scales) with all coefficients set to 0 initially. We denote a set of coefficients at scale i by S_i, $i \in [0, 3]$ (where $S_{i-1} \subset S_i$). The first stage of optimisation finds the coefficients at the coarsest scale, S_0, followed by stages for increasingly higher scales: S_1, S_2 and S_3.

4.2.3.1 The ABCA Test

Optimisation of the free parameters, in particular μ and ν to minimise residual displacements in the ABCA test resulted in poorer performance compared to the other registration algorithm. In particular, Spherical Wavelets-based registration struggled to achieve significant improvements in the separate registration steps especially when large global rotations were required. Figure 4.10 gives a summary of the ABCA test results comparable to those shown in Figure 4.6 - top.

The average vertex displacements from the B to C and C to A registrations were similar to the average residual displacement, while the A to B registration (incorrectly) found no displacements. Therefore, in the A to B registration, no improvement in maximising the SFun (Equation 4.4) was observed. This is most likely due to convergence problems with the variable metric approach when optimising fine-scale wavelet coefficients, possibly due to the highly non-linear similarity function. The conjugate gradient optimisation method yielded similar results.

Summary In summary, both presented methods performed well on the artificial ABCA test suggesting high registration consistency and indicating that the chosen set of features is informative enough to drive the registration. However, registration

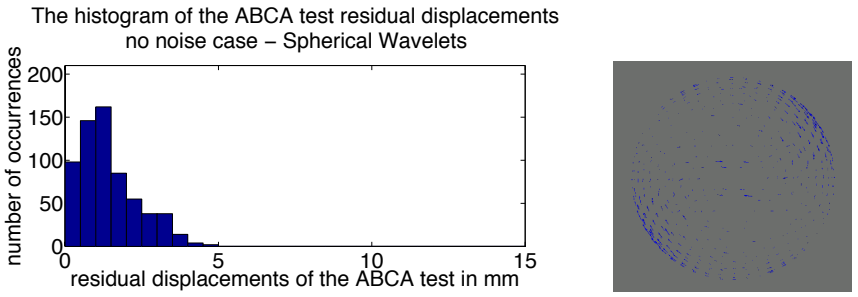

Figure 4.10: The analysis of residual displacements after performing the ABCA test for the registration based on Spherical Wavelets. The displacements are calculated as the Euclidean distance between the corresponding vertices of the input mesh (A) and the input mesh after application of the $\mathcal{W}_C \circ \mathcal{W}_B \circ \mathcal{W}_A$ composition of the deformation fields. **(left)** The histogram of residual displacements after the ABCA test. The mean residual displacement is ~1.5 mm and maximal recorded is ~5 mm. However, the A to B registration incorrectly found no displacements. **(right)** The displacement field (residual displacement vectors shown in blue) corresponding to the histogram on the left.

using spherical wavelets did not always achieve good convergence and optimisation of the associated cost function. We believe this to be due to the difficult non-linear optimisation step at the highest scale (aligning fine details) where the variable metric optimisation algorithm was often unable to properly converge. This highlighted the importance of adjusting the optimisation method, as well as the multi-scale approach, to the specific nature of this registration problem.

4.2.3.2 Comparison on an AD Data Set

After analysing artificial data, we compared two registration methods by co-registering real data (human thalami from the Alzheimer's disease and control groups). The aim of the test was to check whether with real data both methods produce similar deformation fields (despite the differences in the optimisation methods and the way the deformation fields are calculated).

Thus, both methods were tested on the thalami from a dataset comprising age- and gender-matched AD and control groups (12 + 12 subjects - details on data acquisition are given in Chapter 3, section 3.4.2.1). Volumetric Structural

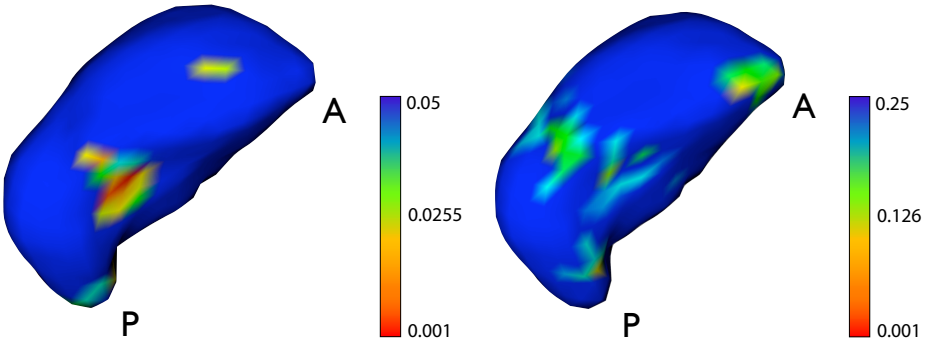

Figure 4.11: Exploratory analysis of residual within-surface displacements between control and AD groups for the left thalamus (medial side is shown). Colourbars represent uncorrected p-values. **(left)** Results for the free deformation-based registration (main FACS/FASS algorithm). **(right)** Results for registration using spherical wavelets. A - anterior, P - posterior.

Connectivity Descriptors were used (section 4.2.1). An AD subject (the most typical of all 24 subjects (Smith et al., 2006)) was chosen as template and all others were registered to this using proposed methods. The differences in the deformation fields between the AD and the control groups were examined using the multivariate Hotelling's T^2 test on vectors encoding vertex positions after co-registration. The F-statistic was formed for each vertex and the p-value on differences in mean positions of that vertex in the two groups was computed (see Chapter 3, section 3.4.2). For both methods the F-test showed group differences in residual within-surface displacements predominantly in medial dorsal parts of the left thalamus (Figure 4.11). Right thalamus showed negligible differences that considerably depended on the choice of the template subject. Spatial positions of changes found in the left thalamus did vary with the choice of template, but were always constrained to the medial dorsal part. Further, more detailed analysis of the applications of the connectivity-driven registration to the investigation of thalamic differences in AD is given in sections 4.2.3.2 and 4.3.

The AD and control data sets were further analysed using the free deformation method and VSCDs (section 4.2.1). To account for the dependency of the results on the template, FASS registration was performed for 8 different templates (control subjects) and the results were combined as explained in section 4.3.2. These results

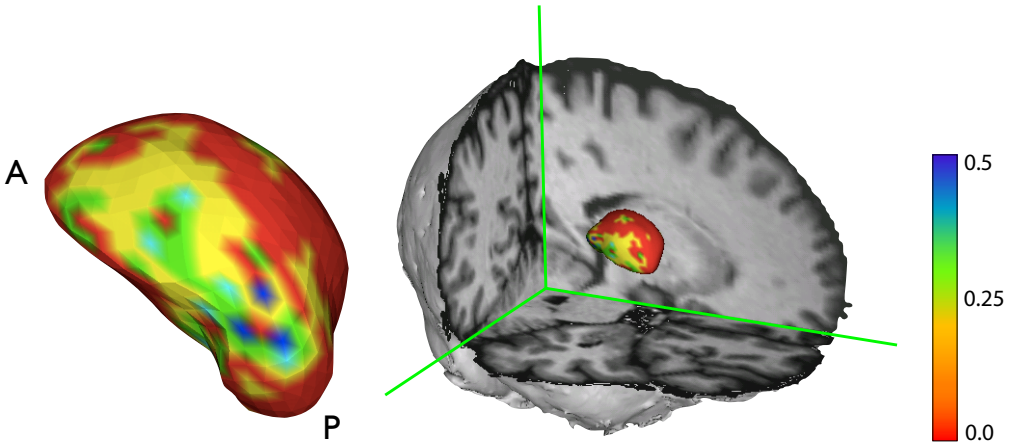

Figure 4.12: **(left)** Exploratory analysis of residual within-surface displacements between Alzheimers and control groups for the left thalamus. FASS results were combined for 8 different template (control) subjects. VSCDs were used to drive registration. Group differences are observed in the mediodorsal parts of the thalamus. Results are comparable to those obtained with registration using SCFVs - Figure 4.15. **(right)** Position of the thalamus within a template brain. Green axes are added for reference. A - anterior, P - posterior. Colourbar represents scaled mean thresholded ($p < 0.05$) maps across templates. More specifically, value 1 indicates that in all templates that vertex has an uncorrected p-value smaller than 0.05. Conversely, values 0 indicates that at a given vertex no templates have an uncorrected p-value smaller than 0.05. For a detailed description see section 4.3.2.

are presented in Figure 4.12. Highlighted areas are similar (but not identical) to those found in a single subject template (Figure 4.11 - left) case[5]. The analysis of the AD and control data sets was performed for exploratory purposes only and is complemented with the equivalent analysis with FASS driven by SCFVs (section 4.3.2).

Summary In conclusion, the method using spherical wavelets was less sensitive when analysing real data (Figure 4.11). Nonetheless, patterns of differences in thalamic connectivity between AD and healthy controls found by both methods are similar to each other and consistent with histological evidence (Xuereb et al., 1991; Braak and Braak, 1991). This suggests that both methods were able to recover similar deformation fields in support to the notion that the connectivity-

[5]Template dependency is discussed throughout the rest of the Chapter.

driven registration reaches consistent solutions irrespective of the actual optimisation method, regularisation, etc.

4.3 Subcortical Registration Using SCFVs From the Cortical Framework

The Structural Connectivity Feature Vectors used in the cortical framework to drive registration can also be straightforwardly utilised for the alignment of subcortical structures. Compared to VSCDs, SCFVs can be easily and efficiently manipulated (e.g. compared, clustered, etc) due to their low dimensionality. At the same time, as we have seen in the cortical framework, they carry enough information to drive registration.

In Chapter 2, section 2.4.3.1 SCFVs were introduced as connectivity descriptors associated with mesh vertices describing structural connectivity of a particular point to a set of predefined cortical and subcortical targets. The SCFV construction procedure in the subcortical setting differs from that in the cortical in the following ways. Firstly, in the subcortical framework, the seed points for tractography are defined by the positions of vertices of the FIRST-extracted (subcortical) meshes. Secondly, no smoothing was applied as given in Chapter 2, section 2.4.3.1, but a different sort of averaging, described in section 4.2.1, Equation 4.1 is used instead. Finally, should the subcortical structure, which is being registered, also be one of the tractography targets, it is excluded from the targets list. Therefore, in the case of the thalamus, the SCFVs would have 88 entries.

The benefits of using SCFVs over VSCDs primarily lie in the improved initial co-registration of the tractography targets across participating subjects. Construction of SCFVs relies on good initial alignment of cortical surfaces, which is achieved through FreeSurfer cortical geometry-driven registration. In the case of VSCDs, the correspondence in VSCDs across subjects is established by volumetric registration

to the standard space, e.g. MNI template. Arguably, the latter is prone to greater misalignment errors (Fischl et al., 1999c). It is for these reasons that we devote more attention to the analysis of SCFV-driven registration. On the other hand, having low-dimensional feature vectors can potentially compromise their discriminatory power (section 4.4). However, as we have previously discussed (Chapter 2, section 2.5.1.1), a comprehensive analysis of the advantages - in terms of registration - of one type of connectivity descriptors over the other, remains the topic of future research.

4.3.1 SCFV Subcortical Templates

The procedure for the SCFV template construction in the cortical setting was described in Chapter 3, section 3.4.1. An iterative procedure was used to derive unbiased 'prototype' SCFVs from a set of healthy control subjects. Such a template can be formed for the subcortical structures segmented by FIRST. The SCFV template can be used in registrations (as the reference) as well as in examining the connectivity-based parcellations. In the rest of this section we present the SCFV template for several subcortical structures, in particular the thalamus, the putamen, the pallidum and the amygdala. We show how these templates can be clustered and used to perform connectivity-driven co-registration of the healthy controls and the Alzheimer's disease patient groups.

4.3.1.1 SCFV Template Clustering

To investigate the structure of the Template Connectivity Matrices[6] we performed clustering of self-correlated template CMs. The presence of anatomically/functionally meaningful clusters would indicate that the SCFVs carry important information on the substructures and nuclei that are not necessarily identifiable in structural, e.g. T1-weighted, images. As a consequence, such information could be of value in the connectivity-driven registration. An additional goal of clustering template CMs is to

[6]For formation details see Chapter 3, section 3.4.1.

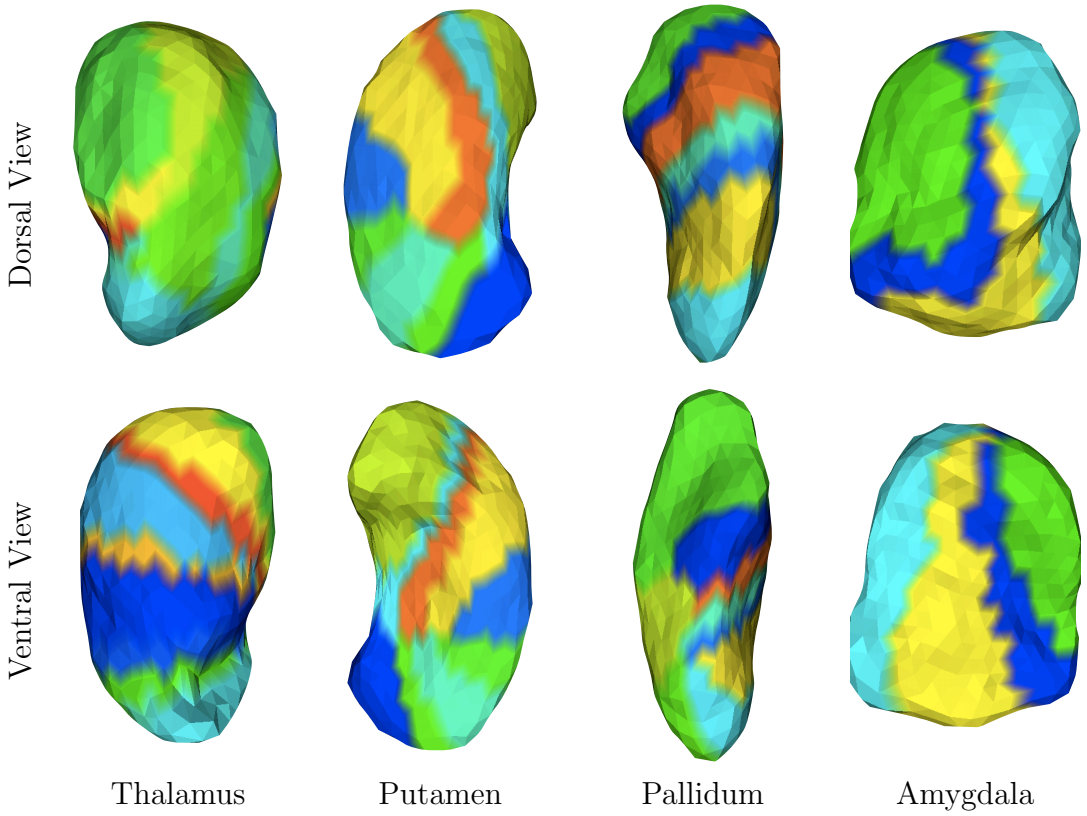

Figure 4.13: Clustering of the template CMs for four subcortical structures. The thalamus was parcellated into 12, the amygdala into 4 and the pallidum and the putamen into 9 clusters (left structures shown). Different clusters are shown in different colours. For all structures, the top corresponds to the anterior of the brain. The upper row shows the dorsal and the lower row the ventral views for each structure. Most clusters are spatially highly consistent suggesting good SCFVs' discriminative power (in the subcortical framework) and successful template formation. Comparable clustering results were observed in structures in the right cortical hemisphere as well.

explore how (spatially) well-defined these clusters are. The presence of consistent (e.g. across different data sets) and coherent clusters also suggests that the connectivity-driven alignment was able to align corresponding parcels across subjects and result in a 'sharper' template (compared to e.g. averaging subject-specific CMs). This is further quantified in section 4.3.1.2.

A data set consisting of 12 healthy control subjects (Chapter 3, section 3.4.2.1) was used to construct template CMs of four subcortical structures - Figure 4.13. Obtained CMs were self-correlated and clustered using the K-means algorithm. The thalamus

was parcellated into 12 clusters as suggested by Wiegell et al. (2003) who also used DW data for identifying thalamic nuclei[7]. However, we should acknowledge that the number of identifiable nuclear groups is subject-specific (Jbabdi et al., 2009b). Still, the aim of this demonstration is not to introduce a method for thalamic parcellation, but to show the existence of the anatomically plausible information within our CMs. The amygdala was parcellated into 4 regions, putatively the lateral, basal + accessory basal, medial and central. The same number of clusters in a similar clustering paradigm was used by Saygin et al. (2010). The putamen and the pallidum were clustered into 9 clusters each. All parcellations, and especially the thalamus and the amygdala, resemble known anatomic substructures (South, 2010; Morel et al., 1997). Observed parcellations of striatum closely resemble areas of differing connectivity observed in histologic analysis (Haber, 2003). Furthermore, although the K-means clustering algorithm did not use any spatial information, all parcels appear highly spatially consistent, sometimes occupying just a one-vertex-thin line (e.g. in the thalamus).

4.3.1.2 Registration Assessment Using Thalamic Clusters

To get a better insight into the outcome of subcortical registration we investigated how the thalamic clusters align before and after FASS. As we have seen in section 4.3.1.1, the thalamus can be clustered into a number of clusters, e.g. 12 or 7 (depending on the method, application, etc) using structural connectivity information. This parcellation resembles the division of the thalamus into thalamic nuclei providing an insight into finer anatomical subdivisions not visible in T1-weighted images. The connectivity-driven registration should improve the co-registration of such parcels with specific anatomical/functional characteristics.

[7]Other authors use a parcellation into 7 clusters such as (Ziyan and Westin, 2008) corresponding to the anterior, medial, midline, intralaminar, centre median, lateral and ventral nuclear groups. However, for the demonstration of the clustering capabilities, e.g. spatial consistency of clusters, we prefer to use a higher number of clusters.

To test this premise we clustered thalami of a group of 24 subjects (Chapter 3, section 3.2.1) as described in section 4.3.1.1. Parcellation was performed for 7 and 12 clusters to examine the robustness of the results to the number of clusters. We chose the thalamus to carry out testing as that is one of the subcortical structures with well examined connectivity-defined clusters (Behrens et al., 2003c). All the thalami have also been registered using FASS to the template CM derived from the same data set. In that way, we could measure the inter-subject cluster correspondence before and after connectivity-driven registration. The cluster correspondence across subjects was imposed in the same way as described in Chapter 3, section 3.2.4 and appendix D.

We should note that the clustering and the connectivity-driven registration are performed using the same connectivity information. Therefore, the assessment of the registration results using clustering is not an entirely independent measure (as would, for example, be using an fMRI data set). Nonetheless, the registration is driven by SCFVs, whereas the defining feature of the parcellations are the borders between clusters. These borders/edges are not explicitly represented in the CM (i.e. the registration cost function) and so the assessment of registration using cluster overlaps is not equivalent to examining the registration cost function before and after the alignment.

Figure 4.14 shows the pairwise Jaccard coefficient for the inter-subject overlap of corresponding clusters before and after FASS. The higher the value of the Jaccard coefficient, the better inter-subject alignment of corresponding clusters. The analysis was repeated for 100 independent initialisations of the K-means clustering algorithm ('clustering iterations'). The results suggest that FASS improves the registration of corresponding thalamic clusters, i.e. putative thalamic nuclei. Therefore, one could speculate that if the thalamic clusters (identified by analysing connectivity patterns) represent functionally (or cytoarchitectonically) distinct areas, then FASS enhances the alignment of functionally analogous areas as well. However, to further

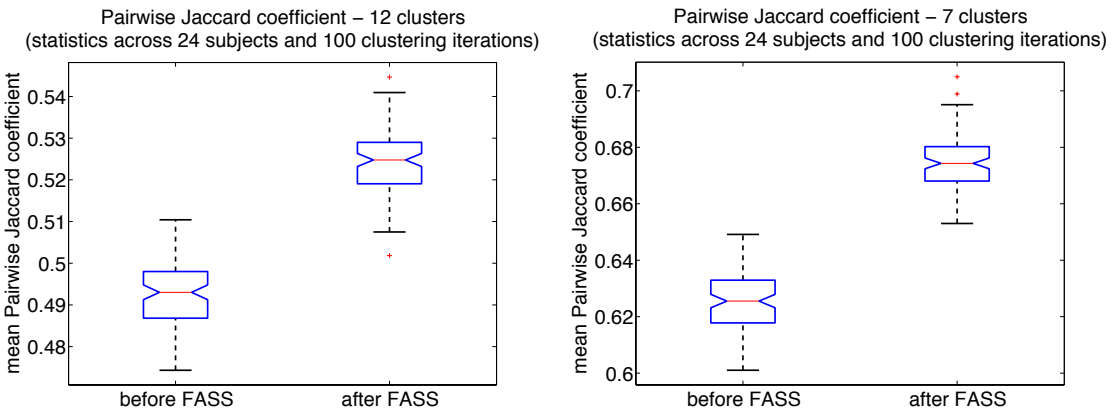

Figure 4.14: The inter-subject overlap of corresponding thalamic parcels (left thalamus) before and after FASS (measured using the extended pairwise Jaccard coefficient). **(left)** Parcellation into 12 clusters. **(right)** Parcellation into 7 clusters. In both cases FASS significantly improves cluster alignment over the original FIRST-imposed vertex-to-vertex correspondence.

investigate this hypothesis, one would need an entirely independent data set, such as fMRI activations within the thalamus, to explicitly test for the improvements in the functional alignment. This is certainly one of the directions for future research.

4.3.2 AD vs. Controls - FASS Results

In Chapter 3, section 3.4.2, we examined the differences in FACS-generated deformation fields between the cortical hemispheres of healthy controls and the Alzheimer's disease patients. Here we repeat the same type of analysis, but for the set of subcortical structures. We would like to examine, using FASS, whether changes in cortical connectivity caused by AD are reflected in the FASS-generated deformation fields.

Global cortical connectivity measured using DW MRI has, so far, been rarely used to assess between-group differences in connectivity patterns of subcortical structures. In particular, it has been difficult to assess connectivity differences in subparts of the deep grey matter structures mostly because they are difficult to successfully segment (progress in parcellating some subcortical structures like hippocampus has been recently made, e.g. (Van Leemput et al., 2009)). We would like to examine

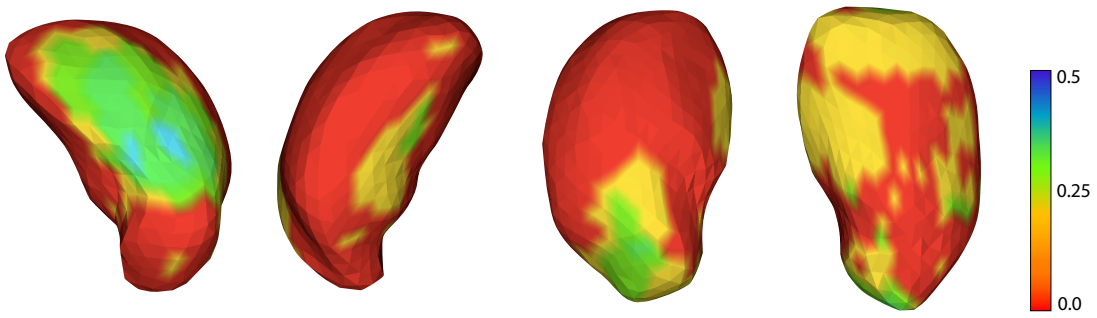

Left - medial Left - lateral Right - dorsal Right - ventral

Figure 4.15: Exploratory analysis of residual within-surface displacements between control and AD groups for the left and right thalami. FASS registration results for 12 different templates (control group subjects) were combined as explained in this section. Colourbar represents scaled mean thresholded ($p < 0.05$) maps across subjects (see main text). More specifically, value 1 indicates that in all templates that vertex has an uncorrected p-value smaller than 0.05. Conversely, values 0 indicates that at a given vertex no templates have an uncorrected p-value smaller than 0.05. In all images, the top corresponds to the anterior of the brain. The central medial aspect of the left thalamus shows changes in deformation fields between two groups (compare this to Figure 4.12). Changes are observed in the posterior part of the right thalamus as well.

whether FASS can provide an insight into the differences in connectivity between different points in the same structure (across subjects) and elucidate with greater spatial specificity which areas are affected by the pathology. The analysis to follow therefore has an exploratory character and we shall not endeavour to interpret the AD pathology or draw conclusions about the accompanying changes in subcortical connectivity.

A group of 24 subjects (12 controls and 12 AD patients) has been co-registered in two ways using FASS. Left and right thalamus, putamen, pallidum and amygdala were analysed. First, all subjects (subcortical structures thereof) were registered to the common templates (template CMs derived as explained in section 4.3.1) derived from the 12 control subjects from the same data set. In that way, 24 deformation fields were generated and between-group differences were analysed using Hotelling's T^2 test (see Chapter 3, section 3.4.2). The second type of analysis used individual control subjects for templates. In that way, each of the control subjects acted as a

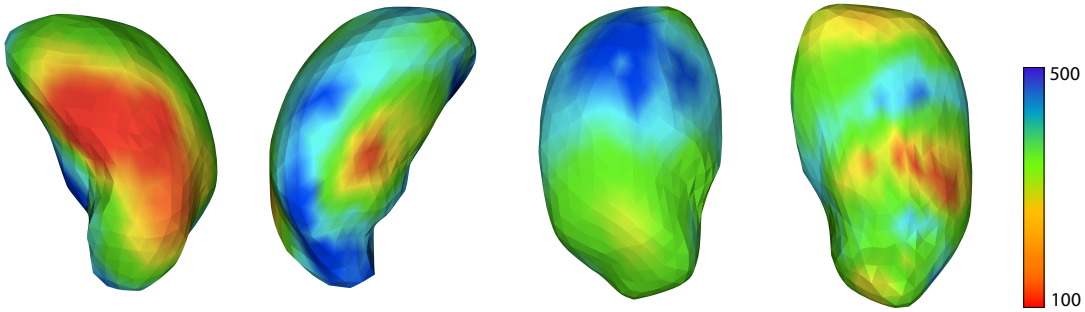

Figure 4.16: Group rank analysis of the uncorrected p-values for different templates (see main text for details) for the left and right thalami. P-values with consistently low rank (i.e. p-values consistently low compared to the others from the same analysis) across templates show up as areas of low mean rank (yellow to red colour). These areas correspond to the ones in Figures 4.12, 4.15 and 4.17 indicating group differences in deformation fields. This suggests that for different templates, low p-values, although not below a statistically acceptable threshold, occur at spatially similar positions. In all images, the top corresponds to the anterior of the brain. The colourbar represents mean (across templates) p-value rank, theoretically ranging from 1 to 642 (number of vertices in a subcortical mesh).

template and 23×12 registrations were performed. The idea behind this was to test for sensitivity of the analysis to the choice of templates.

The first type of analysis (FASS to template CMs) resulted in no statistical differences between the two groups in any of the analysed structures (neither for the left nor for the right hemispheres). Results for the left and right thalami are shown in Figure 4.17. Note that the regions with lower p-values correspond to those in Figure 4.15 where p-values were consistently lower than 0.05 across single subject templates.

In the second type of analysis, registration to some of the single subject templates generated differences between the groups ($p < 0.05$, uncorrected). To represent these findings, we 'pooled together' results from all single subject templates. Since FASS establishes inter-subject vertex-to-vertex correspondence, it is possible to determine (for every vertex) the number of templates for which $p < 0.05$ in a given vertex. In other words, for each template t, a vector P^t: $P^t(i),\ i \in [1, N]$ was generated, where

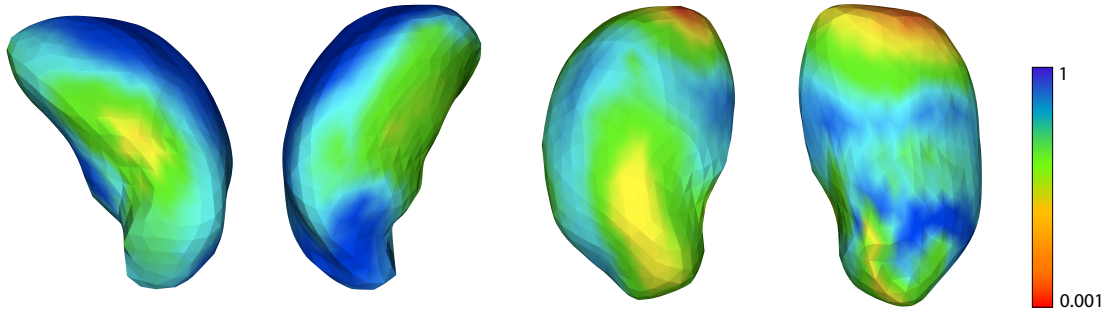

Figure 4.17: Exploratory analysis of residual within-surface displacements between control and AD groups for the left and right thalami - FASS registration to the Template CM. Colourbar represents uncorrected p-values. In all images, the top corresponds to the anterior of the brain. Although no p-values are smaller than 0.05, the regions of lower p-values resemble those in Figure 4.15 where the analyses of group differences were combined for 12 single subject templates. One could speculate that the template CMs proved too 'smooth' to capture subtle differences that exist between groups. Conversely, one-subject templates might be too specific resulting in low robustness of the findings with respect to the template.

N is the number of vertices in a structure and $P^t(i)$ is a binary indicator having value 1 for $p < 0.05$ for vertex i and value 0 otherwise. The results across templates were combined assuming the correspondence in vertex labels across subjects (i.e. due to FASS, vertex i in one template corresponds to vertex i in the other). Thus, the 'pooled' results for a given structure were generated by summing and scaling vectors P^t across different templates: $P(i) = \frac{1}{T}\sum_t P^t(i)$, where T is the total number of templates. These combined results are shown in Figures 4.15, 4.18 and 4.20.

Figure 4.15 shows the combined results for the left and the right thalami. The group differences in residual within-surface displacements are predominantly found in medial parts of the left thalamus and the posterior (putatively pulvinar, which is know to be highly connected to the hippocampus (Zarei et al., 2010)) area of the right thalamus. Note that the results for the left thalamus are similar to those derived using VSCDs (Figure 4.11). These findings also differ from what can be associated with structural thalamic shape changes in AD (Zarei et al., 2010)[8] indicating that

[8]Shape changes were mostly observed in the dorsal and ventral parts of both left and right

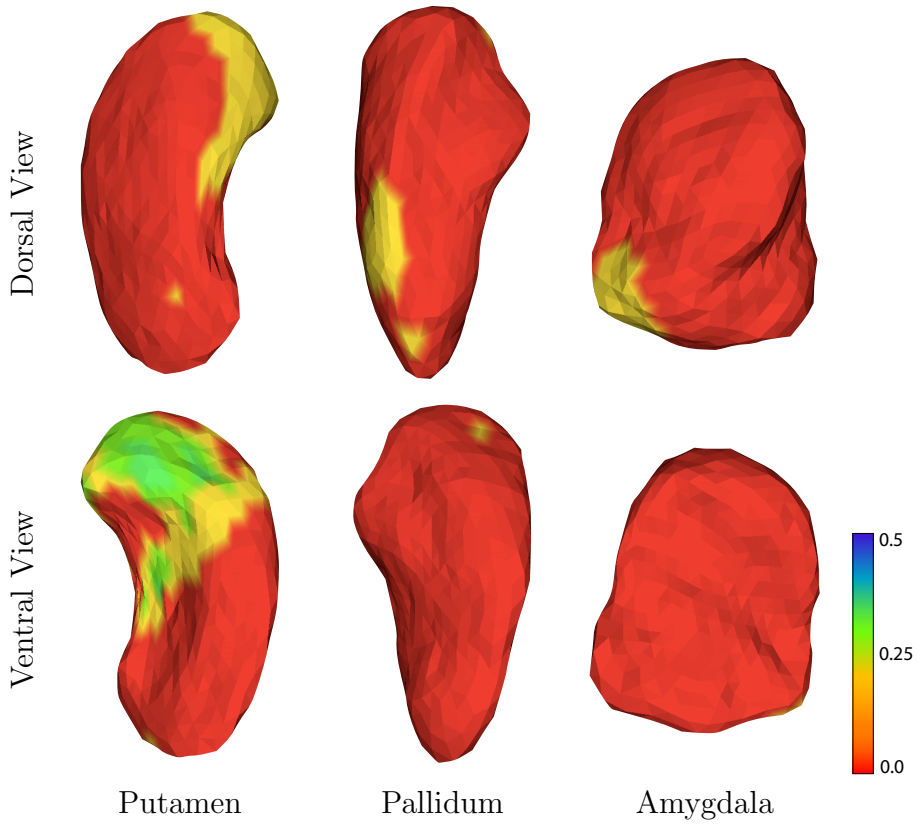

Figure 4.18: Exploratory analysis of residual within-surface displacements between control and AD groups for the left putamen, pallidum and amygdala. FASS registration results for 12 different templates were combined as explained in this section. Colourbar represents scaled mean thresholded ($p < 0.05$) maps across subjects. Value 1 indicates that in all templates that vertex has an uncorrected p-value smaller than 0.05. Conversely, values 0 indicates that at a given vertex no templates have an uncorrected p-value smaller than 0.05. In all images, the top corresponds to the anterior of the brain.

some of the observed effects are not necessarily related to structural shape changes e.g. atrophy.

Figure 4.18 shows an equivalent analysis for the left putamen, pallidum and the amygdala. The left putamen shows group differences in the deformation fields in the anterior ventral aspect of the structure (see Figure 4.22 for the individual single subject template results). No such effect is observed for the right putamen (Figure 4.20). Almost no differences are seen for the left pallidum and the left

thalami.

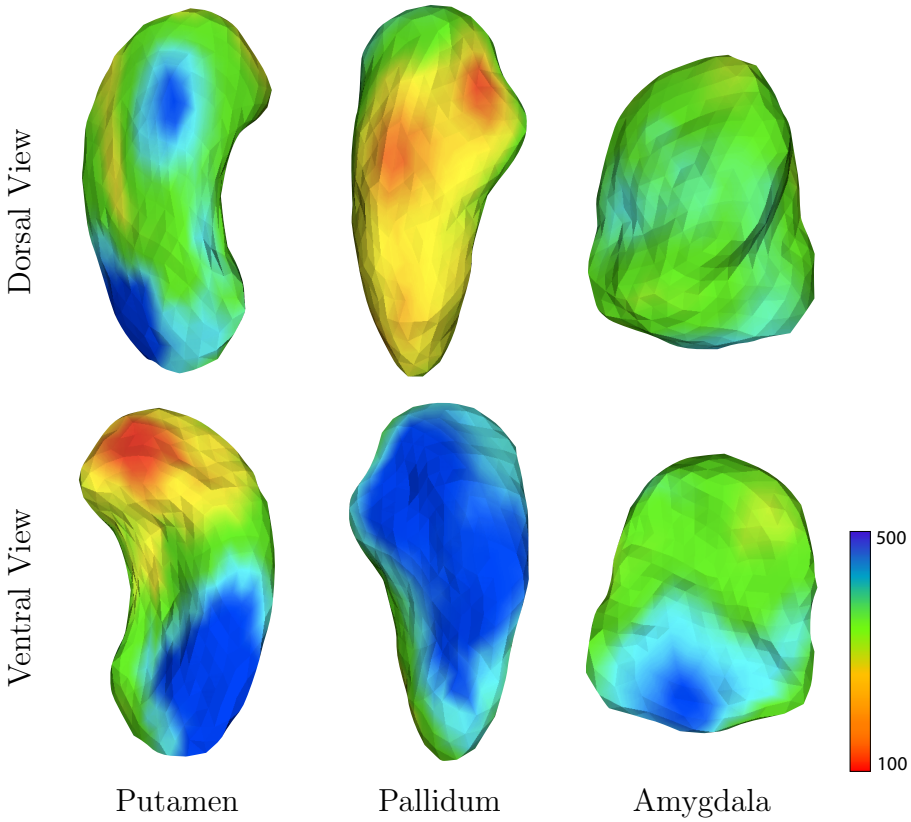

Figure 4.19: Group rank analysis of the uncorrected p-values for different templates for the left putamen, pallidum and amygdala. Areas of consistently low rank across templates are similar to the ones in Figure 4.18 indicating group differences in deformation fields. This suggests that for different templates, low p-values, although not below a statistically acceptable threshold, occur at spatially similar positions. In all images, the top corresponds to the anterior of the brain. The colourbar represents mean (across templates) p-value rank.

amygdala (only one template yielded areas with $p < 0.05$). The right putamen and pallidum show almost no differences whereas some differences exist in the lateral dorsal and ventral aspects of the right amygdala (Figure 4.20). In summary, the observed effects in all structures have pronounced non-bilaterality and are sensitive to the choice of template. None of the effects survive the correction for multiple comparison.

To further investigate the dependency of the results on templates we used p-value *rank* to assess whether low p-values occur across different templates in a spatially consistent manner. In this case, for each template t, $R^t(i)$, $i \in [1, N]$ (N is the

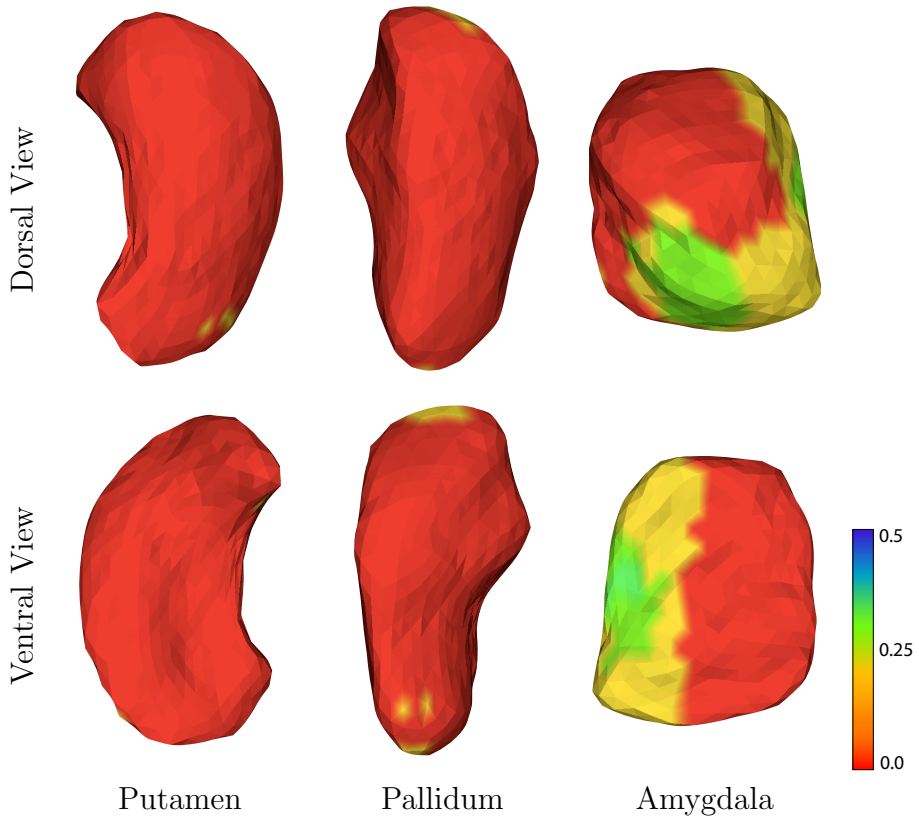

Figure 4.20: Exploratory analysis of residual within-surface displacements between control and AD groups for the right putamen, pallidum and amygdala. FASS registration results for 12 different templates were combined as explained in this section. Colourbar represents scaled mean thresholded ($p < 0.05$) maps across subjects. Value 1 indicates that in all templates that vertex has an uncorrected p-value smaller than 0.05. Conversely, values 0 indicates that at a given vertex no templates have an uncorrected p-value smaller than 0.05. In all images, the top corresponds to the anterior of the brain.

number of vertices, i.e. 642) represents the rank of p-value of vertex i in the analysis of template t. The rank quantifies the ordering of p-values from the lowest to the highest[9]. The combined ranking (across subjects) for a given structure was generated by summing and scaling vectors R^t across different templates: $R(i) = \frac{1}{T}\sum_t R^t(i)$, where T is the total number of templates. Therefore, $R(i) = 1$ indicates that for vertex i all templates have a minimum p-value associated with it. Figures 4.16,

[9]The lowest rank, 1, of a vertex i means that vertex i has the lowest p-value in the analysis for template t. The highest rank, N, means that vertex i has the largest p-value in the analysis for template t.

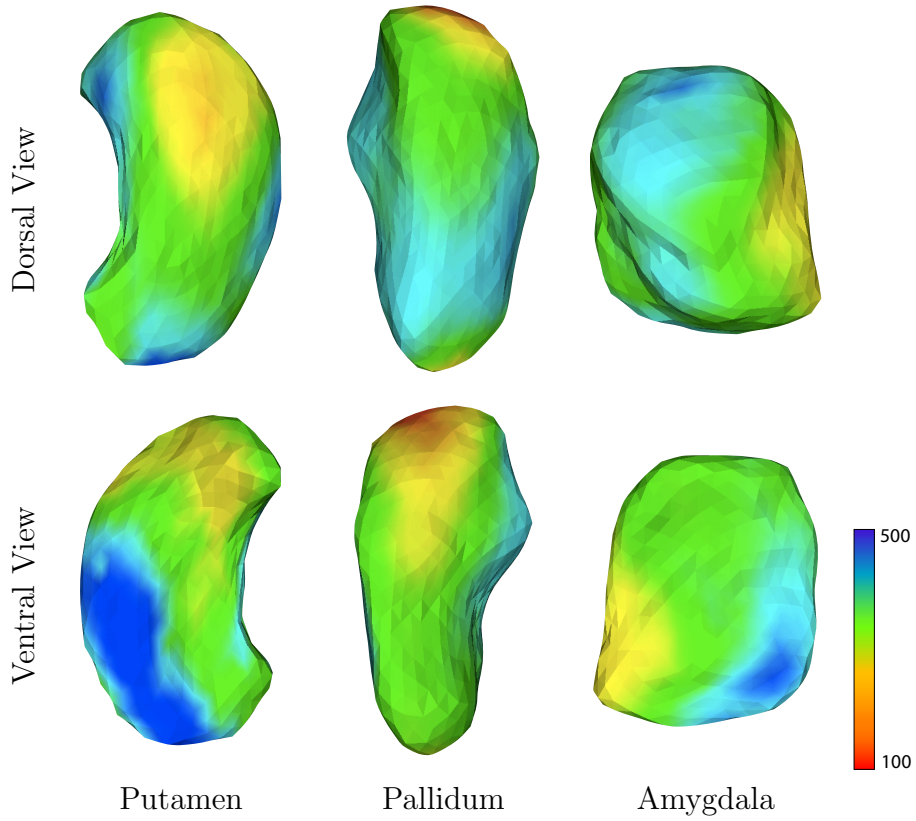

Figure 4.21: Group rank analysis of the uncorrected p-values for different templates for the right putamen, pallidum and amygdala. Areas of consistently low rank across templates are similar to the ones in Figure 4.20 indicating group differences in deformation fields. This suggests that for different templates, low p-values, although not below a statistically acceptable threshold, occur at spatially similar positions. In all images, the top corresponds to the anterior of the brain. The colourbar represents mean (across templates) p-value rank.

4.19 and 4.21 show the rank analysis for the thalamus, the putamen, the pallidum and the amygdala. For all analysed structures, and especially for the left thalamus, left putamen and the right amygdala, the areas of consistently low (p-value) rank resemble the areas highlighted in the combined analysis of thresholded p-value maps (Figures 4.15, 4.18 and 4.20). This suggests that low p-values (although they might not be below the 0.05 threshold) consistently appear in spatially similar areas across different templates. This is encouraging as it could imply that our tests might be underpowered due to a small number of subjects. We should also acknowledge some discrepancies between the rank analysis maps and the combined thresholded p-value

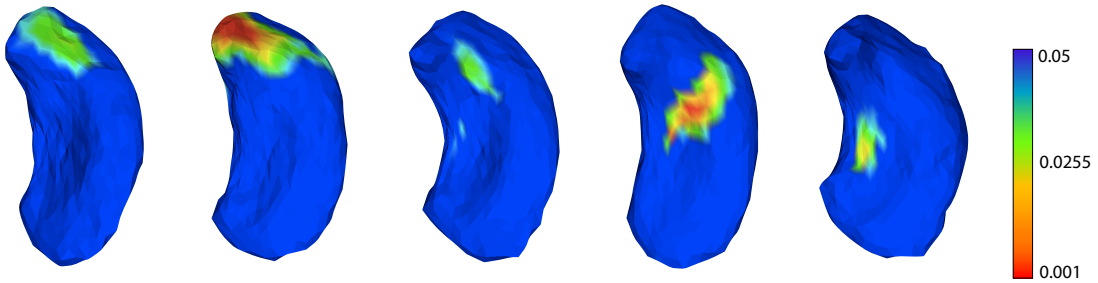

Figure 4.22: Analysis of residual within-surface displacements between control and AD groups for the left putamen (ventral aspect is shown). Colourbar represents uncorrected p-values. Results for 5 different (control) single-subject templates are shown in their own native space (hence the difference in shape and size). In all images, the top corresponds to the anterior of the brain.

maps (see the left pallidum and the left amygdala). These exist due to outliers that are more likely to be seen in the combined thresholded p-value analysis (i.e. spatially small areas with $p < 0.05$ in just one template). On the other hand, in the rank analysis, the outliers are heavily downweighted and likely to make lesser impact.

Analysis Ideally, the connectivity-driven registration should reveal areas within subcortical structures that have differing connectivity patterns in the two groups. Such results could point to the parts of the structures whose connectivity should be additionally examined. Questions like what changes in connectivity are driving the group differences in deformation fields, could be asked. Are the tracts originating in those areas degenerating, or connecting to different places (due to pathology) compared to the control group?

Therefore, the analysis presented here has exploratory character and as such has implied certain areas (such as the medial portion of the thalamus) whose involvement in AD can be expected. Nevertheless, additional analysis is needed to establish how the observed areas relate to the AD-associated connectivity changes. Unfortunately, the task of ascribing the observed areas a direct and unequivocal interpretation has proved challenging. The results heavily depend on the choice of the template. Nevertheless, we have shown that potentially interesting areas (with low p-values)

occur at spatially similar locations for different templates (Including the Template CM).

4.3.3 HDPM Clustering of Thalamic SCFVs

So far we have used similarities between full SCFVs to drive FASS. In other words, the Similarity Matrix (see Chapter 2, section 2.9) had real values from the $[0, 1]$ interval. Every vertex (of the input subject) had a unique SCFV/VSCD descriptor attached to it which was matched to the reference. In the subcortical framework we have shown how certain structures (Figure 4.13) can be parcellated into clusters resembling anatomical/functional subdivisions. These clusters can also be identified across individuals especially in the case of the thalamus (Behrens et al., 2003c; Jbabdi et al., 2009a). In this section we examine whether FASS can be driven by the cluster correspondences instead of full SCFVs. Could we use real thalamic clusterings (with inter-subject cluster correspondence) to construct SMs? The idea behind this is to use real data (clustering of subcortical structures - in particular, the thalamus) to construct registration tasks. In other words, we use class labels/cluster boundaries (not full SCFVs) derived from real data to form a similarity matrix. We then match corresponding *clusters* across individuals and examine the registration outcomes in different groups (e.g. controls and AD patients). However, the proposed registration scheme hinges on accurate and consistent (across subjects) parcellation of subcortical structures of interest. For this reason we use a hierarchical clustering method capable of combining data from different subjects and automatically imposing inter-subject cluster correspondence.

Hierarchical Dirichlet Process Mixture Models (HDPM) Hierarchical infinite mixture modelling was first used by Jbabdi et al. (2009a) to carry out connectivity-based parcellation of the cortex and subcortical structures. The approach models voxel-wise anatomical connectivity as an infinite mixture of

multivariate Gaussian distributions, with a Dirichlet process prior on the cluster parameters. This allows for automatic inference of the number of clusters and also combines data from different subjects to account for inter-subject variability and at the same time produce more robust estimates of the individual parcellations. Multiple subject clustering is performed via a hierarchical mixture of Dirichlet processes and is the feature that we shall use to automatically establish inter-subject cluster correspondence. We should also stress that HDPM-based clustering does not make use of inter-subject registrations and the correspondence among cluster labels is established solely according to the data matrix fed into the HDPM pipeline. A single subject data matrix consists of vertex indices, information about closest neighbours for each vertex and of SCFVs associated with every vertex. The full data matrix fed into HDPM pipeline is a concatenation of single subject data matrices.

4.3.3.1 Using Inter-subject Thalamic Cluster Correspondence to Drive Registration

The HDPM was used to hierarchically cluster left and right thalami (surface meshes thereof) of a set of 24 subjects (12 controls and 12 AD patients). Infinite mixture modelling was not used as we found that the automatically estimated number of clusters regularly exceeded 30 with many clusters consisting of single vertices. Therefore, for this particular data set, we fixed the number of clusters to 7, which proved a good compromise between capturing anatomically plausible divisions and achieving reasonable inter-subject consistency. For greater number of clusters, i.e. 9 or 12, the HDPM clustering often resulted in several clusters consisting of one vertex only or even failed to identify full number of clusters in all subjects. CMs (formed using SCFVs) were fed into the HDPM processing pipeline with all clustering parameters set to their default values (Jbabdi et al., 2009a). The template CM (section 4.3.1) entered the clustering cohort in the same way as the single subject data. Each row of the template CM is a template SCFV associated with a specific vertex. Therefore, the

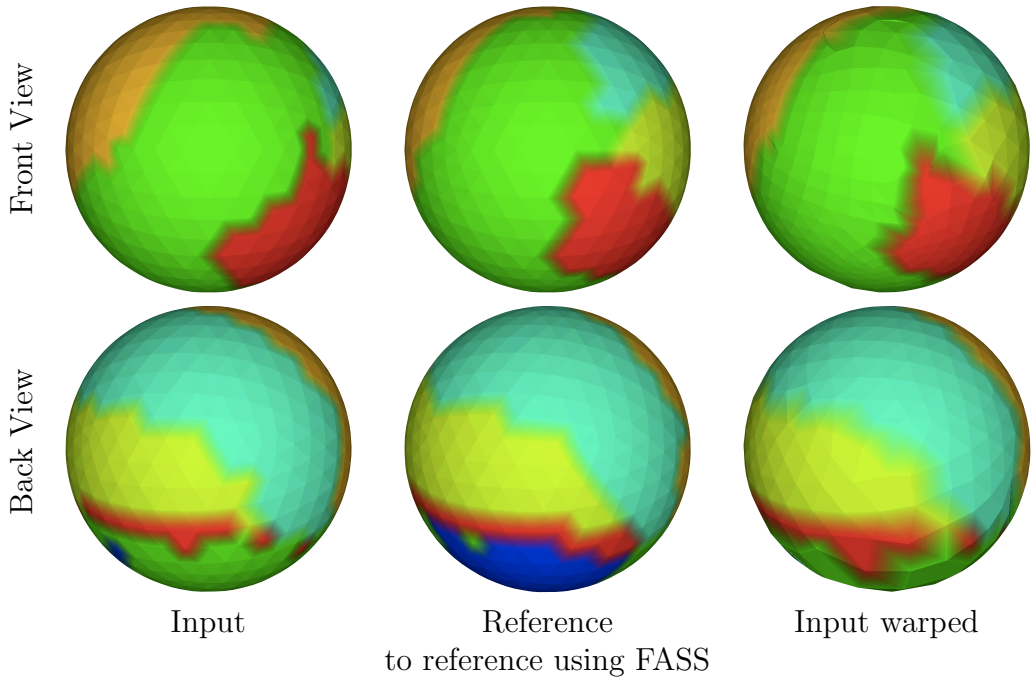

Figure 4.23: An example of a FASS registration driven by correspondences in thalamic clusters. Thalamic parcellations, for the input and the reference subjects, are represented on registration spheres. Thalami were parcellated into 7 clusters (6 visible) each imposing inter-subject cluster correspondence using HDPM. Different colours correspond to different clusters. The figure shows how the input subject is matched to the reference to maximise the overlap of corresponding clusters.

template CM can be clustered in the same way as single subject data by assigning a cluster label to every vertex. We call this the clustering of the template CM or the 'prototype clustering'. The inclusion of the template is necessary as it needs to be clustered as all the other subjects, i.e. cluster label correspondence has to be established between every subject and the template. The template clustering is then used as the references for FASS, avoiding use of a (potentially biased) single subject-based template. In other words, the SMs are calculated between the template clustering and the clusterings of all other subjects (which are in correspondence by construction). The HDPM clustering was performed separately for the left and right thalami.

The HDPM clustering assigns a unique label to a specific cluster (e.g. the one corresponding to the medial dorsal thalamic nucleus) in all subjects in both the control

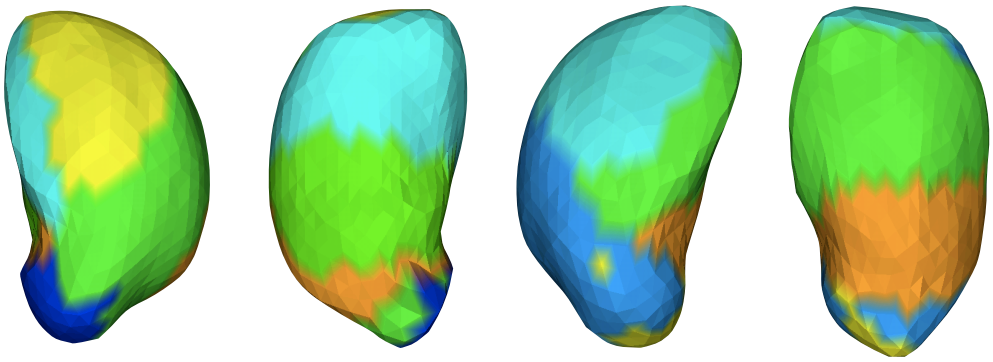

Left Dorsal Left Ventral Right Dorsal Right Ventral

Figure 4.24: HDPM clustering of the Template CM of the left and right thalami as part of the hierarchical clustering cohort - a so called 'prototype' or reference clustering. By clustering the template in this manner, we establish by construction the correspondence in cluster labels between the template/reference and every other single subject. SMs encode cluster correspondences between the template and the input (single) subjects. Different colours represent different clusters (7 in total). Note that clusters are labelled/coloured differently in the left and right thalami.

and AD groups. In the next step we construct SMs using the cluster correspondences. The SMs in this case are binary matrices:

$$\mathrm{SM}(i,j) = \begin{cases} 1, & L_I(i) = L_R(j) \\ 0, & \text{otherwise} \end{cases}, \tag{4.5}$$

where $L_I(i)$ is the cluster label associated with the vertex i of the input subject and $L_R(j)$ is the cluster label associated with the vertex j of the reference/template subject. In the analysis throughout this section, the reference clustering L_R is defined by the clustering of the template CM (Figure 4.24). In this way, 'patches' with the same labels in the input and the reference/template subjects are 'perfectly correlated' as indicated by Equation (4.5) whereas all the other entries in the SM are set to zero.

The SMs (24 in total) become the inputs into the FASS algorithm. The registration outcome for one particular instance is presented in Figure 4.23. After FASS, the cluster correspondence between the input and the reference meshes is markedly better. The quantitative assessment of the improvement in the inter-subject cluster overlap was measured using extended pairwise Jaccard coefficient

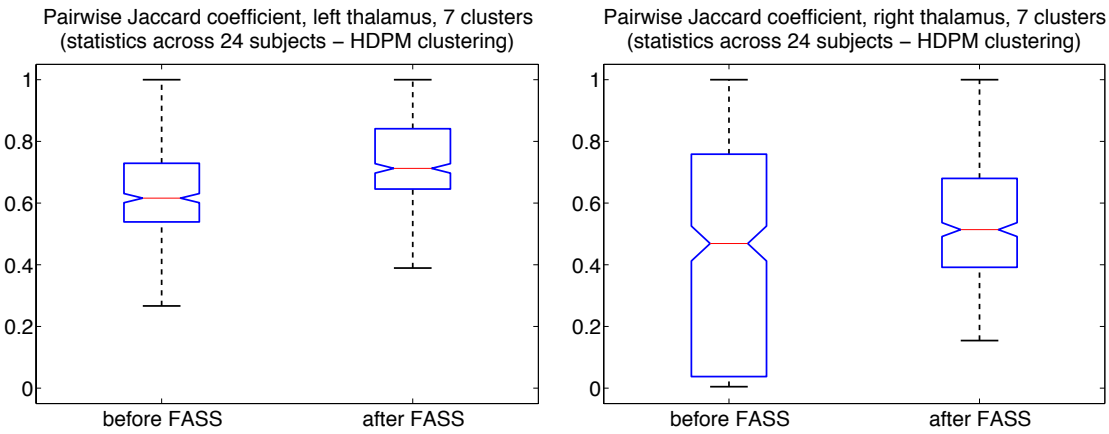

Figure 4.25: The inter-subject overlap of corresponding thalamic parcels before and after FASS (measured using the extended pairwise Jaccard coefficient). **(left)** Left thalamus. **(right)** Right thalamus. In both cases, the mean pairwise Jaccard coefficient was higher after FASS compared to before ($p < 10^{-10}$). Analysis is similar to that from Figure 4.14, however, in this case, a different clustering/cluster correspondence paradigm was used as well as binarised SMs (to drive registration).

(Chapter 3, section 3.3.2) with 24×24 comparisons made in total. The summary of those results is shown in Figure 4.25. For both the left and the right thalami, the mean pairwise Jaccard coefficient was significantly higher after FASS compared to before ($p < 10^{-10}$)[10].

4.3.3.2 AD vs. Controls - Analysis of Cluster Displacements

The analysis of cluster displacements before and after FASS has the potential to inform us of which areas (clusters) underwent greatest deformations to match the template. Similarly, when analysing two different groups of subjects, one could ask which areas remain inconsistently clustered once the registration has been performed. This type of analysis has an advantage that detected changes can be associated with certain thalamic clusters that potentially bear the anatomical or functional connotations.

Figure 4.26 shows the effects of registration on the alignment of corresponding clusters across the whole group of subjects (AD patients and controls analysed

[10]The higher the Jaccard coefficient, the more similar the clusterings are.

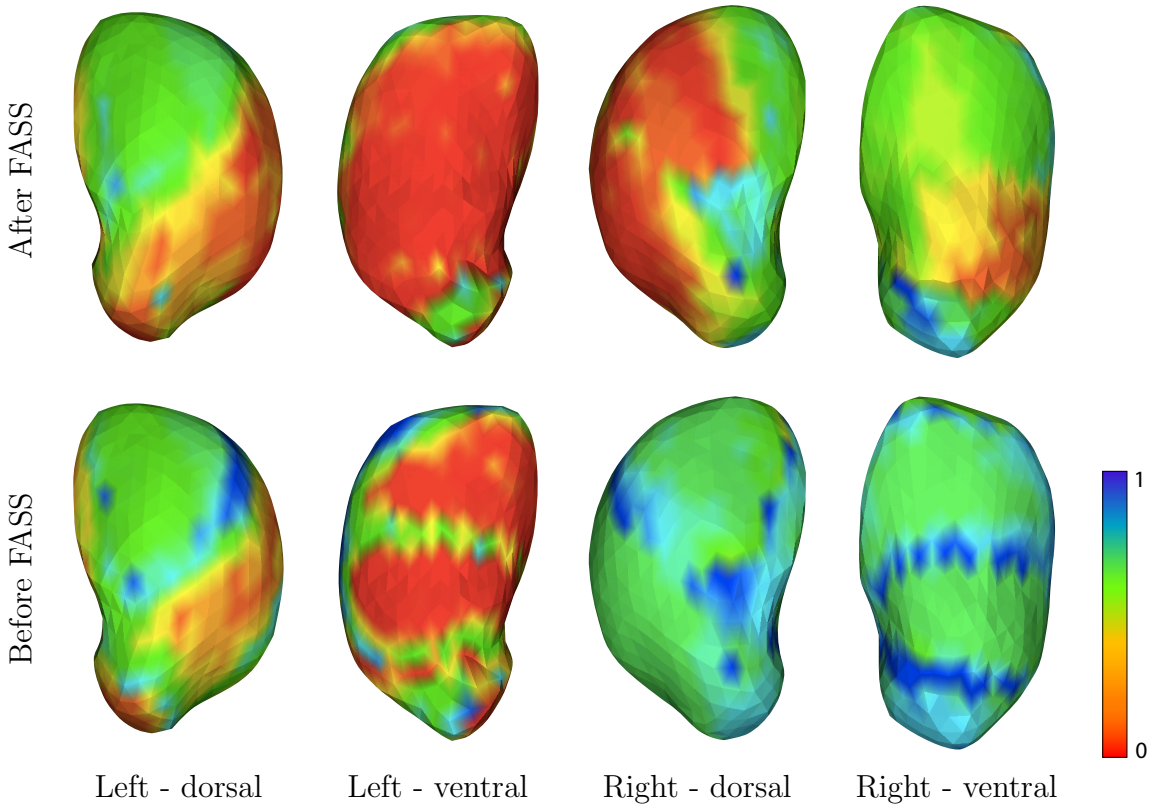

Figure 4.26: Demonstration of the registration effects on the inter-subject overlap of corresponding clusters. Left and right thalami are shown. The top of each image corresponds to the anterior of the brain. The colourbar indicates the quality of the inter-subject overlap of corresponding clusters. Low values (red) denote good overlap whereas high values (blue) denote poor overlap. The extent of 'red' areas increases for both left and right thalami after connectivity-driven registration supporting findings from Figure 4.25.

together). The prototype/template labelling (Figure 4.24) was compared to the cluster labelling of individual subjects and the number of 'incorrect labellings'[11] across subjects was recorded for each vertex. This number was divided by the total number of subjects to produce a *non-overlap coefficient* scaled to the $[0, 1]$ interval. Therefore, the value of 1 indicates that no subjects have the prototype cluster label associated with that vertex. Conversely, a value of 0 means that all subjects have the 'correct'/prototype cluster label at that point/vertex. Therefore, low coefficient

[11] An incorrect labelling at a given vertex occurs when the prototype (cluster) label at that vertex is not the same as the individual subject's (cluster) label at the same vertex.

values indicate good inter-subject alignment of corresponding clusters. Prototype cluster labelling, as well as the corresponding non-overlap coefficient, were separately calculated for the data before and after registration.

Spatial maps showing the non-overlap coefficient for the left and right thalami before and after registration are presented in Figure 4.26. As expected, the figure shows improvements in inter-subject cluster overlap after FASS compared to before. This is especially evident for the dorsal aspect of the right thalamus. The ventral part of the left thalamus also shows similar improvements especially around the likely cluster boundaries. However, from this analysis we cannot conclude anything about the possible differences between the controls and the AD patients in the alignment of corresponding clusters. To gain a further insight into the differences between the two groups we repeated the analysis carried out in Figure 4.26 for the control and AD groups separately.

Figure 4.27 shows the results of the analysis for the left, and Figure 4.28 for the right thalamus, with the AD and control groups analysed separately. Prototype clusterings were calculated for each group, before and after FASS and were used to obtain the non-overlap coefficient. Figure 4.27 demonstrates how the similarity between clustering patterns of the AD and those of the control groups increases after FASS compared to before FASS. From the same figure, one can also infer what areas (in the AD and control groups) underwent greatest changes to match the template.

After a hypothetical ideal registration, the AD and control groups would have had the same clustering patterns (with 0 non-overlap coefficient across the whole structure). However, due to imperfect alignment, the residual differences (reflected in clustering consistency) exist. These residual differences can provide an additional insight into distinctions between two groups not accounted for by registration. In the case of the left thalamus, these residual differences are predominantly localised in the (anterior and central) dorsal aspects of the structure as well as in the posterior ventral

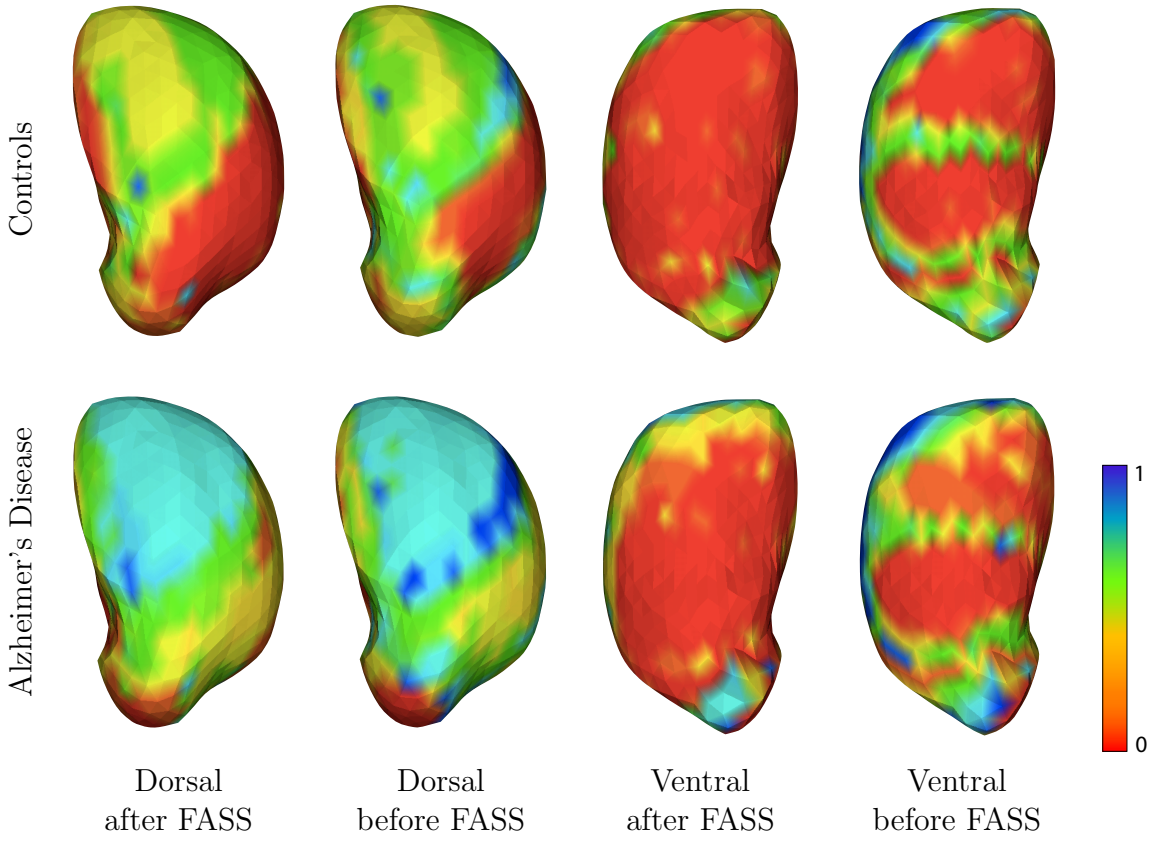

Figure 4.27: The effects of FASS cluster matching were assessed separately for the AD and control groups. The maps of non-overlap coefficient values are shown for the left thalamus. The top of each image corresponds to the anterior of the brain. The colourbar represents non-overlap coefficient values - red colour indicates good inter-subject overlap of corresponding clusters, whereas blue indicates areas of poor overlap.

part (Figure 4.27). Therefore, even after FASS matching corresponding clusters, certain areas exhibit different clustering consistency patterns in the AD and control groups. The same regions are also found to be heavily involved in AD: the dorsal nuclei are the primary site of cell loss and neurofibrillary tangle formation in the thalamus whereas the posterior thalamus is primarily connected to the hippocampus - also involved in AD (Xuereb et al., 1991; Braak and Braak, 1991; Karas et al., 2003b).

Figure 4.28 shows an equivalent analysis for the right thalamus. Clustering consistency markedly improves in the whole dorsal aspect of the thalamus for both

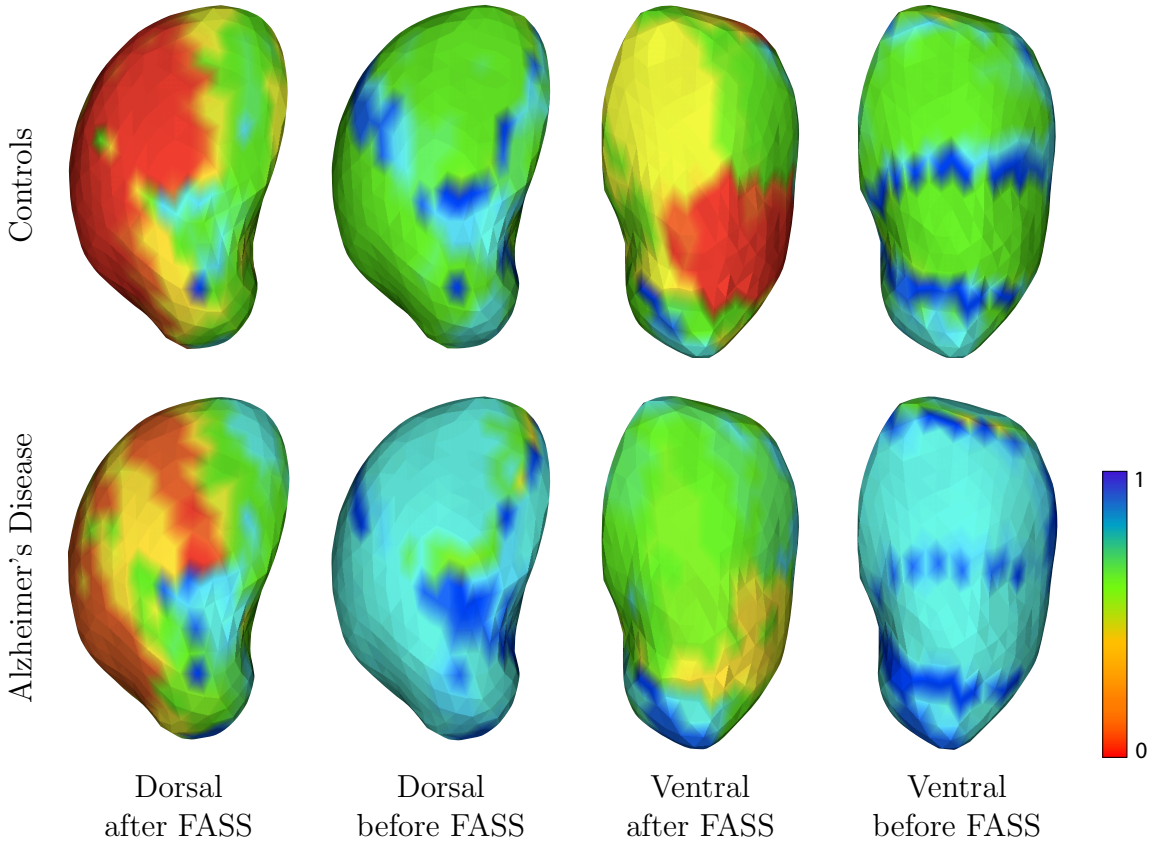

Figure 4.28: The effects of FASS cluster matching were assessed separately for the AD and control groups. The maps of non-overlap coefficient values are shown for the right thalamus. The top of each image corresponds to the anterior of the brain. The colourbar represents non-overlap coefficient values - red colour indicates good inter-subject overlap of corresponding clusters, whereas blue indicates areas of poor overlap.

the AD and control groups[12]. Unlike in the left thalamus where group differences in clustering consistency were concentrated primarily in the dorsal aspect, in the right thalamus the group differences seem to have much wider spread.

Analysis In this section we demonstrated how thalamic (surface) parcellation can be fed into the FASS processing pipeline and used for matching corresponding clusters across subjects. Parcellation was performed using connectivity data (SCFVs) and cluster correspondence across individuals was enforced through the HDPM clustering

[12]One could speculate that a large improvement in cluster correspondence may be due to poor initial FIRST-imposed alignment.

scheme. If thalamic clusters bear anatomical/functional meaning (e.g. correspond to thalamic nuclei), then cluster matching enhances the anatomical/functional alignment. Therefore, the registration results can be associated with specific clusters potentially increasing their interpretability.

The cluster matching scheme was presented as a demonstration of an another way of examining the registration results, complementary to that introduced in section 4.3.2. We have also shown that FASS can be driven by binary SMs, similar to those of the artificial examples, but generated using real data. Finally, the cluster matching was used to explore the differences in cluster alignment after registration in the AD and control groups. The results for the left thalamus indicate the differences in the anterior and medial dorsal aspect of the thalamus, whereas almost no differences were detected for the right thalamus. Although these findings are partially supported by histologic evidence, the absence of bilateral symmetry suggests the asymmetric (changes in) connectivity profiles of the left and right thalami in AD. Similar asymmetry has already been reported in section 4.3.2 (Figure 4.15). Moreover, the left-right asymmetries of diffusion-based measures in the thalamus, such as the Apparent Diffusion Coefficient (ADC), have been observed even in healthy controls (Kantarci et al., 2001). However, further research is needed to understand how the observed asymmetries relate to thalamic connectivity.

Successful cluster matching crucially depends on reliable thalamic parcellation and consistent cluster labelling across subjects. Therefore, the performance of such registration is limited by HDPM. In particular, we have observed that the automatic estimation of the number of clusters can result in difficulties in identifying corresponding clusters across individuals, e.g. a large cluster is identified in one subject, but 'a corresponding' cluster occupies just a couple of vertices in the other subject (or even does not exist). Finally, when matching binary SMs, the areas within clusters are matched so as to only minimise the regularisation part of the

cost function since the intensity cost function is flat in such areas (non-informative data). On the other hand, FASS registration using full SCFVs (not clusters) does not explicitly encounter these problems, but has results that are potentially more difficult to interpret. Therefore, it depends on the actual application as to how FASS (and the SMs) should be manipulated to maximise registration effectiveness and increase the interpretability of results.

4.4 Discussion

In this Chapter we introduced a framework for connectivity-driven registration of surfaces of subcortical structures segmented by FIRST. The connectivity-driven alignment was presented as the enhancement of vertex-to-vertex correspondences established by FIRST (section 4.1.2). We constructed an artificial test data set and showed how it can be used (the ABCA test) to measure registration consistency and estimate free parameters (section 4.2.2.1). Although the presented parameter estimation method does not provide a framework for automatic parameter inference, it nevertheless offers a systematic way for identifying parameter values that can subsequently be fine-tuned to a specific application (e.g., Chapter 2, section 2.4.4.5).

To better examine the influence of the actual registration algorithm on the final results we developed and applied a novel surface registration method based on spherical wavelets. The tests of both methods on artificial and real data resulted in similar deformation fields suggesting that the connectivity-driven registration task converges to a common solution irrespective of the optimisation framework (section 4.2.3). We also introduced the concept of VSCDs, which unlike SCFVs, encode connectivity to the whole brain and commented on potential benefits of using one descriptor over the other (section 4.3).

In the next step we used FASS to explore the differences in deformation fields between two groups of subjects - AD patients and healthy controls (section 4.3.2).

Four different subcortical structures were examined with the special attention devoted to the thalamus. Although the observed differences in deformation fields in the left thalamus are in accordance with histologic evidence, the analysis of real data was laced with several challenges:

- High dependency of the results on the template;
- Low statistical power of the test examining differences in deformation fields (no analysis survived the correction for multiple comparisons);
- Difficult interpretation of the registration results.

The common template derived from the control subjects (template CM) did not help reduce the subject-specific bias. In fact, no differences could be seen between the groups when the template CM was used as reference (No p-values were lower than 0.1. However, the areas of low p-values resemble those detected with the other templates - compare Figures 4.17 and 4.15). This might indicate high variability (i.e. of connectivity profiles) of thalami of different subjects, which presents a problem for statistical inference. Nevertheless, we have shown that irrespective of the template, low p-values[13] (areas of potential group differences in the deformation fields) occur at spatially consistent locations possibly suggesting an underpowered test. In that sense, additional data could boost registration consistency as well as the power to discriminate between the groups. In other words, low statistical power of the test examining differences in deformation fields might be related to this intrinsic data variability and the relatively small number of subjects. Although we are looking for the subtle changes in AD, e.g. complementary to thalamic atrophy, our study has just about half the number of subjects of a typical Volumetric-Based Morphometry (VBM) study investigating volumetric changes in AD, e.g. (Karas et al., 2003a).

[13]Low compared to the other p-values from the same group/analysis.

In the ideal case, the observed differences in the deformation fields would point out areas with differing white matter connectivity patterns. Further tests could then be carried out to examine where exactly those areas project to and what feature of white matter fibres is different between groups. However, before examining differing white matter connectivity patterns, one needs to disassociate potential white matter changes from atrophy, segmentation errors and measurement noise. Future work should examine in detail how all of these factors contribute to the discovered group differences, as well as how they relate to the usage of different connectivity descriptors (SCFVs or VSCDs).

To further explore the AD data set and try to increase the interpretability of group results we performed FASS with SMs derived using correspondences in thalamic clusters (section 4.3.3.1). This analysis also served as yet another demonstration of FASS capabilities to work with different SMs (so far, these were constructed using SCFVs, RSNFVs, VSCDs, as well as direct cluster correspondences). The revealed differences in the inter-subject cluster correspondence between the AD and control groups can now be associated with individual thalamic clusters (putative nuclei) opening up opportunities for the analysis of cluster size/area, deformation patterns, etc. Although associated with difficulties related to thalamic clustering and cluster correspondence, this alignment approach has the potential to increase the interpretability of registration results. In summary, the FASS registration pipeline could incorporate the clustering procedure as part of the postprocessing stream by automatically identifying subcortical clusters/nuclei and relating them to the areas of suspected connectivity changes.

Future work should, among other things, investigate the potential of resting state fMRI in the registration of subcortical structures. One could start by answering questions like: can subcortical structures be parcellated using resting state (RS)

data (e.g. by measuring resting state functional connectivity) in a similar way as using structural connectivity information? Do these two parcellation methods yield similar solutions? Should such analysis reveal potential of RS fMRI for subcortical parcellation, one could then think about constructing suitable RS feature vectors that can drive the alignment of subcortical structures. Some recent research speaks in favour of such hypothesis. Mezer et al. (2009); Zhang et al. (2008) have used RS fMRI timeseries to parcellate the cortex and the thalamus and they report thalamic clusters that resemble anatomical subdivisions. Therefore, it would be of interest to investigate whether functional connectivity of different parts of the thalamus (or other subcortical structures) to a set of targets can be used to construct functional connectivity profiles comparable across subjects and informative/discriminative enough to drive registration.

Chapter 5

Conclusions and Future Work

5.1 Conclusions

Starting from a premise that structural connectivity reflects functional specialisation of the brain, we developed a set of tools that use brain connectivity information to achieve and explore connectivity-based image alignment. These tools are able to calculate, compare and utilise connectivity profiles to achieve better inter-subject alignment of certain functional brain areas. In addition, the concept of connectivity-driven registration of cortical and subcortical surfaces has been carefully validated using artificial test data as well as real data, i.e. a modality different to the one used for driving the registration. Finally, we presented potential applications for the connectivity-driven registration framework, such as the comparative analysis of the control and Alzheimer's patients groups.

5.1.1 Summary of Major Contributions

The major contributions of this thesis can be divided into three groups:

1. Development of methodologies for calculation and inter-subject comparison of brain connectivity;

2. Development of surface registration methods using information derived in 1;

3. Validation of proposed registration pipeline using a novel functional parcellation scheme.

This thesis proposed, to the best of our knowledge, for the first time how probabilistic tractography can be used for inter-subject registration. To achieve this, new methods for constructing and comparing long-range connectivity descriptors were developed. Moreover, we proposed how *both* structural and functional connectivity descriptors can be formed. We further suggested a mechanism for comparing them across subjects using a connectivity 'similarity matrix'. This matrix forms the core of the cost function driving the registration, which is one of the main novelties of this work. Currently, the formation of connectivity descriptors and the similarity matrix is done using FreeSurfer and FSL software, but the developed framework can be implemented in any software package that offers similar functionalities.

Registration methods were developed for both cortical and subcortical frameworks. This thesis argues that subcortical structures, just like the cortex, can be processed and registered as stand-alone entities, and for the first time offers a solution how connectivity-driven registration of subcortical structures, such as the thalamus, can help alignment and inter-subject identification of fine-grain structures such as thalamic nuclei. Although the core registration algorithms use already existing techniques, we also proposed a new registration method based on spherical wavelets and showed how it compares to the free deformation registration framework.

Particular merit of this work is in validating the proposed registration pipeline using an independent imaging modality, i.e. functional MRI. New methods for cortical parcellation using Resting State Networks were proposed and used to demonstrate that the alignment of functional homologues is improved after structural connectivity-driven registration. This result also has neuroscientific value in testing the relationship between structural connectivity and functional segregation.

5.1.2 Registration Methods

From its inception, the connectivity-driven registration method was designed to work with surface data: cortical surface meshes of FreeSurfer and subcortical meshes of FIRST/FSL. To achieve this, we modified already existing surface registration algorithms based on free deformation models to account for the idiosyncrasies of (structural) connectivity information obtained using probabilistic tractography tracking in DW images. These changes primarily consist in the formulation of the cost function which is now driven by similarities in connectivity profiles that can often be high-dimensional. The similarities in connectivity profiles (CPs) are therefore calculated 'off-line' reducing the overall computational burden of the main registration step.

The modified formulation of the surface registration problem required careful testing and validation on an artificial data set. The analysis of the artificial data has shown that similarities in CPs are informative enough to drive registration. Moreover, we used artificial data to test registration consistency and identify free parameters. In all these tests, the developed registration algorithms satisfied preset requirements such as the one that the accumulated registration error (per vertex) in the ABCA consistency test is smaller than the mean distance between two neighbouring mesh vertices. To further explore the properties of our registration algorithm we developed another novel surface registration framework based on spherical wavelets. The two frameworks fundamentally differ in the way the deformation field is calculated and in how the multi-scale optimisation and the regularisation are conducted. Despite these differences, we have shown that the restored deformation fields look similar in both cases which speaks in favour of a unique solution to the connectivity-driven registration problem.

We have also paid particular attention to the construction of CPs, or in case of white matter connectivity, the Structural Connectivity Feature Vectors. SCFVs need to be comparable across subjects and need to be able to locally encode subject-specific connectivity. In that sense, to the best of our knowledge, this work for the first time explores the possibilities of driving (sub)cortical registration using probabilistic tractography. Moreover, we have shown how proposed SCFVs can reveal functional boundaries not identifiable in T1-weighted images and how these boundaries can be used to (drive or) validate registration.

In order to eliminate biases that artificially increase similarities in SCFVs we constrained the tract length to at least 30 mm effectively counting only medium- and long-range connections. SCFVs can therefore be manipulated to sensitise registration to e.g. specific tract lengths or to connectivity to a pre-defined set of brain regions. We have also demonstrated how volumetric connectivity descriptors as well as binary cluster labels can be constructed and used instead of SCFVs. Finally, the developed registration methodology is equally suitable to take as inputs other types of connectivity such as functional connectivity. We have proposed a way to use Resting State Networks instead of white matter connectivity to drive registration.

Overall, the proposed connectivity-driven registration pipeline is intuitive, highly automated and brings together several cutting-edge brain MRI processing tools: FreeSurfer surface reconstruction and registration, FSL probabilistic tractography and FSL subcortical structure segmentation. We hope that (parts of) the connectivity-driven registration pipeline will become integrated into some of these tools and open up new possibilities for the comparison and the alignment of connectivity profiles.

Challenges The investigations into connectivity-driven registration resulted in many challenges some of which remain unsolved in this thesis. The SCFVs are

derived based on the initial correspondences between cortical hemispheres. In other words, the targets for probabilistic tractography are, in the current framework, calculated once at the beginning of each registration. However, during registration, new inter-subject correspondence between surfaces is imposed which also changes the correspondence between original targets. Ideally, SCFVs would be re-derived after every registration iteration. Unfortunately, such a scheme is prohibitive due to its excessive computational, storage, and time demands. Therefore, it remains for the future work to carefully examine how (or whether) one should iterate the registration-SCFV derivation process to enhance the sought-after functional correspondence. In particular, it remains unknown what advantages (if any) re-deriving SCFVs in every iteration might bring. In our experience (judging from the tests on the alignment of (sub)cortical clusters), already after the first registration pass, the correspondences in functional/anatomical areas are improved.

In the current registration algorithm, free parameters are optimised empirically through a set of artificial examples devised to test registration consistency. Ideally however, these parameters should be fine-tuned for the input Similarity Matrix constructed using different features (e.g. using RSN features and using SCFVs). We would therefore like to have a scheme where the regularisation and the intensity cost function contributions are determined automatically from the data. Additional research into how this registration pipeline can be represented as a generative model and infer free parameters e.g. in a Bayesian fashion, could expand the applicability of proposed registration methods.

In the development of the registration pipeline we have tried to make the registration process fast and computationally efficient. Nevertheless, the joint running of FreeSurfer, probabilistic tractography, construction of similarity matrices and the registration itself is a time consuming process that can take up to 40 hours of processing time on a Intel Core 2 Duo 2.4GHz processor. Memory requirements

for storing and processing VSCDs are also significant. Therefore, all validations, especially with the real data, are time consuming, which potentially limits the extent and the diversity of the tests and can thus slow down the adoption of connectivity-driven alignment as a common registration procedure. It is of great importance to continue research in improving the computational effectiveness of the connectivity-driven registration pipeline, especially by optimising the registration algorithm as well as the SCFV construction. Along the same lines, it remains an open question what exactly the SCFVs should look like - i.e. should they be low- or high-dimensional, should they encode local or global connectivity or should all of these choices be made in an application specific way. Although we have already discussed these questions and proposed solutions, a careful quantitative evaluation of the relative merits of different SCFV compositions is yet to be done.

5.1.3 Cortical Framework

The main application of connectivity-driven registration is in the alignment of cortical surfaces. Functional specialisation is related to structural brain connectivity and connectivity-driven registration should therefore improve the alignment of corresponding functional areas (which are often delineated and analysed on the cortical surface). This, in turn, can help functional neuroimaging studies by increasing the sensitivity of group analysis due to better alignment.

Therefore, we have validated FACS by assessing how well it aligns functionally homologous areas in different subjects. Connectivity-driven registration uses DW images to form the cost function that determines the final deformation field. To validate obtained deformations we used another imaging modality, functional MRI, to measure the overlaps across corresponding functional areas. In particular, we have shown that resting state fMRI can reveal boundaries between functional areas especially through the analysis of the Resting State Networks. The inter-subject

overlap of these boundaries (or the regions between them) was measured before and after FACS.

We used RSNs to parcellate the Medial Frontal Cortex into putative SMA and pre-SMA. The introduced segmentation procedure makes a novel use of RSNs and performs better than the seed-based functional connectivity. The SMA/pre-SMA was chosen as it has been shown that the SMA/pre-SMA boundary can be successfully retrieved both from the diffusion as well as the resting state data. The analysis of the inter-subject overlap of the SMA and pre-SMA areas (extracted from RSN data) before and after FACS demonstrates the improvements in cluster alignment brought by the connectivity-driven registration. FACS outperforms FreeSurfer geometry-based registration in aligning the SMA and pre-SMA clusters across subjects.

Furthermore, we have also proposed a novel way to generate maps of putative boundaries between functional areas (using a set of RSNs identifiable in every subject). Cortex-wide maps of local similarities in the RSN profiles can be spatially differentiated to reveal boundaries that resemble those between functional regions. The pre- and post-FACS alignment of these boundaries within the MFC supports previously drawn conclusions. Moreover, we have presented some evidence that FACS improves the cortex-wide inter-subject correspondence of functional areas.

The cingulate cortex is another area where resting state and diffusion data point to the same internal divisions. We therefore used the ICA-based RSN clustering of the cingulate to assess the performance of FACS versus FreeSurfer. In addition to that, we analysed the inter-subject overlap of boundaries within the cingulate. In the case of boundary analysis, the results show superiority of FACS over FreeSurfer and are robust with respect to several overlap measures. However, the inter-subject overlap of extracted clusters did not produce such unequivocal results. Namely, although

FACS consistently outperforms FreeSurfer for 'full' RSN feature-based assessment, the results of the other evaluated registration methods differ with respect to the number of clusters into which the cingulate is parcellated.

Apart from validating FACS alignment using RSNs we constructed, examined and used the SCFV template in co-registering groups of Alzheimer's patients and control subjects. The SCFV template was obtained by an iterative refinement through a series of FACS applications. Interestingly, the decomposition of the SCFV template revealed formations that highly resemble structural and functional cortical areas that were not used in any of the FACS processing steps (therefore are unlikely to occur as consequences of potential biases e.g. in formation of tractography targets). This suggests that the comparison of templates (or certain parts thereof) derived from different groups might reveal areas of differing connectivity and open up doors to construction of connectivity-based imaging biomarkers.

In this sense, we compared the deformation fields of the AD and control groups (when registered to the common template). The areas found to be different between groups project to brain regions known to be involved in AD pathology. This analysis suggests that FACS has potential to reveal areas of differing connectivity, possibly helping the identification and assessment of pathology. However, just as in the subcortical framework, FACS results (e.g. deformation fields) need to be further analysed to reveal causes of group differences and establish which connectivity features differ between groups (e.g. tract length, consistency, the shape of the pathway, etc).

Challenges The validation of FACS relies on measuring the inter-subject alignment of corresponding functional areas. In this thesis, such areas were identified through the analysis of resting state fMRI data. Ideally, functional localisers would be used for identification in addition to the indirect RSN-based methods and many more areas

(instead of only two) would be probed. However, we did not have at hand a data set comprising functional localisers, diffusion- and T1-weighted images. This prevents the generalisation of our conclusions to the whole cortical surface. In other words, validation was performed for two areas only, the MFC and the cingulate, as in these two regions resting state fMRI-based parcellation produces results consistent with the underlying anatomy. On the other hand, the attempted whole-brain analysis did not yield conclusive results.

In addition, as we have already mentioned, the relation between functional segregation and structural connectivity is still being explored, but it is certain that functional specialisation is not defined only by direct structural connections. Therefore, it is not clear in which cortical areas one should neuroscientifically *a priori* expect improvements in alignment. Validation of connectivity-driven registration can thus be seen as application-specific: in case of improving the sensitivity of group fMRI, validation using functional localisers is an appropriate one. However, in cases where the assessment of connections themselves is of major importance (e.g. differences in connectivity between groups) other validation schemes can be devised.

Finally, the analysis of an AD data set did not result in any group differences surviving correction for multiple comparisons. A larger data set and possibly cluster-based correction procedures could increase statistical interpretability of results. Many of the recent publications on the structural analysis use datasets of more than 20 (often about 60) subjects per group (de Jong et al., 2008; Giorgio et al., 2010). In spite of that, the interpretation of differences in deformation fields (or in the associated SCFVs) is not straightforward and requires postprocessing to achieve feasible conclusions (such as e.g. the differences are due to degeneration of a given fibre tract). This is the case with many other structural analysis methods such as Volumetric Based Morphometry (VBM). Although VBM can detect changes in

grey matter, further postprocessing is often required to disentangle whether they are folding-related or thickness-related (or, indeed, even alignment related). Nevertheless, the analysis presented in this thesis primarily aimed to demonstrate how connectivity in pathology can be evaluated in an exploratory manner.

5.1.4 Subcortical Framework

We applied connectivity-driven registration to the alignment of subcortical structures segmented by FIRST. The primary aim was to improve vertex-to-vertex correspondence and align functionally homologous parts of subcortical structures. First, we have shown, in the case of the human thalamus, that when registration is driven by similarities in SCFVs, the correspondence in thalamic clusters improves. Second, the group differences (AD vs. control group) in the deformation fields were analysed for four structures: the thalamus, putamen, pallidum and amygdala. This was carried out in exploratory fashion and in case of the (left) thalamus, changes were observed in the mediodorsal aspect which connects to the prefrontal cortex and the limbic system.

To further demonstrate the applicability of the FASS algorithm, the Similarity Matrix was constructed using thalamic cluster labels. The thalami of the AD and control groups (SCFVs thereof) were clustered using hierarchical Dirichlet process mixture models so that the inter-subject correspondence between cluster labels was automatically established. Using cluster labels only, binary Similarity Matrices were created and fed into the FASS pipeline. In this way we were able to assess not just the deformation fields (as shown in previous analyses), but more importantly, also cluster-specific deformations and the pre- and post-FASS inter-subject consistency. After FASS, the differences in cluster correspondence between the AD and control groups are highly concentrated in the central and medial dorsal aspects of the left thalamus - the areas known to be involved in AD (from histological findings).

Challenges The greatest challenges encountered in the evaluation of FASS, especially in the analysis of an AD data set, were the interpretability of results and poor robustness with respect to the templates. Moreover, just like in the cortical framework, no changes survived the FDR correction for multiple comparisons. There can be many reasons for such findings: small number of subjects, highly differing connectivity patterns across subjects, etc. The latter appears particularly important. Are we perhaps trying to compare and align highly dissimilar objects (SCFVs)? Are even simpler and more robust (across subjects) connectivity profiles more suitable for driving registration? Answering these questions could help explain high dependency of our registration results on the choice of the template/reference. One of the criteria for SCFV construction could, for example, be the robustness of the registration results with respect to different templates. Although we have attempted to shed some light on these questions, more detailed investigation is left for future work. Nevertheless, the aim of this thesis was to develop a connectivity-driven registration pipeline, examine its applicability and limitations and identify areas for further improvements. This work has for the first time proposed how tractography can be used to drive registration.

Another important point not evaluated in this thesis is the sensitivity of FASS pipeline to large pathological changes such as AD-induced cortical atrophy. Brain atrophy affects FIRST segmentation results and therefore potentially biases calculation of SCFVs. We tried to minimise the influences of possible segmentation errors by averaging SCFVs sampled along surface normals. Still, validations of this procedure with regard to cortical atrophy are needed. For example, one could compare (in atrophied brains) averaged SCFVs obtained through the automatic, FIRST segmentation to the ones sampled on the manually defined surface/vertices (e.g. by an expert neuroanatomist), and optimise the averaging scheme to maximise SCFV similarities in the two cases. Moreover, the atrophy affects FreeSurfer cortical

reconstructions (and therefore tractography targets) as well, and proposed tests could potentially also shed some light on those influences.

Finally, we have not examined (because of the difficulties in assembling suitable data sets) how diffusion-weighted data quality influences SCFVs and therefore registration, and consequently such assessments are left for future work. Increased data quality, especially of DW images (e.g. more acquisition directions, better correction for image distortions and head motion, etc), has the potential to improve registration results especially when it comes to analysing group differences. Conversely, better quality of functional data (e.g. higher resolution, better signal to noise ratio) can be expected to improve validation. The robustness of proposed methods to data quality is also of great importance. To examine this in the simplest case, two cortical surfaces could be registered using DW data of differing quality (e.g. different number of acquisition directions) and the obtained deformation fields could be compared to test for robustness with respect to the acquisition parameters. Also, one could, in a similar fashion, perform an analogous test with identical acquisition protocols at two different time points. These tests would certainly tell us more about the robustness of connectivity-driven registration as presented here and possibly highlight areas where further improvements could be made. Ideally, only once the critical robustness issues (with respect to the template, pathology, acquisition parameters, etc) have been resolved, should we fully devote effort to increasing the interpretability of obtained results - possibly by developing additional postprocessing tools. However, the feasibility of such comprehensive analysis first needs to be assessed.

5.2 Future Work

The future work can be divided into three major directions:

- Further validation of surface-based structural connectivity-driven registration and the improvement of the interpretability of results (cortical and subcortical frameworks);
- Development of functional connectivity-based registration protocols and exploration of functional vs. structural connectivity (joint registration framework);
- Development of volumetric connectivity-driven registration protocols.

The first two items have both already been mentioned and explored to a certain extent. For example, we proposed validating FACS (driven by DW images) using fMRI-based cortical parcellation. However, our analysis was constrained to a limited number of functional regions due to lack of appropriate data. Nevertheless, we have shown that the alignment of functional boundaries extracted using functional connectivity gets improved after structural connectivity-driven registration. We have also proposed a method for using RSNs to drive registration. Although not further validated, it opens up possibilities for using resting state fMRI in image registration. Connectivity-driven alignment can therefore play an important role in understanding the relationship and interactions between functional and structural connectivity and contribute towards the identification of the human 'connectome'.

Connectome The Human Connectome Project (HCP) is currently one of the largest studies aimed at understanding brain connectivity and its relation to behaviour and genetics. In a comprehensive MRI study 1200 healthy controls will be scanned using several MRI modalities (including High Angular Resolution Diffusion Imaging, resting state fMRI and task-based fMRI) and 100 out of the 1200 will undergo high-field (7T) MRI scanning as well. Drawing conclusions about such a large cohort will require sophisticated registration methods able to compare (and align) brains

primarily on the basis of their connectivity. In this sense, techniques developed in this thesis present the starting ground for the development of methods capable of integrating both structural and functional connectivity in the registration setting.

An exciting thought is that functional connectivity-driven registration can be additionally informed by the existence of structural connections (or vice versa) and the registration could aim to align not only brain as a whole, but specific functional circuits. Furthermore, such a joint registration concept could be extended by the introduction of descriptors of brain (functional) networks such as RSNs. In other words, structural-functional connectivity registration could also include connection strengths and directionality of a set of cortical networks. Techniques like Dynamic Causal Modelling (DCM) (Friston et al., 2003) or the analysis of RSNs (and other brain networks) (Smith et al., 2010) could further inform registration. For example, the registration cost function could include a term measuring the 'goodness of fit' of DCM network models (which in turn depends on registration) between two subjects.

The third item on the Future Work list was the development of 3D connectivity-driven registration protocols. So far, we presented results of the surface-based registration. However, volumetric registration techniques have could provide an insight not only into the alignment of surfaces, but of full 3D structures. This could yield great benefits to the applications crucially dependent on registration such as group fMRI analysis. Many of the fMRI activations cannot be represented on the surface (e.g. they occur within the thalamus) and the group analysis would not benefit from surface-based alignment. Efforts have been made to combine surface and volumetric registration (CVS) (Postelnicu et al., 2009) and it has been shown that such framework results in better alignment of DW images also (Zollei et al., 2010). Therefore, one could use FACS results directly to inform the surface registration component of CVS in the hope that this will be an even better constraint for the

alignment of fibre tracts. This extension of CVS is straightforward indeed as CVS, just like FACS, takes as its inputs FreeSurfer-aligned grey/white matter surface meshes.

We believe that one could improve the alignment of structural and functional homologues, in techniques like CVS, by informing volumetric registration with long-range connectivity information. We have conducted preliminary investigations to assess how well two state-of-the-art volumetric registration methods, FNIRT and CVS, align features of subcortical structures extracted using connectivity information. These preliminary investigations are described in greater detail in the following section.

5.2.1 Assessment of Potential of Volumetric Connectivity-driven Registration

The aim of the research presented in this section was to investigate whether the inclusion of connectivity features into volumetric registration has the potential to improve brain alignment. To achieve this, we used connectivity features, SCFVs, to parcellate a number of deep grey matter structures, namely the thalamus and putamen, and compare how well corresponding parcels get aligned across subjects with two different volumetric registration methods. In this way we could quantify using an independent imaging modality (i.e. DW images) how well T1-image registration methods perform and find out whether there are any possibilities for further improvements. Should both methods align parcels of subcortical structures perfectly, the inclusion of connectivity information used for parcellation cannot improve alignment. However, if this is not the case, the inclusion of SCFVs could result in better inter-subject volumetric alignment of structurally and functionally homologous regions. Moreover, by comparing two state-of-the-art registration methods in this way, one could infer which one is closer to the 'ideal' registration case.

5.2.1.1 Methods

In order to demonstrate our hypothesis, we evaluated differences between the Combined Volumetric and Surface-based registration (CVS) algorithm and FNIRT (FMRIB's non-linear image registration tool) by examining the post-registration overlap of connectivity-based segmentation labels across a group of control subjects. We began by creating parcellations of the thalamus and putamen for all subjects using Hierarchical Dirichlet Process Mixture Models (HDPM) with spatial constraints (see Chapter 4, section 4.3.3). Both structures were initially segmented by FreeSurfer and then all subjects were spatially aligned to a randomly selected reference for the subsequent analysis.

In order to calculate a parcellation using structural connectivity, every voxel in the region of interest (ROI) has to be described with a feature vector encoding measures of connectivity to other brain areas. A common way to do this would be to parcellate neocortex into its major lobes that would then serve as targets for probabilistic tractography (Behrens et al., 2003a). This process results in low-dimensional feature vectors or connectivity descriptors. In order to boost sensitivity for discriminating between areas of varying connectivity, as well as to reduce the dependency of the feature vector construction on gross volumetric registration, we decided to use SCFVs (as constructed to be utilised in FACS - Chapter 2, section 2.4.3.1) associated not with mesh vertices, but voxels of the ROIs instead. In this case however, we used 160 cortical targets - i.e. 80 in each hemisphere. Such SCFVs were assembled into connectivity matrices as previously described.

Therefore, the final output of probabilistic tractography for every ROI is an $M \times N$ connectivity matrix where M is the number of voxels in the ROI as segmented by FreeSurfer and N is the number of targets, which in our current implementation is set to 160 (80 targets for each hemisphere). Typically, connectivity matrices are very sparse having 90-95% of all values zero. Each of these matrices was preprocessed and

grouped into 9 clusters using a clustering algorithm with spatial constraints based on HDPMs (Jbabdi et al., 2009b). Preprocessing consisted of PCA data dimensionality reduction to size $M \times 20$, typically retaining more than 95% of explained variance. We decided not to use infinite mixture models in our analysis as they regularly produced more than 40 clusters, many of which contained just a few voxels. By varying the number of clusters for a small subset of subjects we found that 9 clusters is a good compromise between capturing anatomically plausible divisions and achieving reasonable inter-subject consistency (for this data set). As we have previously pointed out, HDPM-based clustering does not make use of inter-subject registrations and the correspondence among cluster labels is established solely according to the data matrix fed into the HDPM pipeline.

After the clustering step, all subjects were non-linearly registered to a template (which was part of the hierarchical clustering cohort) using two methods: CVS and FNIRT. We decided not to use a standard space (such as the MNI152) for our analysis, as we test for the alignment of very subtle and subject specific parcellations and are thus in need for a template image with similar detail in it. We believe that if the T1-weighted intensities within our ROIs are not enough to discriminate between their subregions, then the alignment of other cortical areas would constrain the deformations, maximising the alignment of functionally distinct areas. Finally, all clustering results were warped using two registration methods. We then computed the Jaccard overlap measure to test the alignment accuracy among the HDPM clustering of the connectivity descriptors in the ROIs.

CVS Combined Volumetric and Surface-based Registration is a brain image registration method that maximises the alignment of both cortical and subcortical structures. It consists of three image processing steps. First, a surface registration algorithm finds correspondences between the input surfaces from two brain scans and

these correspondences are transformed into a sparse displacement field in Euclidean space. This morph is then diffused into a dense displacement field in the volume using a nonlinear elastic model. Finally, a nonlinear volumetric registration algorithm refines the alignment, bringing subcortical and ventricular structures, which are not near the surfaces, into accurate alignment. This technique has been shown to produce state-of-the-art alignment of cortical folding patterns, architectonics and subcortical structures (Postelnicu et al., 2009).

FNIRT The nonlinear FNIRT registration tool (Andersson et al., 2007b) uses a B-splines representation of the registration warp field and optimises the sum of squared differences in conjunction with an intensity 'bias field' model as its objective function. It is not typically run in a subject-to-subject setting as it was initially optimised for subject-to-template registrations. For our project, together with the original developers of the code, we established a particular set of parameters that achieve high quality subject-to-subject correspondence.

5.2.1.2 Data Description

The experiments described below were run on data provided to us through our collaboration with Dr. Randy Gollub and the Medical Investigation of Neurodevelopmental Disorders (MIND) Institute. Forty-one data sets were selected for this study all of which have been acquired by our collaborators using an identical MRI sequence on a Siemens scanner. The structural data is of $256 \times 256 \times 256$ size with 1mm^3 voxel resolution and TR=12 ms, TE=4.76 ms, TI=4.76 ms, flip angle=20°. The diffusion data scans use single shot echo planar imaging, and a twice-refocused spin echo pulse sequence, optimised to minimise eddy current-induced image distortions (TR/TE=7400/89 ms, b=700 s/mm^2, 256×256 mm FOV, 128×128 matrix, 2 mm (0 mm gap) slice thickness, 10 T2 + 60 DWI, total acquisition time 8 min 38 sec). Sixty-four slices were acquired in the AC-PC plane. The 60

THALAMUS	FNIRT	CVS	PUTAMEN	FNIRT	CVS
Left	**0.19**	0.14	Left	0.20	0.18
Right	0.16	0.13	Right	0.21	0.20

Table 5.1: Mean Jaccard coefficient between the clustering of the template and all the other thalami/putamen (40 comparisons per table entry, see Figures 5.2 and 5.1). Statistically significant difference at $p < 0.01$ level is indicated in bold.

diffusion-encoding gradient directions were determined using the electrostatic shell method, and result in a high signal-to-noise diffusion volume.

5.2.1.3 Results

Comparison between the FNIRT and CVS registration methods was performed for thalamic and putamen clusterings. Table 5.1 summarises the results of compound (calculated across all labels in parallel) Jaccard coefficient measures. Using this measure every subject's clustering was compared to the template clustering. In summary, FNIRT performs consistently better than CVS in the left thalamus.

We also calculated how well both methods perform when each ROI is considered as a single label, without any clustering. In both the left and right hemispheres, FNIRT performed statistically better ($p < 0.001$) over CVS. However, if this effect is regressed out of the findings presented in Table 5.1, the statistical conclusions from Table 5.1 do not change (FNIRT performs better in the left thalamus with $p < 0.05$). Figure 5.1 shows the Jaccard overlap measure between the template and every other subject for FNIRT (left) and CVS (right), for the left thalamus. Jaccard coefficients in case of CVS registration are lower compared to FNIRT.

We also calculated the pairwise compound Jaccard coefficient (compound Jaccard coefficient calculated between each pair of subjects, 40×40 comparisons in total) for the whole dataset excluding the template. These results are shown in Table 5.2 and demonstrated for the left thalamus in Figure 5.2. Figure 5.2 shows the pairwise Jaccard coefficient between every two subjects (for the left thalamus). Overlap values

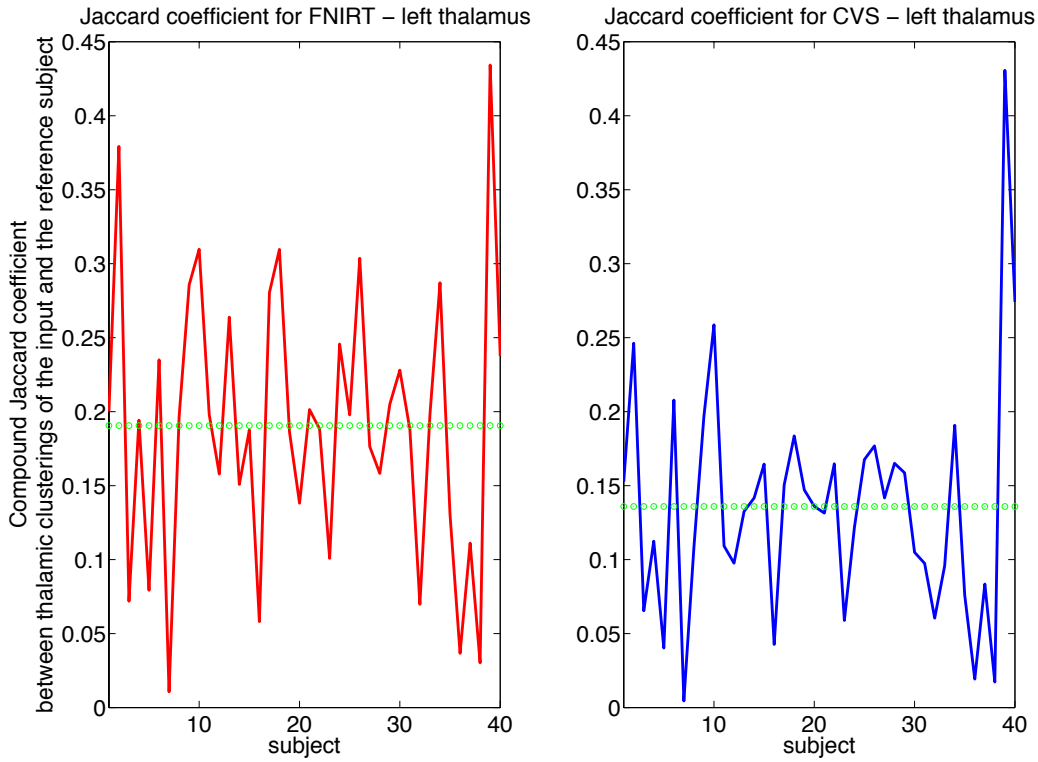

Figure 5.1: Compound Jaccard coefficient for the left thalamus for FNIRT and CVS algorithms. Overlap coefficients were calculated between the *template* and every other subject (40 in total). Mean (compound) Jaccard coefficient value is shown with a green line (see Table 5.1). Subjects are numbered from 1 to 40.

in case of FNIRT (left) are overall higher, i.e. closer to the red end of the colour scale, than that of CVS (right). Similar to previous analyses, we computed the overlap metric with the ROIs treated as a single label. In both the left and right hemispheres, FNIRT performed better ($p < 10^{-100}$). When the effect of the global alignment is regressed out of the findings from Table 5.2, FNIRT still outperformed CVS for both the left and right putamen and for the left thalamus ($p < 0.01$).

THALAMUS	FNIRT	CVS	PUTAMEN	FNIRT	CVS
Left	**0.28**	0.23	Left	**0.26**	0.22
Right	**0.27**	0.20	Right	**0.18**	0.15

Table 5.2: Mean pairwise compound Jaccard coefficient among the clusterings of all ROIs excluding the template (40×40 comparisons per table entry see Figure 5.1). Statistically significant difference at $p < 0.01$ level is given in bold.

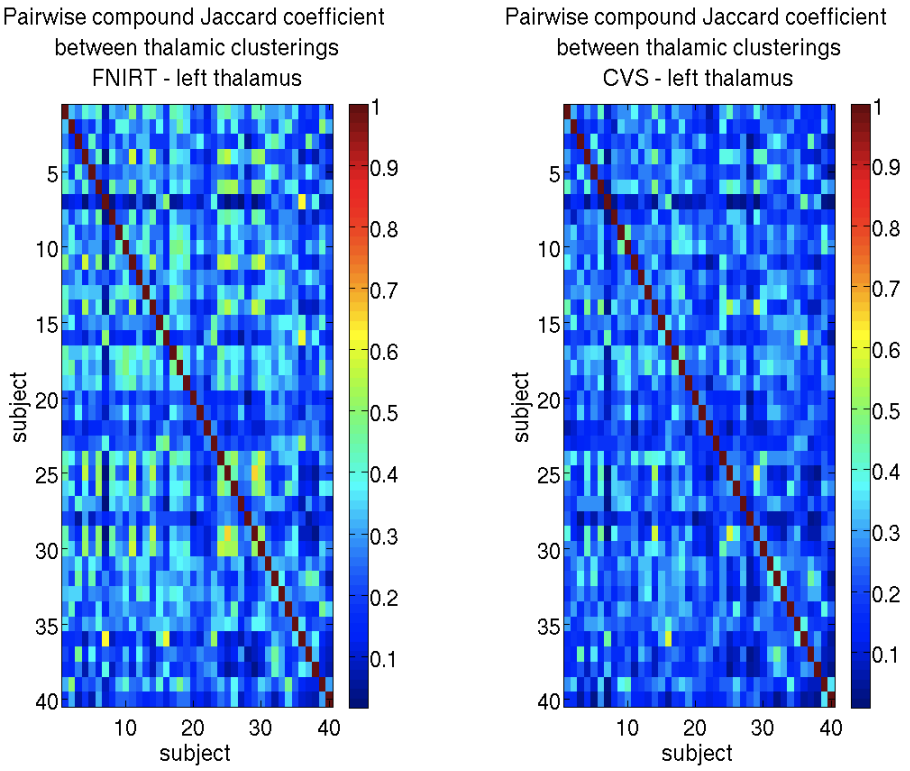

Figure 5.2: Pairwise compound Jaccard coefficient for the clusterings of the left thalamus for FNIRT and CVS algorithms (see Table 5.2). In total, 40×40 comparisons are made and results are visualised as 40×40 matrices/images where the colour reflects the overlap coefficient. Subjects are numbered from 1 to 40.

5.2.1.4 Discussion

The results show that it is possible to make highly local distinctions between competing registration methods possibly giving a new insight into how and where they can be improved. Neither registration method achieves 'perfect' matching of thalamic and putamen clusters. On the contrary, Jaccard similarity measures between two (co-registered) subjects never exceeded 0.5 indicating that the inclusion of SCFVs, i.e. long-range connectivity features, in either of the two evaluated registration methods could yield significant improvements[1]. Moreover, we were able to show that FNIRT outperforms CVS in this evaluation and might therefore be more suitable for the inclusion of SCFVs. However, we should also acknowledge that the proposed

[1]Note that both the thalamus and the putamen are structures whose general shape is almost perfectly matched by both methods.

assessment method can be further enhanced by improving the automatic segmentation of subcortical structures, i.e. ideally one would have expert, manual segmentations. Furthermore, the improvements in robustness of the hierarchical clustering algorithm can be expected to result in more accurate comparisons. Along the same line, questions such as what is the best way to impose inter-subject cluster correspondence and how it affects final findings also need further investigation.

Nevertheless, this research prompted us to think how SCFVs can be included into the FNIRT registration pipeline. For example, SCFVs used in this section can be calculated for a set of subcortical structures (voxels thereof) extracted by e.g. FIRST. This can be done for both the input and the reference subject. If FNIRT can be upgraded to take in multidimensional features instead of scalar values, one could directly include SCFVs associated with subcortical structures into the registration pipeline. In the ideal case, SCFVs should be calculated for every brain voxel, but this is likely to highly increase computational and optimisation requirements. Therefore, by limiting to a set of subcortical structures (e.g. up to 10^4 voxels) we reduce problem dimensionality but hopefully retain enough connectivity information to improve and better constrain registration. FNIRT-based connectivity-driven registration can either be seen as a registration step independent of T1-weighted image registration or as coupled with it. In the latter case, T1-weighted image vs. connectivity weighting would need to be determined to balance out the influences of the two modalities.

In summary, we have proposed a possible upgrade of a volumetric registration method (FNIRT) by including connectivity information as found to be effective in surface-based registration introduced in this thesis. This is certainly one of the primary future research directions and we hope that the research presented here will provide both motivation and a starting point.

Appendix A

Linear Interpolation

Linear interpolation acts on planar triangles (basic elements of tessellated spheres). If vertices of the triangle have coordinates $(x_i, y_i),\ i \in [0, 2]$ in the local Cartesian coordinate system, we seek a linear function of the form $f(x, y) = Ax + By + C$ where A, B and C are to be determined and (x, y) is the point in which we interpolate. If the triangle vertex (x_i, y_i) has an associated value v_i, interpolation coefficients are found by solving the system

$$Ax_i + By_i + C = v_i,\ \ i \in [0, 2]. \tag{A.1}$$

In terms of the Similarity Matrix interpolation, such a scheme naturally bounds the maximum possible interpolated value to the $\max(v_i),\ i \in [0, 2]$.

Appendix B

Regularisation Cost Function

Here we show the similarity between the Thin Plate Spline (TPS) energy functional and the regularisation cost function corresponding to the force given in Equation (2.10). Without any loss of generality let us consider the Thin Plate Spline energy functional (Bookstein, 1989) in 2-dimensional space (e.g. the tangential plane to the S^2). The warp of point x is given by

$$x_i' = x_i + w(x_i), \tag{B.1}$$

where the warp field $w : \mathbb{R}^2 \to \mathbb{R}^2$ has components (w_x, w_y). Thin Plate Spline energy of the warp field w is given by:

$$E_{TPS} = \left(\frac{\partial^2 w}{\partial x^2}\right)^2 + \left(\frac{\partial^2 w}{\partial y^2}\right)^2 + 2\left(\frac{\partial^2 w}{\partial x \partial y}\right)^2 = \tag{B.2}$$

$$= \left(\frac{\partial^2 w_x}{\partial x^2}\right)^2 + \left(\frac{\partial^2 w_x}{\partial y^2}\right)^2 + 2\left(\frac{\partial^2 w_x}{\partial x \partial y}\right)^2 + \tag{B.3}$$

$$+ \left(\frac{\partial^2 w_y}{\partial x^2}\right)^2 + \left(\frac{\partial^2 w_y}{\partial y^2}\right)^2 + 2\left(\frac{\partial^2 w_y}{\partial x \partial y}\right)^2. \tag{B.4}$$

This is the form of the TPS energy that we compare to the regularisation cost function (Equation (2.9)) to show the similarity between the two. Therefore, to obtain the regularisation cost function we start from the Regularisation Force given

in Equation (2.10) and integrate it w.r.t. $w(x_i)$:

$$J_d(x_i') = \left(\frac{1}{n}\left(\sum_{k\in\mathcal{R}_n(i)} w_x(x_k)\right) - \frac{w_x(x_i)}{2}\right) w_x(x_i) + \tag{B.5}$$

$$+ \left(\frac{1}{n}\left(\sum_{k\in\mathcal{R}_n(i)} w_y(x_k)\right) - \frac{w_y(x_i)}{2}\right) w_y(x_i) + const, \tag{B.6}$$

where $\mathcal{R}_n(i)$ is a set of n (immediate) neighbouring vertices of vertex $i \in V_{M_I}$ in V_{M_I} as before. The free constant can be chosen as

$$const = -\frac{1}{2}\left(\left(\frac{1}{n}\sum_{k\in\mathcal{R}_n(i)} w_x(x_k)\right)^2 + \left(\frac{1}{n}\sum_{k\in\mathcal{R}_n(i)} w_y(x_k)\right)^2\right), \tag{B.7}$$

since it does not depend on the centre point $w(x_i)$. For this constant and using that $\frac{1}{n}\left(\sum_{k\in\mathcal{R}_n(i)} w(x_k)\right) - w(x_i)$ is a discrete approximation to the Laplacian $\frac{\partial^2 w}{\partial x^2} + \frac{\partial^2 w}{\partial y^2}$ (valid under the assumption of local mesh regularity), we have

$$J_d(x_i') \approx -\frac{1}{2}\left(\left(\frac{\partial^2 w_x}{\partial x^2} + \frac{\partial^2 w_x}{\partial y^2}\right)^2 + \left(\frac{\partial^2 w_y}{\partial x^2} + \frac{\partial^2 w_y}{\partial y^2}\right)^2\right). \tag{B.8}$$

This equation is similar (but not identical) to Equation (B.2) except in the $\frac{\partial^2 w}{\partial x \partial y}$ and $\frac{\partial^2 w}{\partial x^2}\frac{\partial^2 w}{\partial y^2}$ terms. However, the latter can be neglected in discrete approximation (as it does not depend on the centre point $w(x_i)$ and hence the derivative is zero). The former, on the other hand, introduces rotational symmetry of our energy functional.

Therefore, we have shown the similarity between the TPS energy and the regularisation cost function used in this work. However, the main caveat of the current formulation of the regularisation force is that it associates a zero cost with a saddle point. Nevertheless, the currently used regulariser has the advantages of straightforward computation and a direct geometrical interpretation. Moreover, in Chapter 4, section 4.2.3, the effects of different regularisation paradigms in two different surface registration settings were compared and found to perform similarly on the artificial test data.

Appendix C

Locally Affine Transforms

Figure C.1 shows two corresponding groups of vertices e.g. the group on the left ($\mathcal{A}$) is transformed to the group on the right ($\mathcal{A}'$) due to connectivity-driven surface registration. Every point on the left $P_i \in \mathcal{A}$ has a corresponding point on the right $Q_i \in \mathcal{A}'$ for $i \in [1, n]$. X (shown in blue) is the point spatially belonging to $\mathcal{A}$ but whose mapping to $\mathcal{A}'$ was not explicitly determined by the registration algorithm. We would like to find the projection of that point to $\mathcal{A}'$. n is the size of the vertex group in $\mathcal{A}$ and practically defines n-nearest neighbours of X in the mesh to which X belongs. For our applications n is typically set to 6. In practice, due to high mesh density and low n, this problem is almost planar.

Transformation $\Omega : \mathcal{A} \rightarrow \mathcal{A}'$ can be found by determining the affine matrix between the two groups of vertices. Let matrix A encode coordinates of vertices in $\mathcal{A}$:

$$A_{3n\times 12} = \begin{bmatrix} P_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & P_1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & P_1 \\ P_2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & P_2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & P_2 \\ & & & & \vdots & & & & \\ P_n & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & P_n & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & P_n \end{bmatrix}, \tag{C.1}$$

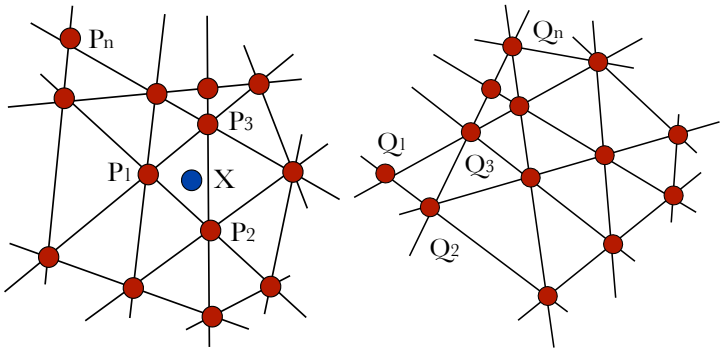

Figure C.1: Two corresponding groups of vertices (shown in red) e.g. the group on the left ($\mathcal{A}$) is mapped to the group on the right ($\mathcal{A}'$) due to connectivity-driven surface registration. Every point on the left P_i has a corresponding point on the right Q_i for $i \in [1, n]$. We seek to find a locally affine mapping between these two groups of vertices/points and apply that mapping to point X (shown in blue). **(left)** The group of n closest neighbours of point X which we seek to project. **(right)** The group of vertices into which the ones from the left are mapped. We aim to find the mapping of point X (left) to this group.

where $P_i = [p_x^i \ p_y^i \ p_z^i \ 1]$ is a row vector with four values: p_x^i, p_y^i and p_z^i represent x, y and z coordinates respectively of a vertex associated with P_i. Let $\tilde{T}$ encode coordinates of points in $\mathcal{A}'$:

$$\tilde{T}_{3n\times 1} = \begin{bmatrix} Q_1 & Q_2 & \cdots & Q_n \end{bmatrix}^T, \tag{C.2}$$

where $Q_i = [q_x^i \ q_y^i \ q_z^i]^T$ is a column vector with three values: q_x^i, q_y^i and q_z^i denote x, y and z coordinates of a vertex associated with Q_i. Then, for large enough n (to ensure a well-conditioned solution), the affine matrix T can be calculated as:

$$A_{PI} = (A^T A)^{-1} A^T \tag{C.3}$$

$$\tilde{W}_{12\times 1} = A_{PI(12\times 3n)} \tilde{T}_{3n\times 1} \tag{C.4}$$

$$T_{4\times 4} = \begin{bmatrix} \tilde{W}(1,1) & \tilde{W}(2,1) & \tilde{W}(3,1) & \tilde{W}(4,1) \\ \tilde{W}(5,1) & \tilde{W}(6,1) & \tilde{W}(7,1) & \tilde{W}(8,1) \\ \tilde{W}(9,1) & \tilde{W}(10,1) & \tilde{W}(11,1) & \tilde{W}(12,1) \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{C.5}$$

Therefore, point $X(x_x, x_y, x_z)$ can be transformed to $X'(x'_x, x'_y, x'_z)$ using the affine matrix T:

$$[x'_x \ x'_y \ x'_z \ 1]^T = T[x_x \ x_y \ x_z \ 1]^T. \tag{C.6}$$

Appendix D

Imposing Labelling Consistency

The input value, C_{MAT}, is the matrix whose each row is the output of one subject's SMA/pre-SMA clustering. As all subjects are co-registered and represented as $ic5$, each row will have the same size, i.e. 10242 and there would be one-to-one correspondence between the rows. SMA and pre-SMA are labelled with non-zero labels whereas all other vertices (in the MFC mask) have zero labels. Therefore $C_{MAT}(i, j)$ is the label associated with vertex j in subject i. N is the number of iterations, typically set to 10 and n is the number of clusters labelled from 1 to n.

Algorithm D.0.1: RELABELCLUSTERS(C_{MAT}, N, n)

comment: $C_{MAT}(1,:)$ is taken as the initial *prototype* clustering IDX_{PROT}

comment: v-number of vertices in $ic5$, m-number of subjects

$IDX_{PROT} \leftarrow C_{MAT}(1,:)$ **comment:** IDX_{PROT} - 'prototype clustering'

$correspondence \leftarrow$ FINDCORRESPONDENCE(IDX_{PROT}, C_{MAT}, n)
$C_{RL} \leftarrow$ RELABEL$(IDX_{PROT}, C_{MAT}, correspondence, n)$
for $i \leftarrow 1$ **to** N
do $\begin{cases} IDX_{PROT} \leftarrow mode(C_{RL}) \text{ } \textbf{comment:} \text{ } mode \text{ across columns} \\ correspondence \leftarrow \text{FINDCORRESPONDENCE}(IDX_{PROT}, C_{MAT}, n) \\ C_{RL} \leftarrow \text{RELABEL}(IDX_{PROT}, C_{MAT}, correspondence, n) \end{cases}$
return (C_{RL})

procedure FINDCORRESPONDENCE(IDX_{PROT}, C_{MAT}, n)
for $i \leftarrow 1$ **to** n
do $\begin{cases} P \leftarrow IDX_{PROT}(:) = i \\ \textbf{for } j \leftarrow 1 \textbf{ to } m \\ \quad \textbf{do} \begin{cases} Q \leftarrow C_{MAT}(j, p) \\ M \leftarrow mode(Q) \\ correspondence(j, i) \leftarrow M \end{cases} \end{cases}$
return $(correspondence)$

procedure RELABEL$(IDX_{PROT}, C_{MAT}, correspondence, n)$
$C_{RL} \leftarrow C_{MAT}$
for $j \leftarrow 1$ **to** m
do $\begin{cases} \textbf{for } i \leftarrow 1 \textbf{ to } v \\ \quad \textbf{do} \begin{cases} \textbf{for } k \leftarrow 1 \textbf{ to } n \\ \quad \textbf{do} \begin{cases} \textbf{if} \begin{cases} (IDX_{PROT}(i) = k) \quad \textbf{and} \\ (C_{MAT}(j, i) = correspondence(j, k)) \end{cases} \\ \quad \textbf{then } C_{RL}(j, i) = k \end{cases} \end{cases} \end{cases}$
return (C_{RL})

Bibliography

K. Amunts, A. Schleicher, U. Burgel, H. Mohlberg, H. B. Uylings, and K. Zilles. Broca's region revisited: cytoarchitecture and intersubject variability. *J Comp Neurol*, 412(2):319–341, 1999. ISSN 0021-9967 (Print); 0021-9967 (Linking).

K. Amunts, A. Malikovic, H. Mohlberg, T. Schormann, and K. Zilles. Brodmann's areas 17 and 18 brought into stereotaxic space-where and how variable? *Neuroimage*, 11(1):66–84, 2000. ISSN 1053-8119 (Print); 1053-8119 (Linking). doi: 10.1006/nimg.1999.0516.

J. L. R. Andersson, M. Jenkinson, and S. M. Smith. TR07JA1: Non-linear optimisation. Technical report, University of Oxford, FMRIB Centre, 2007a.

J. L. R. Andersson, M. Jenkinson, and S. M. Smith. TR07JA2: Non-linear registration, aka Spatial normalisation. Technical report, University of Oxford, FMRIB Centre, 2007b.

A. Anwander, M. Tittgemeyer, D. Y. von Cramon, A. D. Friederici, and T. R. Knosche. Connectivity-Based Parcellation of Broca's Area. *Cereb Cortex*, 17(4): 816–825, 2007. ISSN 1047-3211 (Print); 1047-3211 (Linking). doi: 10.1093/cercor/bhk034.

J. Ashburner and K. Friston. Spatial transformation of images. In R. Frackowiak, K. Friston, C. Frith, R. Dolan, and J. Mazziotta, editors, *Human Brain Function*, pages 43–58. Academic Press USA, 1997.

Y. Assaf and P. J. Basser. Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain. *Neuroimage*, 27(1):48–58, 2005.

M. A. Audette, F. P. Ferrie, and T. M. Peters. An algorithmic overview of surface registration techniques for medical imaging. *Med Image Anal*, 4(3):201–217, 2000. ISSN 1361-8415 (Print); 1361-8415 (Linking).

P. J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In vivo fiber tractography using dt-mri data. *Magn Reson Med*, 44(4):625–32, 2000.

M. F. Bear, B. W. Connors, and M. A. Paradiso. *Neuroscience: exploring the brain*. Lippincott Williams and Wilkins, Philadelphia, PA, 3rd ed edition, 2007. ISBN 0781776074 (alk. paper).

C. Beaulieu. The basis of anisotropic water diffusion in the nervous system - a technical review. *NMR Biomed*, 15(7-8):435–55, 2002. doi: 10.1002/nbm.782.

C. Beckmann and S. Smith. Probabilistic independent component analysis for functional magnetic resonance imaging. *Medical Imaging, IEEE Transactions on*, 23(2):137 –152, feb. 2004. ISSN 0278-0062. doi: 10.1109/TMI.2003.822821.

C. F. Beckmann, M. DeLuca, J. T. Devlin, and S. M. Smith. Investigations into resting-state connectivity using independent component analysis. *Philos Trans R Soc Lond B Biol Sci*, 360(1457):1001–1013, 2005. ISSN 0962-8436 (Print); 0962-8436 (Linking). doi: 10.1098/rstb.2005.1634.

M. Beckmann, H. Johansen-Berg, and M. F. S. Rushworth. Connectivity-based parcellation of human cingulate cortex and its relation to functional specialization. *J Neurosci*, 29(4):1175–1190, 2009. ISSN 1529-2401 (Electronic); 0270-6474 (Linking). doi: 10.1523/JNEUROSCI.3328-08.2009.

T. E. Behrens, H. Johansen-Berg, M. W. Woolrich, S. M. Smith, C. A. 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. *Nat Neurosci*, 6(7):750–7, 2003a.

T. E. 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. *Magn Reson Med*, 50(5):1077–88, 2003b.

T. E. Behrens, H. J. Berg, S. Jbabdi, M. F. Rushworth, and M. W. Woolrich. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *Neuroimage*, 34(1):144–55, 2007.

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. *Nat Neurosci*, 6(7):750–757, 2003c. ISSN 1097-6256 (Print); 1097-6256 (Linking). doi: 10.1038/nn1075.

Y. Benjamini and Y. Hochberg. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1):pp. 289–300, 1995. ISSN 00359246.

P. J. Besl and N. D. McKay. A Method for Registration of 3-D Shapes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14:239–256, February 1992. ISSN 0162-8828. doi: 10.1109/34.121791.

F. L. Bookstein. Principal Warps: Thin-Plate Splines and the Decomposition of Deformations. *IEEE Trans. Pattern Anal. Mach. Intell.*, 11(6):567–585, 1989.

H. Braak and E. Braak. Alzheimer's disease affects limbic nuclei of the thalamus. *Acta Neuropathol*, 81(3):261–268, 1991. ISSN 0001-6322 (Print).

K. Brodmann. *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Barth, Leipzig, 1909 edition, 1909. 1868-1918ill. ; 23 cm Mit 150 Abbildungen im Text.

J. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 8(6):679–698, 1986. ISSN 0162-8828. doi: http://dx.doi.org/10.1109/TPAMI.1986.4767851.

M. Catani, R. J. Howard, S. Pajevic, and D. K. Jones. Virtual in vivo interactive dissection of white matter fasciculi in the human brain. *Neuroimage*, 17(1):77–94, 2002.

P. Cathier and J.-F. Mangin. Registration of Cortical Connectivity Matrices. In *CVPRW '06: Proceedings of the 2006 Conference on Computer Vision and Pattern Recognition Workshop*, page 66, Washington, DC, USA, 2006. IEEE Computer Society.

Y.-P. Chao, J.-H. Chen, K.-H. Cho, C.-H. Yeh, K.-H. Chou, and C.-P. Lin. A multiple streamline approach to high angular resolution diffusion tractography. *Med Eng Phys*, 30(8):989–96, 2008. doi: 10.1016/j.medengphy.2008.01.010.

G. Christensen. Consistent linear-elastic transformations for image matching. *Information Processing In Medical Imaging, Proceedings*, 1613:224–237, 1999.

G. E. Christensen and H. J. Johnson. Invertibility and transitivity analysis for nonrigid image registration. *Journal of Electronic Imaging*, 12(1):106–117, 2003. doi: 10.1117/1.1526494.

A. L. Cohen, D. A. Fair, N. U. F. Dosenbach, F. M. Miezin, D. Dierker, D. C. Van Essen, B. L. Schlaggar, and S. E. Petersen. Defining functional areas in individual human brains using resting functional connectivity MRI. *Neuroimage*, 41(1):45–57, 2008. ISSN 1053-8119 (Print); 1053-8119 (Linking). doi: 10.1016/j.neuroimage.2008.01.066.

Y. Cointepas, D. Geffroy, N. Souedet, I. Denghien, and D. Rivière. The BrainVISA project: a shared software development infrastructure for biomedical imaging research. In *Proc. 16th HBM*, 2010.

B. Conroy, B. Singer, J. Haxby, and P. Ramadge. fMRI-based inter-subject cortical alignment using functional connectivity. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 378–386. 2009.

P. A. Cook, Y. Bai, S. Nedjati-Gilani, K. K. Seunarine, M. G. Hall, G. J. Parker, and D. C. Alexander. Camino: Open-source diffusion-mri reconstruction and processing. In *In Proceedings of the International Society for Magnetic Resonance in Medicine (ISMRM) Meeting*, page 2759, 2006.

A. M. Dale and M. I. Sereno. Improved Localizadon of Cortical Activity by Combining EEG and MEG with MRI Cortical Surface Reconstruction: A Linear Approach. *Journal of Cognitive Neuroscience*, 5(2):162–176, 1993. doi: 10.1162/jocn.1993.5.2.162.

A. M. Dale, B. Fischl, and M. I. Sereno. Cortical surface-based analysis. I. Segmentation and surface reconstruction. *Neuroimage*, 9(2):179–194, 1999. ISSN 1053-8119 (Print). doi: 10.1006/nimg.1998.0395.

J. S. Damoiseaux, S. A. R. B. Rombouts, F. Barkhof, P. Scheltens, C. J. Stam, S. M. Smith, and C. F. Beckmann. Consistent resting-state networks across healthy subjects. *Proc Natl Acad Sci U S A*, 103(37):13848–13853, 2006. ISSN 0027-8424 (Print); 0027-8424 (Linking). doi: 10.1073/pnas.0601417103.

L. W. de Jong, K. van der Hiele, I. M. Veer, J. J. Houwing, R. G. J. Westendorp, E. L. E. M. Bollen, P. W. de Bruin, H. A. M. Middelkoop, M. A. van Buchem, and J. van der Grond. Strongly reduced volumes of putamen and thalamus in Alzheimer's disease: an MRI study. *Brain*, 131(Pt 12):3277–3285, Dec 2008. ISSN 1460-2156 (Electronic); 0006-8950 (Linking). doi: 10.1093/brain/awn278.

M. De Luca, C. F. Beckmann, N. De Stefano, P. M. Matthews, and S. M. Smith. fMRI resting state networks define distinct modes of long-distance interactions in the human brain. *Neuroimage*, 29(4):1359–1367, 2006. ISSN 1053-8119 (Print); 1053-8119 (Linking). doi: 10.1016/j.neuroimage.2005.08.035.

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

H. A. Drury, D. C. Van Essen, C. H. Anderson, C. W. Lee, T. A. Coogan, and J. W. Lewis. Computerized mappings of the cerebral cortex: a multiresolution flattening method and a surface-based coordinate system. *J Cogn Neurosci*, 8(1):1–28, 1996. ISSN 0898-929X (Print); 0898-929X (Linking).

D. J. Felleman and D. C. Van Essen. Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex*, 1(1):1–47, 1991. ISSN 1047-3211 (Print); 1047-3211 (Linking).

N. Filippini, B. J. MacIntosh, M. G. Hough, G. M. Goodwin, G. B. Frisoni, S. M. Smith, P. M. Matthews, C. F. Beckmann, and C. E. Mackay. Distinct patterns of brain activity in young carriers of the APOE-epsilon4 allele. *Proc Natl Acad Sci U S A*, 106(17):7209–7214, 2009. ISSN 1091-6490 (Electronic); 0027-8424 (Linking). doi: 10.1073/pnas.0811879106.

B. Fischl, M. I. Sereno, and A. M. Dale. Cortical surface-based analysis. II: Inflation, flattening, and a surface-based coordinate system. *Neuroimage*, 9(2):195–207, 1999a. ISSN 1053-8119 (Print). doi: 10.1006/nimg.1998.0396.

B. Fischl, M. I. Sereno, R. B. Tootell, and A. M. Dale. High-resolution intersubject averaging and a coordinate system for the cortical surface. *Hum Brain Mapp*, 8(4): 272–84, 1999b.

B. Fischl, M. I. Sereno, R. B. Tootell, and A. M. Dale. High-resolution intersubject averaging and a coordinate system for the cortical surface. *Hum Brain Mapp*, 8(4): 272–284, 1999c. ISSN 1065-9471 (Print); 1065-9471 (Linking).

B. Fischl, D. H. Salat, E. Busa, M. Albert, M. Dieterich, C. Haselgrove, A. van der Kouwe, R. Killiany, D. Kennedy, S. Klaveness, A. Montillo, N. Makris, B. Rosen, and A. M. Dale. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33(3):341–355, 2002. ISSN 0896-6273 (Print); 0896-6273 (Linking).

B. Fischl, A. van der Kouwe, C. Destrieux, E. Halgren, F. Segonne, D. H. Salat, E. Busa, L. J. Seidman, J. Goldstein, D. Kennedy, V. Caviness, N. Makris,

B. Rosen, and A. M. Dale. Automatically parcellating the human cerebral cortex. *Cereb Cortex*, 14(1):11–22, 2004. ISSN 1047-3211 (Print); 1047-3211 (Linking).

B. Fischl, N. Rajendran, E. Busa, J. Augustinack, O. Hinds, B. T. Yeo, H. Mohlberg, K. Amunts, and K. Zilles. Cortical folding patterns and predicting cytoarchitecture. *Cereb Cortex*, 18(8):1973–80, 2008.

M. D. Fox, A. Z. Snyder, J. L. Vincent, M. Corbetta, D. C. Van Essen, and M. E. Raichle. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc Natl Acad Sci U S A*, 102(27):9673–9678, 2005. ISSN 0027-8424 (Print); 0027-8424 (Linking). doi: 10.1073/pnas.0504136102.

K. Friston, L. Harrison, and W. Penny. Dynamic Causal Modelling. *NeuroImage*, 19 (4):1273–1302, 2003.

S. Geyer. The microstructural border between the motor and the cognitive domain in the human cerebral cortex. *Adv Anat Embryol Cell Biol*, 174:I–VIII, 1–89, 2004. ISSN 0301-5556 (Print); 0301-5556 (Linking).

A. Giorgio, L. Santelli, V. Tomassini, R. Bosnell, S. Smith, N. De Stefano, and H. Johansen-Berg. Age-related changes in grey and white matter structure throughout adulthood. *Neuroimage*, 51(3):943–51, Jul 2010. doi: 10.1016/j.neuroimage.2010.03.004.

S. N. Haber. The primate basal ganglia: parallel and integrative networks. *J Chem Neuroanat*, 26(4):317–330, 2003. ISSN 0891-0618 (Print); 0891-0618 (Linking).

P. Hagmann, L. Cammoun, X. Gigandet, R. Meuli, C. J. Honey, V. J. Wedeen, and O. Sporns. Mapping the structural core of human cerebral cortex. *PLoS Biol*, 6(7):e159, 2008. ISSN 1545-7885 (Electronic); 1544-9173 (Linking). doi: 10.1371/journal.pbio.0060159.

D. L. Hill, P. G. Batchelor, M. Holden, and D. J. Hawkes. Medical image registration. *Phys Med Biol*, 46(3):R1–45, 2001. ISSN 0031-9155 (Print); 0031-9155 (Linking).

C. A. R. Hoare. Quicksort. *The Computer Journal*, 5(1):10–16, 1962. doi: 10.1093/comjnl/5.1.10.

S. Hofer, A. Karaus, and J. Frahm. Reconstruction and dissection of the entire human visual pathway using diffusion tensor MRI. *Front Neuroanat*, 4:15, 2010. doi: 10.3389/fnana.2010.00015.

C. J. Honey, O. Sporns, L. Cammoun, X. Gigandet, J. P. Thiran, R. Meuli, and P. Hagmann. Predicting human resting-state functional connectivity from structural connectivity. *Proc Natl Acad Sci U S A*, 106(6):2035–2040, 2009. ISSN 1091-6490 (Electronic); 0027-8424 (Linking). doi: 10.1073/pnas.0811168106.

J. Huang, R. P. Friedland, and A. P. Auchus. Diffusion tensor imaging of normal-appearing white matter in mild cognitive impairment and early Alzheimer disease: preliminary evidence of axonal degeneration in the temporal lobe. *AJNR Am J Neuroradiol*, 28(10):1943–1948, 2007. ISSN 0195-6108 (Print); 0195-6108 (Linking). doi: 10.3174/ajnr.A0700.

P. Jaccard. The distribution of the flora in the Alpine zone. *New Phytologist*, 11: 37–50, 1912.

S. Jbabdi, M. W. Woolrich, J. L. Andersson, and T. E. Behrens. A Bayesian framework for global tractography. *Neuroimage*, 37(1):116–29, 2007.

S. Jbabdi, M. W. Woolrich, and T. E. J. Behrens. Multiple-subjects connectivity-based parcellation using hierarchical Dirichlet process mixture models. *Neuroimage*, 44(2):373–384, 2009a. ISSN 1095-9572 (Electronic); 1053-8119 (Linking). doi: 10.1016/j.neuroimage.2008.08.044.

S. Jbabdi, M. W. Woolrich, and T. E. J. Behrens. Multiple-subjects connectivity-based parcellation using hierarchical Dirichlet process mixture models. *Neuroimage*, 44(2):373–384, 2009b. ISSN 1095-9572 (Electronic); 1053-8119 (Linking). doi: 10.1016/j.neuroimage.2008.08.044.

M. Jenkinson, P. Bannister, M. Brady, and S. Smith. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage*, 17(2):825–841, 2002. ISSN 1053-8119 (Print); 1053-8119 (Linking).

P. Jezzard, P. M. Matthews, and S. M. Smith. *Functional MRI: an introduction to methods*. Oxford University Press, Oxford, 2001. ISBN 0192630717.

H. J. Jo, Z. S. Saad, W. K. Simmons, L. A. Milbury, and R. W. Cox. Mapping sources of correlation in resting state FMRI, with artifact detection and removal. *Neuroimage*, 52(2):571–582, 2010. ISSN 1095-9572 (Electronic); 1053-8119 (Linking). doi: 10.1016/j.neuroimage.2010.04.246.

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. *Proc Natl Acad Sci U S A*, 101(36):13335–13340, 2004. ISSN 0027-8424 (Print); 0027-8424 (Linking). doi: 10.1073/pnas.0403743101.

H. Johansen-Berg, T. E. 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. *Cereb Cortex*, 15(1):31–9, 2005.

D. K. Jones and C. Pierpaoli. Confidence mapping in diffusion tensor magnetic resonance imaging tractography using a bootstrap approach. *Magn Reson Med*, 53 (5):1143–9, 2005. doi: 10.1002/mrm.20466.

K. Kantarci, C. R. J. Jack, Y. C. Xu, N. G. Campeau, P. C. O'Brien, G. E. Smith, R. J. Ivnik, B. F. Boeve, E. Kokmen, E. G. Tangalos, and R. C. Petersen. Mild cognitive impairment and Alzheimer disease: regional diffusivity of water. *Radiology*, 219(1):101–107, 2001. ISSN 0033-8419 (Print); 0033-8419 (Linking).

G. B. Karas, E. J. Burton, S. A. R. B. Rombouts, R. A. van Schijndel, J. T. O'Brien, P. h. Scheltens, I. G. McKeith, D. Williams, C. Ballard, and F. Barkhof. A comprehensive study of gray matter loss in patients with Alzheimer's disease using optimized voxel-based morphometry. *Neuroimage*, 18(4):895–907, 2003a. ISSN 1053-8119 (Print); 1053-8119 (Linking).

G. B. Karas, E. J. Burton, S. A. R. B. Rombouts, R. A. van Schijndel, J. T. O'Brien, P. h. Scheltens, I. G. McKeith, D. Williams, C. Ballard, and F. Barkhof. A comprehensive study of gray matter loss in patients with Alzheimer's disease using optimized voxel-based morphometry. *Neuroimage*, 18(4):895–907, 2003b. ISSN 1053-8119 (Print).

L. Kaufman and P. J. Rousseeuw. *Finding groups in data: an introduction to cluster analysis*. Wiley, New York, 1990. ISBN 0471878766.

J.-H. Kim, J.-M. Lee, H. J. Jo, S. H. Kim, J. H. Lee, S. T. Kim, S. W. Seo, R. W. Cox, D. L. Na, S. I. Kim, and Z. S. Saad. Defining functional SMA and pre-SMA subregions in human MFC using resting state fMRI: functional connectivity-based parcellation method. *Neuroimage*, 49(3):2375–2386, 2010. ISSN 1095-9572 (Electronic); 1053-8119 (Linking). doi: 10.1016/j.neuroimage.2009.10.016.

J. C. Klein, T. E. Behrens, M. D. Robson, C. E. Mackay, D. J. Higham, and H. Johansen-Berg. Connectivity-based parcellation of human cortex using diffusion MRI: Establishing reproducibility, validity and observer independence in BA 44/45 and SMA/pre-SMA. *Neuroimage*, 34(1):204–11, 2007.

M. Lazar and A. L. Alexander. Bootstrap white matter tractography (boot-trac). *Neuroimage*, 24(2):524–32, 2005. doi: 10.1016/j.neuroimage.2004.08.050.

Z.-P. Liang and P. C. Lauterbur. *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*. Wiley-IEEE Press, October 1999. ISBN 0780347234.

P.-M. Lledo, M. Alonso, and M. S. Grubb. Adult neurogenesis and functional plasticity in neuronal circuits. *Nat Rev Neurosci*, 7(3):179–193, 2006. ISSN 1471-003X (Print); 1471-003X (Linking). doi: 10.1038/nrn1867.

M. J. Lowe, M. Dzemidzic, J. T. Lurito, V. P. Mathews, and M. D. Phillips. Correlations in low-frequency BOLD fluctuations reflect cortico-cortical connections. *Neuroimage*, 12(5):582–587, 2000. ISSN 1053-8119 (Print); 1053-8119 (Linking). doi: 10.1006/nimg.2000.0654.

Y. Lu, P.-Y. Liu, P. Xiao, and H.-W. Deng. Hotelling's T2 multivariate profiling for detecting differential expression in microarrays. *Bioinformatics*, 21(14):3105–3113, 2005. ISSN 1367-4803 (Print); 1367-4803 (Linking). doi: 10.1093/bioinformatics/bti496.

O. Lyttelton, M. Boucher, S. Robbins, and A. Evans. An unbiased iterative group registration template for cortical surface analysis. *Neuroimage*, 34(4):1535–1544, 2007. ISSN 1053-8119 (Print); 1053-8119 (Linking). doi: 10.1016/j.neuroimage.2006.10.041.

A. Malikovic, K. Amunts, A. Schleicher, H. Mohlberg, S. B. Eickhoff, M. Wilms, N. Palomero-Gallagher, E. Armstrong, and K. Zilles. Cytoarchitectonic analysis of the human extrastriate cortex in the region of V5/MT+: a probabilistic, stereotaxic map of area hoc5. *Cereb Cortex*, 17(3):562–574, 2007. ISSN 1047-3211 (Print); 1047-3211 (Linking). doi: 10.1093/cercor/bhj181.

J.-F. Mangin, D. Riviere, A. Cachia, E. Duchesnay, Y. Cointepas, D. Papadopoulos-Orfanos, P. Scifo, T. Ochiai, F. Brunelle, and J. Regis. A framework to study the cortical folding patterns. *Neuroimage*, 23 Suppl 1:S129–38, 2004. ISSN 1053-8119 (Print); 1053-8119 (Linking). doi: 10.1016/j.neuroimage.2004.07.019.

D. S. Margulies, A. M. C. Kelly, L. Q. Uddin, B. B. Biswal, F. X. Castellanos, and M. P. Milham. Mapping the functional connectivity of anterior cingulate cortex. *Neuroimage*, 37(2):579–588, 2007. ISSN 1053-8119 (Print); 1053-8119 (Linking). doi: 10.1016/j.neuroimage.2007.05.019.

A. Mezer, Y. Yovel, O. Pasternak, T. Gorfine, and Y. Assaf. Cluster analysis of resting-state fMRI time series. *Neuroimage*, 45(4):1117–1125, 2009. ISSN 1095-9572 (Electronic); 1053-8119 (Linking). doi: 10.1016/j.neuroimage.2008.12.015.

A. Morel, M. Magnin, and D. Jeanmonod. Multiarchitectonic and stereotactic atlas of the human thalamus. *J Comp Neurol*, 387(4):588–630, 1997. ISSN 0021-9967 (Print); 0021-9967 (Linking).

S. Mori. *Introduction to diffusion tensor imaging*. Elsevier, Amsterdam, 2006. ISBN 9780444528285 (alk. paper).

D. Nain, S. Haker, A. F. Bobick, and A. Tannenbaum. Multiscale 3-D Shape Representation and Segmentation Using Spherical Wavelets. *IEEE Trans. Med. Imaging*, 26(4):598–618, 2007.

S. Ogawa, R. S. Menon, D. W. Tank, S. G. Kim, H. Merkle, J. M. Ellermann, and K. Ugurbil. Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging. A comparison of signal characteristics with a biophysical model. *Biophys J*, 64(3):803–12, Mar 1993. doi: 10.1016/S0006-3495(93)81441-3.

I. Oguz, M. Niethammer, J. Cates, R. Whitaker, T. Fletcher, C. Vachet, and M. Styner. Cortical correspondence with probabilistic fiber connectivity. *Inf Process Med Imaging*, 21:651–663, 2009. ISSN 1011-2499 (Print); 1011-2499 (Linking).

M. Ono, S. Kubik, and C. D. Abernathey. *Atlas of the cerebral sulci*. G. Thieme Verlag, Stuttgart, 1990. ISBN 0865773629 (Thieme Medical Publishers).

S. Pajevic and P. J. Basser. Parametric and non-parametric statistical analysis of dt-mri data. *J Magn Reson*, 161(1):1–14, 2003.

G. J. M. Parker and D. C. Alexander. Probabilistic monte carlo based mapping of cerebral connections utilising whole-brain crossing fibre information. *Inf Process Med Imaging*, 18:684–95, 2003.

G. J. M. Parker and D. C. Alexander. Probabilistic anatomical connectivity derived from the microscopic persistent angular structure of cerebral tissue. *Philos Trans R Soc Lond B Biol Sci*, 360(1457):893–902, 2005. doi: 10.1098/rstb.2005.1639.

G. J. M. Parker, H. A. Haroon, and C. A. M. Wheeler-Kingshott. A framework for a streamline-based probabilistic index of connectivity (PICo) using a structural interpretation of MRI diffusion measurements. *J Magn Reson Imaging*, 18(2):242–54, Aug 2003. doi: 10.1002/jmri.10350.

B. Patenaude. *Bayesian Statistical Models of Shape and Appearance for Subcortical Brain Segmentation*. PhD thesis, University of Oxford, 2007.

T. Paus, F. Tomaiuolo, N. Otaky, D. MacDonald, M. Petrides, J. Atlas, R. Morris, and A. C. Evans. Human cingulate and paracingulate sulci: pattern, variability, asymmetry, and probabilistic map. *Cereb Cortex*, 6(2):207–214, 1996. ISSN 1047-3211 (Print); 1047-3211 (Linking).

M. Perrin, Y. Cointepas, A. Cachia, C. Poupon, B. Thirion, D. Riviere, P. Cathier, V. El Kouby, A. Constantinesco, D. Le Bihan, and J.-F. Mangin. Connectivity-based parcellation of the cortical mantle using q-ball diffusion imaging. *Int J Biomed Imaging*, 2008:368406, 2008. ISSN 1687-4188 (Print); 1687-4188 (Linking). doi: 10.1155/2008/368406.

A. Petrovic, S. M. Smith, R. A. Menke, and M. Jenkinson. Methods for tractography-driven surface registration of brain structures. *Med Image Comput Comput Assist Interv*, 12(Pt 1):705–712, 2009.

A. Petrovic, S. M. Smith, and M. Jenkinson. Functional Parcellation of SMA and pre-SMA in Human MFC Using Resting State Networks: Comparison With Seed-Based Functional Connectivity Parcellation Method. In *Human Brain Mapping Conference*, 2010.

G. Postelnicu, L. Zollei, and B. Fischl. Combined volumetric and surface registration. *IEEE Trans Med Imaging*, 28(4):508–522, 2009. ISSN 1558-0062 (Electronic); 0278-0062 (Linking). doi: 10.1109/TMI.2008.2004426.

W. K. Pratt. *Digital image processing: PIKS Scientific inside*. Wiley-Interscience, Hoboken, N.J., 4th ed., newly updated and rev. ed edition, 2007. ISBN 0471767778.

W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C*. Cambridge University Press, second edition, 1995.

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. *Proc Natl Acad Sci U S A*, 98 (2):676–682, 2001. ISSN 0027-8424 (Print); 0027-8424 (Linking). doi: 10.1073/pnas.98.2.676.

P. Roca, D. Riviere, P. Guevara, C. Poupon, and J.-F. Mangin. Tractography-based parcellation of the cortex using a spatially-informed dimension reduction of the

connectivity matrix. *Med Image Comput Comput Assist Interv*, 12(Pt 1):935–942, 2009.

S. E. Rose, F. Chen, J. B. Chalk, F. O. Zelaya, W. E. Strugnell, M. Benson, J. Semple, and D. M. Doddrell. Loss of connectivity in Alzheimer's disease: an evaluation of white matter tract integrity with colour coded MR diffusion tensor imaging. *J Neurol Neurosurg Psychiatry*, 69(4):528–530, 2000. ISSN 0022-3050 (Print); 0022-3050 (Linking).

S. E. Rose, A. L. Janke, and J. B. Chalk. Gray and white matter changes in Alzheimer's disease: a diffusion tensor imaging study. *J Magn Reson Imaging*, 27(1):20–26, 2008. ISSN 1053-1807 (Print). doi: 10.1002/jmri.21231.

M. F. Rushworth, T. E. Behrens, and H. Johansen-Berg. Connection patterns distinguish 3 regions of human parietal cortex. *Cereb Cortex*, 16(10):1418–30, 2006.

M. R. Sabuncu, B. D. Singer, B. Conroy, R. E. Bryan, P. J. Ramadge, and J. V. Haxby. Function-based intersubject alignment of human cortical anatomy. *Cereb Cortex*, 20(1):130–40, 2010.

D. H. Salat, D. S. Tuch, A. J. W. van der Kouwe, D. N. Greve, V. Pappu, S. Y. Lee, N. D. Hevelone, A. K. Zaleta, J. H. Growdon, S. Corkin, B. Fischl, and H. D. Rosas. White matter pathology isolates the hippocampal formation in Alzheimer's disease. *Neurobiol Aging*, 31(2):244–256, 2010. ISSN 1558-1497 (Electronic); 0197-4580 (Linking). doi: 10.1016/j.neurobiolaging.2008.03.013.

Z. Saygin, D. Osher, A. van der Kouwe, and G. John. Connectivity-based segmentation of human amygdala nuclei using probabilistic tractography. In *Human Brain Mapping Conference*, 2010.

J. D. Schmahmann and D. N. Pandya. *Fiber pathways of the brain*. Oxford University Press, Oxford, 2006. ISBN 0195104234 (alk. paper).

J. Scholz, M. C. Klein, T. E. Behrens, and H. Johansen-Berg. Training induces changes in white-matter architecture. *Nat Neurosci*, 12(11):1370–1, 2009.

P. Schröder and W. Sweldens. Spherical Wavelets: Texture Processing. In *Rendering Techniques*, pages 252–263, 1995a.

P. Schröder and W. Sweldens. Spherical wavelets: efficiently representing functions on the sphere. In *SIGGRAPH*, pages 161–172, 1995b.

F. Segonne, E. Grimson, and B. Fischl. A genetic algorithm for the topology correction of cortical surfaces. *Inf Process Med Imaging*, 19:393–405, 2005.

S. M. Smith. Fast robust automated brain extraction. *Hum Brain Mapp*, 17(3):143–155, 2002. ISSN 1065-9471 (Print); 1065-9471 (Linking). doi: 10.1002/hbm.10062.

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. *Neuroimage*, 31(4):1487–1505, 2006. ISSN 1053-8119 (Print); 1053-8119 (Linking). doi: 10.1016/j.neuroimage.2006.02.024.

S. M. Smith, P. T. Fox, K. L. Miller, D. C. Glahn, P. M. Fox, C. E. Mackay, N. Filippini, K. E. Watkins, R. Toro, A. R. Laird, and C. F. Beckmann. Correspondence of the brain's functional architecture during activation and rest. *Proc Natl Acad Sci U S A*, 106(31):13040–13045, 2009. ISSN 1091-6490 (Electronic); 0027-8424 (Linking). doi: 10.1073/pnas.0905267106.

S. M. Smith, K. L. Miller, G. Salimi-Khorshidi, M. Webster, C. F. Beckmann, T. E. Nichols, J. D. Ramsey, and M. W. Woolrich. Network modelling methods for FMRI. *NeuroImage*, In Press, Corrected Proof:–, 2010. ISSN 1053-8119. doi: DOI:10.1016/j.neuroimage.2010.08.063.

S. N. Sotiropoulos. *Processing of diffusion MR images of the brain: from crossing fibres to distributed tractography.* PhD thesis, University of Nottingham, 2010.

M. South. The amygdala in a nutshell: A comprehensive and current view The human amygdala. Paul J. Whalen and Elizabeth A. Phelps (Eds.). *J Clin Exp Neuropsychol*, 32(8):909–910, 2010. ISSN 1744-411X (Electronic); 1380-3395 (Linking). doi: 10.1080/13803391003625345.

O. Sporns, G. Tononi, and R. Kotter. The human connectome: A structural description of the human brain. *PLoS Comput Biol*, 1(4):e42, 2005. ISSN 1553-7358 (Electronic). doi: 10.1371/journal.pcbi.0010042.

J. Talairach and P. Tornoux. *Co-planar stereotaxic atlas of the human brain: 3-dimensional proportional system : an approach to cerebral imaging.* Georg Thieme, Stuttgart, 1988. ISBN 0865772932.

P. Thompson and A. W. Toga. A surface-based technique for warping three-dimensional images of the brain. *IEEE Trans Med Imaging*, 15(4):402–417, 1996. ISSN 0278-0062 (Print); 0278-0062 (Linking). doi: 10.1109/42.511745.

V. Tomassini, S. Jbabdi, J. C. Klein, T. E. Behrens, C. Pozzilli, P. M. Matthews, M. F. Rushworth, and H. Johansen-Berg. Diffusion-weighted imaging tractography-based parcellation of the human lateral premotor cortex identifies dorsal and ventral subregions with anatomical and functional specializations. *J Neurosci*, 27(38): 10259–69, 2007a.

V. Tomassini, S. Jbabdi, J. C. Klein, T. E. J. Behrens, C. Pozzilli, P. M. Matthews, M. F. S. Rushworth, and H. Johansen-Berg. Diffusion-weighted imaging tractography-based parcellation of the human lateral premotor cortex identifies dorsal and ventral subregions with anatomical and functional specializations. *J*

Neurosci, 27(38):10259–10269, 2007b. ISSN 1529-2401 (Electronic); 0270-6474 (Linking). doi: 10.1523/JNEUROSCI.2144-07.2007.

D. S. Tuch. Q-ball imaging. *Magn Reson Med*, 52(6):1358–72, 2004.

D. C. Van Essen. A tension-based theory of morphogenesis and compact wiring in the central nervous system. *Nature*, 385(6614):313–318, 1997. ISSN 0028-0836 (Print); 0028-0836 (Linking). doi: 10.1038/385313a0.

D. C. Van Essen and D. L. Dierker. Surface-based and probabilistic atlases of primate cerebral cortex. *Neuron*, 56(2):209–225, 2007. ISSN 0896-6273 (Print); 0896-6273 (Linking). doi: 10.1016/j.neuron.2007.10.015.

D. C. Van Essen, H. A. Drury, S. Joshi, and M. I. Miller. Functional and structural mapping of human cerebral cortex: solutions are in the surfaces. *Proc Natl Acad Sci U S A*, 95(3):788–795, 1998. ISSN 0027-8424 (Print); 0027-8424 (Linking).

D. C. Van Essen, H. A. Drury, J. Dickson, J. Harwell, D. Hanlon, and C. H. Anderson. An integrated software suite for surface-based analyses of cerebral cortex. *J Am Med Inform Assoc*, 8(5):443–59, 2001.

K. Van Leemput, A. Bakkour, T. Benner, G. Wiggins, L. L. Wald, J. Augustinack, B. C. Dickerson, P. Golland, and B. Fischl. Automated segmentation of hippocampal subfields from ultra-high resolution in vivo MRI. *Hippocampus*, 19(6):549–557, 2009. ISSN 1098-1063 (Electronic); 1050-9631 (Linking). doi: 10.1002/hipo.20615.

L. Vincent and P. Soille. Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 13(6):583 –598, jun. 1991. ISSN 0162-8828. doi: 10.1109/34.87344.

B. A. Vogt. *Cingulate neurobiology and disease*. Oxford University Press, Oxford, 2009. ISBN 9780198566960 (hbk.).

S. Wakana, H. Jiang, L. M. Nagae-Poetscher, P. C. M. van Zijl, and S. Mori. Fiber tract-based atlas of human white matter anatomy. *Radiology*, 230(1):77–87, Jan 2004. doi: 10.1148/radiol.2301021640.

V. J. Wedeen, R. P. Wang, J. D. Schmahmann, T. Benner, W. Y. Tseng, G. Dai, D. N. Pandya, P. Hagmann, H. D'Arceuil, and A. J. de Crespigny. Diffusion spectrum magnetic resonance imaging (DSI) tractography of crossing fibers. *Neuroimage*, 41 (4):1267–77, 2008.

W. M. r. Wells, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis. Multi-modal volume registration by maximization of mutual information. *Med Image Anal*, 1 (1):35–51, 1996. ISSN 1361-8415 (Print); 1361-8415 (Linking).

J. S. Werner and L. M. Chalupa. *The visual neurosciences*. MIT Press, Cambridge, Mass., 2004. ISBN 0262033089.

M. R. Wiegell, D. S. Tuch, H. B. W. Larsson, and V. J. Wedeen. Automatic segmentation of thalamic nuclei from diffusion tensor magnetic resonance imaging. *Neuroimage*, 19(2 Pt 1):391–401, 2003. ISSN 1053-8119 (Print); 1053-8119 (Linking).

M. Wilms, S. B. Eickhoff, K. Specht, K. Amunts, N. J. Shah, A. Malikovic, and G. R. Fink. Human V5/MT+: comparison of functional and cytoarchitectonic data. *Anat Embryol (Berl)*, 210(5-6):485–495, 2005. ISSN 0340-2061 (Print); 0340-2061 (Linking). doi: 10.1007/s00429-005-0064-y.

R. P. Woods, J. C. Mazziotta, and S. R. Cherry. MRI-PET registration with automated algorithm. *J Comput Assist Tomogr*, 17(4):536–546, 1993. ISSN 0363-8715 (Print); 0363-8715 (Linking).

S. Xie, J. X. Xiao, G. L. Gong, Y. F. Zang, Y. H. Wang, H. K. Wu, and X. X. Jiang. Voxel-based detection of white matter abnormalities in mild Alzheimer disease. *Neurology*, 66(12):1845–1849, 2006. ISSN 1526-632X (Electronic); 0028-3878 (Linking). doi: 10.1212/01.wnl.0000219625.77625.aa.

J. H. Xuereb, R. H. Perry, J. M. Candy, E. K. Perry, E. Marshall, and J. R. Bonham. Nerve cell loss in the thalamus in Alzheimer's disease and Parkinson's disease. *Brain*, 114 (Pt 3):1363–1379, 1991. ISSN 0006-8950 (Print).

B. T. Yeo, M. Sabuncu, P. Golland, and B. Fischl. Task-optimal registration cost functions. *Med Image Comput Comput Assist Interv*, 12(Pt 1):598–606, 2009.

B. T. Yeo, M. R. Sabuncu, T. Vercauteren, N. Ayache, B. Fischl, and P. Golland. Spherical demons: fast diffeomorphic landmark-free surface registration. *IEEE Trans Med Imaging*, 29(3):650–68, 2010a.

B. T. Yeo, M. R. Sabuncu, T. Vercauteren, D. J. Holt, K. Amunts, K. Zilles, P. Golland, and B. Fischl. Learning task-optimal registration cost functions for localizing cytoarchitecture and function in the cerebral cortex. *IEEE Trans Med Imaging*, 29(7):1424–41, 2010b.

P. Yu, P. E. Grant, Y. Qi, X. Han, F. Ségonne, R. Pienaar, E. Busa, J. Pacheco, N. Makris, R. L. Buckner, P. Golland, and B. Fischl. Cortical Surface Shape Analysis Based on Spherical Wavelets. *IEEE Trans. Med. Imaging*, 26(4):582–597, 2007.

M. Zarei, B. Patenaude, J. Damoiseaux, C. Morgese, S. Smith, P. M. Matthews, F. Barkhof, S. A. R. B. Rombouts, E. Sanz-Arigita, and M. Jenkinson. Combining shape and connectivity analysis: an MRI study of thalamic degeneration in Alzheimer's disease. *Neuroimage*, 49(1):1–8, 2010. ISSN 1095-9572 (Electronic); 1053-8119 (Linking). doi: 10.1016/j.neuroimage.2009.09.001.

D. Zhang, A. Z. Snyder, M. D. Fox, M. W. Sansbury, J. S. Shimony, and M. E. Raichle. Intrinsic functional relations between human cerebral cortex and thalamus. *J Neurophysiol*, 100(4):1740–1748, 2008. ISSN 0022-3077 (Print); 0022-3077 (Linking). doi: 10.1152/jn.90463.2008.

Z. Zhang. Iterative point matching for registration of free-form curves and surfaces. *International Journal of Computer Vision*, 13(2):119–152, OCT 1994.

B. Zitova and J. Flusser. Image registration methods: a survey. *Image and Vision Computing*, 21(11):977–1000, OCT 1 2003. doi: DOI10.1016/S0262-8856(03)00137-9.

U. Ziyan and C.-F. Westin. Joint segmentation of thalamic nuclei from a population of diffusion tensor MR images. *Med Image Comput Comput Assist Interv*, 11(Pt 1):279–286, 2008.

L. Zollei, A. Stevens, K. Huber, S. Kakunoori, and B. Fischl. Improved tractography alignment using combined volumetric and surface registration. *Neuroimage*, 51(1):206–213, 2010. ISSN 1095-9572 (Electronic); 1053-8119 (Linking). doi: 10.1016/j.neuroimage.2010.01.101.