Imperial College of Science, Technology and Medicine
Department of Computing

Connectivity-Driven Parcellation Methods for the Human Cerebral Cortex

Salim Arslan

A dissertation submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy
of
Imperial College London

October 2017

Abstract

The macro connectome elucidates the pathways through which brain regions are structurally connected or functionally coupled to perform cognitive functions. It embodies the notion of representing, analysing, and understanding all connections within the brain as a network, while the subdivision of the brain into interacting cortical units is inherent in its architecture. As a result, the definition of network nodes is one of the most critical steps in connectivity network analysis. Parcellations derived from anatomical brain atlases or random parcellations are traditionally used for node identification, however these approaches do not always fully reflect the functional/structural organisation of the brain. Connectivity-driven methods have arisen only recently, aiming to delineate parcellations that are more faithful to the underlying connectivity. Such parcellation methods face several challenges, including but not limited to poor signal-to-noise ratio, the curse of dimensionality, and functional/structural variations inherent in individual brains, which are only limitedly addressed by the current state of the art.

In this thesis, we present robust and fully-automated methods for the subdivision of the entire human cerebral cortex based on connectivity information. Our contributions are four-fold: First, we propose a clustering approach to delineate a cortical parcellation that provides a reliable abstraction of the brain's functional organisation. Second, we cast the parcellation problem as a feature reduction problem and make use of manifold learning and image segmentation techniques to identify cortical regions with distinct structural connectivity patterns. Third, we present a multi-layer graphical model that combines within- and between-subject connectivity, which is then decomposed into a cortical parcellation that can represent the whole population, while accounting for the variability across subjects. Finally, we conduct a large-scale, systematic comparison of existing parcellation methods, with a focus on providing some insight into the reliability of brain parcellations in terms of reflecting the underlying connectivity, as well as, revealing their impact on network analysis.

We evaluate the proposed parcellation methods on publicly available data from the Human Connectome Project and a plethora of quantitative and qualitative evaluation techniques investigated in the literature. Experiments across multiple resolutions demonstrate the accuracy of the presented methods at both subject and group levels with regards to reproducibility and fidelity to the data. The neuro-biological interpretation of the proposed parcellations is also investigated by comparing parcel boundaries with well-structured properties of the cerebral

cortex. Results show the advantage of connectivity-driven parcellations over traditional approaches in terms of better fitting the underlying connectivity. However, the benefit of using connectivity to parcellate the brain is not always as clear regarding the agreement with other modalities and simple network analysis tasks carried out across healthy subjects. Nonetheless, we believe the proposed methods, along with the systematic evaluation of existing techniques, offer an important contribution to the field of brain parcellation, advancing our understanding of how the human cerebral cortex is organised at the macroscale.

Acknowledgements

I would like to first of all thank my supervisor Daniel Rueckert for giving me the opportunity to undertake my PhD in the BioMedIA group and for his support, guidance, and inspiration throughout this journey. I would like to express my gratitude to Ben Glocker for his comments and feedback during the preparation of this thesis, as well as, Yi-ke Guo and Georg Langs for their time and commitment in the examination process. Many thanks to my colleagues at Imperial College London for providing a great atmosphere and research environment, in particular to Sofia Ira Ktena and Sarah Parisot for their valuable inputs and contributions. Special thanks go to Amani El-Kholy for always being there to sort things out.

I am also extremely grateful to my family for all the support and encouragement, especially during hard times. Finally, I thank my beloved wife, Dilara, for *everything* - nothing would have been possible without her being by my side. She is the real mvp...

Declaration of Originality

I, Salim Arslan, hereby declare that the work described in this thesis is my own, except where specifically acknowledged.

© The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives license. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the license terms of this work.

To Dilara, my amazing wife

Acronyms

2D Two-dimensional

3D Three-dimensional

dMRI Diffusion Magnetic Resonance Imaging

fMRI Functional Magnetic Resonance Imaging

rs-fMRI Resting-State Functional Magnetic Resonance Imaging

t-fMRI Task Functional Magnetic Resonance Imaging

ACC Anterior Cingular Cortex

ARI Adjusted Rand Index

BA Brodmann's Area

BIC Bayesian Information Criterion

BOLD Blood Oxygenation Level Dependent

CDP Connectivity-Driven Parcellation

DTI Diffusion Tensor Imaging

DWI Diffusion Weighted Imaging

EEG Electroencephalography

EPI Echo-planar Imaging

GM Gray Matter

HARDI High Angular Resolution Diffusion Imaging

HCP Human Connectome Project

ICA Independent Component Analysis

MEG Magnetoencephalography

MFC Medial Frontal Cortex

MNI Montreal Neurological Institute

MRF Markov Random Field

MRI Magnetic Resonance Imaging

PCA Principle Component Analysis

PET Positron Emission Tomography

ROI Region of Interest

RF Radio Frequency

RSFC Resting-State Functional Connectivity

RSN Resting-State Network

SAD Sum of Absolute Differences

SC Silhouette Coefficient

SMA Supplementary Motor Area

SNR Signal-to-Noise Ratio

WM White Matter

Contents

List of Tables **xv**

List of Figures **xvii**

1 Introduction **1**

1.1 Motivation 1

1.2 Research Contributions and Thesis Outline 4

2 Studying the Human Brain **7**

2.1 Introduction 7

2.2 The Human Brain 9

2.2.1 The Cerebral Cortex 10

2.3 Mapping the Brain 11

2.4 Imaging the Macro Connectome 13

2.4.1 Magnetic Resonance Imaging 13

2.4.2 Diffusion Magnetic Resonance Imaging 15

2.4.3 Functional Magnetic Resonance Imaging 17

2.5 Structural and Functional Connectivity 22

2.5.1 Estimating Structural Connectivity 22

2.5.2 Estimating Functional Connectivity 25

2.6 Imaging Data 27

2.6.1 The Human Connectome Project 27

2.6.2 HCP Minimal Preprocessing Pipelines 28

2.6.3 Datasets 33

3 Parcellating the Human Brain 36

3.1 Introduction 36

3.2 Brain Parcellation: Aim and Scope 37

3.2.1 Historical Foundations of Brain Parcellation 37

3.2.2 Neuroimaging for Brain Parcellation 40

3.2.3 Parcellation for Connectome Analysis 41

3.3 Connectivity-Driven Parcellation 43

3.3.1 Methods for Connectivity-Driven Parcellation 44

3.3.2 What Defines A Reliable Connectivity-Driven Parcellation? 58

3.3.3 Challenges in Connectivity-Driven Parcellation 60

4 Single Subject Parcellation of the Cerebral Cortex 62

4.1 Introduction 63

4.2 Methodology 65

4.2.1 Data . . . 65

4.2.2 A Two-Level Parcellation Framework . . . 66

4.3 Experiments . . . 71

4.3.1 Performance Measures . . . 71

4.3.2 Qualitative Assessment of Parcellations with Other Modalities . . . 77

4.3.3 Inter-Subject Variability Across Proposed Parcellations . . . 78

4.3.4 Comparison Methods . . . 78

4.3.5 Experimental Setup . . . 82

4.4 Results . . . 83

4.4.1 Parameter Selection . . . 83

4.4.2 First-Level Clustering Results . . . 85

4.4.3 Subject-Level Parcellation Results . . . 87

4.4.4 Inter-Modality Assessment of the Proposed Parcellations . . . 93

4.5 Discussion . . . 94

5 Cortical Boundary Mapping through Manifold Learning 101

5.1 Introduction . . . 102

5.2 Methodology . . . 105

5.2.1 Data . . . 105

5.2.2 Method . . . 105

5.3 Experiments . . . 117

5.3.1 Quantitative and Qualitative Evaluation Measures . . . 117

5.3.2 Comparison Methods . . . 119

5.4 Results . . . 120

5.4.1 Parameter Selection . . . 120

5.4.2 Quantitative Assessment of the Parcellation Quality . . . 122

5.4.3 Visual and Inter-Modality Assessment . . . 123

5.5 Discussion . . . 126

6 Joint Spectral Decomposition of the Cerebral Cortex 133

6.1 Introduction . . . 134

6.2 Methodology . . . 136

6.2.1 Data . . . 136

6.2.2 Joint Spectral Decomposition . . . 137

6.3 Experiments . . . 143

6.3.1 Evaluation . . . 143

6.3.2 Comparison Methods . . . 144

6.3.3 Experimental Setup . . . 146

6.3.4 Parameter Selection . . . 147

6.4 Results . . . 148

6.4.1 Parameter Analysis . . . 152

6.5 Discussion . . . 154

7 A Systematic Comparison of Parcellation Methods **157**

7.1 Introduction . . . 158

7.2 Materials and Methods . . . 159

7.2.1 Data . . . 159

7.2.2 Parcellation Methods . . . 160

7.2.3 Parcellation Evaluation Techniques . . . 165

7.3 Results . . . 174

7.3.1 Experimental Setup . . . 174

7.3.2 Reproducibility Results . . . 175

7.3.3 Cluster Validity Results . . . 176

7.3.4 Multi-Modal Comparisons . . . 180

7.3.5 Network Analysis . . . 184

7.4 Discussion . . . 186

8 Conclusions **196**

8.1 Future Work . . . 198

List of Publications **200**

Bibliography **203**

List of Tables

2.1 Summary of the demographic information of Dataset 1. 34

2.2 Summary of the demographic information of Dataset 2. 35

4.1 Summary of the parcellation methods used for comparison. 79

6.1 Computational performance of the proposed method. 154

7.1 Computed parcellation methods. 161

7.2 Pre-computed, publicly available parcellation methods. 163

7.3 Parcellation evaluation techniques. 165

List of Figures

1.1 A typical network analysis pipeline driven by a connectivity-based parcellation. 3

1.2 Parcellation of anterior cingulate cortex based on different properties. 4

2.1 An illustration of the human brain and cerebral cortex. 10

2.2 Different types of MR images. 15

2.3 Illustration of a typical dMRI pipeline. 25

2.4 Illustration of a typical fMRI pipeline. 27

2.5 Grayordinate standard coordinate system. 30

3.1 Brodmann's areas of the cerebral cortex. 38

3.2 A parcellation of the left hemisphere obtained with k-means. 46

3.3 Cortical parcellations obtained using spectral clustering. 48

3.4 Cortical parcellations with different resolutions obtained via hierarchical clustering. 50

3.5 Spatial maps obtained via independent component analysis. 52

3.6 Whole-brain cortical parcellation with boundary mapping. 54

3.7 An example of non-linear manifold learning. 56

4.1 The proposed two-level parcellation framework. 66

4.2 Illustration of the mesh graph that represents the cerebral cortex and the supervertex clustering algorithm. 67

4.3 Illustration of the agglomerative hierarchical clustering algorithm. 70

4.4 Illustration of the neighbourhood criterion for a set of supervertices. 71

4.5 Visual outline of the parcellation evaluation pipelines. 82

4.6 Evolution of the boundaries of a supervertex with different values of α. 83

4.7 The number of spatially disconnected vertices with respect to α and N. 84

4.8 Reproducibility and homogeneity as a function of the number of parcels. 84

4.9 Parcellations of the left lateral cortex derived from one subject for 1000 regions. 86

4.10 First-level reproducibility results. 86

4.11 First-level homogeneity values and Silhouette coefficients. 87

4.12 Parcellations of the left lateral cortex derived from one subject for 100 regions. . 88

4.13 Subject-level reproducibility results. 90

4.14 Subject-level homogeneity results. 92

4.15 Subject-level Silhouette analysis results. 92

4.16 Inter-subject consistency maps obtained from all subjects. 94

4.17 Parcellation boundaries compared to task fMRI activation maps. 95

4.18 Parcellation boundaries compared to task fMRI activation maps. 96

4.19 Parcellation boundaries compared to myelin maps and Brodmann areas. 97

5.1 Parcellation pipeline summarising all steps after preprocessing. 106

5.2 Illustration of non-linear feature reduction using Laplacian eigenmaps. 109

5.3 Illustration of eigenvector discretisation and boundary estimation. 115

5.4 Connectivity profiles of different vertices. 116

5.5 Boundary map generation and cortical parcellation. 117

5.6 The number of parcels obtained by the proposed method as a function of the number of eigenvectors. 121

5.7 Homogeneity values and Silhouette coefficients based on structural connectivity. 123

5.8 Homogeneity values and Silhouette coefficients based on functional connectivity. 124

5.9 Parcellations of the left lateral cortex derived from one subject. 125

5.10 Inter-subject consistency and the sulcal depth maps. 126

5.11 Parcellations of four different subjects compared to myelin maps. 127

5.12 Parcellations of four different subjects compared to the Brodmann atlas. 128

5.13 Quantitative assessment of the overlap between the proposed parcellations and Brodmann areas. 129

6.1 Visual representation of the parcellation pipeline. 137

6.2 Visualisation of spectral embedding at different spatial resolutions. 139

6.3 Illustration of how to compute the sum of absolute differences. 144

6.4 Illustration of how to compute a symmetric adjacency matrix from four toy parcellations. 146

6.5 Group-to-group reproducibility results for different number of parcels. 149

6.6 Group-wise whole-brain parcellations obtained by the joint spectral decomposition method. 150

6.7 Homogeneity and Silhouette analysis results for different number of parcels. . . . 151

6.8 Functional consistency results for different number of parcels. 151

6.9 Evolution of different evaluation measurements with respect to the number of (a) inter-subject links and (b) supervertices. 153

7.1 Visual outline of parcellation generation and evaluation steps. 160

7.2 Illustration of the generation of null models. 168

7.3 Cyto- and myelo-architecture of the cerebral cortex. 170

7.4 Individual subject identification procedure. 174

7.5 Group-level reproducibility results. 176

7.6 Group-level homogeneity results. 177

7.7 Homogeneity of the parcellations and their respective 1000 null models. 178

7.8 Difference between the actual homogeneity and the homogeneity distribution of null models. 179

7.9 Group-level Silhouette analysis results. 180

7.10 Quantitative evaluation measures obtained from *K-Means-AVR* after relabelling spatially disjoint regions. 181

7.11 Agreement of group-level parcellations with task activation. 182

7.12 Agreement of group-level parcellations with the Brodmann atlas. 183

7.13 Agreement of group-level parcellations with highly myelinated cortical areas. . . 184

7.14 Gender classification results. 185

7.15 Individual subject prediction results. 186

7.16 The first half of the parcellations considered in this study as projected on the left hemisphere. 194

7.17 The second half of the parcellations considered in this study as projected on the left hemisphere. 195

Chapter 1

Introduction

1.1 Motivation

Understanding the brain's behaviour and function has been a prominent and ongoing research subject for over a century [287, 231]. Neuronal interconnections constitute the primary means of information transmission within the brain and are strongly related to the way the brain functions [193, 202, 227]. These connections constitute a complex network that can be estimated at the macroscale via modern neuroimaging techniques such as Magnetic Resonance Imaging (MRI) [233, 55, 261]. While structural connectivity networks are typically inferred from diffusion MRI (dMRI), functional networks can be mapped using resting-state functional MRI (rs-fMRI) [101, 66]. The former allows estimation of the physical (anatomical) connections, while the latter elucidates putative functional connections between spatially remote brain regions.

Analysing brain connectivity from a network theoretical point of view has shown significant potential for identifying organisational principles in the brain and their connections to cognitive procedures [155, 224, 41, 231] and brain disorders, such as Alzheimer's disease [238, 146], attention-deficit/hyperactivity disorder [275] and schizophrenia [19]. This allows to study the brain and its function from a new perspective that accounts for the complexity of its architecture. One of the critical steps in the construction of brain connectivity networks is the definition

of network nodes [231, 66]. Adopting a vertex- or voxel-based representation yields networks that are very noisy and of extremely high dimensionality, making subsequent network analysis steps often intractable [244]. An alternative approach to node definition is to parcellate the cerebral cortex into a set of distinct regions of interest (ROI), where each ROI (i.e. parcel) corresponds to a node of the connectivity network. This further allows to reduce the complexity of connectivity, an aspect that is highly critical for the study of brain dynamics with whole-brain models [55].

Traditionally, parcellations derived from neuro-anatomical landmarks [254, 61, 74] or microstructural features [38, 269, 175] have been used to define ROIs for network analysis [231]. Whereas such parcellations are of great importance in order to derive neuro-biologically meaningful brain atlases, they might fail to fully reflect the intrinsic organisation of the brain and capture the functional/structural variability inherent in individual subjects, due to brain maturation or pathology [67, 132, 272, 92]. In addition, they are typically generated on a single or small set of individuals, which can make them biased and unable to accurately represent population variability or adapt to new subjects [245, 54]. Alternative approaches include use of random parcellations for node definition [231]; however, this kind of approaches could fail to represent the underlying cortical organisation faithfully and may lead to ill-defined nodes in the constructed network [226].

More recent parcellation methods attempt to overcome these problems by using connectivity information, captured from rs-fMRI or dMRI data, to derive a set of network nodes for connectivity analysis [244, 66] as illustrated in Fig. 1.1. This type of approaches typically casts the parcellation problem as a clustering problem, in which a connectivity profile is first computed for each individual subunit (e.g. vertices or voxels). These connectivity profiles are then submitted to a parcellation scheme for grouping subunits, such that the connectivity is similar for subunits within the same cluster, but different between clusters [66]. Since connectivity-based parcellations are directly obtained from the underlying data, such methods can potentially provide highly homogeneous ROIs and separate regions with different patterns of connectivity more accurately. As a result, they are likely to yield a more reliable set of network nodes, as these nodes are typically represented by a single entity (such as the average connectivity

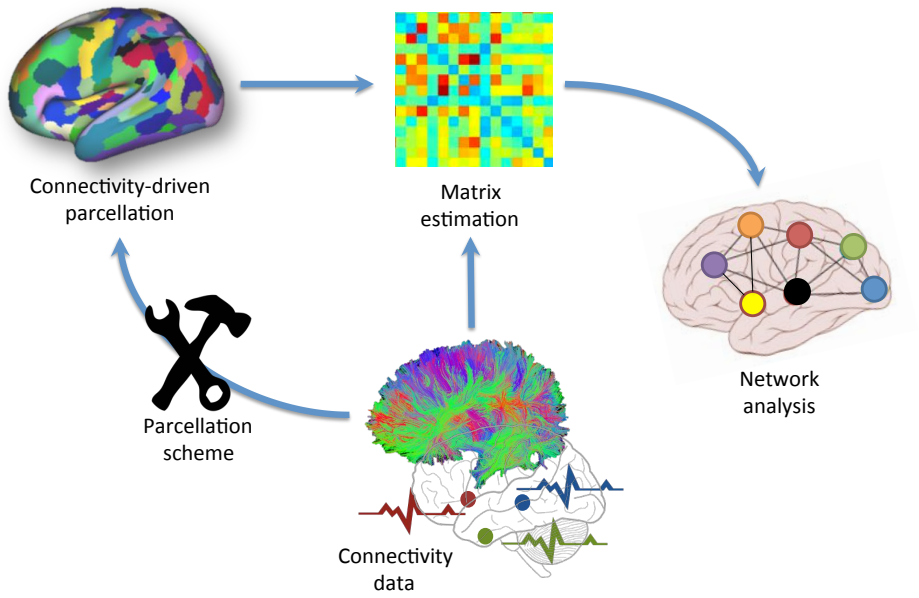
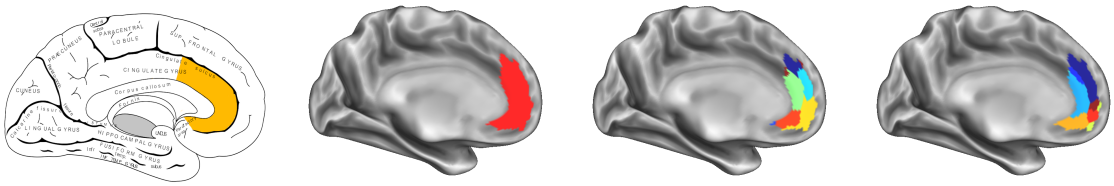


Figure 1.1: A typical network analysis pipeline driven by a connectivity-based parcellation.

profile) in network analysis [218, 93].

The role of connectivity in human brain mapping studies is also crucial, as it provides complementary information for subdividing the cortex into anatomically and functionally distinct regions [193]. In particular, the functional operations performed by a cortical area are thought to simultaneously depend on its local micro-architecture and connectivity [193, 66]. As a result, most of the neuro-anatomical parcellations alone cannot precisely match the degree of functional segregation of the cerebral cortex [175, 86]. For example, the anterior cingulate cortex (ACC), which is involved in certain brain functions such as cognition and emotion [236], is typically represented as a single ROI in anatomical brain atlases [239, 254, 74] (Fig. 1.2); however, it exhibits a great amount of heterogeneity in structural [22] and functional connectivity [157]. As a result, when ACC is localised from a connectivity-driven parcellation, it typically consists of several subregions with varying shape and size, as shown in Fig. 1.2(c-d). In general, connectivity-driven parcellations provide a greater flexibility to study the brain function, as they enable the segregation of the cortex at different scales, as opposed to anatomical atlases with fixed resolutions.

Despite many attempts to parcellate the brain with respect to connectivity, the problem is still



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

Figure 1.2: (a) Medial surface of left cerebral hemisphere, with anterior cingulate cortex (ACC) is highlighted. (b) ACC as delineated by a cortical-folding based anatomical atlas [74]. The cerebral organisation of the same area based on (c) resting-state functional connectivity and (d) dMRI-based structural connectivity, obtained with the connectivity-driven approaches, [13] and [12], respectively. Each colour indicates a distinct region, with homogeneous features (e.g. connectivity).

open to improvements. This is primarily due to the fact that, like all other clustering problems, the parcellation problem is ill-posed, thus, obtaining accurate subdivisions of the cortex depends on the proposed method's fidelity to the underlying data, as well as its capability to encapsulate valuable information from the naturally complex, noisy, and high-dimensional connectivity patterns in the brain [187, 66]. The heterogeneity of the population under investigation, i.e. inter-subject variability, also possesses an additional challenge, especially regarding the identification of shared patterns of connectivity and the delineation of spatially coherent cortical regions across different subjects [32, 131, 273].

1.2 Research Contributions and Thesis Outline

The aim of this thesis is to develop robust and reliable methods for subdividing the cerebral cortex into spatially contiguous, non-overlapping, and distinct regions with respect to underlying connectivity. Parcellations obtained by the proposed methods can provide high-level abstractions of the functional specialisation and segregation in the cerebral cortex. Such parcellations can further be used to define the network nodes in connectivity analysis, thus, help better understand how connectivity changes through development, ageing, and neurological disorders. In addition, our parcellations allow to study the brain's cortical organisation from a multi-scale perspective through the subdivision of the cerebral cortex at varying levels of granularity.

Chapter 2 provides a brief overview of the neuro-biological basis of the brain with a special emphasis on the cerebral cortex, which is responsible for high-level brain functions such as language and memory. It continues with an introduction to imaging of the brain at the macroscale, particularly focusing on rs-fMRI and dMRI, two mostly-used techniques to capture connectivity non-invasively. After summarising the quantitative methods for estimating structural and functional connectivity from the MRI data, we describe the data used to develop and evaluate the proposed parcellation methods throughout the thesis.

Chapter 3 briefly covers the historical foundations of brain parcellation and reviews techniques that allow the segregation of the cerebral cortex based on information obtained from different neuro-anatomical properties, with a particular focus on the connectivity-driven approaches. Both Chapter 2 and 3 describe the challenges for obtaining reliable connectivity-driven parcellations, such as low signal-to-noise ratio (SNR) and high dimensionality, as well as, limitations induced by imaging techniques and pre-processing pipelines prior to parcellation.

Chapter 4 presents a novel method for the subject-specific parcellation of the cerebral cortex using resting-state functional connectivity (RSFC). It is based on the idea of combining two different clustering strategies in a parcellation framework that allows subdividing the cortex at varying levels of detail. The chapter also provides an extensive comparison of different subject-level parcellation methods with respect to their fidelity to the underlying connectivity. Most of the statistical techniques used for evaluating parcellations are introduced in this chapter and used throughout the remaining of the thesis.

Chapter 5 is motivated by the idea of identifying local connectivity patterns in the brain through dimensionality reduction, which can then be used to compute whole-brain cortical parcellations on a single subject basis. Contrarily to the preceding chapter, it aims to delineate subdivisions of the brain with respect to structural connectivity estimated from dMRI. The proposed method casts the parcellation problem as the localisation of boundaries between cortical regions with distinct connectivity patterns and solves it using image segmentation techniques.

Both Chapter 4 and 5 discuss the limitations of the presented methods, in particular with respect to different connectivity types used, and investigate the variability across individuals

from a cortical parcellation point of view. In addition, multi-modal comparisons are also provided to assess the degree of alignment between connectivity-driven parcellations and well-defined neuro-anatomical properties of the cerebral cortex.

Chapter 6 proposes a robust group-wise parcellation framework that is simultaneously driven by the within- and inter-subject variability in connectivity. A joint graphical model is formed that can effectively capture the fundamental properties of connectivity within the population, while still preserving individual subject characteristics. The method is presented as an alternative approach to compute a group-wise parcellation, which is typically obtained from either average datasets or *a priori* delineated subject-level parcellations, and shown to produce a more robust and accurate representation of a group of subjects in terms of reflecting the underlying connectivity.

Chapter 7 provides a large-scale, systematic comparison of existing parcellation methods using publicly available sources. Experiments consist of quantitative assessments of 24 group-level parcellations (including connectivity-driven, random and anatomical parcellations) for different resolutions. Several criteria are simultaneously considered to evaluate parcellations, including (1) reproducibility across different groups, (2) fidelity to the underlying connectivity data estimated from rs-fMRI, (3) agreement with functional activation and well-known properties of the cerebral cortex, and (4) two simple network analysis tasks. This extensive empirical study highlights the strengths and shortcomings of the various methods and aims to provide a guideline for the choice of parcellation technique and resolution according to the task at hand.

Chapter 8 concludes this work with an emphasis on the thesis achievements and possible future directions, including the further use of connectivity-driven parcellations in subject- and group-level connectivity studies.

Chapter 2

Studying the Human Brain

Abstract

This chapter provides an overview of the background required to understand the connectivity-driven parcellation methods presented in this thesis. We first survey the neuro-biological basis of the human brain with an emphasis on the cerebral cortex, the folded surface of the brain, that is responsible for high-level functions. After giving some insight towards the importance of connectivity in the context of brain mapping, we summarise the fundamentals of magnetic resonance imaging (MRI) techniques used for capturing the connectivity that underlie the brain function and anatomy. We then elaborately explain the most commonly used approaches for estimating the functional and structural connectivity from the MRI data. The chapter ends with an overview of the data used to conduct experiments in this thesis.

2.1 Introduction

The human brain is the 'headquarters' of the human nervous system that allows carrying out a variety of operations, ranging from relatively primitive actions such as executing movement, to more complex functions such as thinking and speaking, as well as many different cognitive

processes that separate the humans from other animals [119, 46]. Although the neuro-biological foundations of the brain are mostly revealed through advances in neuroscience, the relationship between the brain and cognitive functions that constitute the human behaviour is still not completely uncovered [47, 119]. How is the brain involved in certain cognitive operations that underlie the basis of thought, memory, perception, and act? Going as back as to the ancient Greek times, philosophers and scientists have been endeavouring to answer this question; yet to this time, mapping the brain's function and anatomy still stands as one of the greatest challenges in the field of modern neuroscience [47, 119, 20].

In order to understand how the brain functions, one should first know its elemental units and their interactions with each other [233]. To this end, the following section explains the fundamentals of the brain anatomy, focusing on the main structures that constitute the brain and their roles in the nervous system. We exclusively cover the cerebral cortex, since many cognitive functions and mental operations take place in this convoluted layer of neural tissue.

The input and output connectivity of a cortical region is considered to play a critical role to determine its functionality [193]. Studying connectivity is therefore important to better understand the link between function and anatomy. Towards this end, Section 2.3 gives the rationale behind brain mapping studies that explore connectivity at the micro-, meso-, and macro-scales. The latter is of particular interest, as connectivity analysis at the macroscale allows exploring functional interactions and anatomical pathways between different brain regions.

In order to provide prior knowledge about connectivity, Section 2.4 covers the fundamentals of MRI, the most-widely used imaging technology for the *in-vivo* connectivity studies. After briefly explaining the physics behind MRI, the section proceeds with a detailed coverage of diffusion and functional MRI, the main imaging techniques used for capturing the structural and functional connectivity, respectively. In this section, we further provide information about the general drawbacks and problems associated with each technique, as well as the standard preprocessing pipelines necessary to bring the imaging data to an analysable basis.

Given the basics of brain imaging at the macroscale, Section 2.5 covers the common approaches for estimating functional and structural connectivity. We first summarise different tractography

algorithms used to delineate anatomical pathways with respect to dMRI, and then give a brief overview of the most widely employed statistical techniques for modelling the functional interactions between different cortical regions captured with fMRI.

In Section 2.6, we wrap up the chapter with an overview of the imaging data used throughout this thesis. We conduct our experiments on the publicly available data collected and distributed by the Human Connectome Project [261]. After a brief introduction to the project, we summarise the image acquisition and preprocessing pipelines, and provide details about the two datasets used in different parts of this thesis.

2.2 The Human Brain

The human brain controls the human nervous system and facilitates mental operations with its highly complex structure [176]. It receives input from the environment via sensory organs, processes these signals through a serial -as well as parallel- set of sophisticated pipelines and generates complex responses to coordinate the human body. The integration and interaction of these signals constitute the 'mind', a set of operations that leads to observed human behaviour [119] and distinguishes humans from each other, despite the fact that the neuro-anatomical structure of the brain is highly similar across individuals.

Like those of all vertebrates (and most of the other animals), the human brain is located in the head and protected by the skull. It consists of three distinct parts: the cerebrum, the cerebellum, and the brain stem, as shown in Fig. 2.1. The largest of these is the cerebrum, which consists of two approximately symmetric cerebral hemispheres, interconnected by a bundle of nerve fibers. It is covered with a thin layer of neural tissue, the cerebral cortex, and contains several subcortical structures including the basal ganglia, the hippocampus, and the amygdala. The cerebral hemispheres are associated with motor and sensory functions, as well as involved with aspects of different cognitive functions, such as memory, language, and emotion [119, 20]. They are connected to the rest of the brain via the diencephalon, which consists of two substructures, the thalamus and hypothalamus. The former processes motor and sensory information

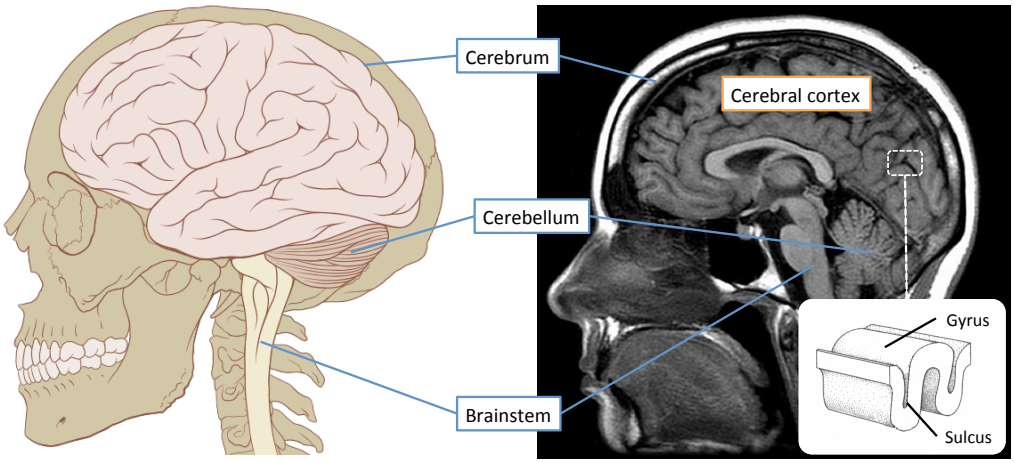


Figure 2.1: (*Left*) An illustration of the human brain and skull, showing three distinct structures (image from Wikimedia Commons [281]). (*Right*) Sagittal view of an MR image of the living human brain, in which the cerebral cortex and some other major structures are clearly visible (image is taken from FlickR [76]). Inlay shows a gyrus and sulcus, the characteristic infoldings that give the cerebral cortex its highly convoluted shape.

transmitted to the cerebral cortex, while the hypothalamus is involved with the regulation of autonomic and endocrine functions.

Underneath the cerebrum lies the brain stem that facilitates the brain's integration with the central nervous system by attaching it to the spinal cord. The brain stem is primarily responsible for connecting the motor and sensory systems in the brain to the rest of the body. It is also involved in several vital autonomic functions, including breathing, maintaining consciousness and controlling the heart rate [178].

The cerebellum is located at the back of the brain, behind the brain stem. It contributes to motor planning and control, including the regulation of movement, maintenance of balance and learning of motor skills. It is also involved in certain cognitive functions, such as language processing [35].

2.2.1 The Cerebral Cortex

The cerebral cortex is a thin layer of neural tissue that overlays the cerebral hemispheres (Fig 2.1). It is primarily associated with most of the mental operations that lead to the observed

human behaviour, including but not limited to, complex thought, memory, emotion, learning, planning, acting, control of movement, sensation, vision, and auditory processing [119]. As a result, it naturally constitutes the focus of many brain studies and is also targeted in this thesis for parcellation purposes.

The cerebral cortex is mainly composed of two different types of tissue, commonly referred to as 'gray matter' (GM) and 'white matter' (WM). GM consists of neuronal cell bodies, glial (non-neuronal) cells, and unmyelinated axons. The name 'gray' comes from its very light gray appearance that is caused by the blood vessels and neuronal cell bodies [127]. WM is composed of bundles of long-range myelinated axons that interconnect different areas of GM and transmit nerve impulses between neurons. It is refereed to as WM due to the white fatty substance (myelin) that surrounds the axons, which acts as an electrical insulation and facilitates the transmission of nerve impulses [125].

Like in all other mammals, the human cerebral cortex has a highly convoluted shape, consisting of several characteristic deep infoldings. The ridges in this convoluted structure are called *gyri*, while the grooves that make these convolutions are called *sulci*. A gyrus and sulcus are shown in Fig. 2.1. The neuro-developmental process that leads to cortical folding, i.e. gyrification, starts at approximately 15-20 weeks gestational age [142] and continues even later after birth [84]. Although the precise reason behind this convoluted shape is not known, it is thought to be an evolutionary strategy to accommodate more GM within the skull's limited volume, and hence, providing a higher brain capacity for processing information and performing cognitive functions [119, 153, 289].

2.3 Mapping the Brain

Understanding the role of connectivity in brain function is the key to reveal the neural mechanisms facilitating the observed human behaviour [233]. The foundations of brain mapping were set in the nineteenth and twentieth centuries, by neuroscientists like Ramon y Cajal and Carl Wernicke, who emphasised the importance of neuronal circuits in understanding the func-

tional organisation of the brain and provided references towards underpinning the anatomical wiring of the brain systems [287, 120]. Invasive methods of brain research, such as histological staining and tracking probes, provided avenues for early structural mapping efforts and new insights towards understanding the circuitry of neurons [120]. With the development of MRI, it became possible to investigate connectivity within the living brain, as MRI allows to map the anatomical (structural) pathways that underlie human brain function at near millimetre resolution, as well as to capture functional interactions between different cortical areas.

These advances have collaboratively led to the emerging field of *connectomics*, which can be defined as the scientific efforts for capturing, mapping, and analysing connectivity in the brain. The ultimate aim of connectomics is to produce the 'human connectome', a comprehensive map of the brain's neural organisation [233] that can be produced and studied at different scales, ranging from single neurons (microscale) to populations of neuronal units (mesoscale) to high-level systems such as distinct brain regions and functional networks (macroscale).

At the microscale, the connectome refers to a complete map of all neural units and synapses. However, a neuron-by-neuron mapping of the human brain is currently not feasible, due to the amount of neural units comprising the brain. The human cerebral cortex alone contains an estimated number of 20 billion neurons, each of which is connected to thousands of other cells, yielding an immensely complex network with trillions of synaptic connections [14]. Considering current limitations in imaging technologies and analysis tools, the microscale connectome is not likely to be derived in the near future.

At the mesoscale, the connectome is composed of neuronal populations, the so-called 'local processing units', that are formed by linking hundreds or thousands of individual neural cells [120]. Several recent animal studies, with a particular focus on mice [182] and rats [37], have revealed how the brain is wired at this scale [233]. However, such methods currently rely on invasive probing techniques, and hence, cannot be used for mapping the human brain.

The macroscale connectome refers to functional interactions between brain regions and the anatomical pathways that connect them to each other. Among different modalities used to visualise the macro connectome, MRI is by far the most common technique due to its widespread

availability, safety and spatial resolution [55]. Diffusion MRI and functional MRI are widely used to capture the brain's structural and functional organisation, respectively. The former allows the reconstruction of anatomical pathways that interconnect different cortical areas in the brain, while the latter enables the delineation of functional connections between spatially remote brain regions (either at rest or while the subject is performing a task).

Analysis of structural and functional connectivity provides two complementary views of brain mapping and enables the identification of brain's functional and anatomical circuitry [231]. Studying their evolution through ageing and their variation across individuals can enable the discovery of biomarkers for various neurological diseases, such as autism and Alzheimer's disease, and may ultimately help develop more effective diagnosis and treatment techniques for brain disorders [77]. Such investigations are made possible through network analysis, in which the connectome is represented as a network (or graph) [231]. At the macroscale, the network nodes correspond to the distinct cortical areas (or regions of interest, ROIs) and the edges represent the connections (axonal projections or functional interactions) between them. Definition of the network nodes through parcellation of the cerebral cortex lies at the heart of the macro connectomics, as errors at this stage can propagate into the subsequent stages and consequently reduce the reliability of network analysis. This motivates the generation of connectivity-driven parcellations for node identification, which constitutes the primary focus of this thesis. Fundamentals of brain parcellation and different methods to derive connectivity-based ROIs for network analysis will further be discussed in the following chapter.

2.4 Imaging the Macro Connectome

2.4.1 Magnetic Resonance Imaging

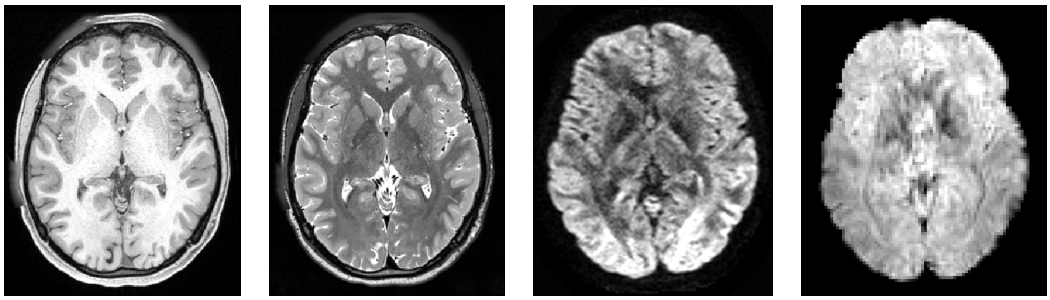
MRI makes use of strong magnetic fields, radio waves, and field gradients to generate detailed images of internal body structures, including but not limited to the heart, the liver, breasts, as well as the brain and other organs. A major advantage of MRI over other imaging techniques,

such as Computer Tomography, Positron Emission Tomography (PET) or X-Ray Imaging, is that it does not involve ionising radiation and is therefore relatively safe and allows repeatable scans. On the downside, it may cause discomfort, as MRI scans require patients to stand still inside a narrow tube for a fairly long period of time and to wear headphones for reducing the loud high-pitched noise caused by the scanner.

MRI is based on the physical phenomenon of Nuclear Magnetic Resonance, in which atomic nuclei absorb and emit the energy from radio frequency waves if undergo an external magnetisation. The water and fat molecules in the human body naturally consist of hydrogen atoms, which have their own magnetic fields and inherently spin around their axis at random orientations. An MRI scanner generates a strong magnetic field B_0 that causes protons in hydrogen nuclei to align with the direction of the B_0 field, precessing around their axis and generating a longitudinal magnetisation. A pulse of radio frequency (RF) is directed at these precessing protons to temporarily create a second magnetic field perpendicular to B_0. This process excites the nuclear spin, bringing some of the nuclei into transverse phase coherence with each other.

Once the RF pulse is removed, the nuclei lose their magnetisation and reach to an equilibrium, that is, recover their initial state with respect to the magnetic field B_0. The energy transition in the process of realignment (i.e. relaxation) can be measured with a coil and used to generate images. Intensity values in these images result from the different concentration of hydrogen protons in different types of tissue and are characterised by two relaxation factors: T1 (longitudinal) relaxation, which is the realignment of nuclei with the direction of the external magnetic field B_0, and $T2$ (transverse) relaxation, which can be defined as the lose of coherence among the nuclei, resulting in decay of transverse magnetization.

By changing the parameters of the MRI pulse sequence, different contrast images can be generated based on the T1 and T2 relaxation factors, e.g. T1-weighted (T1w) and T2-weighted (T2w) images, which can distinguish between different types of tissue. Example T1w and T2w images are presented in Fig. 2.2.



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

Figure 2.2: Different types of MR images. (a) T1w image, (b) T2w image, (c) diffusion weighted image, and (d) functional MR image. All images are acquired from a healthy adult brain.

2.4.2 Diffusion Magnetic Resonance Imaging

Diffusion Magnetic Resonance Imaging (dMRI) or Diffusion Weighted Imaging (DWI) for short, is a form of MRI that relies upon the diffusivity of water molecules for generating contrast in MR images. Water molecules tend to move with random replacements (i.e. microscopic movements) in a free (unrestricted) environment. However, molecular diffusion in biological tissues is naturally impeded by physical boundaries, such as cell membranes, fibers, and macromolecules. For example, in the brain, the motion of water molecules is restricted within the neural axons. Water tends to diffuse more rapidly in the direction aligned with the axon's fibrous structure and more slowly in other directions. These anisotropic patterns of diffusion can reveal the microscopic architecture of the underlying tissue, and with the help of computational modelling, can be used to delineate white matter fibers that interconnect different regions in the brain.

In order to capture the displacements of water molecules in the tissue, MRI must be sensitised to diffusion. The principles of diffusion weighting were introduced by Stejskal and Tanner in 1965 [235], in which the homogeneity of the magnetic field is linearly altered by using two gradient pulses with the same magnitude but opposite gradient direction, yielding nuclei to precess at different rates. Assuming that the nuclei have not moved in between the pulses, the MR scanner will measure the same signal, as the two pulses will cancel each other. For example this can be observed within the ventricles, where the diffusion is typically unrestricted. On the other hand, if the diffusion is hindered, for example by the myelin sheath, the position of the nuclei will likely be different, consequently leading to a signal loss. This reduced signal

is reflected in the DW image as an estimate of the water displacement within each voxel.

One critical parameter of DWI is the b-value, which determines the degree of diffusion weighting on the generated MR image. The choice of b-value is usually application-dependent and needs to be tuned based on the task at hand, as it involves a trade-off between sensitivity and signal-to-noise ratio (SNR) (i.e. higher b-values yield sharper contrast diffusion images, but also lower SNR) [117].

The structure of myelinated axonal fibers restricts the motion of water molecules, leading to anisotropic diffusion patterns that are not equal in all directions. Applying diffusion gradients in at least 6 non-collinear directions enables the definition of a diffusion tensor, from which local diffusion patterns can be characterised for each voxel. A diffusion tensor provides quantitative information about the degree of diffusion anisotropy and directions of the fibers through spectral decomposition, and hence, facilitates the delineation of axonal pathways with the help of computational techniques, such as tractography (see Section 2.5.1).

One major drawback of the diffusion tensor imaging (DTI) model is its assumption that the fiber orientation is uniform at each voxel. However, axonal fibers are known to cross each other (i.e. crossing fibers), temporarily merge with one other, or disperse (i.e. fan out) as they approach their destinations, which may lead to heterogeneity in single voxels that can not be accounted for by simple fiber orientation models, such as the diffusion tensor [55]. Current estimates show that up to 90% of the voxels in WM may contain 'crossing fibers', suggesting that DTI may not be able to detect axonal pathways in the majority of WM [113]. More sophisticated diffusion modelling techniques can account for this heterogeneity by using higher order models [165], commonly referred to as High Angular Resolution Diffusion Imaging (HARDI) [252]. HARDI approaches are able to resolve multiple orientations [252] by relying on a more complex fiber orientation model, e.g. many more gradient directions, compared to the diffusion tensor. As a result, they provide more precise modelling for the heterogeneous distributions of crossing-fiber orientations [278, 55]. Two well-known HARDI methods are Q-ball imaging [251] and Diffusion Spectrum Imaging [277], which are currently being used by the major human connnectome projects to generate more accurate mappings of the human brain [249, 261].

Preprocessing for dMRI

The most commonly observed artefacts in DW images are caused by spatial distortions related to motion (e.g. head motion, respiration, cardiac pulsation etc.), magnetic field inhomogeneities, and eddy currents [138]. The latter emerges from rapid switching of gradient pulses during acquisition, yielding to different field gradients than the ones initially programmed. Spatial distortions introduced by eddy currents can be minimised on the acquisition time, for example by using 'self-shielded' gradient coils [138], or by intentionally altering the shape of the currents sent to the gradient hardware to account for expected eddy-current distortions [186].

The local variations in the magnetic field, usually emerging from the interactions between areas with different magnetic susceptibility (e.g. air-tissue interfaces) can also contribute to spatial distortions [55]. Impact of such artefacts can be reduced with the use of parallel imaging techniques or corrected via mathematical modelling of the field variations [7, 55].

Patient motion typically leads to ghosting effects and large signal variations in the images [138]. Although various techniques can be used to minimise the impact of motion in dMRI, its complete elimination is not possible without anaesthesia or using an MRI sequence that is less prone to motion artefacts [138]. This is achieved by echo-planar imaging (EPI) [201], which has now become a golden standard in dMRI due to its ability to effectively 'freeze' the motion during acquisition [138]. However, EPI often provides less spatial resolution than conventional sequences due to hardware limitations and may introduce its own artefacts, e.g. blurry images [116]. In general, a post-acquisition solution to correct diffusion artefacts involves registering the distorted images to the b_0 image, i.e. the image with no diffusion weighting [228].

2.4.3 Functional Magnetic Resonance Imaging

Functional Magnetic Resonance Imaging (fMRI) is a non-invasive, *in-vivo* imaging technique for capturing dynamic changes in neurocognitive activity associated with cerebral blood flow [104]. The technique is based upon the fact that the amount of blood that is circulated within the cerebral cortex and neural activity regulating mental processes are closely linked. That is,

when a brain area is involved in neural activity, blood flow towards the active area increases in response to high energy demand.

Among others, the most popular form of fMRI uses Blood Oxygenation Level Dependent (BOLD) contrast, which measures the inhomogeneities in MRI signal due to the level of oxygen in blood [180]. BOLD reflects the change in heamodynamic response, which facilitates the rapid delivery of blood to the active neuronal tissue to supply their demand for oxygen and other nutrients, such as glucose [181, 44]. Haemoglobin is a protein molecule in red blood cells that is responsible for carrying oxygen from the respiratory organs to the tissues and shows different magnetic susceptibility depending on its bond with oxygen. While oxygen-carrying haemoglobin (oxyhaemoglobin) behaves as a diamagnetic substance, its oxygen-free form is typically paramagnetic. When there is a demand for high energy, for example in case of neural firing, oxygen is transferred from oxyhaemoglobin to neurons, making the blood de-oxygenated, and hence, increasing the level of deoxyhaemoglobin. In order to compensate for the greater amount of energy demand, vascular system increases the blood flow, changing the relative level of oxyhaemoglobin and deoxyhaemoglobin in favour of the former. The difference in magnetic susceptibility of haemoglobin molecules facilitates the detection of the increase in blood flow via MRI, as areas with high concentration of oxyhaemoglobin produces a higher signal than areas with low concentration [5].

Depending on the problem under investigation, two types of fMRI can be employed, each of which provides information about different aspects of brain activation and functional connectivity. The first one is task fMRI (t-fMRI), in which neural activity is recorded while the subject is performing a variety of cognitive tasks designed to activate different brain regions [183]. The primary motivation of t-fMRI studies is to examine the brain under a controlled setting, whose limits are delineated by a psychological paradigm. These paradigms may target different aspects of cognition, such as primary sensory processing, information processing, decision making, problem solving, and many more [16]. Data collection in t-fMRI is typically followed by a statistical analysis stage in which image intensities are compared to task paradigm in order to reveal functional activation [222, 5].

While localisation of function constitutes the primary purpose of t-fMRI studies, more recent approaches also attempt to determine the mappings between certain mental operations and patterns of neural activity [179]. These so called 'mind-reading' methods combine data representation techniques with pattern recognition tools in different applications, including but not limited to lie detection [57], object recognition [183], and human behaviour prediction [162].

The second type of fMRI is known as resting-state fMRI (rs-fMRI), as it captures neurocognitive activity while the subject is 'at rest', i.e. while the brain is not driven by any external stimuli. In other words, the subject is not requested to perform a neuropsychological task, but told to relax with eyes are shut or fixated to an object. In 1995, Biswal et al. showed that the brain is still active in the absence of an externally prompted stimulus, hence spontaneous fluctuations in BOLD signals can be used to identify functional interactions between different brain regions [30]. Since its discovery, rs-fMRI has been the primary tool to explore the brain's functional organisation [150, 212, 79, 56, 177, 152, 227] as well as to study changes/alterations in functional connectivity due to neurological and psychiatric disorders [234, 275, 78, 77].

It is noteworthy that rs-fMRI allows to explore the brain's functional organisation as a whole, without being biased by a neuropsychological paradigm. On the other hand, t-fMRI only targets certain brain areas in order to investigate their activation in response to external stimuli, but ignores the activation from non-target areas. With regards to this, we focus on rs-fMRI to estimate functional connectivity in the context of brain parcellation, while information derived from t-fMRI is used as a complimentary measure to validate the location of functional areas identified by rs-fMRI based parcellations.

Preprocessing for fMRI

BOLD fMRI signals are inherently confounded by several potential noise sources, such as head movement, cardiac/respiratory pulsation, or scanner-induced artefacts [97, 51, 200]. This may lead to a relatively low SNR and pose major challenges for the usability of fMRI data and the interpretability of subsequent analysis. As a result, several preprocessing stages are required to remove noise and correct for artefacts in the acquired BOLD signals prior to any attempt to

analyse the data. Noise removal is particularly important for functional connectivity studies, since spurious correlations induced by structured noise may severely increase the amount of falsely identified connections [51]. Standard preprocessing steps generally include slice-timing correction, head motion correction, distortion correction, temporal filtering, coregistration, spatial normalisation, and spatial smoothing.

Slice-timing correction is used to realign individual slices in an fMRI brain volume to a reference slice based on their relative timing, as each slice typically records activity at a slightly different time point, leading to a between-slice temporal offset [220].

Head motion is one of the most commonly conferred sources of noise in fMRI. Artefacts induced by motion can lead to mislocalisation of function, activation being detected outside the brain volume, or artefactual fMRI signals [55]. It is typically corrected by applying rigid-body registration to a reference volume (such as the first volume) or by identifying structured noise related to motion with the use of statistical analysis techniques, such as independent component analysis [200].

Artefacts related to distortion emerges from inhomogeneities in the magnetic field and can be corrected by aligning the functional image with a structural image or by acquiring two fMRI images with different echo time, in which distortion correction is facilitated by mapping the spatial distribution of non-uniform areas in the magnetic field [106]. Alternatively, bias field estimation techniques can be utilised to correct for distortion, in case the magnetic field distribution is not known *a priori* [95].

Temporal filtering is the removal of frequencies from the signal that are not of interest, which may be induced by respiratory and cardiac pulsation. BOLD signals related to neuronal activity are typically exist in a particular frequency range, usually located between 0.01 and 0.1 Hz [248]. A band-pass filtering is hence applied to surpass spurious frequencies and only focus on the neurobiologically meaningful range.

Prior to analysis of fMRI data, BOLD signals are registered to a structural MR image of higher resolution of the same subject (i.e. coregistration), so that functionally distinct regions or

activated voxels can be mapped into anatomical space. This is generally followed by spatial normalisation to a common space, which could either refer to a volumetric brain template, e.g. Montreal Neurological Institute (MNI) template [159], or a surface template, e.g. FreeSurfer's fsavereage surface [74]. This brings individual subjects to a standard anatomical basis, and hence, facilitates the integration of results across multiple subjects, populations, as well as different analysis pipelines.

Even by assuming a perfect coregistration between the anatomy and function, a reliable functional alignment across subjects cannot be guaranteed [243]. A typical solution to this problem is to sacrifice spatial resolution for group-level analysis, i.e. applying spatial smoothing to BOLD signals, which is generally achieved through convolution with a Gaussian filter. Such a process not only alleviates the impact of misalignment across subjects (e.g. inter-subject variability), but also improves SNR in individual subject data, if the spatial extent of activation matches the width of the filter used [144]. On the downside, smoothing typically reduces the overall resolution in fMRI and may introduce blurring in the group data [243, 144]. As a result, activation in relatively small areas may be mislocalised or may even completely disappear depending on the amount of spatial smoothing [144].

Regardless of these artefacts, BOLD fMRI constitutes the current state of the art for imaging the brain activation, with its non-invasive nature and relatively high spatial resolution over alternative electro-physiological recordings, such as electroencephalography (EEG) and magnetoencephalography (MEG), which, despite offering a high temporal resolution (typically at the order of milliseconds), suffer from poor spatial resolution and lack of spatial localisation [147]. Another limitation of BOLD fMRI is the fact that it only provides an indirect measure of the neural activity (i.e. through change in blood flow), while EEG and MEG directly record neural activity in the form of electrical impulses from the brain. In addition, the true biological meaning of BOLD signals is still under investigation [66] and its lower temporal resolution (typically at the order of seconds) compared to electro-physiological recordings is also a limiting factor for capturing high-frequency patterns.

2.5 Structural and Functional Connectivity

2.5.1 Estimating Structural Connectivity

Structural connectivity describes the anatomical pathways linking distinct brain areas, which generally refers to WM projections interconnecting cortical and subcortical regions [232]. These connections form the biological basis for information transfer between remote brain regions and are therefore fundamental for mapping the structural human connectome, and relatedly, a better understanding of brain function [233, 248].

Reconstructing Pathways through Tractography

As described in the previous section, dMRI facilitates the mapping of the diffusion process of water molecules, and hence, provides information about the location, orientation, and anisotropy of axonal fibers through different diffusion models. With this information available, tractography (fiber tracking) techniques can be utilised to reconstruct trajectories of the major pathways in the brain, from which structural connectivity patterns can be quantified [167, 24]. Tractography estimates the trajectories by tracking the streamlines through a 3D vector field, in which each vector represents the diffusion orientation measured by dMRI [165]. The majority of the tractography methods can be broadly divided into two categories, depending on how they reconstruct white matter tracts from the underlying diffusion model: i.e. deterministic tractography and probabilistic tractography.

Deterministic tractography relies on the local diffusion information for integrating the streamlines on a step-by-step basis [24]. Initiated from a seed (a voxel or a region of interest), streamlines are reconstructed with respect to the primary diffusion direction at each voxel. Due to the fact that the directional information is encoded on an imaging grid, each voxel provides only one measurement of orientation. Hence, interpolation methods are essential to transform these discrete measurements into a continuous coordinate system, allowing to estimate the fiber orientation at locations away from the voxel centres [24]. This could be achieved, for example,

by applying each voxel's measurement over the entire voxel [166] or using a weighted interpolation of the fiber orientation by incorporating measurements from neighbouring voxels [18, 137]. Regardless of the interpolation technique, tracking algorithms continue to delineate streamlines until a termination criterion is reached. Different approaches to terminate this process may include using a white matter mask to constrain the search space of the tracking algorithm or to define heuristic rules to stop the streamline from progressing, for example when the computed anisotropy at a voxel falls beneath a user-defined threshold [24].

The primary intuition behind relying on such a heuristic to stop the tracking algorithm is due to the fact that deterministic tractography does not provide a means to model uncertainties inherently associated with each voxel along a streamline. *Probabilistic tractography*, by contrast, aims to handle the uncertainties so that tractography can keep tracking streamlines through regions of high uncertainty, in which deterministic techniques would be forced to terminate. [24]. Errors propagating as tractography progress, for example due to imaging noise or crossing fibers, have necessitated the use of probabilistic tractography in order to achieve a more accurate reconstruction of axonal pathways [23, 24]. Probabilistic tractography techniques typically start with characterising the uncertainty of the fiber orientation at each voxel with a probability density function [25]. This function allows any streamline to follow an infinite number of orientations, each of which is assigned a different level of probability. Hence, once a voxel is reached, the next orientation is sampled from the probability distribution associated with that voxel. By drawing many such samples (each time providing slightly different orientations), an entire tract connecting two points can be reconstructed [24].

A major benefit of probabilistic tractography, apart from its ability to quantify uncertainty, is its robustness to noise [25, 24]. Paths progressing towards wrong routes due to noisy voxels are likely to produce low probability values, and hence, are quickly dispersed [25, 24]. Although probabilistic tractography does not require the definition of a termination criterion, a large curvature threshold is typically employed to prevent streamlines tracing back to where they start, a necessary mechanism to avoid an artificial increase in probability values and generation of implausible pathways [24]. One major drawback of probabilistic tractography techniques is the computational cost, especially when used to estimate whole-brain connectivity. However, recent

advances in GPU-based parallelisation techniques have drastically reduced the computational time required to perform probabilistic tractography for a typical spatial resolution [99].

A major challenge for tractography techniques, either deterministic or probabilistic, emerges from their inability to detect fibers in their entirety [55]. Although, tracking algorithms can successfully estimate the location of fibers within the white matter, the diffusion orientation is typically highly ambiguous around WM-GM interface, resulting in the termination of tractography before reaching the GM [172] and making it impossible to accurately identify the origin or termination of axonal tracts [172, 55].

Quantifying Structural Connectivity

Tractography techniques facilitate the delineation of trajectories connecting cortical areas through WM and primarily provide information regarding the density and location of connections in the brain [55]. However, reconstructed tracts do no inherently allow to quantify the 'connection strength' between two cortical regions in the context of brain connectivity [110]. Rather than, estimates of connectivity can be acquired through some quantitative measurements by utilising different aspects of the diffusion data and reconstructed tracts. While connectivity estimates based on the latter include the number [96] or the total length [250] of reconstructed streamlines between pairs of brain regions, measures driven from diffusion properties of the underlying data usually rely on the computed anisotropy along white matter tracts [110]. Steps leading to structural connectivity analysis through a typical dMRI pipeline are illustrated in Fig. 2.3.

Although probabilistic tractography provides a quantitative means for the uncertainty in streamline trajectories, it is still not possible to acquire a true quantitative measurement of the connectivity [24]. Structural connectivity from probabilistic tractography is typically estimated by counting the number of streamlines passing through a region and dividing it by the total number of streamlines [24]. However, the reliability of connections obtained in this context is limited by several confounding factors. For example, it is typically assumed that brain regions connected via a major bundle should have a noticeable trace in the diffusion data, and hence, low uncertainty (high confidence) in their trajectories; but, probabilistic tractography is likely

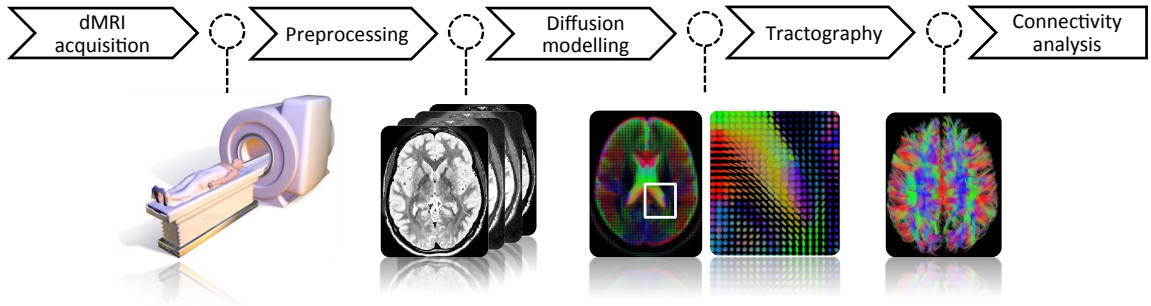


Figure 2.3: Illustration of a typical dMRI pipeline. Images of the MRI scanner and diffusion tensor model are from Wikimedia Commons, referenced as [283] and [282], respectively.

to assign a higher confidence to a locally non-dominant fiber bundle than such a major bundle, in case it is crossed by other fibers [55]. Similarly, uncertainty along a streamline's path typically increases with the tract length. Other non-interesting factors, such as imaging noise, modelling errors, partial volume effects, and voxel size also affect uncertainty, and hence, may influence the estimation of structural connections [24, 55].

Despite all these limitations and many other aforementioned confounding factors affecting the acquisition and modelling of diffusion data, dMRI and tractography are indispensable for brain mapping, as together they provide the only non-invasive way of measuring (or estimating) the structural connectivity within the brain.

2.5.2 Estimating Functional Connectivity

Functional connectivity is conventionally defined as the temporal dependency of neurophysiological events between spatially remote brain areas [83]. It is used to identify co-activation between different regions that share similar functional characteristics and/or work together to perform cognitive processes. Compared to its structural counterpart, functional connectivity is not necessarily supported by physical connections through WM. Similarly, anatomically connected regions (as defined via tractography) also do not need to reflect functional dependency [96].

Several mathematical modelling techniques can be used to quantify functional connectivity, each providing a different perspective on the temporal interactions between BOLD signals measured

in brain areas [55]. In the simplest form, bivariate tests, such as Pearson's correlation [30] or coherence [237], can be employed between every pair of timeseries to measure their statistical dependence. A limitation of bivariate approaches emerges from the fact that they do no account information from multiple regions [55], and hence, cannot distinguish direct from indirect connections that may be mediated by third-party regions [226, 66]. Whereas such distinction would not be of great importance when correlation is used to measure the distance (similarity) between two BOLD signals, for example in the context of connectivity-driven brain parcellation, it may be a critical consideration point for network analysis [227]. To this end, partial correlation can correctly estimate the conditional linear dependency between two regions, as it typically accounts for interaction with every single region of interest [158]. While the reliability of partial correlations decreases when the number of observations exceeds the sample size (the number of regions) [196], such as in the case of typically long fMRI acquisitions, regularisation techniques (e.g. graphical lasso [80]) can overcome this limitation [227].

Although more complex techniques exist to measure functional connectivity, such as methods based on higher-order statistics or lag-based approaches [227], results derived from large-scale simulations indicate that correlation-based approaches generally yield more accurate connectivity estimates [227]. Various data-driven approaches can also be used to study functional connectivity, including but not limited to principle component analysis (PCA), independent component analysis (ICA), and clustering techniques. However such methods in general are more appropriate for identifying spatially distributed networks that are functionally connected during rest (i.e. resting state networks) or deriving nodes in network analysis (for parcellating the brain), and therefore, will be extensively covered in the next chapter. Steps leading to functional connectivity analysis through a typical fMRI pipeline are illustrated in Fig. 2.4.

While functional connectivity can generally estimate existing connections between brain regions with good accuracy, it typically suffers from false positives (interestingly, structural connectivity is, by contrast, more prone to false negatives) [66]. False positives usually emerge due to the low SNR inherent in fMRI data, induced by imaging errors, head motion, and physiological noise [55]. Special care is taken to reduce the impact of such artefacts via a series of preprocessing steps prior to data analysis as previously discussed (see Section 2.4.3); yet, it is not possible

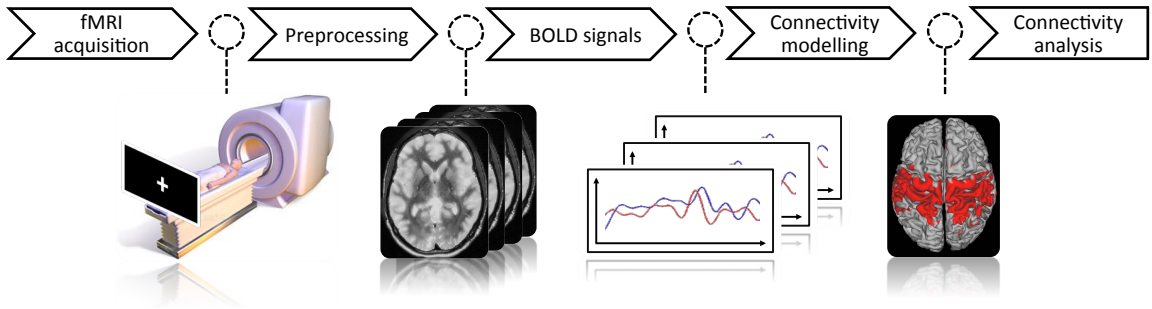


Figure 2.4: Illustration of a typical fMRI pipeline. MRI scanner image is from Wikimedia Commons [283].

to completely eliminate the inaccurate BOLD signal fluctuations. Some techniques aim to alleviate the impact of false positives by using thresholding mechanisms, with the assumption that negative or weak correlations correspond to spurious activity or noise [255, 245, 55, 132].

2.6 Imaging Data

In our experiments throughout this thesis, we rely on the imaging data provided by the Human Connectome Project (HCP) [261]. In the remainder of this section, we first introduce the HCP and then briefly explain the HCP data acquisition and preprocessing pipelines [87]. We finally summarise the two datasets used in different chapters of this study.

2.6.1 The Human Connectome Project

HCP (`https://www.humanconnectome.org/`) is one of the recent scientific efforts to map the human brain [261]. Led by Washington University, University of Minnesota, and Oxford University (the WU-Minn HCP consortium), it aims to chart the neural circuitry of the human connectome in healthy adults by using non-invasive imaging technologies. HCP provides invaluable information to explore neural mechanisms that underlie the brain function and behaviour, which in turn, contributes to our understanding of the human mind.

Since its first data release, the HCP datasets of almost 900 healthy adults have been made freely available to the scientific community via the HCP Database, `https://db.humanconnectome.`

org/. The project focuses on four imaging modalities to acquire data with high spatial and temporal resolution [156]. Resting-state fMRI and diffusion MRI respectively provide information about functional and structural brain connectivity and constitute our primary data sources for developing connectivity-driven parcellation algorithms. Task-evoked fMRI reveals much about brain function with respect to various cognitive tasks designed to activate different brain regions and is used as a complementary information source to evaluate the proposed parcellations. Structural MRI captures the shape of the cerebral cortex and subcortical brain areas as well as provides the anatomical (both volumetric and surface-based) templates, allowing the integration of function and anatomy and analysis of results across multiple subjects.

Prior to HCP, several other projects, including but not limited to the 1000 Functional Connectomes Project [31] and the Human Connectome Project led by the UCLA-Harvard consortium [249] also worked with a similar motivation and enriched the understanding of the human brain with their contributions to the field of macro connectomics. However, HCP has taken the flag one step further by producing MRI data in several different modalities that are preprocessed with novel pipelines [87], and thus, offering ready-to-analyse high-resolution structural, functional, and diffusion datasets for studying the human brain from many different aspects. Besides these attempts to map the adult human connectome, the Developing Human Connectome Project, which aims to create the first connectome of early life (from 20 to 44 weeks post-conceptional age), is still at its preliminary stage and is expected to produce its first data release in 2017 [1].

2.6.2 HCP Minimal Preprocessing Pipelines

HCP brings multiple MRI modalities together in a common framework, allowing to perform cross-subject comparisons and multi-modal analysis of brain architecture, connectivity, and function [87]. This is achieved through a series of newly developed preprocessing methods, tailored to many modality-specific challenges faced in both the acquisition and analysis of structural, functional, and diffusion images. Throughout this section, we briefly summarise the key points in the HCP preprocessing pipelines, particularly putting emphasis on the HCP

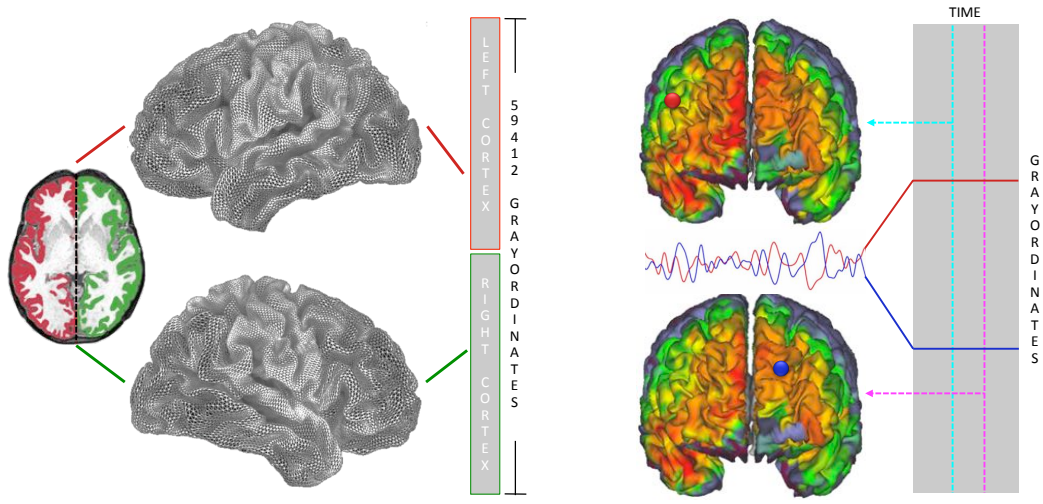
standard cortical space, i.e. the so-called 'grayordinate' spatial coordinate system, as well as data acquisition and preprocessing steps for the structural, fMRI and dMRI data. It is worth noting that, this section is compiled from the HCP minimal preprocessing pipelines paper [87] and is up-to-date with the reference manual of HCP 900 subjects data release (S900) [105].

Grayordinate Standard Coordinate System

Since we attempt to parcellate the cerebral cortex, it is beneficial to use cortical neuroimaging data and surface-constrained methods, rather than relying on volume-based data. Due to the fact that the highly convoluted cortical sheet can be easily analysed as a 2D surface, we can further utilise geodesic distances along the surface, which can be more relevant than 3D Euclidean distances within the volume [259, 227, 87, 93]. Relying on the geometry of the cerebral cortex further allows a more effective spatial smoothing and inter-subject registration with greater overlap [87].

HCP datasets are provided in a standard coordinate space obtained by mapping the gray matter voxels onto cortical sheet, that allows combining the left and right cerebral hemispheres into a single file[1]. This grayordinate space does not only reduce the computational and storage demands of high resolution MRI data, but also enables a more compact data representation. All gray matter voxels in the cerebral cortex apart from the medial wall are represented at a 2 mm resolution using 59412 grayordinates, including 29696 and 29716 cortical vertices from the left and right hemispheres, respectively (Fig. 2.5). Data files in the grayordinate system typically contain a single 2D matrix, in which the x axis always represents indices of the standard set of grayordinates, while the y axis may represent, for instance, timepoints as illustrated in Fig. 2.5. The grayordinate space achieves spatial correspondence across different subjects, and thus, facilitates straightforward cross-subject comparisons.

[1]It is worth noting that, the original grayordinate space also includes gray matter data from subcortical regions, but since we focus on the cerebral cortex only, we provide information for a simplified version of the grayordinate space with an emphasis on the cerebral hemispheres and omit details from elsewhere in the brain.



(a) (b)

Figure 2.5: (a) Grayordinate standard coordinate system. Each cerebral hemisphere roughly contributes about 30k cortical vertices, achieving a total of 59412 grayordinates corresponding to all of the gray matter voxels apart from the non-cortical medial wall. (b) A standard timeseries file in the grayordinate space. Each row represents a grayordinate and the columns record fMRI time points. Each row corresponds to a timeseries mapped to a cortical vertex, as shown in red and blue. Reading along a column produces a spatial pattern, which corresponds to the recorded BOLD signals at a given timepoint across the whole cerebral cortex. Spatial patterns at different timepoints (columns) are shown in cyan and magenta.

Structural Pipelines

The HCP structural acquisitions provide T1w and T2w images at 0.7 mm isotropic resolution. This high resolution facilitates the creation of more accurate cortical surfaces and myelin maps compared to conventionally lower resolution MRI data. Please see [105] for a list of all imaging parameters used in structural MRI acquisitions.

The structural pipelines consist of three stages. The first pipeline produces a 'native' structural volume space for each subject, which represents the best approximation of the subject's physical brain. After aligning T1w and T2w images, each native volume is registered to MNI space. The second structural pipeline includes 1) anatomical brain segmentation, 2) reconstruction of white and pial cortical surfaces, and 3) cortical-folding based inter-subject registration to FreeSurfer's standard surface atlas (fsaverage) [72] using a novel multimodal surface matching algorithm [205]. Finally, the third structural pipeline produces all of the volume and surface files

required for visualisation purposes, along with 1) applying the cortical folding-based registration to the Conte69 surface template [259], 2) downsampling registered surfaces to 32k standard cortical mesh, in which connectivity analysis can be performed, 3) creating the final brain mask, and 4) creating myelin maps, which are computed as the ratio of T1w and T2w intensities [88].

Functional Pipelines

Resting-state fMRI data for each subject was acquired in two sessions, divided into four runs of approximately 15 minutes (1200 timepoints) each. The sessions were held on different days and EPI phase encoding was applied in a right-to-left (R-L) direction in one run and in a left-to-right (L-R) direction in the other run for minimising distortion and blurring. During the scans, the subjects were presented a fixation cross-hair, projected against a dark background, which prevented them from falling asleep. The functional data was acquired at a spatial resolution of 2 mm isotropic, allowing an accurate mapping of gray matter BOLD signals onto the cortical sheet [87]. Please see [105] for a list of all imaging parameters used in rs-fMRI acquisitions.

The pipeline for the rs-fMRI data consists of two stages, which are referred to as 'fMRIVolume' and 'fMRISurface' in the HCP minimal preprocessing paper [87]. fMRIVolume briefly involves the following steps: 1) spatial distortion correction, 2) realignment of the timeseries to correct for subject motion, 3) EPI distortion correction using FSL's 'topup' tool [7] and registration to the T1w image, 4) bias field correction, 5) normalisation of BOLD signals to a global mean of 10000, and 6) masking the data with the final brain mask. The second stage, fMRISurface, is employed to bring the 4D volume data to the grayordinate standard space. This is achieved by first mapping the gray matter voxels onto the native cortical surface and then transforming them onto the 32k standard triangulated mesh (Conte69) [259], using the cortical folding-driven registration's deformation field. The outcome of this pipeline is a standard set of timeseries for every subject with spatial correspondence, at a spatial resolution of 2 mm average surface vertex spacing. Excluding the non-cortical medial wall vertices, each cortical hemisphere is represented by around 30k vertices (the same number in each subject) in the grayordinate space as illustrated in Fig. 2.5(b). The fMRI timeseries are then slightly smoothed (2 mm FWHM)

on the surface to match the vertex spacing of the standard 32k mesh. The preprocessed data is cleaned of structured noise with ICA-FIX [21, 211], an ICA-based tool that automatically removes artefactual components from rs-fMRI data. Following these preprocessing and denoising steps, every single timeseries is temporally normalised to zero-mean and unit-variance.

Following completion of rs-fMRI acquisitions in each of the two scanning sessions, subjects were asked to complete a battery of tasks designed to activate different brain regions. Seven major domains were assessed in order to cover a wide range of neural systems, including 1) visual, motion, somato-sensory, and motor systems (MOTOR), 2) category specific representations (GAMBLING), 3) working memory, cognitive control systems (WM), 4) language processing (LANGUAGE), 5) social cognition (SOCIAL), 6) relational processing (RELATIONAL), and 7) emotion processing (EMOTION). These physiological paradigms are described in detail in [16]. Two fMRI scans were collected for each task. Similar to rs-fMRI acquisition, one scan was acquired with R-L phase encoding direction and the other with the opposite phase encoding direction (L-R).

Task activation maps for each subject/task were computed using FSL's FEAT, a data-analysis tool based on general linear modelling [16]. Given the experimental design, FEAT creates a model that best fits the BOLD signals, and provides a statistical map that shows the activated brain areas in response to the stimuli. The analysis is carried out across sessions to obtain activation maps for each subject, which in turn, are used for evaluation purposes throughout this thesis.

Diffusion Pipelines and Tractography

The whole set of DW volumes for each subject was acquired in six separate series (runs) and grouped into three pairs, each representing a different gradient scheme. The paired two runs include the same DW directions, but with reversed phase-encoding (i.e. alternating between R-L and L-R directions in consecutive runs). Each gradient scheme included approximately 90 diffusion weighting directions and six b_0 images. Diffusion weighting consisted of three shells of $b = 1000$, 2000, and 3000 s/mm^2. Approximately the same number of images were acquired for

each $b-$value within each run. DW images were collected with a spin-echo EPI sequence at a 1.25 mm isotropic resolution, using a Stejskal-Tanner (monopolar) diffusion-encoding scheme. Please see [105] for a list of all imaging parameters used in dMRI acquisitions.

The diffusion pipeline briefly consists of the following steps: 1) b_0 intensity normalisation, 2) estimation of distortion field by feeding the R-L and L-R b_0 images to the FSL's 'topup' tool [7], 3) estimation of eddy-current distortions and head motion for each image volume using FSL's 'eddy' tool [6], 4) distortion correction and registration of b_0 image to the T1w image, 5) resampling the distortion-corrected diffusion data from eddy into 1.25 mm native structural space, and 6) masking the data with the final brain mask to reduce the file size. The diffusion gradient vectors are also accordingly rotated and registered to the native structural space [87].

Following these preprocessing steps, diffusion data is fed into FSL's multi-shell spherical deconvolution toolbox (bedpostX) [25, 23, 110] for fiber orientation estimation. The output files generated by bedpostX are used by another FSL tool, probtrackX, to perform probabilistic tractography with respect to the estimated fiber orientations. Probabilistic tracking is performed on the native mesh (before registration) representing the grey/white matter interface (specified by the white and pial surfaces). The algorithm generates probability distributions from user-defined seed vertices. In our case, 5000 streamlines are seeded from all cortical vertices (i.e. each vertex is considered as a seed). The output tractography matrix contains a so-called 'connectivity' value for each vertex, which corresponds to the number of streamlines that pass through that vertex and reach to the rest of the mesh. This matrix is finally used to estimate structural connectivity in our attempts to generate cortical brain parcellations.

2.6.3 Datasets

Our experiments are based on 200 subjects, acquired from the HCP S900 data release [105] and grouped into two datasets of 100 subjects each, which will be hereinafter referred to as Dataset 1 and Dataset 2. All subjects had their scans successfully completed for all imaging modalities considered by the HCP.

Dataset 1 consists of 100 subjects (54 female, 46 male adults, age 22-35), whose demographic characteristics are summarised in Table. 2.1. This dataset corresponds to the 'Unrelated 100 Subjects' in the HCP database. We use Dataset 1 in particular for the development of connectivity-driven parcellations in Chapters 4, 5, 6, and 7.

Dataset 2 consists of 100 randomly selected subjects (50 female, 50 male adults, age 22-35). Summary of the demographic information for all subjects is given in Table. 2.2. We use Dataset

Table 2.1: Summary of the demographic information of Dataset 1.

Subject	Gender	Age	Subject	Gender	Age	Subject	Gender	Age
100307	F	26-30	128127	F	26-30	163129	M	26-30
100408	M	31-35	128632	F	31-35	176542	F	26-30
101107	M	22-25	129028	M	26-30	178950	M	31-35
101309	M	26-30	130013	M	26-30	188347	F	26-30
101915	F	31-35	130316	F	26-30	189450	M	22-25
103111	M	26-30	131217	F	26-30	190031	F	26-30
103414	F	22-25	131722	F	26-30	192540	F	26-30
103818	F	31-35	133019	F	26-30	196750	F	26-30
105014	F	26-30	133928	M	26-30	198451	M	26-30
105115	M	31-35	135225	F	22-25	199655	M	31-35
106016	F	31-35	135932	F	26-30	201111	M	26-30
108828	M	31-35	136833	M	31-35	208226	F	22-25
110411	M	31-35	138534	M	22-25	211417	M	36+
111312	F	31-35	139637	F	31-35	211720	M	31-35
111716	F	31-35	140925	F	22-25	212318	F	31-35
113619	F	31-35	144832	F	22-25	214423	M	22-25
113922	M	31-35	146432	M	31-35	221319	F	26-30
114419	M	31-35	147737	F	31-35	239944	F	31-35
115320	F	31-35	148335	F	31-35	245333	F	31-35
116524	M	26-30	148840	F	22-25	280739	F	31-35
117122	F	26-30	149337	M	31-35	298051	F	31-35
118528	F	26-30	149539	M	22-25	366446	M	22-25
118730	M	22-25	149741	M	26-30	397760	M	31-35
118932	M	26-30	151223	F	26-30	414229	M	31-35
120111	F	26-30	151526	M	26-30	499566	F	26-30
122317	M	31-35	151627	F	26-30	654754	F	31-35
122620	M	26-30	153025	F	22-25	672756	M	31-35
123117	M	26-30	154734	M	31-35	751348	M	26-30
123925	F	26-30	156637	M	31-35	756055	M	22-25
124422	F	31-35	159340	M	26-30	792564	F	31-35
125525	F	31-35	160123	M	31-35	856766	F	31-35
126325	F	26-30	161731	F	31-35	857263	M	22-25
127630	F	22-25	162733	F	26-30	899885	F	26-30
127933	M	31-35						

2 primarily for evaluation purposes in Chapter 7; that is, group-level parcellations are computed from Dataset 1, but evaluated on Dataset 2, so as to reduce the possible bias towards the computed parcellations with respect to the publicly available ones.

Table 2.2: Summary of the demographic information of Dataset 2.

Subject	Gender	Age	Subject	Gender	Age	Subject	Gender	Age
100206	M	26-30	131823	M	26-30	173738	M	31-35
100610	M	26-30	134728	M	26-30	174841	M	22-25
101006	F	31-35	135528	F	26-30	176239	M	31-35
102311	F	26-30	135730	M	22-25	176441	M	26-30
102513	M	26-30	137229	M	22-25	177241	M	26-30
102816	F	26-30	138837	M	22-25	177645	F	26-30
103515	F	26-30	139839	M	26-30	178243	M	22-25
104012	F	26-30	141119	M	26-30	178647	F	31-35
104416	F	31-35	145127	M	22-25	179245	F	31-35
105923	F	31-35	146533	M	26-30	181232	F	26-30
106521	F	26-30	150928	M	22-25	181636	F	31-35
108121	F	26-30	151425	F	26-30	185947	F	26-30
108222	M	31-35	154532	M	26-30	187547	F	31-35
108323	F	26-30	154835	F	31-35	188549	M	22-25
110613	M	26-30	155938	M	26-30	188751	M	26-30
111009	F	26-30	156334	F	26-30	189349	F	26-30
112112	M	26-30	156435	M	22-25	191942	M	26-30
112516	F	31-35	156536	M	22-25	195041	F	31-35
112920	M	31-35	158035	F	26-30	196144	F	31-35
114621	M	26-30	158338	M	22-25	200008	F	26-30
115017	F	31-35	158843	M	26-30	200614	F	31-35
115825	M	22-25	159441	F	26-30	236130	F	26-30
116726	M	26-30	159744	M	31-35	250427	F	31-35
118225	M	26-30	162026	F	31-35	263436	F	31-35
120212	F	31-35	162935	M	22-25	287248	F	26-30
121416	M	26-30	164636	M	22-25	297655	F	31-35
121921	M	31-35	165638	M	26-30	304727	F	31-35
123521	M	22-25	168240	M	26-30	321323	F	31-35
124826	F	31-35	169444	F	26-30	358144	F	26-30
128935	M	31-35	169949	M	26-30	391748	F	31-35
129331	F	26-30	171330	M	26-30	525541	F	26-30
129634	M	22-25	172332	F	26-30	709551	F	31-35
130417	M	26-30	172433	M	22-25	885975	F	26-30
130619	F	26-30						

Chapter 3

Parcellating the Human Brain

Abstract

This chapter defines the concept of brain parcellation and answers the questions such as 'What is a parcellation?' and 'In what terms a parcellation is of use?'. In the light of this, we briefly cover the historical foundations of brain parcellation and review widely-used techniques to subdivide the cerebral cortex (and subcortical nuclei) with respect to different neuro-anatomical properties, with a particular focus on functional and structural connectivity. We then extensively review connectivity-driven parcellation approaches by emphasising on their limitations and advantages over each other. We wrap-up the chapter by describing the challenges for obtaining reliable connectivity-driven parcellations and briefly summarise how we aim to overcome each challenge throughout this thesis.

3.1 Introduction

Subdivision of the brain into anatomically and functionally distinct regions (cortical areas or subcortical nuclei) has been one of the ultimate challenges in the field of neuroscience, as it constitutes a major milestone towards understanding the brain. Cortical areas can be separated

from their neighbours based on different properties such as micro-architecture (i.e. cyto-, myle- and receptor-architectonics) [291, 64], anatomy (e.g. formations of sulci/gyri) [254, 61, 72], function (e.g. involvement in cognitive operations) [249, 243, 244], and connectivity (e.g. a set of extrinsic inputs and outputs) [193, 25, 121, 66]. Brain parcellation has many important roles in neuroscience and particularly in macro connectomics, as it provides 1) a reference map of the brain, enabling the comparison of results across subjects, groups, and studies, 2) a common language for researchers to discuss their findings, 3) a snapshot of the functional and structural organisation of the brain at macroscale, and 4) a means to reduce the complexity of connectivity, an aspect that is highly critical for the study of brain dynamics with whole-brain models [249, 55, 218, 86].

While the first chapter has briefly introduced the rationale behind parcellation, this chapter provides a more detailed explanation of the brain parcellation concept and its implications for neuroscience. Although we refer to different types of neuro-anatomical properties and techniques to segregate the brain, the majority of the chapter focuses on connectivity-driven parcellation. We provide insight towards understanding the role of connectivity in order to define cortical parcellations and identify network nodes for connectome analysis. We further review many different methods developed for parcellating the brain, particularly the cerebral cortex, by using functional and structural connectivity information, with an emphasis on the key points to define a reliable parcellation. Finally, we summarise the challenges that are likely to emerge in the development process of connectivity-driven parcellations.

3.2 Brain Parcellation: Aim and Scope

3.2.1 Historical Foundations of Brain Parcellation

The concept of parcellation spans more than a century in the field of neuroscience. In 1909, Korbinian Brodmann produced one of the first parcellations of the brain by dividing the cerebral cortex into 52 functionally and anatomically distinct regions (Fig. 3.1) [38]. However, prior to

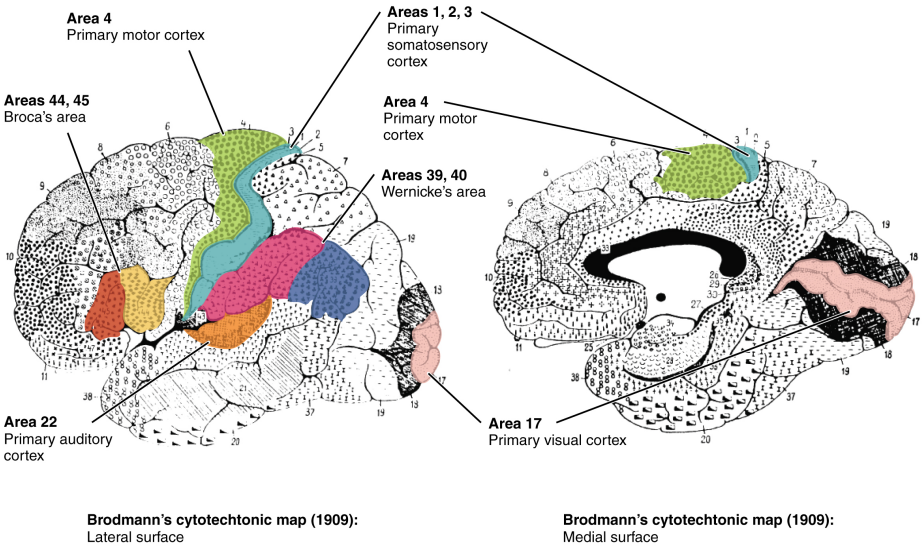


Figure 3.1: Brodmann's areas of the cerebral cortex, with an emphasis on regions specialised in distinct functionality. Image from OpenStax [184].

his seminal work, the foundations of brain parcellation and functional localisation were already set. As back as 1786, Vicq D'Azry had described the convolution characteristics of the cerebral cortex, emphasizing the differences in its morphology compared to the other animals [195]. In the early nineteenth century, Franz Joseph Gall had published his revolutionary ideas on the functional localisation in the brain, arguing that particular regions of the cerebral cortex were responsible for specific functions [119], while Louis Pierre Gratiolet had divided the cerebral hemispheres into the frontal, temporal, parietal, and occipital lobes, with respect to the overlaying skull bones [195]. In 1861, Pier Paul Broca had discovered that the posterior region of the frontal lobe (now referred to as Broca's area) was involved in speech production. Inspired by Broca's work, in 1876 Karl Wernicke had identified another cortical region (in the posterior part of the temporal lobe, now called Wenicke's area) that was responsible for understanding of speech [119]. By the late 1800s, Camillio Golgi and Santiago Ramon y Cajal had revealed the first detailed descriptions of nerve cells by using staining techniques.

Nevertheless, it was Brodmann, who came up with a cortical parcellation by delineating different regions in the cerebral cortex based on the variations in cellular decomposition of gray matter tissues. Brodmann not only influenced the field of neuroscience with a reference map of the cortex, but also provided insight towards the basis of correlation between anatomy and

function [288]. As a result, Brodmann's areas are still widely used today for the identification of activated foci, attributing its popularity to the modern imaging techniques, which allow *in-vivo* mapping of the architectonic data into a spatial reference system, hence enabling comparison of function with anatomical areas [287].

The clinical relevance of the brain's cellular organisation was already a major topic in the early twentieth century. Following Brodmann's pioneering studies, Cecile Vogt and Oscar Vogt subdivided the cerebral cortex based on the features derived from myeloarchitectonic features. The Vogts focused on the local variations in size and density of fiber bundles in neural tissue and achieved a cortical parcellation with a finer granularity (composed of around 200 distinct regions). Other post-mortem studies on the cellular level, such as by Constantin Freiherr von Economo and Georg Koskinas [270], contributed to the field of cytoarchitectonics by defining more cortical areas with homogeneous cellular decomposition. Micro-structural organisation of the cerebral cortex was later illuminated by means of receptor-architectonic mapping, an alternative architectonic technique that relies on the analysis of signals obtained from neurotransmitter receptors after being stimulated by chemical messengers [288, 291, 290].

Until recently, functional localisation and parcellation of the human cortex mostly involved post-mortem examination of the brain, including the aforementioned architectonic techniques as well as observations on lesion-induced brain areas in patients with particular disorders (inability in speech and language processing being the most popular). It was possible to localise function in the living brain, but only via invasive techniques such as using electrodes to stimulate the cortex during surgery [119]. Although, all these endeavours extensively contributed to our understanding of the human brain, their extent to reveal the link between anatomy and function was still limited. First of all, the majority of early studies were only based on a single brain or explored very small groups of subjects, hence they were not able to take the inter-subject variability in the organisation of the cerebral cortex into consideration [288]. Second, most of the studies were investigator-dependent, as it was not possible to replicate the experiments in a systematic manner [287]. In addition, studies targeting the living brain, on top of being invasive, were mostly performed on patients with brain disorders, making the generalisation of observations/findings to healthy people questionable [119].

3.2.2 Neuroimaging for Brain Parcellation

Starting in the early 1980s, structural and functional imaging technologies have revolutionised the study of organisational principles of the brain. MR and PET imaging have opened up new horizons in the field of neuroscience by allowing non-invasive exploration of the living brain from various aspects. Functional MRI has quickly become the most commonly used tool for functional localisation studies with its ability to capture the human brain while the subject is performing a pre-defined task. Parallel advancements in medical image analysis have led to development of various brain atlases and statistical analysis techniques, allowing to map activated foci with anatomically-defined cortical regions in the brain [81, 144, 222, 287, 197], hence, finally linking function with anatomy.

The Talairach and Tournoux atlas [239] has become particularly popular as it is one of the first attempts to provide a 3D spatial coordinate system, where cortical areas are labelled according to surface anatomy (sulci/gyri) and Brodmann's architectonic map; but only based on a single brain. In order to account for anatomical variability across subjects, probabilistic brain atlases (including a reference template and labelled map of distinct brain regions) have been introduced by co-registering and averaging multiple structural and functional images [52, 160, 159]. As the imaging techniques progressed further, these attempts have been followed by many other cortical and subcortical brain parcellations [130, 254, 74, 64, 61, 286]. Such subdivisions of the brain mainly follow the architectonic boundaries and/or anatomical landmarks in the cerebral cortex and are still being extensively used in neuroimaging studies.

By the beginning of the new millennium, enough evidence was accumulated to consider the brain as a mosaic of distinct modules that are specialised in certain functions, such that specific facets of cognition, emotion and behaviour can be anatomically localised [82]. However, it has been also revealed that, no single brain region could perform a particular function by itself, nor would be a one-to-one mapping between brain regions and mental operations [82, 249]. This is further supported by numerous studies targeting the brain at rest (for an extensive review see [257]), which have identified several anatomically separated, but functionally connected cortical regions (commonly referred to as resting-state networks or RSNs) that are associated with a variety

of functions, such as sensory/motor, visual processing, auditory processing, memory, and the so-called default mode network (DMN), which gets activated 'by default' when subject is not involved in a task [40].

Most of RSNs are consistently identified across studies, despite differences in the experimental setting, data acquisition and analysis techniques [56, 224, 257]. Besides, it has been also reported that functionally linked brain regions composing these large-scale networks continuously exchange information via anatomical connections [256]. As a result, the dynamic interactions and information transmission through functional and structural connections between anatomically distinct regions (such as defined by their local micro-structure) are considered to facilitate performing a particular cognitive, sensory, or motor operation [193, 249].

Put together, the last few decades have seen a tremendous effort to provide a high-level abstraction of the fundamental organisation of the brain at macroscopic scales. This can be easily attributed to the advances in image acquisition techniques, which have facilitated the multi-scale subdivision of the brain using various modalities and methods. This has also brought the scope of parcellation to a much broader extent. As a matter of fact, the term 'parcellation' does not refer to a unique map or tool, but rather define a spectrum of subdivisions that encapsulate fundamental neuro-biological information about cortical organisation and allow the mapping of brain function and anatomy with respect to different aspects. In this context, the role of structural and functional connectivity is particularly essential, as they can provide complimentary information towards parcellating the brain into neuro-anatomically distinct areas that are involved in different functions.

3.2.3 Parcellation for Connectome Analysis

In addition to its role in the functional segregation and integration of the cerebral cortex, parcellation constitutes one of the core steps to explore the system organisation of the human connectome at the the macroscale. As previously discussed in Section 2.3 *Mapping the Brain*, the macro connectome is typically modelled as a network (or graph) of interactions between different regions of interest (ROIs) located in the cerebral cortex, where each ROI constitutes a

node in the network. Attempts to perform connectivity analysis at the vertex level (i.e. considering each vertex as a distinct region) are generally not feasible due to the high dimensionality of connectomic data, which may yield numerous problems, including computational cost, sensitivity to noise, and lack of interpretability of the resulting networks [249]. In fact, a higher level of abstraction provides a more neurobiologically meaningful basis for constructing a connectome, considering the aforementioned integration of neural subunits into spatially localised clusters in order to perform a function [151, 98, 55, 244]. As a result, connectivity analysis is generally carried out with respect to a parcellation, which may cover the whole cortex, i.e. whole-brain connectivity analysis, or target a local area of interest, i.e. ROI-based analysis, depending on the study under investigation.

ROI-based analysis is carried out by *a priori* selecting a region or seed, typically based on a traditional task experiment and/or using biological knowledge. The correlations between this predefined ROI and the rest of the cortex (or other ROIs) are then explored via statistical analysis techniques [174]. One major limitation of ROI-based analysis emerges from the fact that it cannot account for the joint interactions among multiple brain regions, as the signal outside the target area is not taken under consideration. The results, therefore, are likely to depend on the choice of the ROI(s), which may be prone to expert's error [27]. Besides, such ROIs may not be generalised to new subjects or data [244].

Whole-brain connectivity analysis targets the entire brain and generally relies on a set of nodes acquired from a cortical parcellation. Parcellations derived from anatomical/cytoarchitectonic brain atlases [254, 239, 61, 74] or random parcellations [96, 101, 231] are traditionally used for node identification, however these approaches might constitute several major limitations [55, 244]. Despite many widely-used brain atlases, there exists no consistency across parcellations defined by different atlases [33], which can be attributed to the fact that each atlas represents a labelling of brain according to the accumulated knowledge at the time of atlas creation [244]. More importantly, ROIs derived from brain atlases or random parcellations are likely to comprise a mixture of signals, since such parcellations are not defined from the underlying data [226, 54, 218]. One typical example of this is the aforementioned anterior cingulate cortex (ACC), which exhibits a great amount of heterogeneity regarding structural [22] and

functional connectivity [157], despite the fact that it is typically represented as a single ROI in traditional brain atlases [254]. Similarly, many anatomical atlases divide the frontal lobe into relatively large regions, hence a post-parcellation is typically carried out to achieve a finer granularity, which may better fit the data [244].

Using ill-defined ROIs (ROIs that do not match well the actual structures of interest) may have a great impact on the estimation of network structure. According to large-scale simulations carried out by Smith et al. [226], incorrect ROI delineation for especially whole-brain connectivity analysis would severely affect the network modelling, as signals from unrelated parts of the brain are mixed to each other. In addition to network modelling, topological architecture of brain functional networks can also depend on the selection of the underlying parcellation, and thus, network statistics such as small-worldness can vary between networks derived from different atlases [274]. Parcellation is, therefore, of great importance for the subsequent stages in connectivity analysis, which fundamentally rely on the network components. In this case, connectivity-driven parcellations (or more broadly, data-driven parcellations) provide a more reliable means for the definition of network nodes -hence a better model of the signal-, as they are typically learned from the same connectomic data used to conduct the connectivity analysis [243, 75, 135, 55, 249, 244]. Although, connectivity-driven parcellations are not likely to reflect the neuro-biological organisation of the brain, they can be typically associated with known brain structures for a more clinically relevant interpretation of the subsequent analysis [244, 71, 93].

3.3 Connectivity-Driven Parcellation

Connectivity-driven parcellation (CDP) aims at dividing the cerebral cortex (or a targeted area) into subregions that differ from their neighbours in connectivity, which can be in the form of long-range anatomical connections or functional interactions with other areas. The idea stems from the fact that each vertex or ROI can be described with a connectivity profile that shows how it is connected with the rest of the cerebral cortex [193]. The CDP problem

is then cast as a clustering problem, where the objective is to subdivide the cortical surface into a set of clusters in a way that connectivity is similar for vertices in the same cluster but different between clusters [66]. This consequently allows to represent each parcellated region with a single, consistent connectivity profile in network analysis. CDP can, therefore, provide a well-suited model to the underlying connectivity structure at a lower dimensionality [244].

3.3.1 Methods for Connectivity-Driven Parcellation

Connectivity has been one of the major sources of information for brain parcellation, and consequently, many CDP methods have been proposed, usually in association with a clustering algorithm driven by the similarity between connectivity profiles or fMRI timeseries. Among them k-means, spectral clustering, and hierarchical clustering are generally the most commonly employed approaches in CDP studies, while various other techniques, including but not limited to independent component analysis, boundary mapping, non-linear manifold learning, region growing, and mixture models have also been alternatively used, each with its own limitations and advantages. Below we will review the most prominent approaches to obtain connectivity-based parcellations, with an emphasis on the differences between various methods.

K-means Clustering

K-means clustering constitutes one of the most extensively used methods for CDP studies [250, 170, 161, 28]. It typically subdivides a ROI into a predefined number of subregions, which is specified by k as an external parameter. After selecting the initial cluster centroids (k seeds), k-means clustering algorithm is driven by two alternating steps: (1) associating data samples (e.g. timeseries or connectivity profiles) with clusters (assignment step) and (2) the estimation of new cluster centroids (update step). At each iteration, the cluster assignments take place in a way that minimises the sum of squared differences between the data samples and their associated cluster centroids.

The method has been employed in varying aspects of the CDP problem. For example, Mezer

et al. [161] applied k-means to rs-fMRI BOLD timeseries to obtain whole-brain parcellations including subcortical areas and thalamus, while Kahn et al. [118] relied on a similar approach in order to subdivide the orbitofrontal cortex. Functional connectivity captured at rest is further used to parcellate the medial frontal cortex into supplementary motor areas (SMA) and pre-SMA subregions [121]. Anwander et al. [8] and Tomassini et al. [250] used k-means in conjunction with tractography-driven anatomical connectivity to parcellate Broca's area and lateral premotor cortex, respectively. An analogous approach was followed by Nanetti et al. [170] to segment the SMA and the insula. K-means applied to connectomic data has further revealed the heterogeneous connectivity patterns within the cingulate cortex [22]. A resting-state fMRI-driven parcellation of the left hemisphere obtained with k-means is shown in Fig. 3.2.

While k-means provides a 'hard' clustering (i.e. every vertex can only be assigned to a single cluster), its fuzzy counterpart c-means can alternatively be used to obtain a 'soft' clustering, where each each vertex is represented by a weighted label vector [29]. For example, by applying c-means to the rs-fMRI data, Lee and colleagues [139] have subdivided the cortical and subcorbital gray matter into several regions, each is associated with a different RSN.

Although being extensively used in the context of CDP, several points should be taken into consideration before adapting a k-means based approach. First of all, the algorithm highly depends on the initialisation of cluster centroids, which can be random or based on *a priori* knowledge, such as anatomical landmarks [66]. Partially due to this randomness, but mostly because of the internal dynamics of the algorithm, it is not likely to obtain the same parcellation every time the algorithm is employed. It is, thus, important to run the algorithm for many times before deciding on the final parcellation [170]. It is also worth noting that, repeating the clustering process for different values of k yields different solutions and is not expected to obtain a hierarchy of nested regions that are linked to each other in a child-parent relationship [66].

k-means could become computationally prohibitive especially when the dimensionality of the data is extremely high. This is typically the case when the whole cortex is targeted instead of a relatively small ROI and the parcellation is carried out with a very large k. To overcome the

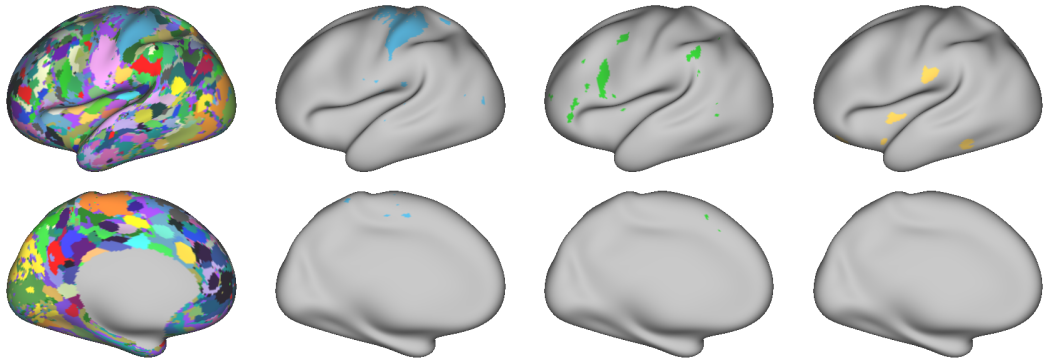


Figure 3.2: A parcellation of the left hemisphere obtained with k-means using resting-state fMRI connectivity of a single subject ($k = 100$). 3 arbitrarily chosen regions are shown individually alongside the parcellation to emphasise the fact that k-means does not impose any spatial constraints.

impact of computational limitations, Thirion et al. [244] suggest to utilise principal component analysis (PCA) as a preprocessing stage to reduce the dimensionality of the data. Another limitation of k-means is that it naturally does not provide spatially contiguous parcellations, as the spatial structure of the data is not accounted for. For example, such disconnected parcels can be seen in Fig. 3.2. A typical procedure to acquire spatially more coherent regions is to incorporate external constraints (for example anatomical coordinates of vertices) into the clustering algorithm [75, 250]. A modified k-means approach driven by a similar principle will be presented in the next chapter.

Spectral Clustering

Another widely-used clustering approach that has been frequently employed to CDP is spectral clustering [114, 124, 255, 54, 217, 218, 188], owing its popularity to the fact that the similarity of connectivity profiles between any pair of vertices can be naturally represented as a connectivity matrix [66], where nodes correspond to vertices and edges represent the connections between the nodes. Structures of interest in this matrix can be revealed through spectral decomposition and further used to obtain subdivisions of the underlying data. This is typically achieved by computing the first eigenvectors of the normalised graph Laplacian of the similarity matrix [219].

These eigenvectors provide a feature matrix in a lower dimensional space and can be submitted to a conventional clustering algorithm for parcellation, preferably after being reordered.

Spectral reordering modifies the structural organisation of the connectivity matrix with respect to the information provided by the eigenvectors, such that vertices with higher similarity are positioned together within the matrix [17]. For example, Johansen-Berg et al. [114] used spectral reordering to parcellate the medial frontal cortex into two subregions (SMA and pre-SMA), in which the number of clusters was not predefined, but estimated by visual inspection [48]. However, when there is no clear separation between the data points in the reordered connectivity matrix, traditional clustering algorithms, such as k-means, are generally employed to obtain the final parcellation borders [8, 124]. This is further beneficial to reduce the inspector bias and obtain parcellations in a more data-driven manner [48].

Alternatively to spectral reordering, eigenvectors can be directly fed into a clustering algorithm to compute a parcellation with a predefined number of regions. For example, Thirion et al [244] applied spectral clustering to task-based fMRI data in conjunction with c-means to obtain ROIs that can be used for functional connectivity analysis. Van den Heuvel et al. [255] modelled the whole-brain CDP problem as a normalised cut (NCut) graph partitioning problem [219] and used spectral clustering to identify seven RSNs from a combined connectivity graph of several healthy subjects. Craddock et al. [54] extended the idea of NCut spectral clustering to obtain spatially contiguous cortical parcellations that can be used as nodes in connectome analysis, by imposing spatial constraints to the input similarity matrix, i.e. only retaining the edges constructed between anatomically adjacent vertices and discarding the others. In such cases, final parcellations can be typically obtained with k-means. However, as discussed above, k-means heavily depends on initialisation and may yield different parcellations on a run-to-run basis. For a more robust clustering, it can be replaced by an alternative technique proposed in [285], that includes learning a rotation to discretise the eigenvector representation, such that the discrete eigenvectors reflect an optimal partitioning of the graph [255, 54, 244]. Several cortical parcellations obtained using spectral clustering with normalised cuts at different levels of detail are shown in Fig. 3.3.

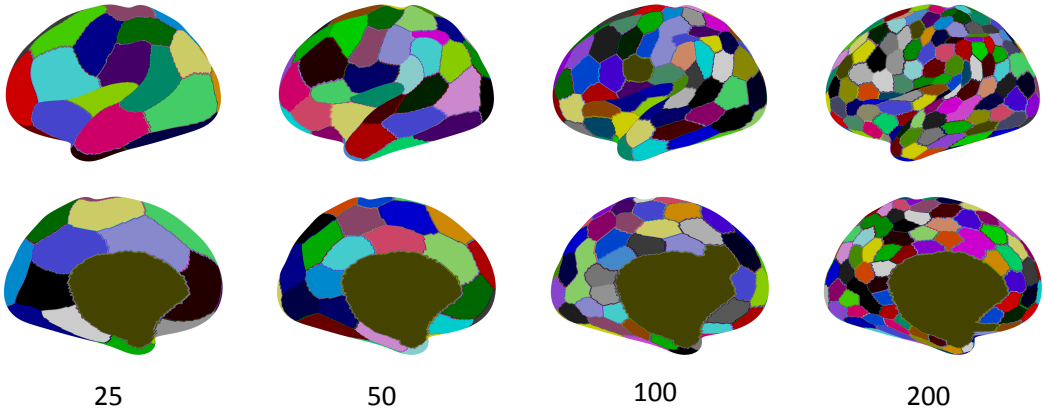


Figure 3.3: Cortical parcellations of the left hemisphere obtained with the normalised cuts spectral clustering algorithm using resting-state fMRI connectivity of a subject. The numbers indicate the predefined resolution of each parcellation.

There are at least two drawbacks associated with NCut-based spectral clustering, particularly when spatial constraints are in use. NCut graph partitioning tends to yield a 'balanced' subdivision of the input graph. However, in cortical areas where a reliable separation does not exist (such as areas with relatively weak connectivity or low SNR), it is likely to generate parcellations with uniformly sized clusters, which can bee seen in Fig. 3.3. This relatively simple model may thus fail to reflect the complex structures generally present in the brain [66]. Due to spatial constraints imposed to the similarity matrix, NCut is driven to a large extent by the underlying structure of the graphical model, which consequently leads to capturing the spatial organisation of the brain rather than the underlying connectivity characteristics [54, 244]. Similar to k-means and in contrast with hierarchical clustering, spectral clustering does not yield a hierarchy of nested regions when the algorithm is employed for different number of clusters, however there may still be some sort of inherently-captured similarity between parcellations of different granularity [66].

An important consideration point in spectral clustering is how to construct a similarity matrix that could summarise the connectivity characteristics shared across individual subjects when a group-wise study is devised. While popular solutions include averaging individual subject connectivity matrices [194] or computing an intermediate graph representation from subject-level parcellations [255, 54], more recent techniques construct a joint model of connectivity

patterns [218, 187, 188]. A novel method to obtain a group-level graphical representation of a population will be provided in Chapter 6.

Hierarchical Clustering

Hierarchical clustering is a technique for grouping data points into a set of nested clusters that can be represented as a hierarchical tree, which is often called a dendrogram [115]. Cutting this tree at different levels of depth produces parcellations with the desired precision (Fig. 3.4). Strategies for hierarchical clustering generally fall into two groups, depending on the direction towards which the algorithm is operated, i.e. bottom-up (agglomerative) and top-down (divisive). The former approach starts with each data point as a singleton cluster and the most similar clusters are merged as the algorithm progresses up the hierarchy. The divisive approach initially assigns all data points into a single cluster and progressively divides the least homogeneous cluster at each iteration. However, this strategy is very rarely used for brain parcellation, and hence is not covered here.

Agglomerative hierarchical clustering, on the other hand, is widely adapted as a parcellation technique in numerous CDP studies [53, 212, 28, 169, 112, 65, 32, 164]. Its popularity can be attributed to the fact that, in contrast with the aforementioned clustering approaches, hierarchical clustering provides a spectrum of parcellations that can be explored at different coarseness levels [66] as shown in Fig. 3.4. Such a clustering model has a higher tendency to reflect the hierarchical organisation of the human connectome; however, it is worth noting that a hierarchical parcellation does not necessarily provide a neurologically meaningful segregation of the brain [66].

Hierarchical clustering techniques have been adapted to investigate various facets of brain parcellation with regards to both functional and structural connectivity, though the former being more extensively studied. Cordes et al. [53] made one of the first attempts to use hierarchical clustering for the identification of functional connectivity patterns in resting state data. Quite a few studies have been carried out to identify consistent RSNs across healthy groups of subjects by using hierarchical clustering at different processing stages [212, 28, 169]. For example, Sal-

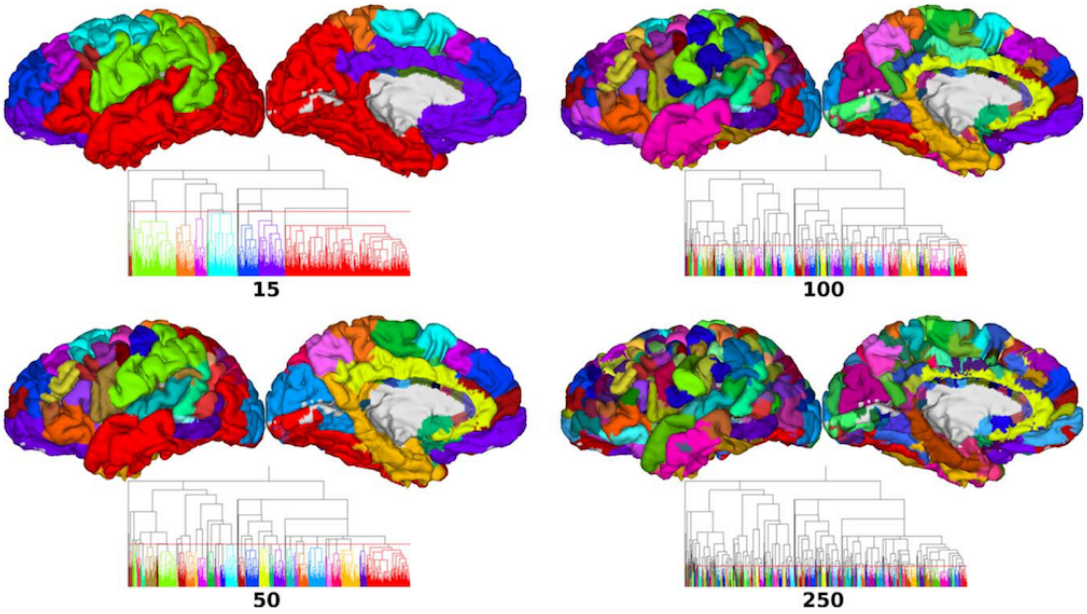


Figure 3.4: Cortical parcellations of the left hemisphere acquired by cutting a hierarchical tree at different levels, denoted by the red horizontal lines. The numbers indicate the predefined resolution of each parcellation. Image is adapted from [164].

vador et al. [212] revealed six major systems through a hierarchical parcellation of a connectivity network defined from an anatomical template. Bellec et al. [28] followed a similar procedure for RSN identification, but instead of relying on an anatomical template for network construction, a data-driven parcellation is acquired through a multi-level approach that combines region growing, k-means clustering, and hierarchical clustering. Brain reading studies that aim to predict the subject's behaviour during a t-fMRI session have also utilised hierarchical techniques to obtain structured patterns of activation from task-based data [112, 162]. Finally, hierarchical clustering has been applied to both rs-fMRI and dMRI to obtain whole-brain cortical parcellations that can be used to substitute the network nodes in network analysis [32, 164].

One critical consideration point in these studies is the selection of a 'linkage rule' that determines which clusters to merge with respect to a similarity measure. Among many linkage rules, average (mean) link [212, 28, 169, 164] and Ward's method [227, 32, 65, 162, 112] seem to dominate the literature. The former is driven by the similarity measured in terms of the mean connectivity between vertices within each pair [229], whereas Ward's method merges a pair if the resulting cluster leads to minimum increase in the total within-cluster variance [276]. It is

important to note that, hierarchical clustering algorithms can yield very imbalanced parcellations depending on the merging criterion [66]. As there is no best linkage rule, it needs to be selected experimentally and/or based on the study under investigation. It may also be critical to impose constraints to the merging process, (e.g. two clusters can be merged if and only if they are adjacent to each other), in order to obtain spatially contiguous, and hence, biologically more plausible parcellations [32, 244, 164].

Although hierarchical approaches generate a nested multi-level parcellation, only few levels of granularity are explored in practice [66]. While these levels are typically determined arbitrarily, but still accounting for a wide range of resolutions [32, 164], other methods may rely on more sophisticated techniques to decide where to cut the dendrogram [169]. One critical drawback of hierarchical clustering is its sensitivity to noise, which may lead to erroneous (false) boundaries detected at higher levels to progress down the hierarchical tree and subsequently reduce the reliability of parcellations at coarser levels. In other words, once a parcel is formed by merging two clusters, it cannot be undone in later iterations [164]. As a result, an optimisation based clustering strategy like k-means may be more appropriate for parcellation when dealing with small datasets/ROIs or only a single resolution is desired [164].

Independent Component Analysis

Independent Component analysis (ICA) is currently considered as the reference model for identifying networks that underlie the intrinsic functional organisation of the brain [21]. The key idea behind ICA is to decompose fMRI data into a set of components that are maximally independent from each other. Each component is then interpreted as a spatially distributed connectivity network and represented in the form of a weighted set of vertices (i.e. a spatial map) [21]. Several spatial maps obtained using ICA are shown in Fig. 3.5. In this figure, each map is thresholded to identify the cortical regions with the strongest connectivity within each component. ICA-based analysis techniques [122, 21, 56, 224, 152] have been widely used to subdivide the human connectome into robust functional networks across resting [21, 56] and task-based [224] experimental settings, as well as across different groups of subjects [249]. The

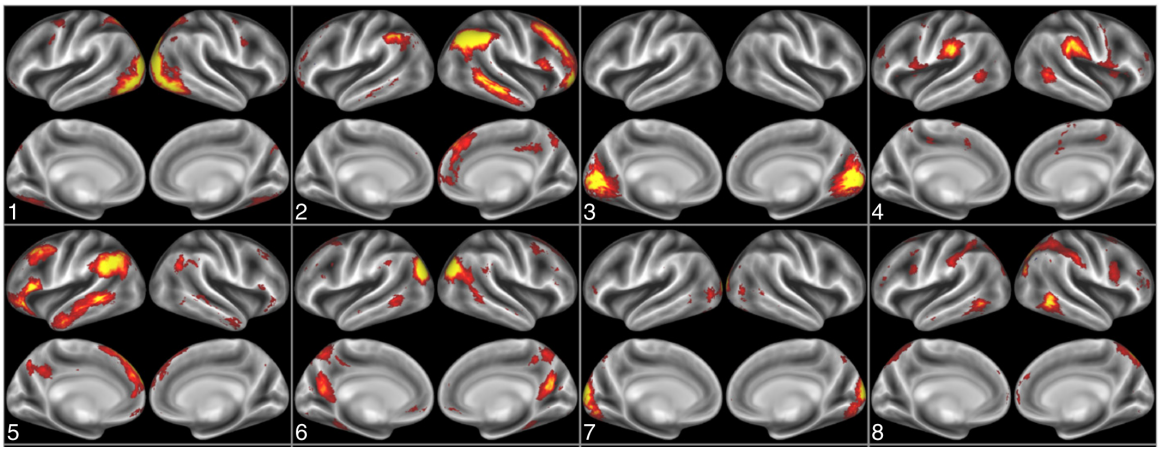


Figure 3.5: Spatial maps obtained by applying ICA to resting-state fMRI connectivity data of 20 subjects. Overlays correspond to group-derived z-score maps thresholded at $z > 5$. Image is adapted from [227].

multivariate signal that forms the basis of ICA can also be formulated in the dictionary learning framework [263, 2]. Different strategies can be used to combine fMRI data from subjects, such as concatenating individual timeseries [122, 21] or using more sophisticated methods that take the inter-subject variability into consideration [264, 263, 2].

ICA may capture different components of the signal, including non-neural components and actual functional networks [2]. On one hand, this could allow extraction of artefactual patterns from the data, which is not possible with the aforementioned clustering techniques [66]. On the other hand, it constitutes a limitation, as manual selection of components may be required [218]. Spatial maps obtained via ICA may potentially have several separate regions scattered across any given map and overlap with each other, hence, do not provide spatially contiguous and non-overlapping parcellations. As a result, ICA-based techniques may not be preferred by some researchers [223, 66]. However, it might be possible to achieve spatial coherence by imposing additional post-processing steps, such as (1) increasing the number of components until spatial contiguity is ensured [123], (2) using sparsity-induced regularisation [2], or (3) subdividing components that are not spatially contiguous [227]. Since ICA inherently assigns each component a distinct timeseries, subsequent network modelling using these components as nodes is guaranteed not to be rank deficient, whereas a hard parcellation is likely to have multiple parcels with very similar timeseries, which may potentially affect the network modelling [223].

Boundary Mapping

Boundary mapping is an alternative CDP approach that casts the clustering problem as the identification of abrupt changes between connectivity patterns captured at rest and solves it using image analysis techniques [49, 279, 280, 93]. The key idea behind boundary mapping is to identify locations on the cerebral cortex where the functional connectivity pattern changes rapidly, potentially representing boundaries between functionally localised cortical areas [49]. Since the cerebral cortex is treated as a 2D cortical sheet, edge detection algorithms can be used to compute gradient maps from the resting-state correlations. Such maps provide information about the transitional zones in connectivity across the cortex, thus, can be fed into an image segmentation tool (such as watershed algorithm [267]) for delineating cortical parcellations. Boundary mapping has been successfully used to obtain areal segregations of the cortical surface from resting-state functional connectivity, targeting both specific cortical regions (such as the frontal cortex [49] and the lateral parietal cortex [171]), and the entire cerebral cortex [279, 280, 136, 93, 92]. A whole-brain boundary map and the resulting cortical parcellations are provided in Fig. 3.6.

One major drawback of the boundary mapping techniques is that they do not allow model selection, such that the granularity of the parcellation cannot be predefined, which is in contrast with most of the previously mentioned clustering techniques [244]. In addition, boundary mapping is very sensitivity to spurious gradients captured from connectivity patterns, especially when the whole cortex is under consideration. As a result, a considerable amount of data is required to reduce the possible impact of falsely identified edges, which makes the technique more suitable for application to averaged-datasets of large populations [93]. On the other hand, a subject-level parcellation at the whole-brain scale can still be achieved, for example by scanning a single subject many times [136]. However, it should be noted that the high dimensionality may still constrain the detection of robust connectivity patterns at the subject level.

Considering the fact that detected edges can be represented as a probabilistic map, some researchers position boundary mapping in between hard and soft clustering methods [66]. While

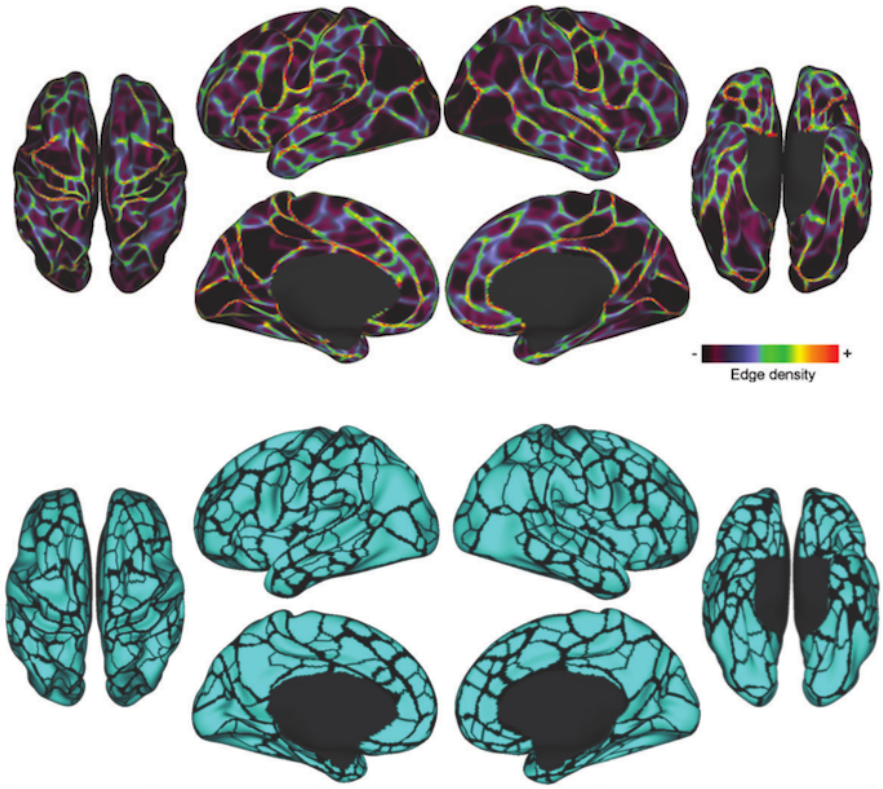


Figure 3.6: (*Top*) A boundary map obtained from resting-state functional connectivity patterns of 120 healthy young adult subjects. Hot colours indicate the locations where connectivity abruptly changes. (*Bottom*) Cortical parcellations derived from the boundary map. Both images are adapted from [93].

boundary mapping has been a very popular approach for functional CDP, to the best of our knowledge it has not been applied to structural connectivity estimated from dMRI; though, Chapter 5 will cover a newly developed method to fill this gap.

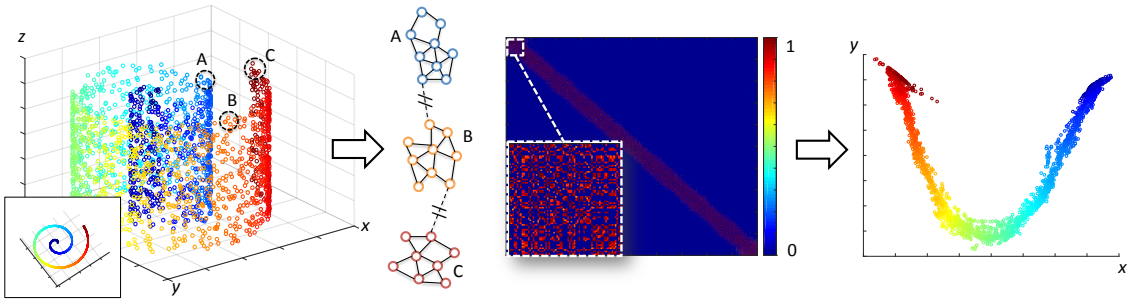
Non-linear Manifold Learning

This group of statistical techniques reconceptualises the CDP problem as a non-linear feature reduction problem and proposes using manifold learning to capture patterns of interest within the inherently high dimensional connectivity space [216, 242, 133, 132, 131, 134]. Contrarily to linear methods such as PCA and ICA, non-linear feature reduction techniques, including but not limited to ISOMAP [241], locally linear embedding (LLE) [209], Laplacian eigenmaps (i.e. spectral embedding) [26], and diffusion maps [50], seek a low-dimensional manifold that

is assumed to represent the intrinsic geometry of the underlying data. Such methods encapsulate local proximity information within a nearest-neighbourhood graph, which is then used to compute embedding coordinates in a low dimensional space by applying spectral analysis techniques to certain matrix forms associated with the underlying high dimensional graph [216]

ISOMAP (short for Isometric Mapping) computes a lower-dimensional embedding by first approximating geodesic distances between all points and then applying the classical multi-dimensional scaling (MDS) [36] to the distance matrix. LLE assumes that the manifold in the data is locally linear and finds a non-linear embedding according to the optimal local weights, which are computed for each data point in the form of a linear combination of their nearest neighbours. Laplacian eigenmaps relies on the assumption that the data lies in a low-dimensional manifold in a high-dimensional space [26] and uses the graph Laplacian matrix to perform dimensionality reduction. Diffusion maps, while is very similar to Laplacian eigenmaps in terms of underlying graph representation, takes a probabilistic perspective to the dimensionality reduction problem and defines a Markov chain from a dataset, in which the distance between pairs of data points is characterised by their transition probability [216]. An illustrative example featuring the well-known Swiss roll dataset is provided in Fig. 3.7 to demonstrate typical steps involved in non-linear manifold learning, in which Laplacian eigenmaps is used to compute a 2-dimensional embedding from a three dimensional dataset.

Various non-linear manifold learning approaches, diffusion maps [216, 132, 131] and Laplacian eigenmaps [242] being the most popular, as well as a broader class of spectral decomposition techniques [216, 243, 54, 218] are used for learning low-dimensional embeddings from the connectivity data (see [132] for other related work). The common point in these studies is that the interactions between remote brain regions are represented with a graph in the high dimensional space, which is then decomposed into a feature matrix with lower dimensionality that can be used to reveal, for example, temporal patterns of functional activity [242, 216] or further fed into a clustering algorithm to identify cortical parcellations [54, 132]. Novel CDP approaches that make use of non-linear manifold learning in the form of Laplacian eigenmaps [26] and spectral decomposition techniques will be proposed in Chapters 5 and 6, respectively.



(a) (b) (c)

Figure 3.7: An example of non-linear manifold learning: (a) The well-known Swiss roll dataset generated from 2000 data points, inset showing the bird-view of the space from the z axis. (b) The local proximity information in the 3D space is represented with a graphical model (i.e. an adjacency matrix), in which only the closest data points are connected to each other. The graph is partly illustrated for three arbitrary neighbourhoods (A, B, and C). (c) The 2-dimensional embedding obtained using Laplacian eigenmaps, in which locally connected data points in the 3D space remain nearby in the manifold.

Alternative Methods for CDP

With the increasing attention in macro connectomics, last decade has seen a tremendous progress towards CDP. While it is not feasible to cover every single CDP approach in as much detail as the aforementioned methods, we provide an overview of alternative techniques that approach the CDP problem from different points of view.

Region growing constitutes a good example to the parcellation techniques that impose a spatial constraint to the clustering process and has been used for many times in CDP studies [151, 98, 27, 32]. The general idea is to iteratively merge vertices with their most similar neighbour until a stopping criterion is satisfied. Region growing, therefore, inherently maximises the homogeneity within each parcel [98, 27]. However, this may pose a limitation especially in whole-brain CDP, i.e. imbalanced parcellations with many small but few large parcels. This is typically avoided by capping the parcel size with a predefined threshold [27] or driving the algorithm from an initial set of seed vertices that approximately cover the entire cortex [32]. The former leads to parcellations with evenly sized clusters, which may consequently show similar properties as in spectral clustering. Region growing tends to segregate the cortical surface into many, but highly homogeneous clusters, some of which may only consist of a single vertex. As a result, further

processing (such as a second level clustering) may be required to obtain more neuro-biologically plausible parcellations [27, 28, 32].

A different family of CDP methods consists of probabilistic modelling techniques, which describe the underlying data as a set of probability distributions, such as Gaussian mixture models [90, 89, 253, 217, 132]. Alternatively, von Mises-Fisher distribution is adapted for CDP with regard to its usability for clustering high dimensional data [135, 210]. For example, Yeo et al. [245] used such a probabilistic model [135] to subdivide the cerebral cortex into several RSNs based on functional connectivity data collected from 1000 healthy subjects. Langs et al. [132] combined the functional characteristics of several subjects obtained from a language fMRI study into a joint probabilistic distribution, from which a generative model is learned and used to identify a functional atlas, without relying on the spatial distribution of the contributing functional units (i.e. no explicit spatial constraints imposed to the data) [132]. A similar stream of processing is later instrumentalised to study the intrinsic functional organisation of the brain at the individual subject level using rs-fMRI [134].

In probabilistic modelling techniques, cluster labels and other parameters are usually inferred via an iterative expectation maximisation procedure [210, 90, 132]. While parametric models typically require fixing the number of clusters *a priori*, it may also be uncovered automatically by incorporating a label cost to the model [210]. Alternatively, using a mixture model within a non-parametric setting (such as Dirichlet processes) may allow inferring the number of clusters directly from the data [111]. Non-parametric models are particularly popular tools in the context of whole-brain CDP, as they do not rely on any assumptions on the granularity of parcellation [15, 109].

In general, mixture models may be highly sensitive to initialisation and require multiple runs to achieve a reliable parcellation [90] or to infer the right number of clusters [111]. Spatial regularisation in this context is generally enforced through Markov Random Field (MRF) priors [111, 210], which apply a penalty whenever neighbouring regions are assigned to different labels. However, this may not ensure specially coherent parcellations, as spatial contiguity is a global property that cannot be enforced locally [102]. To overcome this, Honnorat et al. [102]

suggest using a convexity prior that guarantees connectedness within parcels; however, such a constraint yields parcels with regular (convex) shapes and comes at the expense of computational inefficiency depending on the spatial dimensionality (i.e. number of nodes). While this method is applied to resting-state fMRI data to compute whole-brain cortical parcellations, another MRF-based approach tackles the same problem, but with application to dMRI-based structural connectivity [192], which is initialised from random parcellations and iteratively updates parcel boundaries in a coarse-to-fine multi-resolution procedure.

3.3.2 What Defines A Reliable Connectivity-Driven Parcellation?

As extensively reviewed in the previous section, there exit many techniques that attack the parcellation problem from a connectivity-driven perspective. Put together, previous work has enhanced our capacity to understand the brain's functional and structural organisation at multiple scales, ranging from high-level spatially distributed structures of connectivity networks to whole-brain segregations of the cerebral cortex. In this thesis, we position ourselves next to the latter, and propose novel solutions to subdivide the cortex into spatially contiguous, non-overlapping, and homogeneous regions, which can subsequently be used to (1) map the cortical organisation of the brain with respect to connectivity and (2) to define the network nodes in connectome analysis. Towards this end, we follow several notable criteria defined by the CDP literature to evaluate the 'reliability' of a whole-brain parcellation in the context of cortical segregation and connectome analysis:

- Subdivided regions should be highly homogeneous, in particular, comprise vertices with similar timeseries or connectivity profiles, since network nodes derived from parcellations are typically represented by a single feature vector (such as the average timeseries) [243, 54, 218, 93].

- Each region should specify a functionally and anatomically distinct area that has a different timeseries or connectivity profile from its neighbours, so as the entire parcellation can provide an abstraction of the functional specialisation and segregation in the

cerebral cortex, and can provide a more accurate network modelling for brain mapping [226, 32, 66, 93].

- ROIs should be spatially contiguous and non-overlapping to distinguish homogeneous network nodes (used to construct graphical models for connectivity analysis) from large-scale spatially-distributed networks (heterogeneous regions, usually spanning across the brain) [243, 27, 224]. Spatial connectedness further yields a physiologically more interpretable connectivity analysis, since enforcing spatial contiguity allows producing anatomically localised regions [54, 218, 109].

- Parcellations should be reproducible to a certain extent, meaning that, individual subject parcellations obtained from different scans of the same subject as well as group-wise parcellations computed from different groups of people who share some defining characteristics (e.g. healthy young adults) should exhibit some amount of functional and structural similarity [54, 32, 218, 102, 188].

- A CDP method should ideally provide cortical parcellations at different levels of detail, such that it would be possible to study the brain's functional/structural organisation in a multi-scale manner [54, 32, 218, 102, 188].

In practice, it may not be possible to address all these criteria simultaneously, as there may typically exist trade-offs between different points, e.g. in order to achieve higher fidelity one may need to sacrifice reproducibility. As a result, clustering algorithms may *a priori* make various assumptions or introduce implicit/explicit constraints, depending on the parcellation problem under consideration. As a general note, it is highly important to realise that each assumption and processing decision made by a clustering techniques comes with different consequences and inevitably bias the resulting parcellations in different aspects, including the shape, number, size, and spatial contiguity of the parcels [66].

For example, ICA assumes that the fMRI data consists of a mixture of statistically independent components and that spatially distributed functional networks can be effectively separated from signals of non-neural (e.g. artefactual) origin. With a similar objective, but from a differ-

ent perspective, nonlinear manifold learning techniques rely on the assumption that structures of interest in the connectivity data live in a low dimensional embedding, which can be captured using spectral decomposition [242, 216, 132]. Other techniques alter the structure of the underlying data to obtain more robust parcellations, for instance, by applying thresholding to suppress negative and weak correlations, assuming that correlations under a threshold correspond to spurious connections [255, 202, 54, 11]. It is also common to rely on spatial constraints [54] or regularisation [2] for computing what is expected to be physiologically more plausible parcellations.

3.3.3 Challenges in Connectivity-Driven Parcellation

Despite many promising results in the literature, the CDP problem is still open to improvements. Majority of the methodologies we have discussed so far either target relatively small cortical regions (e.g. the insula or SMA) or attempt to reveal the intrinsic connectivity that underlie the brain function (e.g. RSNs). However, the parcellation of the entire cortex is a more challenging task and the aforementioned whole-brain CDP methods only partly satisfy the above criteria. The main challenges that may have a critical impact on the reliability of a parcellation and how we overcome them throughout this thesis are summarised as follows:

- Connectivity data is typically very noisy because of the previously discussed (Chapter 2, *Imaging the Macro Connectome*) limitations associated with the imaging modalities used to capture connectivity. In particular, fMRI mainly suffers from low SNR, while dMRI and tractography lack sensitivity towards delineating long-range connections. Considering the 'garbage-in garbage-out' principle, it is highly critical to clean the data before carrying out any parcellation. We conduct our experiments on the HCP datasets, which have been subject to novel pre-processing pipelines, providing relatively high quality data [87, 261].

- The inherently high dimensionality of the connectivity data from which parcellations are acquired poses another challenge, particularly when subdividing the entire cerebral cortex. Today, MRI captures the brain at millimetre scale and allows longer scanning time than

ever. This is obviously beneficial towards obtaining images in greater detail; however, as the dimensionality increases, the computational requirements for analysing the data also grows rapidly. It may as well be more challenging to separate structures of interest (or true signal) from noise. Towards this end, we provide effective solutions to deal with both spatial and temporal dimensionality, where the former is particularly important, as it helps improve the SNR in rs-fMRI timeseries.

- Inter-subject variability stands as a challenge especially for group-level studies that aim to identify shared patterns of connectivity across subjects. As connectivity-driven parcellations are directly learnt from the underlying data, they can better capture the variability across subjects compared to traditional brain atlases [188, 3]. However, it is still important to pay particular attention towards incorporating individual functional/structural characteristics into the group-wise CDP computation [218, 131, 173]. From a different perspective, inter-subject variability constitutes a critical source of information towards patient-centred medicine and understanding the neural basis of variation in the brain function and human behaviour [272, 240]. Therefore, it is critical to study the brain's functional and structural organisation at the subject level. Taking both views into consideration, we propose methodologies to compute different cortical parcellations that can be effectively (1) tailored to individual subjects and (2) used to summarise connectional characteristics shared across a group of subjects.

- Evaluation in CDP is a challenge in itself, since there exists no widely-accepted parcellation that can be used as the 'golden standard'. It is, therefore, highly critical to validate computed parcellations considering different facets of the CDP problem [244, 66]. Given the aforementioned criteria that define the reliability in the context of CDP, we mainly perform evaluation through measuring reproducibility across subjects/groups, homogeneity within parcels, and the degree of separation between neighbouring regions, as well as carry out multi-modal comparisons with well-known neuro-anatomical patterns (such as Brodmann's cortical regions). More on evaluation will be discussed in the next chapter.

Chapter 4

Single Subject Parcellation of the Cerebral Cortex

This chapter is based on:

S. Arslan, S. I. Ktena, A. Makropoulos, E. C. Robinson, D. Rueckert, and S. Parisot, *Human Brain Mapping: A Systematic Comparison of Parcellation Methods for the Human Cerebral Cortex*, NeuroImage, 2017. (*In Press*)

S. Arslan and D. Rueckert, *Multi-Level Parcellation of the Cerebral Cortex Using Resting-State fMRI*. International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI), vol. 9351 of LNCS. Springer, pp. 47-54, 2015.

S. Arslan, S. Parisot, and D. Rueckert, *Supervertex Clustering of the Cerebral Cortex Using Resting-State fMRI*, Organization for Human Brain Mapping (OHBM), Honolulu, 2015.

Abstract

This chapter describes a method for the subject-specific parcellation of the cerebral cortex. We propose a two-level parcellation framework which deploys a different clustering strategy at each level. We adapt a k-means clustering approach to group vertices into a relatively large number

of highly homogeneous regions with respect to functional and spatial similarities between them. These clusters constitute a robust, high-level representation of the cerebral cortex, which, at the second level, are combined into a spectrum of parcellations with different resolutions using a spatially-constrained hierarchical clustering approach. Using data from 100 healthy subjects, we show that our algorithm segregates the cerebral cortex into distinct parcels at different levels of granularity with higher reproducibility and functional homogeneity, compared to a state-of-the-art two-level approach. We also discuss the pros and cons of our approach over other data-driven clustering techniques, including but not limited to k-means, hierarchical clustering, and spectral clustering. Finally, we show that functional connectivity at the individual subject level can be more effectively represented by connectivity-driven parcellation techniques, rather than random or anatomical parcellations.

4.1 Introduction

In this chapter, our aim is to compute robust and homogeneous parcellations for individual subjects, which would provide a reliable abstraction of the brain's functional organisation, and thus, enable generation of graphical models for a more effective analysis of the brain connectivity. Such models are currently derived from anatomical landmarks [254, 61], cytoarchitectonic features [38], or rely on random parcellations [231], thus may fail to fully reflect the underlying function in the brain. We propose a method driven by rs-fMRI, which records neurocognitive activity by measuring the fluctuations in the BOLD signals while the subject is at rest. Since the brain is still active in the absence of external stimuli, resting-state functional connectivity (RSFC) estimated from BOLD timeseries can be used to identify the brain's functional organisation [30, 150].

RSFC-based parcellation techniques generally focus on identifying patterns of brain connectivity that are shared across individuals [202, 245]. Studies replicated under different experimental settings and across many groups of people have discovered that the human brain is organised into several large-scale functional networks that are linked to brain cognitive functions [224,

258] and human behaviour [129]. Although these studies are typically conducted on RSFC datasets averaged across many subjects, similar functional systems are also detected at the subject level [215]. However a great amount of individual variability is also observed [168, 273], especially in the spatial distribution of these networks [273]. In addition, evidence suggests that the human connectome possesses connectional traits that are unique to each subject [168, 16, 92]. A recent study [71] has further shown that RSFC can be used to derive distinct features to successfully distinguish one individual from another. These features, however, may not be observed in group-averaged datasets [92]. Therefore, parcellating the cerebral cortex on a single subject basis can provide a natural start point for detecting such features, and help better understand how connectivity changes across individual brains.

With this motivation, we propose a two-level parcellation framework that could be applied to individual subjects for subdividing the cerebral cortex in its entirety. At the first level, we utilise a k-means clustering approach to group the cortical vertices into relatively large number of homogeneous regions. This stage reduces the high dimensionality of the data at the vertex level and improves the poor SNR in single subject data. One critical issue that needs to be addressed in order to ensure the success of such a clustering approach is the definition of the distance function that drives the k-means clustering. To this end, we propose a hybrid distance function based on rs-fMRI correlations and geodesic distance, which allows to cluster functionally homogeneous vertices, while enforcing spatial contiguity within a cluster. At the second level, we build a hierarchical tree on top of the pre-segmented regions to obtain individual parcellations reflecting the functional organisation of the cortex without losing the spatial integrity within the parcels. Rather than parcellating the cortex into a fixed number of regions, our framework provides parcellations within a range of 50 to 500 parcels, nested within a hierarchical tree, and thus, allowing the analysis of functional connectivity at different levels of detail.

The most closely related work to our approach is another two-level parcellation method composed of region growing and hierarchical clustering [32]. The major difference between two methods emerges from the first parcellation stage. We show that the proposed framework can provide a more reliable abstraction of the individual brain organisation, achieving higher repro-

ducibility and functional consistency than the other approach. In addition, we expand our experiments by including other RSFC-driven parcellations that are obtained by some widely-used clustering algorithms such as k-means, agglomerative hierarchical clustering [276] and spectral clustering with normalised cuts [54] as well as another region growing approach [27]. We show that connectivity-driven parcellations have the potential to provide a more reliable representation of the underlying data compared to traditional parcellations obtained from anatomical landmarks [61, 74] and random parcellations [214, 244].

The remainder of this chapter is organised as follows: In Section 4.2, we summarise the proposed parcellation method, with a particular emphasis on the initial parcellation and hierarchical clustering stages. In Section 4.3, we define several measures for the evaluation of parcellations and summarise other clustering techniques used for comparison. In Section 4.4 we show the performance of different approaches with respect to visual and quantitative results. Finally, in Section 4.5, we discuss the pros and cons of the proposed method with respect to the other clustering approaches, and provide some insight towards future research directions.

4.2 Methodology

4.2.1 Data

We conducted our experiments on Dataset 1, which contains rs-fMRI scans from 100 subjects. The details of the dataset are provided in Chapter 2 and briefly summarised here. The data for each subject was acquired in two sessions that were held on different days and divided into four runs of approximately 15 minutes each. All data was preprocessed and denoised by the HCP minimal preprocessing pipelines [87]. Following the preprocessing steps, all timeseries were normalised to zero-mean and unit-variance. For each subject, we temporally concatenated the 15-minute scans acquired on the same day into two 30-minute rs-fMRI datasets and used them to evaluate our approach.

4.2.2 A Two-Level Parcellation Framework

The proposed method starts with subdividing the cerebral cortex into highly homogeneous and relatively small regions using a modified k-means algorithm applied to the BOLD timeseries. This high-resolution parcellation is then submitted to a hierarchical clustering algorithm to obtain a multi-scale, nested set of subject-level parcellations. A visual representation of the parcellation framework is given in Fig. 4.1.

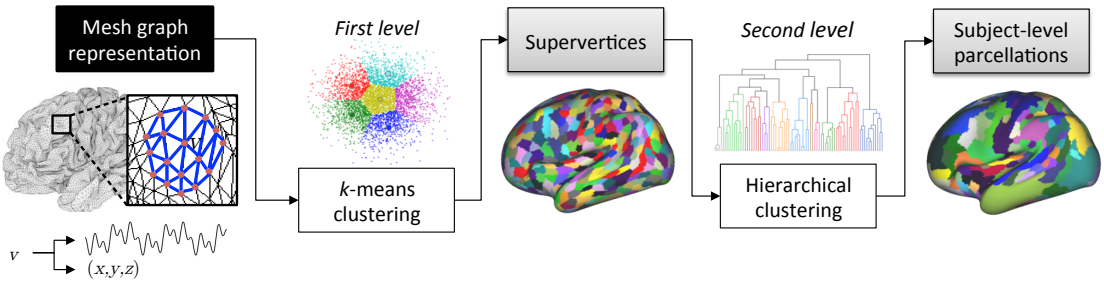


Figure 4.1: The two-level parcellation framework. *Initial representation:* The cerebral cortex is represented as a standard triangulated mesh, in which each vertex v is associated with a rs-fMRI timeseries and a set of anatomical coordinates (x, y, z) in the 3D space. *First level:* The cortex is parcellated into a relatively large number of supervertices using a k-means clustering variant. *Second level:* Subject-specific parcellations at different resolutions are obtained by merging supervertices with hierarchical clustering.

Initial Parcellation via Supervertex Clustering

The parcellation process begins with clustering the cortical vertices into a set of functionally uniform regions. Inspired by the superpixel approaches used in image segmentation [4], we develop a similar method to construct the initial clusters, i.e. *supervertices*. Our approach is a variation of the k-means clustering algorithm, but distinguishes itself from the other variants in two aspects: (1) We limit the search space of each supervertex to the expected average cluster size, which greatly reduces the number of distance calculations, thus improves the computational performance. (2) We define a hybrid distance function to measure the functional similarity between two vertices. This distance function is capable of assigning highly correlated vertices into same clusters, while enforcing spatial contiguity within a cluster.

The algorithm is designed to operate on the cerebral cortex, where the functional activity can be identified. We represent the cerebral surface as a smooth, triangulated mesh with no topological defects (Fig. 4.2). The mesh vertices and their associations are modelled as a weighted graph $G = (V, E)$, in which V is the set of vertices (nodes) and E is the set of edges connecting neighbour vertices. Each edge $e \in E$ is associated with a weight w that measures the Euclidean distance between two vertices with respect to their anatomical coordinates (x, y, z), i.e. for vertices v_i and v_j connected via e_{ij}, $w_{ij} = \|v_i - v_j\|_2 = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2 + (z_i - z_j)^2}$.

Initially, k seed vertices are selected as the supervertex centroids from V by uniformly subsampling the mesh graph using the iso2mesh library [69]. For every supervertex centroid S_i we compute its distance to the vertices which fall within a predefined range R of the supervertex centroid. Limiting the search space significantly increases the efficiency of the algorithm without affecting the final parcellations, since the cortical vertices are likely to be assigned to clusters within their local neighbourhoods due to spatial constraints imposed by the distance function.

The algorithm iteratively assigns each vertex to a supervertex by computing their dissimilarity with a hybrid distance function. At the end of each iteration, the geometric centres of newly-formed supervertices are chosen as the new centroids. The algorithm converges when all supervertices remain unchanged between two consecutive iterations. These steps are algorithmically summarised in Algorithm 1 and an illustration of the approach is given in Fig. 4.2.

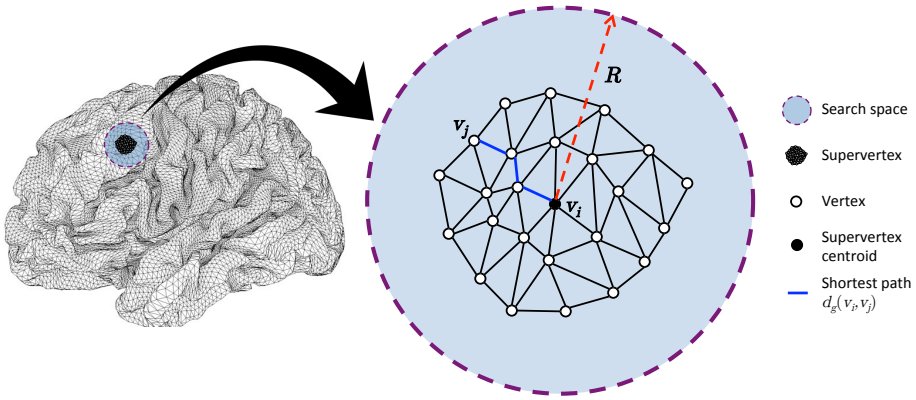


Figure 4.2: Illustration of the mesh graph that represents the cerebral cortex (*left*) and the supervertex clustering algorithm, emphasizing on a single supervertex (*right*). Note that the search space is not necessarily a circle, but displayed here as one for simplicity.

Algorithm 1: Supervertex Clustering

```
/* k initial supervertex centroids are selected by uniform sampling.   */
foreach vertex v do
    labels(v) ← 0
    distances(v) ← ∞
repeat
    changed ← false
    foreach supervertex centroid S_i do
        /* Distance calculated only for vertices within a range.       */
        foreach vertex v within a range of R around S_i do
            D = distance between S_i and v
            if D < distances(v) then
                distances(v) ← D
                labels(v) ← i
                changed ← true
    Compute the new supervertex centroids
until changed ≠ true
```

The supervertex clustering algorithm takes both the functional similarity and spatial proximity of the vertices into account while assigning cluster memberships. In order to avoid inconsistencies in clustering behaviour, the distance function should bring these different measures into a common space. We address this issue by defining separate distance functions for the functional and spatial domains, normalised by their maximal values within a cluster, and then combine them in a 2D Euclidean distance function.

In the functional domain, a vertex v_i is represented by its time series t_i. Functional similarity between two vertices v_i and v_j is measured by the Pearson's correlation coefficient r defined as follows:

$$r(t_i, t_j) = \frac{\sum_{l=0}^{L}(t_{il} - \mu_i)(t_{jl} - \mu_j)}{(L-1)\sigma_i\sigma_j} \tag{4.1}$$

where l $(0 \leq l \leq L)$ corresponds to a sample of a timeseries of length L, μ_i (μ_j) is the mean and σ_i (σ_j) is the sample standard deviation of t_i (t_j). r always produces values within the range [-1, 1]. The extrema at both ends indicate strongly (anti-)correlated time series. In order to

convert Pearson's correlation into a distance measure, we define the following transformation:

$$d_c(v_i, v_j) = 1 - r(t_i, t_j) \tag{4.2}$$

where d_c is referred as the Pearson's distance and yields values $0 \leq d_c \leq 2$. This transformation ensures the distance between highly correlated vertices being close to zero, thus increases their likelihood of being assigned to the same cluster.

In the spatial domain, instead of computing the distance directly from the spatial coordinates of two vertices, we propose to use the geodesic distance along the cortical surface, since it more naturally reflects the spatial geometry of the cerebral cortex. Given that each edge is subject to a weight, the geodesic distance d_g between v_i and v_j can be approximated as the sum of edge weights along the shortest path that connect v_i to v_j in the mesh graph, as illustrated in Fig. 4.2.

To combine two measures into a hybrid function, we introduce the functional and spatial normalisation factors, N_c and N_g, respectively. These normalisation factors are set to their corresponding maximal distance values in a cluster. We straightforwardly set N_c to 2, since the Pearson's distance values fall within the range $[0, 2]$. Similarly, N_g is set to the predefined local search limit, R, since the maximum distance within a cluster cannot exceed it. Finally, we combine these normalised measures in a Euclidean distance function D as follows:

$$D = \sqrt{\alpha {d_c}^2 + (1 - \alpha) {d_g}^2} \tag{4.3}$$

Here, α is introduced as a weighting parameter ($0 \leq \alpha \leq 1$) that controls the influence of functional dissimilarity over spatial closeness. We further investigate its effect on the parcellation performance in the following section.

Subject-Level Parcellation via Hierarchical Clustering

The supervertex clustering stage parcellates the cerebral cortex into a relatively large (1000 to 3000) number of highly homogeneous, non-overlapping regions. Although parcellations at this scale can be used for applications when accuracy is of great importance [54, 244], a second stage clustering must be performed to obtain more neuro-biologically interpretable sub-divisions of the cortex [54, 32].

Towards this end, we adapt an agglomerative hierarchical clustering algorithm [115] to join supervertices into non-overlapping parcels, without impairing their functional uniformity. This approach starts with each supervertex as singleton clusters and builds a hierarchy of clusters using a bottom-up strategy in which pairs of supervertices are merged into a single cluster with respect to a linkage rule. An illustration of the clustering algorithm is provided in Fig. 4.3.

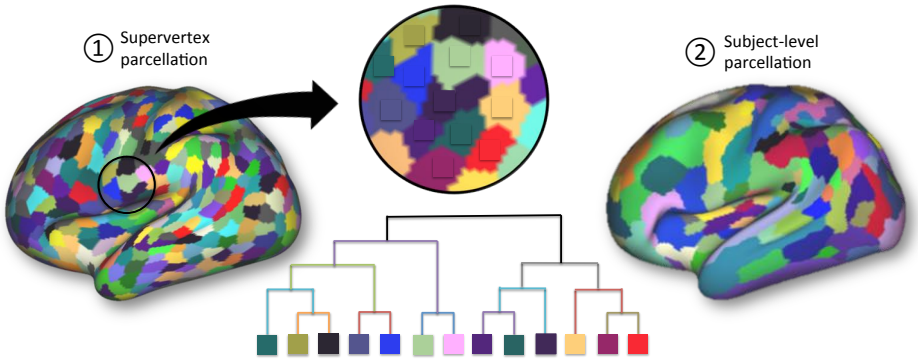


Figure 4.3: Illustration of the agglomerative hierarchical clustering algorithm.

In order to ensure the spatial coherency throughout the parcellation process, a connectivity constraint is imposed, so that only neighbouring clusters are merged into a higher level. Here, two clusters S_a and S_b are considered as neighbours, or adjacent, if vertices $v_i \in S_a$ and $v_j \in S_b$ are connected by a direct edge, e_{ij}, in the cortical mesh. An illustration is provided in Fig. 4.4 to make the neighbouring criterion more clear. For instance, the yellow, black, and blue supervertices in the figure are all adjacent to each other, since there is at least one edge that directly connects one to another. On the other hand, the pink and blue supervertices are only connected through a chain of neighbouring vertices, hence, are not considered as neighbours.

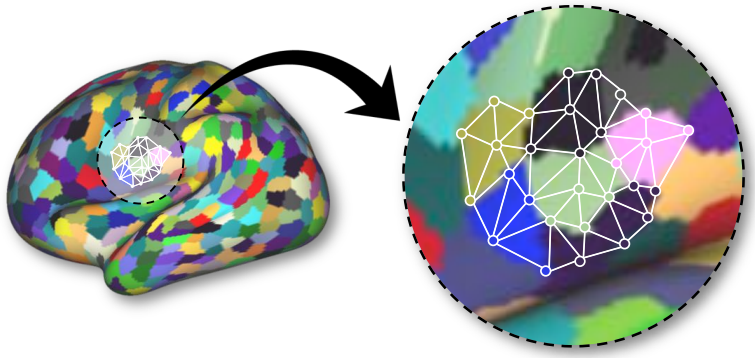


Figure 4.4: Illustration of the neighbourhood criterion for a set of supervertices. Two supervertices S_a and S_b are considered as neighbours, if vertices $v_i \in S_a$ and $v_j \in S_b$ are immediately connected in the cortical mesh.

The algorithm is driven by Ward's minimum variance method [276] and the similarity between pairing clusters is computed by the Euclidean distance. Among other linkage rules, we use Ward's method as we have observed that it yields the best quantitative results. Similar observations are reported in [32] and a wide range of fMRI-based parcellation methods in the literature are driven by Ward's linkage rule as well [94, 112, 227, 244]. The proposed hierarchical clustering technique produces a dendrogram [115], in which the leaves represent the supervertices and the root represents an entire hemisphere. Cutting this tree at different levels of depth produces individual subject parcellations with the desired precision. We investigate the effects of different granularities on the parcel reproducibility and functional reliability in the following section.

4.3 Experiments

4.3.1 Performance Measures

Evaluating the quality of parcellation methods is a challenging task since there is no ground-truth parcellation of the cerebral cortex. Considering this, we gather the most commonly used evaluation techniques from the literature to evaluate parcellations with respect to varying resolutions and quantitatively compare different clustering techniques to each other. These

techniques can be separated into three categories with regards to the parcellation aspects they assess: (1) reproducibility, (2) homogeneity, and (3) Silhouette analysis. These techniques are explained in detail in this chapter and used throughout the thesis for measuring the clustering quality of parcellations.

We also provide a qualitative assessment of the parcellations provided by the proposed two-level framework in terms of its agreement with task activation, the cytoarchitecture of the cerebral cortex, and well-defined structures of myelination.

Reproducibility

Reproducibility is a widely-accepted technique for evaluating the robustness of a parcellation method with respect to the underlying data/subjects. It measures the extent of alignment in parcellation boundaries between different parcellations. It can be used to evaluate parcellations obtained from a) different subjects (inter-subject reproducibility), b) the same subject but different rs-fMRI acquisitions (within-subject, or scan-to-scan reproducibility), and c) different groups (group-to-group reproducibility). Due to the high inter-subject variability within a population, it is not expected to obtain high reproducibility between different subjects. Nevertheless, a robust parcellation method should yield very similar parcellations for the same subject with different acquisitions. A high reproducibility should be expected of group level parcellations, assuming the group size is large enough.

We measure the reproducibility of parcellations using two well-known techniques, Dice coefficient [62] and adjusted Rand index (ARI) [103], each of which provides a means of assessing reproducibility from a different aspect. The former measures the amount of overlap at the parcel level, while ARI provides a more direct evaluation of the similarity between two parcellations by considering each as a whole.

Dice coefficient. The Dice coefficient [62] is a very popular measure of overlap between two labelled areas. It has been extensively used for evaluating brain parcellations [245, 54, 102, 32,

13, 188]. Given two parcels X and Y, the Dice coefficient is calculated as:

$$Dice = \frac{2|X \bigcap Y|}{|X| + |Y|} \tag{4.4}$$

where $|\cdot|$ indicates the number of vertices in a parcel. In order to obtain a global measure of parcellation reproducibility, we follow the approach proposed in [32]. We first compute Dice coefficients for every pair of parcels and match the ones with the highest overlap. The Dice coefficients of matching parcels are then averaged to obtain a global reproducibility score for the whole parcellation. The matching process is performed in an iterative manner, where matching pairs identified in one iteration cannot be matched with other parcels at the next iterations. This process is repeated until all pairs are identified. A Dice coefficient of 1 implies a perfect match (identical parcellations) and 0 indicates no match.

Low SNR in functional connectivity data or high variability within a group may yield a subdivision of some regions from one parcellation to the next, even when the same algorithm is performed on different acquisitions/subsets. To account for this effect and reduce its impact on reproducibility, we also use a modified version of Dice coefficient that merges the subdivided regions so as to maximise the overlap with the other parcellation as described in [32]. This is done by iteratively matching each parcel in one parcellation to those in the other, if their overlap ratio is ≥ 0.5 (i.e. one parcel comprises at least half of the other parcel). After this process, regions in one parcellation that are matched with the same parcel in the other are merged and the average Dice coefficient is computed between the matched pairs as described above. This measure will be referred to as the *joined* Dice coefficient in the remainder of this chapter.

Adjusted Rand index. ARI [103] is another technique for the evaluation of parcellation reproducibility [244]. In contrast to Dice coefficient, it measures the agreement of two parcellations without the necessity of initially matching parcels. As a result, it can more effectively

measure the agreement between two parcellations with different numbers of clusters [163].

ARI is built upon counting the number of items (in our case, vertices) on which two parcellations agree or disagree [268]. It classifies $\binom{M}{2}$ pairs of vertices into one of the four sets $(M_{11}, M_{00}, M_{01}, M_{10})$, based on their labeling in each parcellation. For parcellations $\mathbf{U}$ and $\mathbf{V}$, M_{11} corresponds to the number of pairs that are assigned to the same parcel in both $\mathbf{U}$ and $\mathbf{V}$, M_{00} corresponds to the number of pairs that are assigned to different clusters in both $\mathbf{U}$ and $\mathbf{V}$, M_{01} corresponds to the number of pairs that are assigned to the same parcel in $\mathbf{U}$, but different parcels in $\mathbf{V}$, and M_{10} corresponds to the number of pairs that are assigned to the same parcel in $\mathbf{V}$, but different parcels in $\mathbf{U}$. Intuitively, M_{00} and M_{11} account for the agreement of parcellations, whereas M_{01} and M_{10} indicate their disagreement [268]. After counting the number of pairs, ARI for parcellations $\mathbf{U}$ and $\mathbf{V}$ is computed as follows:

$$ARI(\mathbf{U}, \mathbf{V}) = \frac{2(M_{00}M_{11} - M_{01}M_{10})}{(M_{00} + M_{01})(M_{01} + M_{11}) + (M_{00} + M_{10})(M_{10} + M_{11})} \tag{4.5}$$

An ARI of 1 indicates a perfect correspondence between parcellations, whereas a value of 0 implies that the parcellations do not agree on any of the labels.

Homogeneity

Homogeneity is a highly popular parcellation evaluation technique [54, 218, 93, 13, 188, 102] that aims to measure the similarity of vertices aggregated in the same parcel, since a good parcellation should have the ability to group vertices with highly similar functional connectivity [54, 93]. A high homogeneity is particularly important for subsequent network analysis where network nodes are typically represented by the average signal (e.g. BOLD timeseries) within each parcel [218, 93]. Given a parcellation $\mathbf{U} = \{U_1, U_2, \ldots U_K\}$, the homogeneity h_k of a parcel U_k is measured by calculating the average similarity between every pair of vertices

assigned to U_k as

$$h_k = \frac{1}{n_k(n_k-1)} \sum_{i,j \in U_k, i \neq j} s(v_i, v_j) \tag{4.6}$$

where, n_k denotes the number of vertices in parcel U_k and $s(v_i, v_j)$ corresponds to the similarity between vertices v_i and v_j. While this similarity can be directly measured by comparing time-series of two vertices, we use an alternative approach and first define 'connectivity fingerprints' for each vertex as suggested in [54, 32]. A connectivity fingerprint (also known as connectivity profile) is a feature vector which indicates the similarity of a vertex with all other vertices, and hence, provides a global measure for how a vertex is connected to the rest of the cerebral cortex. When measuring functional similarity, such maps are computed by correlating timeseries of a vertex with the timeseries of other vertices. Two of these maps can then be compared by using Pearson's correlation, after being subject to Fisher's r-to-z transformation [49].

A global homogeneity value for $\mathbf{U}$ is finally obtained by averaging the homogeneity values across all parcels [54].

Silhouette analysis

Another useful and popular technique to quantify parcellation reliability is Silhouette coefficient (SC) [208], which can be used as an indicator of how well vertices fit in their assigned parcel. For each vertex, it compares the *within-parcel dissimilarity* defined as the average distance to all other vertices in the same parcel, to the *inter-parcel dissimilarity* obtained from those assigned to other parcels [245, 54]. SC not only evaluates the compactness of parcels, but also their degree of separation from each other. It is defined as follows:

$$SC_i = \frac{b_i - a_i}{\max(a_i, b_i)} \tag{4.7}$$

Given a parcellation $\mathbf{U} = \{U_1, U_2, \ldots U_K\}$, a_i and b_i correspond to within-parcel and inter-parcel dissimilarity of vertex $v_i \in U_k$, respectively, and are defined in the following equations.

$$a_i = \frac{1}{n_k - 1} \sum_{j \in U_k, i \neq j} d(v_i, v_j) \tag{4.8}$$

$$b_i = \frac{1}{M} \sum_{j \in \mathbb{N}(U_k)} d(v_i, v_j) \tag{4.9}$$

Here, n_k denotes the number of vertices in U_k, $\mathbb{N}(U_k)$ denotes the set of parcels that are neighbours with U_k, with $M = \sum_{j \in \mathbb{N}(U_k)} n_j$ being the number of vertices within these neighbouring parcels. $d(v_i, v_j)$ stands for the distance measure defined as $1-r$, where r is Pearson's correlation coefficient computed between the connectivity fingerprints of v_i and v_j. Instead of computing the inter-parcel dissimilarity with respect to the vertices in all other parcels, we restrict the computations to the neighbouring parcels. This is because (1) it is unlikely for a vertex to be assigned to a remote parcel due to spatial constraints imposed on the parcellations, and (2) computing inter-parcel dissimilarity with respect to all vertices outside a parcel quickly yields a bias towards obtaining high Silhouette coefficients, as the inter-parcel dissimilarity tends to be extremely high due to the many vertices with low similarity contributing to its computation.

Due to the fact that we use correlation distance as the dissimilarity measure, SC ranges within $[-1, +1]$. A negative SC implies misclassification of a vertex, while a value close to 1 indicates that the vertex is clustered with a high degree of confidence. If most vertices possess high Silhouette values, the parcellation is considered to be of high quality. A global score is obtained for each parcellation by averaging the Silhouette coefficients across all vertices.

4.3.2 Qualitative Assessment of Parcellations with Other Modalities

The previously proposed measurements assess the performance of a cortical parcellation from a clustering point of view and show how well the underlying connectivity data is effectively represented at a lower dimensionality. However, when defining regions of interest for neuro-anatomical purposes, the agreement of the parcellation boundaries with other neuro-biological properties also constitutes a critical aspect of parcellation quality. To this end, we provide additional visual results showing the agreement of the parcellation boundaries with several other cortical features, including myelin maps, task fMRI activation maps, and cytoarchitectural areas as defined by Brodmann.

fMRI data recorded while the subject is performing a behavioural or psychological task is typically used to discover the activated areas in the cerebral cortex [183]. Task-evoked fMRI images thus provide, although with limited coverage, a means to delineate the functional organisation of the brain. As a result, local agreements are likely to be observed between RSFC parcellation boundaries and highly activated cortical regions [32]. Similarly, a strong alignment has been observed between myelin maps and resting-state fMRI gradients [88]. We should therefore expect the boundaries of RSFC-driven parcellations to align with well-structured myelination patterns. We also assess the agreement of our parcellations with Brodmann's cytoarchitectonic areas [38]. Despite the fact that connectivity obtained from BOLD timeseries does not necessarily need to reflect the cytoarchitecture of the cerebral cortex, several studies report some alignment with certain cytoarchitectonic areas [32, 280, 93], such as the motor and visual cortex.

For our visual assessments, we use the complimentary data provided by the HCP. The Brodmann parcellations contain labels for the primary somato-sensory cortex (BA[3,1,2]), the primary motor cortex (BA4), the pre-motor cortex (BA6), Broca's area (BA[44,45]), the visual cortex (BA17 and MT), and the perirhinal cortex (BA[35,36]). For the comparisons with the task activation, we select the following subset of contrasts from five different task protocols: the RH-AVG and T-AVG contrasts of the MOTOR protocol; the STORY-MATH contrast of the LANGUGAGE protocol; the RANDOM contrast of the SOCIAL protocol, the PUNISH contrast of the GAMBLING protocol; the MATCH contrast of the RELATIONAL protocol.

4.3.3 Inter-Subject Variability Across Proposed Parcellations

The proposed parcellation scheme provides a subdivision of the cerebral cortex that is unique to each single subject. Due to high functional and structural variability across subjects as well as low SNR in the BOLD timeseries, it is not likely to obtain a high similarity across different subject-level parcellations. However, local agreements between parcellation borders are likely to be expected within some cortical areas, depending on the robustness of the registration algorithm used to align subjects.

In order to measure the extent of this agreement across subjects, we make use of the Dice-based overlapping measure previously introduced in Section 4.3.1. To this end, a 'consistency map' is computed for each subject by matching its parcellation with the other subject-level parcellations and each cortical vertex is assigned a score of 1 if it is within a matched parcel, and 0, otherwise. This procedure is repeated for all individuals and a consistency map is obtained on a per subject basis. We then compute a global consistency map for the entire cortex by averaging the maps across subjects. Cortical regions with high consistency scores are likely to be parcellated in a similar way (e.g. having parcels of similar size/shape) for most of the subjects .

4.3.4 Comparison Methods

We compare our method to a set of different parcellation techniques and assess the effectiveness of the two-level parcellation framework separately for each level. Initially, we evaluate the performance of the supervertex clustering scheme by comparing it with two alternative approaches based on region growing [27, 32]. At the subject level, different strategies are applied to the RSFC data to obtain individual subject parcellations, including a state-of-the-art two-level approach that combines region growing with hierarchical clustering [32], as well as popular clustering algorithms adopted for the RSFC data, namely, k-means, Ward's hierarchical clustering [276], and spectral clustering with normalised cuts [255, 54]. Finally, we also include parcellations obtained from non-connectivity data to show the potential of the connectivity-

Name	Reference	Resolution	Description
Blumensath	Blumensath et al. (2013); re-implemented as described in the original paper.	Varying	A two-level method that combines region growing and Ward's hierarchical clustering.
Bellec	Bellec et al. (2006); codes available from www.nitrc.org/projects/niak/	Varying	A competitive region growing approach driven by parcel homogeneity. A size threshold is applied to avoid over-growing of parcels.
K-Means	K-means clustering as described in Thirion et al. (2014); in-house implementation, featuring scikit-learn's KMeans and PCA functions.	Varying	K-means clustering is applied to a low dimensional representation of the BOLD timeseries, obtained by PCA. Only the components that explain 50% of the variance are kept for the sake of efficiency. No spatial constraints are explicitly imposed.
Ward	Ward (1963); in-house implementation, featuring scikit-learn's AgglomerativeClustering function.	Varying	A hierarchical tree is built by merging pairs of clusters according to Ward's linkage rule. Only adjacent clusters are joined into a higher level in order to ensure the spatial contiguity.
N-Cuts	Craddock et al. (2012); in-house implementation of spectral clustering with normalised cuts, featuring discretisation (Yu and Malik 2003).	Varying	Spectral clustering with normalised cuts. An affinity matrix is built by cross-correlating the adjacent vertices with each other. Spectral decomposition is applied to the normalised graph Laplacian and the parcellations are obtained by discretising the eigenvectors.
Destrieux	Fischl et al. (2004); available as part of the HCP datasets (VanEssen et al., 2013).	150 (75 L, 75 R)	A surface-based parcellation that subdivides the cortex according to the limit between the gyral and sulcal regions.
Desikan	Desikan et al. (2006); available as part of the HCP datasets (VanEssen et al., 2013).	70 (35 L, 35 R)	Another surface-based parcellation that subdivides the cortex with respect to anatomical landmarks based on the gyri.
Geometric	Geometric parcellations as described in Thirion et al. (2014).	Varying	K-means clustering is applied to the spatial vertex coordinates. No connectivity information is accounted for.
Random	Random parcellations as described in Schirmer et al. (2015).	Varying	Poisson disk sampling is used to generate regions of approximately equal size by ensuring that two region centres are not closer than a given threshold that controls the number of parcels.

Table 4.1: Summary of the parcellation methods used for comparison.

driven methods for computing more reliable subject level parcellations. These methods include anatomical parcellations [74, 61], random parcellations [214], and geometric parcellations [244]. All methods are described below and summarised in Table 4.1. We rely on the following method descriptions and implementations for the comparison methods throughout the thesis, unless otherwise noted.

Blumensath's region growing [32]. This two-level method combines region growing and Ward's hierarchical clustering to derive subject-level parcellations of the cerebral cortex. At the initial level, each seed vertex is grown into a region (cluster) by an iterative process based on the similarity between the cluster centroids and the adjacent vertices to them. In order for a cluster to obtain a vertex, the correlation between the vertex' timeseries and the cluster's timeseries should exceed p times the maximal correlation between all other regions' timeseries and their associated neighbourhood vertex timeseries [32].

After all vertices have been assigned to a cluster, subject-level parcellations are obtained by merging adjacent clusters using hierarchical clustering with Ward's linkage rule [276] as in our proposed two-level framework. However, this approach differs from our method in three ways. (1) Instead of relying on a fixed number of seed vertices to start the parcellation process, it *a priori* locates potentially homogeneous cortical regions that can be used to compute the seeds. (2) These seeds are used as the cluster centroids, but never updated during the region growing process. (3) The representative timeseries of a cluster is computed by averaging the timeseries within the immediate neighbourhood of the seed vertex, instead of typically including all within-cluster vertices.

Bellec's region growing [27]. This region growing approach is originally proposed to reduce the spatial dimensionality of the RSFC data. The objective function is based on maximising the global homogeneity of the parcellation. Parcel homogeneity in this context is defined as the average correlation between the timeseries associated with any pair of vertices within a region.

Since this method is driven by homogeneity, a size constraint (t) is imposed in order to avoid having non-uniformly evolved parcellations in terms of parcel size (e.g. having many small, highly homogeneous parcels with a few very large parcels across the cortex) [27]. During the region growing process, regions that exceed the size parameter t are validated and excluded from further consideration. This ensures that the size of a validated region cannot exceed $(2t - 2)$, as it is possible that two regions of size $(t - 1)$ can still be merged during the growing process. Different from the original paper, we specify the number of regions beforehand in order to be consistent with the other approaches; hence the algorithm keeps merging the smallest adjacent regions until the desired resolution has been reached.

***K*-means algorithm [244].** As an alternative technique, the well-known k-means clustering algorithm is used to compute subject-level parcellations for varying resolutions. To overcome the computational limitations inherent to *k-means* due to high dimensionality, we follow the instructions in [244] and reduce the dimensionality of the data for each subject via PCA, capturing about 50% of the variance. K-means naturally does not compute spatially contiguous

parcellations, as the spatial structure of the data is not accounted for. Therefore, resulting parcellations can consist of extremely disjoint regions.

Ward's hierarchical clustering [276]. We cluster the RSFC data using another popular clustering algorithm, Ward's agglomerative hierarchical clustering [276]. As opposed to k-means, spatial coherency within parcels is explicitly ensured by only merging adjacent clusters at each iteration. Different from the proposed approach and [32], this method is directly applied to the vertex-level RSFC data, instead of relying on an initial fine-resolution parcellation.

Spectral clustering with normalised cuts [54]. A spectral clustering technique based on normalised cuts is applied to the connectivity data as described in [54]. An affinity matrix W is built by cross-correlating vertices with each other. In order to obtain spatially contiguous parcellations, we only retain the edges constructed between adjacent vertices and discard the others. After ensuring that the resulting affinity matrix is fully-connected and positive-semidefinite (i.e. all $W_{ij} \geq 0$), we apply spectral decomposition to the normalised graph Laplacian, defined as $L = D^{-1/2}(D - W)D^{-1/2}$, where D is a diagonal matrix with each entry $D_{ii} = \sum_j W_{ij}$ representing the degree of vertex i. The eigenvectors corresponding to the K smallest non-zero eigenvalues are clustered via discretisation [285] to obtain the final parcellations.

Anatomical parcellations. We also assess the performance of surface-based anatomical parcellations in order to represent the functional organisation of the brain at the subject level. These atlases are distributed as part of the HCP datasets [61, 74] and tailored to each individual subject with respect to anatomical features, such as cortical folding.

The Destrieux atlas [74] subdivides the cerebral cortex into a fixed number of 75 parcels per hemisphere according to the limit between the gyral and sulcal regions, determined by the curvature of the surface. The Desikan atlas [61] offers a coarser anatomical parcellation (35 parcels per hemisphere) based on gyral landmarks, where a gyrus is defined as the visible part on the pial surface and limited by the adjacent banks of the sulci [61].

Other parcellations. In addition, we include two other approaches to our experiments that do not account for any connectivity information, with the aim of having a baseline towards data-driven approaches. First, we compute random parcellations using Poisson disk sampling as described in [214]. Poisson disk sampling is used to generate regions of approximately equal size by ensuring that two region centres are not closer than a given threshold that controls the number of parcels. Second, we obtain geometric parcellations by applying k-means algorithm to the anatomical vertex coordinates as described in [244]. Since the parcellation is directly driven by spatial proximity between cortical vertices, it is likely to obtain regularly shaped parcels, with roughly similar size.

4.3.5 Experimental Setup

All subjects in our cohort are represented by two 30-minute rs-fMRI datasets (rs-fMRI 1 and rs-fMRI 2). We perform reproducibility analysis for each subject by comparing their parcellations obtained from rs-fMRI 1 and rs-fMRI 2 (i.e. scan-to-scan reproducibility). Clustering accuracy (homogeneity and Silhouette analysis) is computed based on the first parcellation, but using data from rs-fMRI 2, to avoid possible biases that may have emerged during the computation of the parcellations. There is only one random, *Desikan*, and *Destrieux* parcellation is available for each subject, therefore we skip the reproducibility analysis for these methods. A summary of the evaluation pipelines is given in Fig. 4.5.

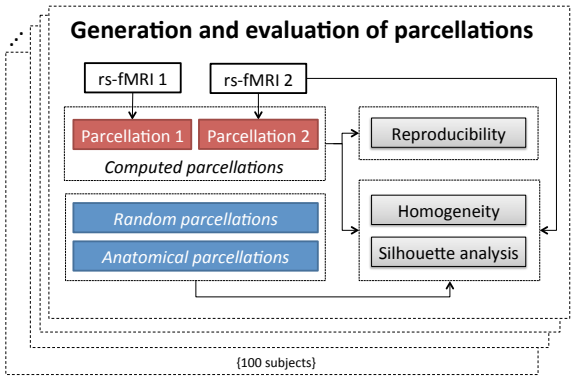


Figure 4.5: Visual outline of the parcellation evaluation pipelines.

4.4 Results

4.4.1 Parameter Selection

Other than the number of parcels K, the proposed algorithm has two external model parameters: (1) the weighting parameter α, and (2) the number of supervertices N.

The weighting parameter α is used to determine the influence of functional connectivity over spatial proximity, and thus, has a direct impact on the supervertices. Whereas values of α closer to 1 tend to group more functionally correlated vertices together, the spatial contiguity of the clusters can be severely affected, reducing the reproducibility of the final parcels. Decreasing α to an extent would produce more regularly-shaped supervertices, but ultimately may lead to lower homogeneity, as the clustering is only driven by the spatial closeness of vertices to supervertex centroids. Given this trade-off, we set α to 0.7, which offers the highest homogeneity without violating the spatial integrity of supervertices. The impact of α over the spatial contiguity of a supervertex can be seen in Fig. 4.6.

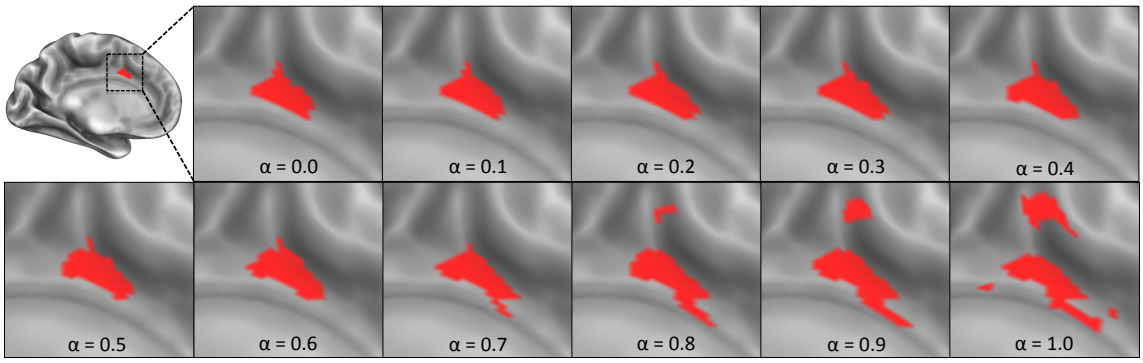


Figure 4.6: Evolution of the boundaries of a supervertex with different values of α.

In order to show that this tendency holds across all supervertices/subjects, we located the disjoint supervertices and counted the number of vertices that are not spatially connected to their assigned supervertex using the neighbouring criterion introduced in Section 4.4. The results obtained for different numbers of supervertices are given in Fig. 4.7. As can be seen in the figure, $\alpha = 0.7$ stands as a cut-off point where supervertices start to become disjoint due to higher influence of functional similarity over spatial constraints.

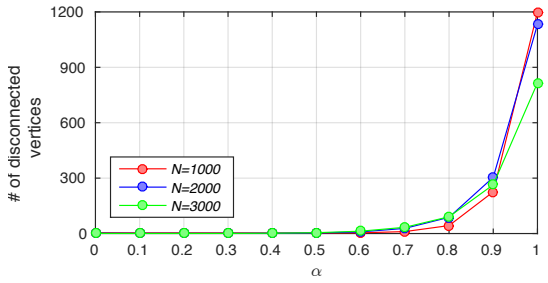


Figure 4.7: The number of vertices that are not spatially linked to the largest connected component in their associated supervertices with respect to different values of α and N.

The number of supervertices, N, is the second external parameter imposed to the method. It provides a means of spatial dimensionality reduction, and therefore, can help alleviate the impact of noise at the vertex level. While selecting a very low N may provide more reproducible parcellations, it may lead to over-smoothing of the underlying functional data (i.e. low homogeneity). Increasing N gradually yields more homogeneous supervertices (and hence parcellations), but at the cost of lower reproducibility. The accuracy (by means of reproducibility and homogeneity) of the parcellations with respect to different number of supervertices is shown in Fig. 4.8. Given this trade-off, we initially parcellate the cortex into approximately 1000 supervertices per hemisphere (i.e. $N = 2000$ initial regions per subject).

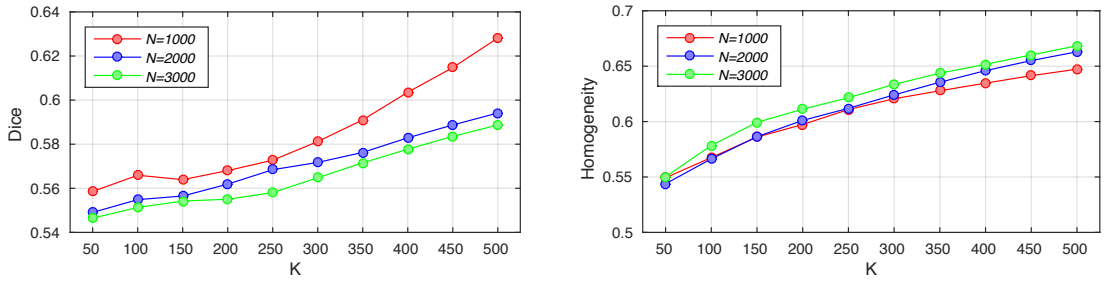


Figure 4.8: Reproducibility (*left*) and homogeneity (*right*) as a function of the number of regions (K) for different number of supervertices (N).

Two other methods, *Blumensath* and *Bellec*, are also driven by external parameters. The p parameter in *Blumensath* determines the degree of similarity between vertices and regions competing for them. Based on our observations and the discussion in the original paper, we set this parameter to 0.9, which significantly increased the computational speed of the approach with only a minimal impact on the final parcellations [32]. The other region growing approach,

Bellec relies on a size constraint t that avoids over-growing of regions [27]. For each supervertex resolution we set t to the expected average parcel size (i.e. number of vertices / number of supervertices).

4.4.2 First-Level Clustering Results

We first provide the evaluation results for the initial parcellation stage (i.e. first level). We assess the performance of the supervertex clustering scheme (denoted as *Arslan*), against two functional connectivity-driven region growing approaches (*Blumensath* and *Bellec*) and geometric (*Geometric*) parcellations, which do not take any connectivity information into account, and hence, provide a baseline for comparison. We test the methods with respect to three different resolutions of approximately N= 1000, 2000, and 3000 regions. We determine the exact number of regions for each subject from the *Blumensath* [32] parcellations, as it is the only method that is driven by seed vertices generated from the underlying data, rather than a set of pre-determined centroids. Using the same initial parcellation resolution for each method facilitates their comparison and avoids unfairness with respect to the number of regions which may have emerged otherwise. Parcellations of the lateral cortex of the left hemisphere derived from one subject for a resolution of $N = 1000$ regions are given in Fig. 4.9 for visual inspection.

Scan-to-scan reproducibility results computed by Dice coefficients and ARI are given in Fig. 4.10. It is immediately clear from the figure that, the proposed supervertex clustering approach (*Arslan*) and *Blumensath* yield the most reproducible parcellations across all resolutions. Although similar Dice results are achieved by both methods, *Arslan* appears to produce more robust regions with respect to ARI. *Bellec* can not generate reproducible regions, most likely due to the fact that it does not rely on any spatial constraints in the region growing process. *Geometric* is purely driven by the spatial proximity of the parcels, but due to randomness of the k-means algorithm, the resulting parcellations tend to be very different.

Homogeneity values and Silhouette coefficients indicating the fidelity of the parcellations to the underlying data are given in Fig. 4.11. Contrarily to reproducibility, stark differences in performance are not observed across methods, which even tend to perform more equivalently with

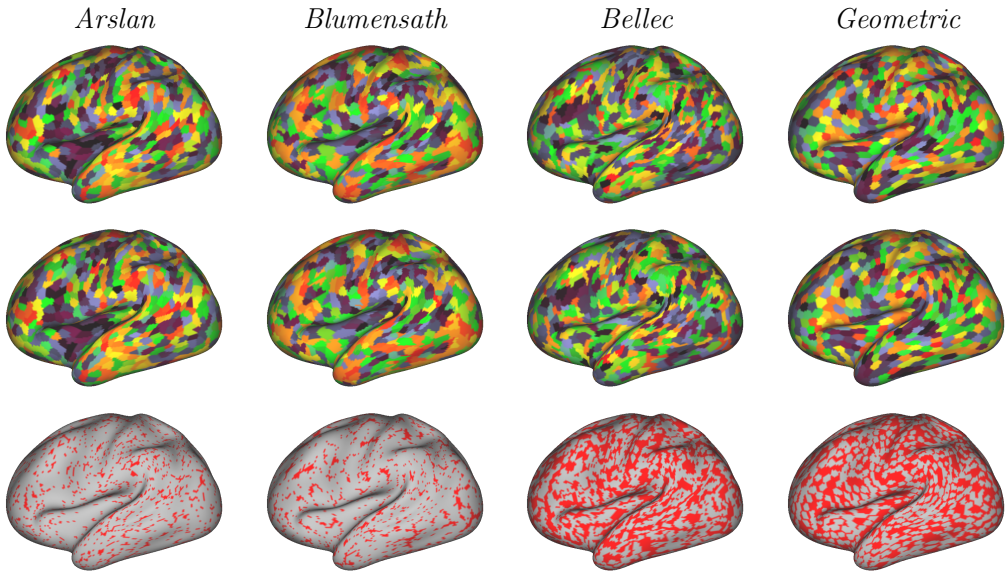


Figure 4.9: Parcellations of the lateral cortex of the left hemisphere derived from one subject using different approaches for $N = 1000$ regions. The parcellations in the first and second row were obtained from different scans of the same subject in order to evaluate scan-to-scan reproducibility. Parcel colours were matched for better visualisation and easier comparison. The last row shows the non-matching regions.

increasing resolution. This is an expected result given the high resolution of the parcellations, which consist of very small regions. Since the BOLD timeseries are smoothed to some extent during preprocessing, this yields an inherent correlation between vertices that are spatially close to each other.

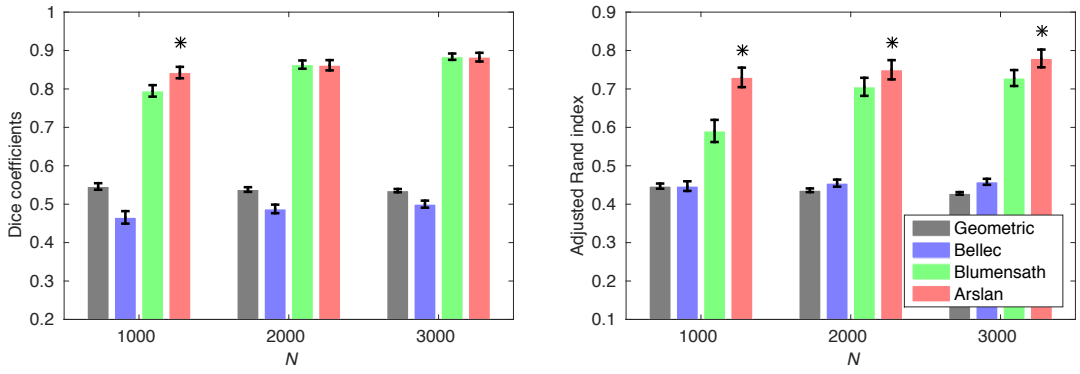


Figure 4.10: First-level reproducibility results obtained using Dice coefficient (left) and adjusted Rand index (right). Whiskers in the bars indicate the variability across repetitions as measured by standard deviation. Stars (*) show statistical significance between the winner and the runner-up based on Wilcoxon signed rank test with $p < 0.01$.

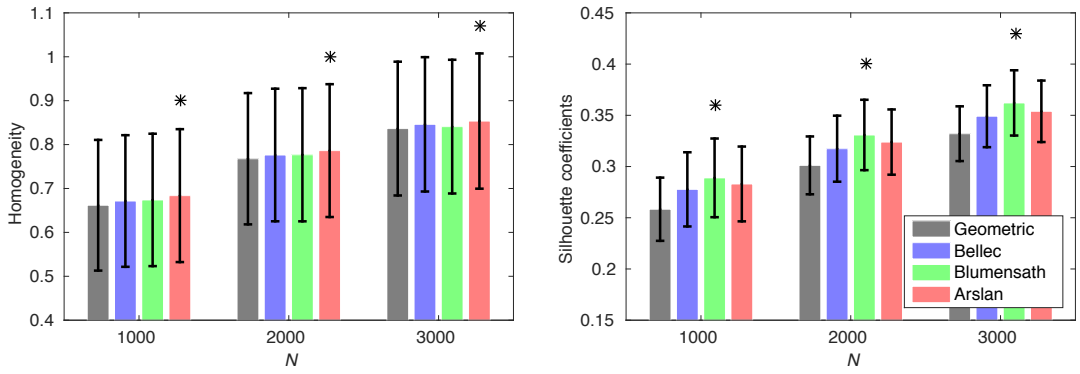


Figure 4.11: First-level homogeneity values (left) and Silhouette coefficients (right). Whiskers in the bars indicate the variability across repetitions as measured by standard deviation. Stars (*) show statistical significance between the winner and the runner-up based on Wilcoxon signed rank test with $p < 0.01$.

Nevertheless, *Arslan* and *Blumensath* consistently perform better than the others in terms of homogeneity and Silhouette analysis, respectively. Although driven by a criterion that promotes homogeneity in regions, *Bellec* performs worse than the other two connectivity-driven approaches, most likely due to the size constraint that avoids over-growing. As expected, *Geometric* shows the worst performance, indicating that parcellations driven by the spatial proximity of cortical vertices do not fit the underlying connectivity well, even at very fine resolutions.

4.4.3 Subject-Level Parcellation Results

We next provide the evaluation results for the second parcellation stage (i.e. subject-level parcellation). For ease of comparison between different methods, we report average evaluation measures in the form of line graphs for all computed resolutions. In order to show the variability across individuals we provide box plots alongside the line graphs, but only for a subset of granularity levels (i.e. for 100, 200, and 300 regions). Parcellations of the lateral cortex of the left hemisphere derived from one subject for a resolution of $K = 100$ regions are given in Fig. 4.12 for visual inspection. Reproducibility results are provided in Fig. 4.13. Cluster validity results, including homogeneity values and Silhouette coefficients, are presented in Fig. 4.14 and 4.15, respectively.

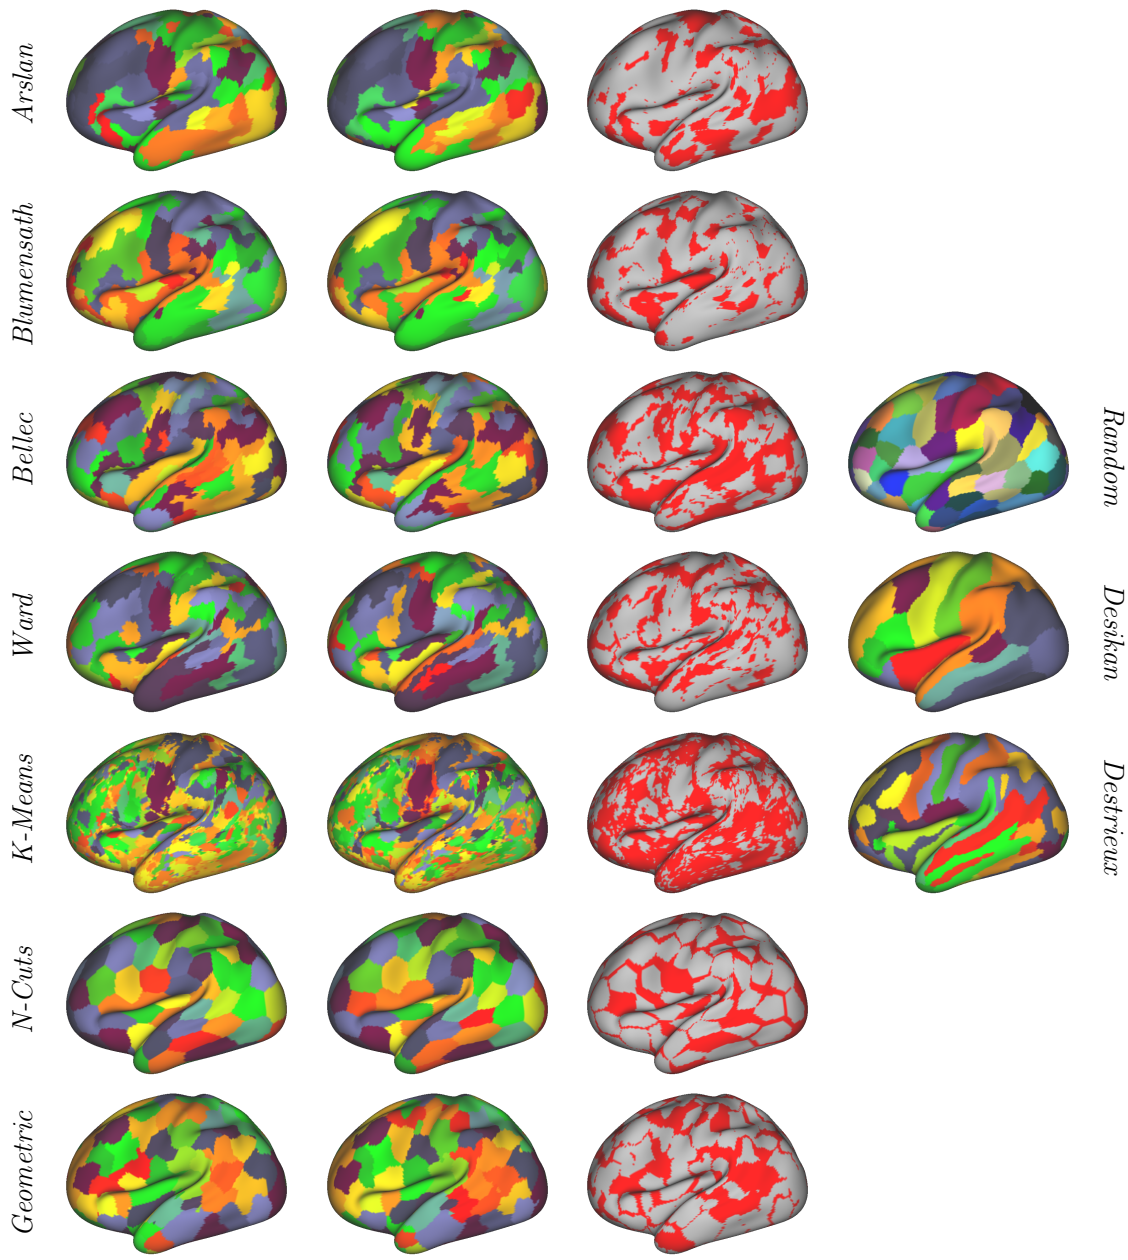


Figure 4.12: Parcellations of the lateral cortex of the left hemisphere derived from one subject for $K = 100$ regions using different approaches. The parcellations in the first and second column are obtained from different scans of the same subject in order to evaluate scan-to-scan reproducibility. Parcel colours are matched for better visualization and easier comparison. The third column shows the non-matching regions. The last column shows the two anatomical parcellations, *Desikan* and *Destrieux*, with 35 and 75 regions, respectively and one random parcellation with the same resolution as the connectivity-driven parcellations.

Reproducibility results computed by the Dice coefficient and ARI indicate that *Geometric* and *N-Cuts* yield the most reproducible results. Although *Geometric* shows a better performance

than *N-Cuts* at relatively low resolutions, this trend shifts towards the favour of *N-Cuts* at higher resolutions. The performance of *N-Cuts* can be explained by the hard spatial constraints imposed to the adjacency matrices that drive the spectral clustering algorithm, which promotes uniformly-sized parcels [54, 32]. Obtaining highly reproducible parcellations for *Geometric* is also expected, as the parcellations of a subject are generated from the same set of spatial coordinates.

Hierarchical clustering applied directly to the underlying BOLD timeseries (*Ward*) shows a poor performance, in particular when compared to the other two hierarchical methods derived from an initial finer parcellation, namely, *Arslan* and *Blumensath*. It is worth noting that Dice overlap measurements indicate a more favourable performance for *Blumensath* with respect to *Arslan*, while a reverse trend is observed in ARI. This can be attributed to the fact that, *Blumensath* parcellations rely on a set of cluster centroids that are learned from the underlying data and are not updated during the region growing process. This inherent bias is propagated to the second level and yields higher Dice scores between pre-matched parcels. On the other hand, higher ARIs indicate a better reproducibility on a cortex-wise scale. In general, results obtained from these two-level approaches suggest that methods initialised with a finer parcellation may be more robust, which can be linked to the fact that the impact of noise is reduced by the initialisation scheme. *Bellec* generally yields low reproducibility scores, but still shows a better performance than *K-Means*. Nevertheless, it should be noted that this method is originally developed to obtain parcellations with much finer resolutions (over 1000 regions per hemisphere) [27], hence, it may not be adapted to this range of resolutions. *K-Means* unsurprisingly yields the least reproducible parcellations, most likely due to the high degree of spatial discontiguity inherent in parcellations obtained via this clustering technique, which is purely driven by the affinity of the underlying BOLD timeseries. Spatially disjoint regions in *K-Means* parcellations can be seen in Fig. 4.12.

As expected, the Dice coefficient is strongly increased by merging subdivided regions. In particular, this process yields more favourable results for the methods based on hierarchical clustering, i.e. *Ward*, *Arslan* and *Blumensath*, for which the improvement is up to 15%. *Blumensath* even surpasses *N-Cuts* and *Geometric* for resolutions with more than 150 parcels, becoming the top

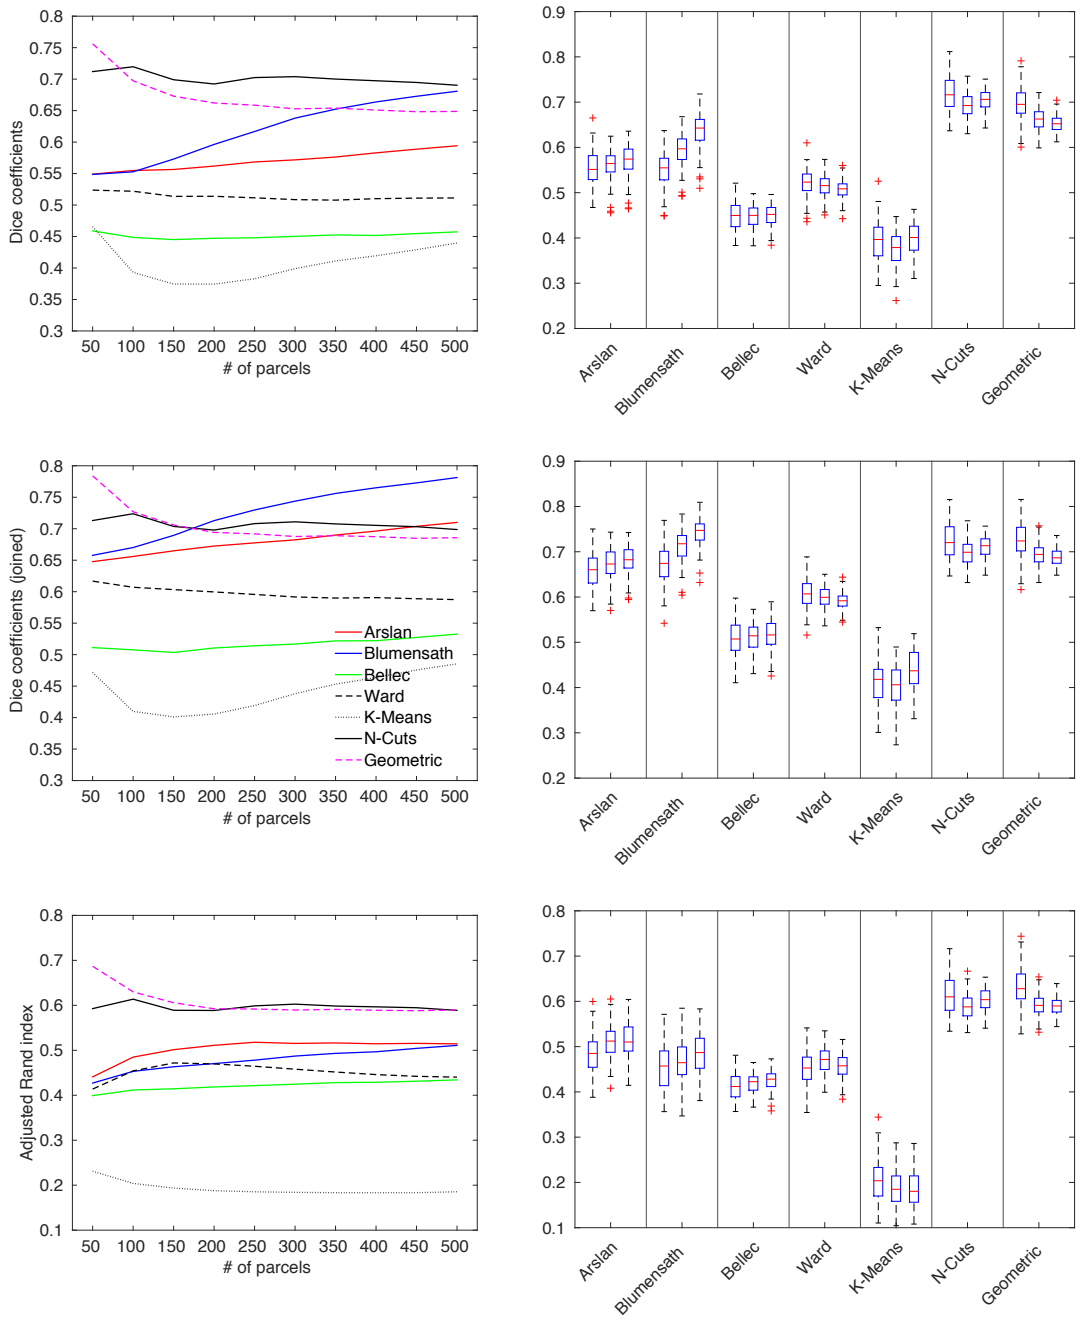


Figure 4.13: Subject-level reproducibility results. *Left*: Average reproducibility values obtained using Dice coefficient (top), joined Dice coefficient (middle), and adjusted Rand index (bottom). *Right*: Box plots indicate the reproducibility distribution across subjects for 100, 200, and 300 regions, from left to right, for each method.

performing method regarding reproducibility. Other approaches tend to have a less significant improvement, mostly at a rate of $5-8\%$, while *N-Cuts* and *Geometric* are minimally affected.

This trend can be attributed to the fact that hierarchical clustering subdivides the cortex with a bottom-up process, where boundaries derived at lower resolutions are propagated to higher levels. Joining over-parcellated regions may therefore increase the similarity between parcellations that subdivide the same regions at different levels of the hierarchical clustering tree.

Cluster validity measurements show a clear tendency in favour of connectivity-driven approaches. The most prominent trend is that, regardless of the parcellation resolution, *K-Means* outperforms all other methods in terms of both homogeneity (Fig. 4.14) and Silhouette analysis (Fig. 4.15). This would indicate that *K-Means* generates the best clustering of the underlying data. It is followed by the hierarchical approaches, each of which performs almost equally regarding Silhouette coefficients, while *Ward* is the best with respect to homogeneity. In particular, *Arslan* consistently generates more homogeneous parcellations than *Blumensath*, which might be attributed to the different techniques used by each method for computing an initial parcellation of the cerebral cortex before applying hierarchical clustering. We represent the cortical surface by functionally uniform supervertices and incorporate a flexible spatial constraint into our distance function, enabling any vertex to be assigned to a cluster if they exhibit high correlation and are spatially close (but not necessarily neighbours) as opposed to the *Blumensath*'s region growing, which is based on stable seed points and a more strict definition of spatial proximity. In general, this initial stage helps obtain parcellations with a slightly higher degree of confidence than *Ward*.

Amongst the connectivity-driven parcellations, *N-Cuts* shows the poorest performance. This can be due to the size bias inherent in this parcellation scheme that could limit the agreement with the underlying data. On the other hand, anatomical parcellations *Desikan* and *Destrieux*, yield the worst measurements and are surpassed by *Geometric* and *Random*. This might suggest that anatomical information alone does not allow to map the brain's functional organisation.

All methods show a performance increasing with the number of parcels computed. This is linked to the fact that evaluation measurements depend on the size of the parcels (e.g. smaller parcels yield better results). It should be noted that this trend may benefit the *K-Means* parcellations, which comprise many small discontinuous parcels.

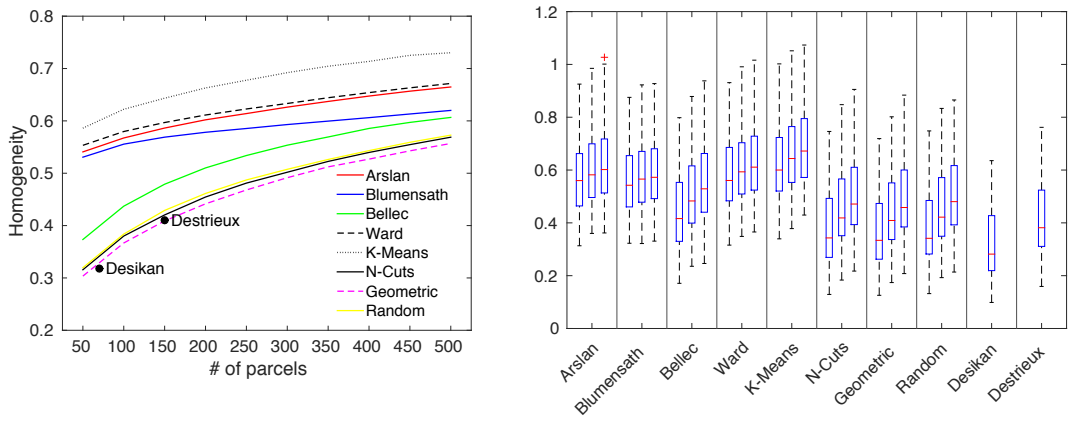


Figure 4.14: Subject-level homogeneity results. *Left*: Lines show homogeneity values for all resolutions, averaged across subjects, whereas black dots correspond to the average homogeneity obtained from the *Desikan* and *Destrieux* atlases, at a fixed resolution of 70 and 150 parcels, respectively. *Right:* Box plots indicate the homogeneity distribution across subjects for K = 100, 200, and 300 parcels, from left to right for each computed method.

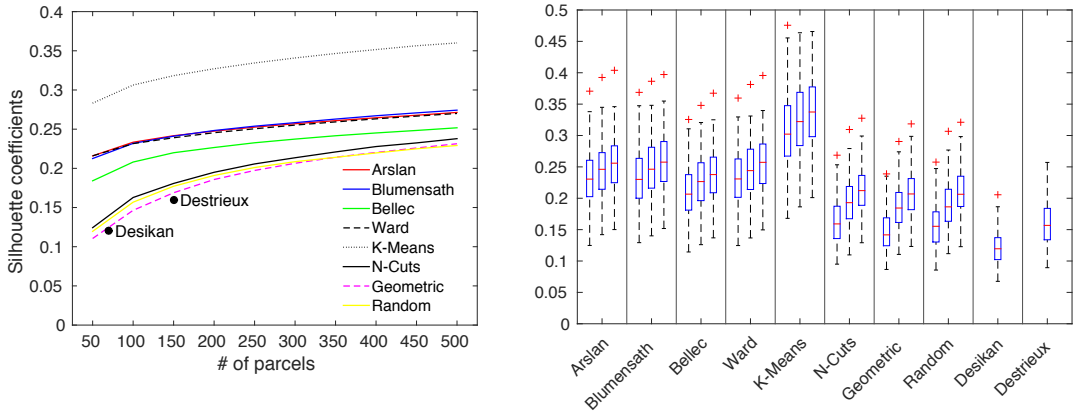


Figure 4.15: Subject-level Silhouette analysis results. *Left*: Lines show Silhouette coefficients (SC) for all resolutions, averaged across subjects, whereas black dots correspond to the average SC obtained from the Desikan and Destrieux atlases, at a fixed resolution of 70 and 150 parcels, respectively. *Right:* Box plots indicate the SC distribution across subjects for K = 100, 200, and 300 parcels, from left to right for each computed method.

Another important observation is the higher inter-subject variability of homogeneity and Silhouette analysis results compared to reproducibility. While one can infer that cluster validity measures are more sensitive than Dice coefficients, this could also be attributed to the fact that reproducibility measures the spatial similarity of parcellations that have been registered onto the same standard cortical surface; as a result, an inherent alignment already exists across subjects. This yields a lower inter-subject variability, especially for the spatially constrained

methods and with respect to increasing resolution. On the other hand, functional organisation of the brain as estimated by RSFC can dramatically change from one subject to the next and even between different acquisitions of the same subject. Combining this with the impact of low SNR inherent to rs-fMRI, it may not be possible to parcellate all subjects with high homogeneity and/or confidence. This can be a critical point for consideration, for example, when a group-level study is devised.

4.4.4 Inter-Modality Assessment of the Proposed Parcellations

While the previously given measurements assess the performance from a clustering point of view, the consistency of the parcellation boundaries with the task-related activated regions as well as other neuro-biological properties, such as cytoarchitecture and myelination, constitute another important aspect of evaluation. It provides a means of assessing the reliability of functional parcellations to represent the cortical organisation of the brain as delineated by other features.

To this end, we present example task activation maps of three different subjects in Figs. 4.17 and 4.18. Myelin maps and cytoarchitectonic regions of the same subjects are given in Fig. 4.19. The white borders in the figures show the parcellation boundaries obtained by the proposed two-level parcellation algorithm at a granularity level of $K = 100$ regions. In order to better interpret the consistency of the agreement between different modalities, we show the consistency maps obtained for the proposed parcellations at the same resolution in Fig. 4.16.

It appears that the proposed parcellations show a good degree of alignment with the task-activated areas and well-structured patterns of myelination. This agreement is more prominent within the primary motor, somato-sensory, and visual cortex as indicated by the arrows. In addition, the boundaries around area MT are also found to align well with the changes in the myelo-architecture of the cortex and some activation maps. Although we only present visual results for several subjects, similar results are likely to occur for other subjects, especially within the motor area, which is typically subdivided into alike regions across subjects as indicated in the inter-subject consistency maps in Fig. 4.16.

Our findings regarding the similarity between the rs-fMRI parcellations and sharp changes in myelin maps agree with the previous literature, where strong alignment has been observed between the highly myelinated cortical areas and resting-state fMRI gradients [88]. These observations might also indicate that the connectivity estimated at rest may truly reflect the functional organisation of the brain, especially within certain areas of the cortex. Although the functional connectivity obtained from BOLD timeseries does not necessarily match the cytoarchitecture of the cerebral cortex, some alignment is observed in Brodmann's areas BA17 (part of visual cortex) and BA[4,6] (motor cortex), for which several other studies have also reported similar observations [32, 280, 93].

4.5 Discussion

In this chapter, we proposed a new two-level clustering approach to parcellate the cerebral cortex using publicly available data provided by the HCP. The generation and evaluation of the parcellations is based on RSFC, which is thought to represent the functional organisation of the brain on a whole-cortex scale. As there is no gold standard parcellation of the cerebral cortex, we evaluated the reliability of parcellations in terms of two widely-accepted criteria [244]: (1) reproducibility (i.e. stability) of a parcellation on a scan-to-scan basis, and (2) fidelity (i.e. accuracy) of a parcellation with respect to the underlying data. We tested the proposed approach with respect to different parcellation schemes and popular clustering methods on

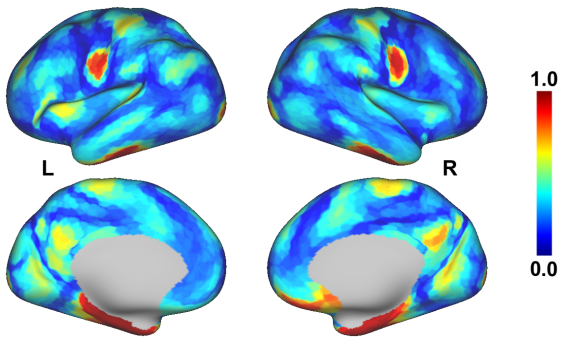


Figure 4.16: Inter-subject consistency maps obtained from all subjects in the dataset for a granularity of 100 regions. Colour maps are normalised to $[0, 1]$ for better visualisation. Hotter colours indicate higher consistency between parcels across subject-level parcellations.

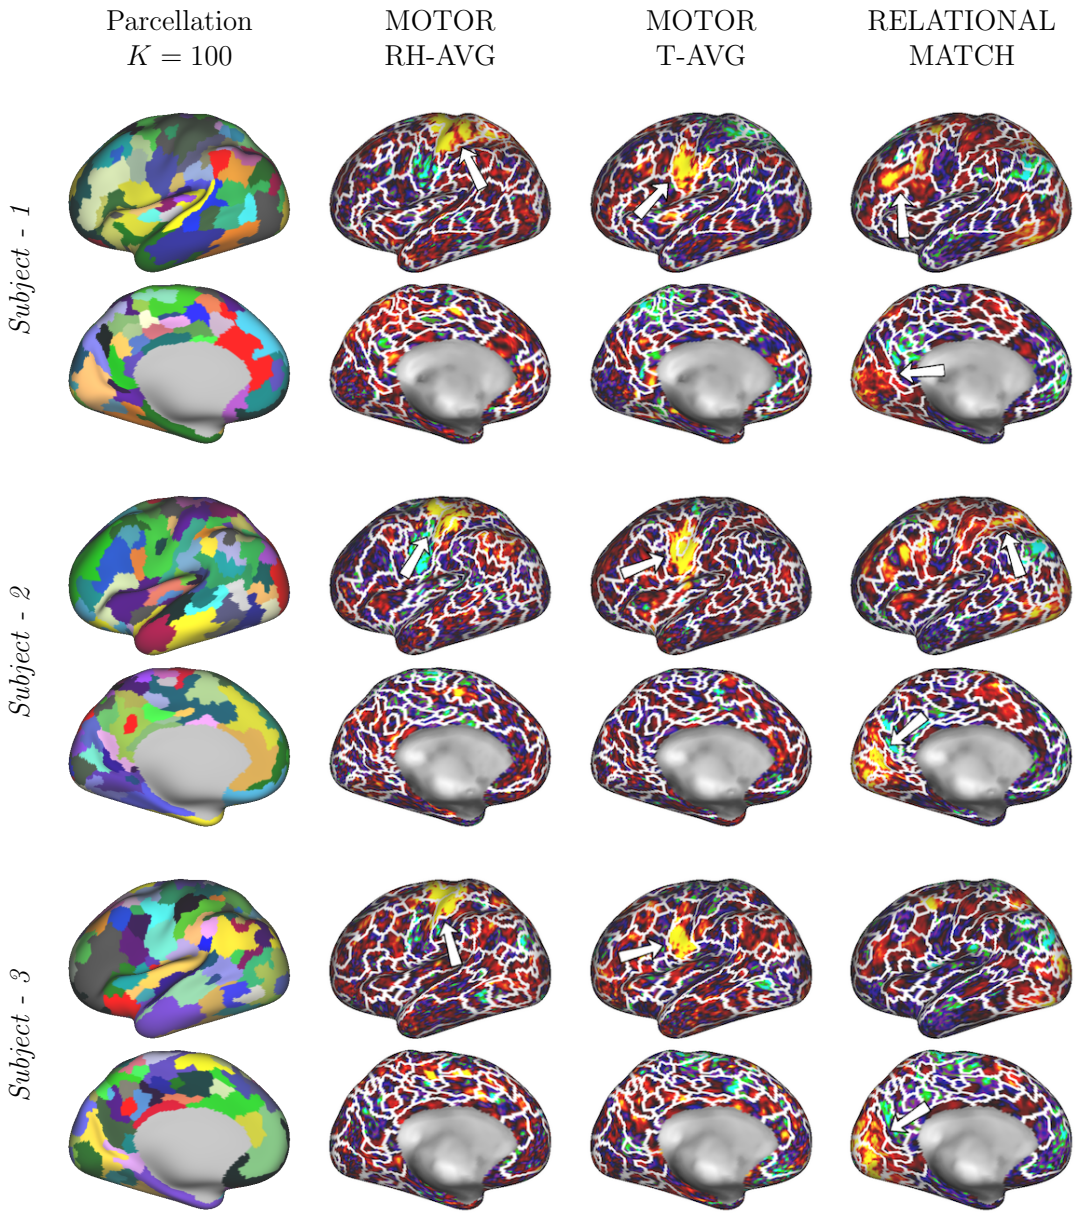


Figure 4.17: Parcellations of the left hemisphere obtained from three different subjects using the proposed approach ($K = 100$). Parcellation borders are superimposed onto the subject-specific task activation maps represented in terms of z scores. Hot colours indicate activation (yellow being higher than red) and white arrows show aligning borders with the activated areas.

a coarse-to-fine multi-resolution basis, considering 1000 to 3000 and 50 to 500 regions per subject. We not only assessed the performance from a clustering point of view, but also showed the extent of agreement between the parcellation boundaries and functionally active regions obtained via task fMRI as well as other cortical features such as the cytoarchitectonic areas and well-structured patterns of myelination.

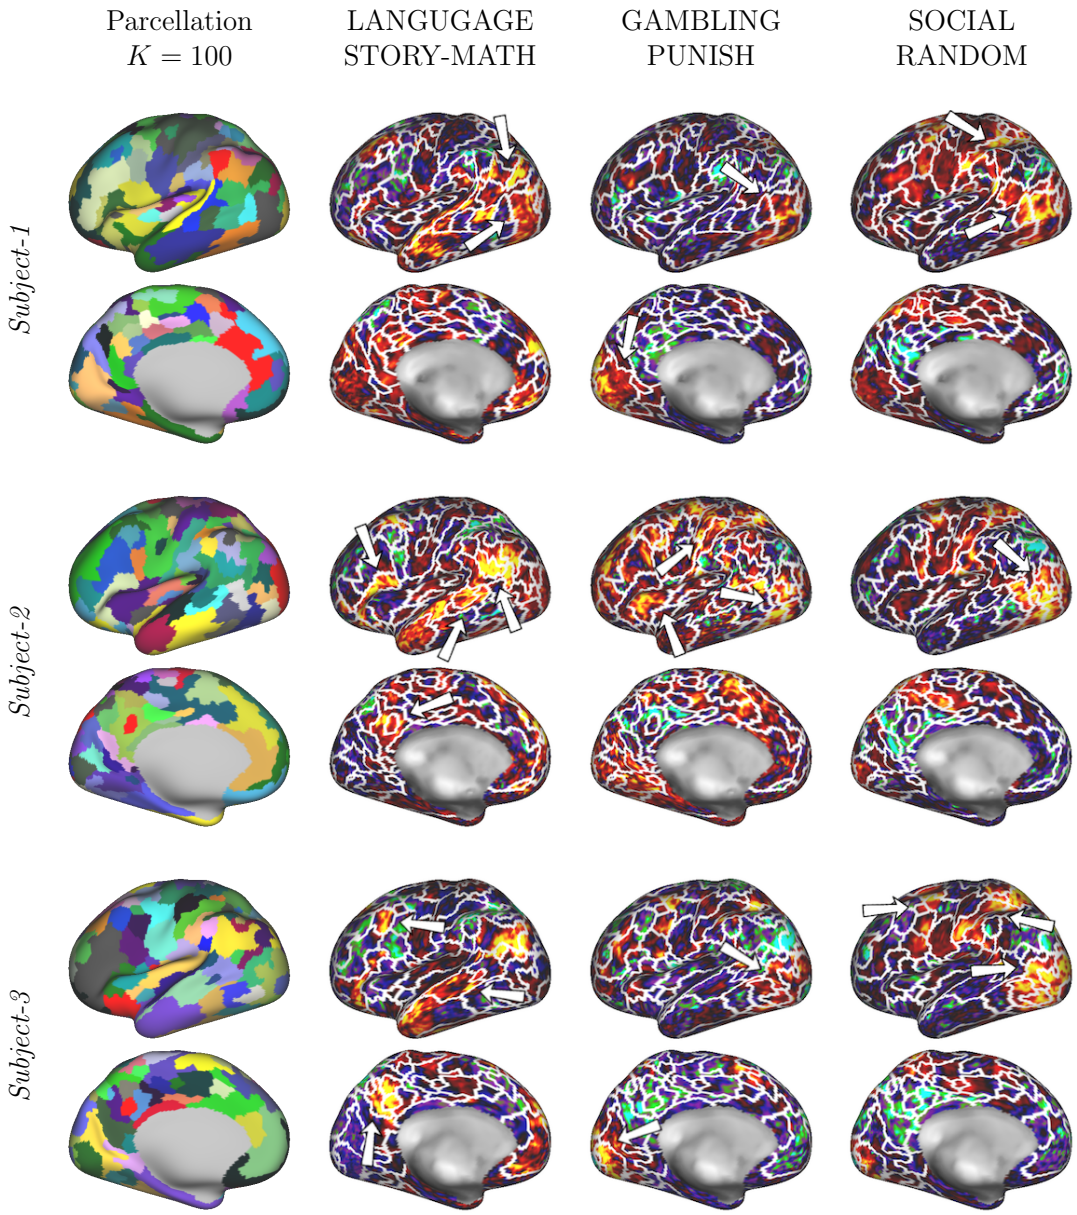


Figure 4.18: Parcellations of the left hemisphere obtained from three different subjects using the proposed approach ($K = 100$). Parcellation borders are superimposed onto the subject-specific task activation maps represented in terms of z scores. Hot colours indicate activation (yellow being higher than red) and white arrows show aligning borders with the activated areas.

We observed that data-driven parcellations in general have a much better agreement with the underlying connectivity compared to anatomical and random parcellations. It is worth noting that, when considering cluster validity measures, there typically exists a trade-off between stability and fidelity to the underlying data [2]. For instance, k-means clustering largely leads in terms of clustering accuracy due to the fact that it is solely driven by functional connectivity. It,

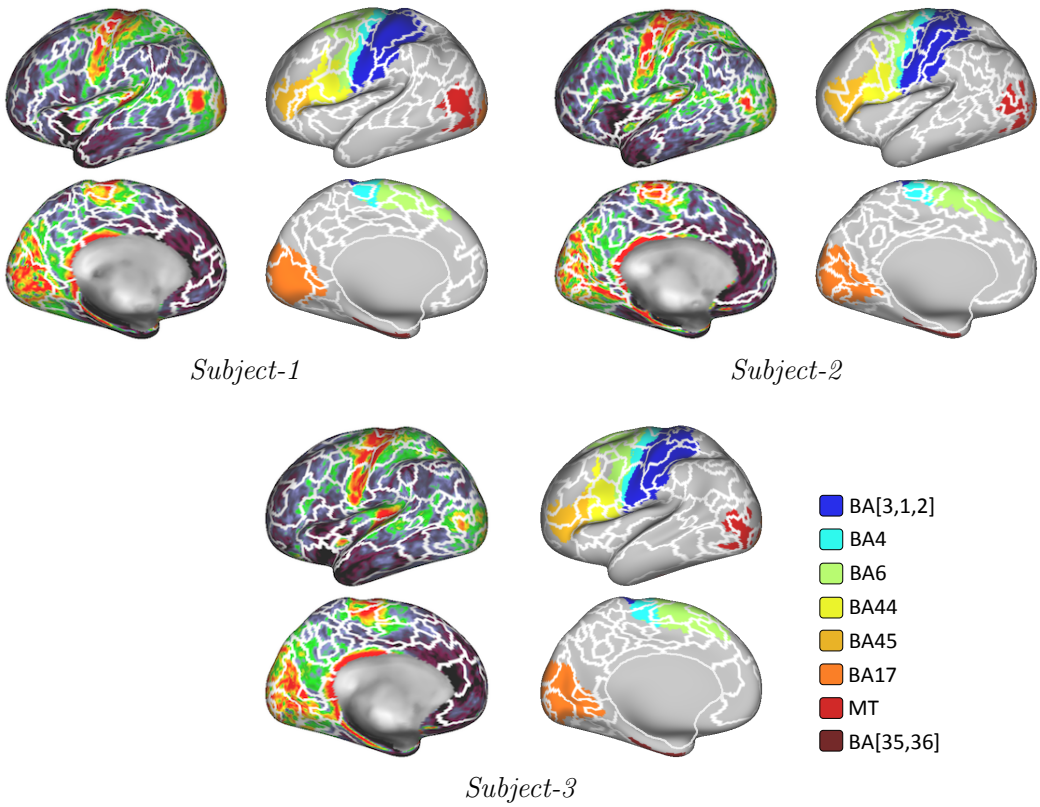


Figure 4.19: Parcellations of the left hemisphere obtained from three different subjects using the proposed approach at a resolution of $K = 100$ regions are overlaid onto subject-specific myelin maps and Brodmann areas.

however, shows an extremely poor performance regarding reproducibility. Contrarily, methods that are dominated by spatial constraints, such as spectral clustering, are likely to capture stable features regarding the geometry of the cortical mesh. This appears to confer a strong advantage for reproducibility, but it also constrains the parcellation task and leads to inaccurate alignment with the brain's underlying functional organisation [66].

Hierarchical clustering yields a performance that is likely to reside in-between: it offers the advantage of generating spatially joint parcels, which can help obtain more reproducible parcellations, while still capturing the functional features with high fidelity. It can even compete with spectral clustering in terms of stability when a finer resolution parcellation is initially used to smooth the underlying data. For instance, the proposed parcellation scheme achieves higher reproducibility than the typical hierarchical clustering with only limited impact on the accuracy of the parcellations. This can be linked to the supervertex clustering stage, which relies

on a modified k-means algorithm that allows injection of flexible spatial constraints into the parcellation process. As a result, it can compute a cortical parcellation that can simultaneously maximise reproducibility and accuracy to a large extent.

Another advantage of the proposed method in the context of brain mapping emerges from the fact that it provides a hierarchy of nested regions that can be explored at different resolutions [32, 66]. It is not necessary to pre-specify a cluster number, as the clustering algorithm inherently provides a multi-scale segregation of the brain. Since not all parts of the cerebral cortex can be subdivided with the same precision, this would allow a fine-to-coarse delineation of cortical areas under investigation without the necessity to use a pre-defined mask. However, this may also confer a disadvantage in terms of propagating errors from one resolution to the next [187].

Qualitative assessment of RSFC parcellations revealed some degree of alignment with task-activated cortical regions and other neuro-anatomically defined areas. However, it is worth noting that this agreement is highly variable across subjects and RSFC-based parcels do not necessarily match an entire task-activated region or cytoarchitectonic area. Rather than, some of the parcellation borders are likely to align with the gradients in cortical maps or the boundaries of anatomically defined regions. While a direct comparison between RSFC parcellations and other modalities can be difficult to interpret, it may still provide additional confidence for a parcellation method's ability to delineate meaningful cortical segregations and assess its reliability from a neuro-biological point of view [32].

Subject-level parcellations are of great importance to study the functional organisation of the brain and its variability within a population. Whereas a high reproducibility can be obtained on a scan-to-scan basis, the functional and structural differences inherent in each individual brain do not allow a subject-to-subject comparison in a similar manner. Instead, we used cross-subject comparisons to evaluate the degree of variability between individuals as shown in Fig. 4.16. Our experiments revealed that certain areas, such as the motor cortex, can be consistently parcellated into similar regions and/or at a similar granularity; however, typically a great amount of variance exists within the rest of the cortex.

Understanding the source of variability across subjects is an important research problem and constitutes an obstacle in itself. While a possible alteration in connectivity could be associated with a brain disorder, it could also be attributed to genetic variations [63], topological differences between subjects [131], varying strengths of connections between brain areas in individuals [92], as well as simply could be due to potential registration errors or low SNR in the underlying data. Interestingly, the fact that we obtain less variance (i.e. more similarity) within the motor and visual cortex can be linked to the HCP's cortical-folding based registration technique [205], which promotes a more consistent registration for these areas [188].

The variability across individual parcellations does not allow constructing spatial correspondence across subjects. However, population-level connectome analysis typically rely on the assumption that vertex (or voxel) level correspondence has been ensured *a priori*. This can be achieved by computing a group-wise parcellation that ideally represents the shared characteristics within a population. Such a group representation can be useful in the context of connectome analysis, for instance, to identify how connectivity changes within a population through ageing. Similarly, they can be used for deriving biomarkers in order to better understand disease-related differences in the brain connectivity. In addition, the inter-subject variability problem can be approached from a different perspective: cortical regions that are most consistent and/or least similar across subjects can be located by comparing individual parcellations to a group representation obtained from the same population.

There are two typical extensions that can be attached to our method for computing group-wise parcellations: i) averaging BOLD timeseries across subjects and submitting it to the proposed parcellation framework, or ii) mounting a third-level clustering stage on top of the individual subject parcellations. Given the accuracy of the proposed parcellations at the subject level, we implemented the latter approach and evaluate its performance in Chapter 7 along with a large set of other group-wise parcellation techniques in the literature. While this chapter focused on RSFC captured in the form of BOLD timeseries for a connectivity-based delineation of the brain's cortical organisation, parcellations can also be obtained from other types of connectivity data, such as structural connectivity estimated from dMRI. Next chapter explores this phenomenon within the context of cortical brain parcellation and proposes a novel method

driven by dMRI-derived structural connectivity.

Chapter 5

Cortical Boundary Mapping through Manifold Learning

This chapter is based on:

S. Arslan, S. Parisot, and D. Rueckert, *Boundary Mapping through Manifold Learning for Connectivity-Based Cortical Parcellation.* International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI), vol. 9900 of LNCS. Springer, pp. 115-122, 2016.

Abstract

This chapter proposes a new parcellation method that can be used to model the brain's cortical organisation with respect to structural connectivity estimated from dMRI. To this end, we learn a manifold from the high-dimensional local connectivity patterns of an individual subject and identify parcellation boundaries as points in this low-dimensional embedding where the connectivity patterns change. We compute spatially contiguous and non-overlapping parcels from these boundaries after projecting them back to the cortical surface. This is achieved by boundary mapping, a technique that separates cortical regions using well-known image segmentation algo-

rithms. Our experiments with a set of 100 healthy subjects show that the proposed method can produce spatially contiguous parcels with distinct patterns of connectivity and a higher degree of homogeneity at varying resolutions compared to the state-of-the-art parcellation methods, hence can potentially provide a more reliable set of network nodes for connectome analysis.

5.1 Introduction

Connectivity-driven parcellation (CDP) techniques have recently gained a lot of attention due to their potential to reveal the functional and structural architecture of the human brain. As these approaches are directly applied to the rs-fMRI and/or dMRI data, they can learn a somewhat optimal representation of connectivity in a lower dimensional space, and therefore, can better reflect the network organisation of the brain compared to traditional parcellations derived from anatomical landmarks or randomly obtained ROIs [231]. CDP methods tend to model the connectivity data in association with clustering algorithms [66]. As we have extensively reviewed in the previous chapters, many clustering approaches such as hierarchical clustering, spectral clustering (normalised-cuts being the most popular), and k-means clustering (as well as their variants) have been repeatedly used for the parcellation of the human cerebral cortex [244]. However, despite promising results, the CDP problem is still open to improvements. This is primarily due to the facts that (1) the problem is not well-posed and (2) there is not a reference parcellation to supervise the model selection. As a result, obtaining accurate parcellations directly depends on the proposed method's fidelity to the underlying data [187] and its capacity to differentiate regions with different connectivity profiles [66].

An alternative cortical parcellation approach known as boundary mapping casts the clustering problem as the identification of transitions between connectivity patterns and solves it using image segmentation techniques [49, 279, 280, 93]. This technique has been successfully used to derive areal parcellations of the cortical surface from resting-state functional connectivity (RSFC) [93]. The key idea emerges from the fact that RSFC patterns are likely to show rapid changes between different cortical locations in the human brain, similarly to the changes in

connectivity that segregate cortical areas in non-human primates [70]. The locations where connectivity is in transition are therefore assumed to represent the parcellation boundaries in the human cerebral cortex [49, 93]. Although boundary mapping has been extensively used to derive local or whole-brain cortical parcellations from RSFC [279, 280, 93], to the best of our knowledge, it has not been applied to structural connectivity estimated from dMRI. In contrast to the indirect estimation of connectivity achieved with rs-fMRI, dMRI can estimate the physical white matter connections in the brain. Parcellations derived from dMRI have, therefore, a more intuitive interpretation and tend to be more robust than RSFC parcellations [188].

One limitation of boundary mapping is its sensitivity to spurious spatial gradients captured from connectivity patterns. A considerable amount of data is required to obtain a robust parcellation from these gradient maps, and hence, it is typically applied to averaged-datasets of large populations [93]. This may obscure patterns of brain organisation specific to individuals [136, 92]. It also poses an extra challenge for the method's adaptability to the individual subjects: a subject-level parcellation may only be achieved if sufficient data is collected, e.g. the subject is scanned repeatedly over time [136]. However, most neuroscience datasets typically provide at most few scans per subject/modality. Another challenge emerges from the correlation/connectivity maps used to obtain the gradients. The high dimensionality of these maps may constrain the detection of robust connectivity patterns at the subject level. To this end, a dimensionality reduction stage may be beneficial for a more robust and interpretable representation of the brain connectivity and may facilitate the use of boundary mapping on a single subject basis.

With this motivation, we introduce a new parcellation method, in which we learn a non-linear manifold from local connectivity characteristics of an individual subject and develop an effective way of computing parcels from this manifold based on boundary mapping. Our approach rests on the assumption that through dimensionality reduction, we can capture the underlying connectivity structure that may not be visible in high-dimensional space [131]. We use a low-dimensional embedding to locate transition points where connectivity patterns change and interpret them as an abstract delineation of the parcellation boundaries. After projecting back to the native cortical space, these boundaries are used to compute non-overlapping and

spatially contiguous parcels. We achieve this with a watershed segmentation technique, which has been effectively used for whole-cortex boundary mapping before [136, 93]. Non-linear manifold learning has been previously used to identify functional networks from fMRI [132, 134] and for surface matching [131], as well as within many other fMRI analysis techniques, such as [243, 216]. Here, we propose to use such technique in association with dMRI-based structural connectivity and boundary mapping, in order to compute cortical parcellations for individual subjects, which can be used to investigate the brain's cortical organisation at the subject level and to derive network nodes for a whole-brain connectome analysis.

Given the lack of a ground truth parcellation of the cortex, we assess the performance of the proposed method from a clustering point of view and evaluate the degree of fidelity to the underlying structural connectivity data using two well-known techniques: (1) parcel homogeneity [218, 244, 93] and (2) Silhouette analysis [245, 54, 66, 190], as introduced in the previous chapter. Besides structural connectivity that is used for model estimation, we evaluate the parcellations with respect to functional connectivity data obtained from resting-state fMRI as a means of external validation [66]. Our method is compared to three well-known clustering algorithms applied to the connectivity data, namely, hierarchical clustering, spectral clustering with normalised cuts, and k-means clustering, as well as one dMRI-based state-of-the-art connectivity-based parcellation technique [187]. In addition, two parcellation schemes which do not take into account any connectivity information [214, 244] are also incorporated to our experiments to form a baseline for comparisons. Besides the quantitative evaluation measures, we provide visual results of the parcellations across different subjects/resolutions and provide complimentary information to understand the inter-subject variability within the population by means of parcellation consistency across subjects. Finally, we show the extent to which our parcellation boundaries agree with well-established patterns of myelination and cytoarchitecture as defined by the Brodmann atlas [38].

The remainder of this chapter is organised as follows: In Section 5.2, we summarise the proposed parcellation method, starting with the estimation of structural connectivity and following with a detailed explanation of the manifold learning technique used for dimensionality reduction. The section then covers the steps for boundary map generation and cortical parcellation. In

Section 5.3, we summarise the quantitative measures and clustering techniques used for evaluation. In Section 5.4 we provide the performance measurements of the proposed method as well as the other approaches and show the visual results for qualitative and inter-modality assessment of the parcellations. Finally, in Section 5.5, we conclude the chapter with a detailed discussion on the advantages and limitations of the proposed method, and provide some insight towards future research directions.

5.2 Methodology

5.2.1 Data

We conduct our experiments on Dataset 1, which comprises 100 subjects (54 female, 46 male healthy adults, ages 22-35). The details of the dataset are provided in Chapter 2 and briefly summarised here. Data was acquired and preprocessed following the HCP minimal preprocessing pipelines [87]. For each subject, the cortical gray-matter voxels were registered onto the 32k standard triangulated mesh with a 2 mm spacing yielding a standard set of cortical vertices per hemisphere. The dMRI data was acquired using a multi-shell approach, in which diffusion weighting consisted of three shells at b-values 1000, 2000, and 3000 s/mm^2 and 90 gradient directions obtained per shell [230].

5.2.2 Method

Our method starts with preprocessing the dMRI data using probabilistic tractography to estimate a structural connectivity network, which is then reduced in dimensionality through non-linear manifold learning with Laplacian eigenmaps [26]. Next, we identify boundaries in this low-dimensional embedding as points where connectivity patterns change. Finally, driven by these boundaries, we make use of watershed segmentation to achieve a whole-brain cortical parcellation. All steps are explained thoroughly in the subsequent sections and the parcellation pipeline following preprocessing is summarised in Fig. 5.1.

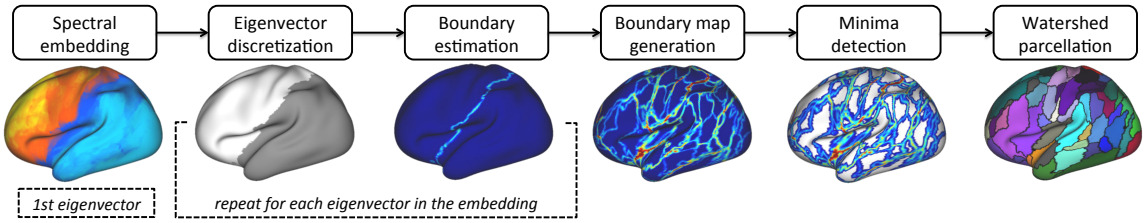


Figure 5.1: Parcellation pipeline summarising all steps after preprocessing.

Estimating Structural Connectivity

The tractography matrix has been obtained from the preprocessed dMRI data by performing probabilistic tractography with FSL's bedpostX and probtrackX methods [23, 110], as described in Chapter 2. bedpostX allows to automatically determine the number of fibers passing through each voxel within the brain and estimate their orientation distribution. probtrackX performs probabilistic tractography with respect to estimated fiber orientations. We seeded 5000 streamlines from each of the n cortical vertices and obtained a tractography matrix that records the number of streamlines originating from each vertex and reaching the rest of the cerebral cortex.

Long-range connections computed with probabilistic tractography cannot be as effectively estimated as the short-range ones, due to the accumulation of errors along the tracts [111, 164, 188]. As a result, the estimated strength of connectivity is likely to be much higher for short-range connections compared to the long-range ones, even though the actual connections may have the same strength. This can potentially influence the connectivity matrix and have an impact on the resulting parcellations. To reduce the bias towards short-range connections, we apply an element-wise log transformation to the tractography matrix as suggested in [111, 164]. The process reduces the dynamic range of fiber counts, and hence, helps alleviate the bias towards short connections while not losing any information [188], as opposed to an alternative solution including thresholding short-range connections [207].

We estimate a structural connectivity matrix $C \in \mathbb{R}^{n \times n}$ for each subject in the dataset by cross-correlating the log-transformed tractography matrix using Pearson's correlation coefficient. Each row in C represents how a vertex is connected to the other vertices in the cortical surface and therefore represents a connectivity profile, showing how a vertex is connected to

the rest of the cerebral cortex. We excluded the medial wall vertices from further processing as they do not possess reliable information for connectivity analysis.

Manifold Learning via Laplacian Eigenmaps

Overview We use *Laplacian eigenmaps*[1] [26] to compute a non-linear embedding from a structural connectivity network. This method can reveal the intrinsic geometry of the underlying connectivity by forming an affinity matrix based on how vertices are connected within their neighbourhoods, hence preserving local proximity information between vertices. Based on the notion of the graph Laplacian obtained from the affinity matrix, a low-dimensional representation of the data can then be computed by using spectral decomposition.

Several key features of the algorithm distinguish it from other non-linear manifold learning techniques and make it a powerful tool to explore structural (or functional) brain connectivity. First of all, the core algorithm is very simple. It only requires to define a nearest-neighbour graph and solves one sparse eigenvector problem. Second, the method exhibits stability with respect to the embedding, that is, regardless of the resolution of the underlying network (i.e. different choices of n), the resulting embeddings will recover the same underlying manifold. Last but not least, Laplacian eigenmaps are closely related to spectral clustering techniques in a sense that the algorithm inherently provides a natural clustering of the data, which can be attributed to the fact that the eigenmaps preserve local proximity information in the embedding. As a result, the resulting low-dimensional embedding can be used for obtaining soft clusters [54, 244]. Furthermore, this locality-preserving characteristic makes the algorithm more robust to outliers and noise [26].

The Algorithm The generic dimensionality reduction problem can be defined as the following. Given a set $x_1, \ldots, x_n$ of n points in $\mathbb{R}^m$, find a new set of n points $y_1, ..., y_n$ in $\mathbb{R}^d$ such that y_i represents x_i in a low dimensional space, where $d << m$. Here, we assume that the data points (i.e. vertices) $x_1, \ldots, x_n \in \mathcal{M}$ and $\mathcal{M}$ is a manifold embedded in $\mathbb{R}^m$. Specific to our

[1] It is worth noting that, this section is compiled mainly from the original Laplacian eigenmaps paper by Belkin and Niyogi [26]. Readers are referred to the original manuscript for more detailed information.

problem, each vertex is represented by an n-dimensional connectivity profile, hence, $m = n$. The procedure that leads to eigenmaps can be outlined in three key steps as follows:

1. **Constructing an affinity matrix**: An affinity matrix W is formed by constructing an edge between vertices i and j if x_i and x_j are 'close' to each other. Among different methods to define this closeness, we use the k nearest neighbour approach ($k \in \mathbb{N}$) and connect vertices i and j by an edge if i is among k nearest neighbours of j or j is among k nearest neighbours of i, forming a symmetric relationship. k is chosen large enough to effectively capture the local connectivity structure and to ensure that the affinity matrix is connected for each subject.

2. **Choosing the weights**: The edges are typically weighted by a heat kernel if i and j are connected

$$W_{ij} = \exp\Big(-\frac{\|x_i - x_j\|^2}{t}\Big), \tag{5.1}$$

   otherwise $W_{ij} = 0$. Here, $\|\cdot\|$ is Euclidean norm and t controls the weight decay. In practice t can be chosen infinite [242], leading to $W_{ij} \in \{0, 1\}$. Alternatively, the heat kernel can be replaced by an appropriate similarity measure, such as Pearson's correlation, together with a thresholding or scaling scheme to discard negative correlations [131, 132]. In this chapter, we choose this approach and transform the connectivity matrix C into a locality-preserving affinity matrix $W \in \mathbb{R}^{+n\times n}$ by only retaining the correlations of the k nearest neighbours of each vertex.

3. **Spectral decomposition**: Given that W is connected, we compute eigenvalues and eigenvectors by solving the generalized eigen decomposition problem

$$L\mathbf{f} = \lambda D\mathbf{f}, \tag{5.2}$$

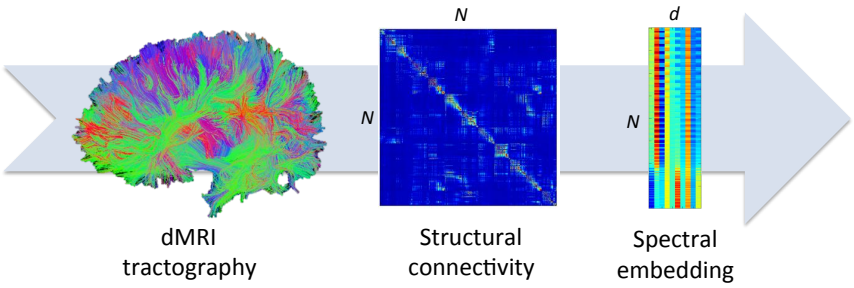


Figure 5.2: Illustration of non-linear feature reduction using Laplacian eigenmaps.

where D is a diagonal matrix with each entry $D_{ii} = \sum_j W_{ij}$ representing the degree of x_i and $L = D - W$ is the graph Laplacian, which is a symmetric and positive semidefinite matrix. The Laplacian matrix can be considered as an operator on functions defined on vertices in W [26]. Solving Eq. 5.2 reveals the eigenvectors $\mathbf{f}_0, \mathbf{f}_1, \ldots, \mathbf{f}_{n-1}$, ordered according to their eigenvalues $0 = \lambda_0 \leq \lambda_1 \leq \ldots \leq \lambda_{n-1}$. After omitting the eigenvector $\mathbf{f}_0$ corresponding to λ_0, we can use the next d eigenvectors to define an embedding that can approximate a low dimensional manifold [26]. Hence, each cortical vertex i in this high-dimensional structural connectivity matrix can now be expressed as a row in the spectral embedding, i.e. $i \mapsto (\mathbf{f}_1(i), \ldots, \mathbf{f}_d(i))$ as illustrated in Fig. 5.2.

An Optimal Embedding The local proximity between pairs of vertices is optimally preserved by the spectral embedding generated using the Laplacian eigenmaps algorithm. The goal of the embedding function $\mathbf{f}$ is to map nearby (or strongly connected) vertices in the high-dimensional space to a d-dimensional Euclidean space and ensure that they still remain nearby.

Let us first consider the problem of mapping the vertices to a line, i.e. where $d = 1$. We are seeking a mapping $\mathbf{y} = (y_1, y_2, \ldots, y_n)^T$, such that the following criterion is minimised:

$$\sum_{i,j} (y_i - y_j)^2 W_{ij} \tag{5.3}$$

Given how the weights W_{ij} are assigned, the objective function incurs a heavy penalty if points

that are close in high-dimensional space are mapped far apart in $\mathbf{y}$. For any $\mathbf{y}$, we can rewrite Eq. 5.3 as

$$\frac{1}{2}\sum_{i,j}(y_i - y_j)^2 W_{ij} = \mathbf{y}^T L \mathbf{y}, \tag{5.4}$$

where L is the Laplacian matrix. Recall that W_{ij} is symmetric and $D_{ii} = \sum_j W_{ij}$. Thus,

$$\begin{aligned}\sum_{i,j}(y_i - y_j)^2 W_{ij} &= \sum_{i,j}(y_i^2 + y_j^2 - 2y_i y_j) W_{ij} \\ &= \sum_i y_i^2 D_{ii} + \sum_j y_j^2 D_{jj} - 2\sum_{i,j} y_i y_j W_{ij} \\ &= 2\mathbf{y}^T L \mathbf{y}.\end{aligned} \tag{5.5}$$

Since all W_{ij} are non-negative, Eq. 5.5 shows that $\mathbf{y}^T L \mathbf{y} \geq 0$, which proves that L is positive semidefinite. The minimisation problem is reduced to finding

$$\underset{\substack{\mathbf{y} \\ \mathbf{y}^T D \mathbf{y} = 1}}{\operatorname{argmin}} \mathbf{y}^T L \mathbf{y}. \tag{5.6}$$

Here, the additional constraint $\mathbf{y}^T D \mathbf{y} = 1$ removes an arbitrary scaling factor in the embedding [26]. The D matrix provides a means to account for the impact of vertices in the graph. A larger D_{ii} implies more edges connected to vertex i, hence increases its importance for the minimisation [26]. Given that L is a symmetric and positive semidefinite matrix, the Rayleigh-Ritz theorem [154] indicates that the solution of the minimisation problem is given by the vector $\mathbf{y}$, which corresponds to the eigenvector associated with the minimum non-zero eigenvalue of L [271].

Considering the more general problem of embedding the graph into d-dimensional Euclidean

space where $d > 1$, we can represent the embedding by the $n \times d$ matrix $Y = [\mathbf{y}_1, \ldots, \mathbf{y}_d]$, in which the ith row corresponds to the embedding coordinates of the ith vertex. Similarly to Eq. 5.3, we need to minimise

$$\sum_{i,j} \|\mathbf{y}^{(i)} - \mathbf{y}^{(j)}\|^2 W_{ij} = \mathrm{tr}(Y^T L Y), \tag{5.7}$$

where $\mathbf{y}^{(i)} = [\mathbf{y}_1(i), \ldots, \mathbf{y}_d(i)]^T$ is the d-dimensional representation of the ith vertex. By introducing a constraint that prevents an embedding from collapsing onto a subspace of fewer than d-1 dimensions, the objective function is reduced to finding

$$\underset{Y^T D Y = I}{\mathrm{argmin}}\ \mathrm{tr}(Y^T L Y). \tag{5.8}$$

As in the one-dimensional case, standard methods [154] show that the solution that minimises the objective function is provided by the matrix of eigenvectors associated with the lowest eigenvalues of the graph Laplacian [26].

Connections to Spectral Clustering The eigenvectors of the graph Laplacian provide an embedding that reflects the intrinsic geometry of the data and interestingly this solution can also be interpreted in the context of spectral clustering [26], which can be used to sub-partition a graph into K pre-defined clusters. Here, we briefly outline the ideas behind spectral clustering with a particular focus on the normalised-cuts (N-Cuts) technique proposed by Shi and Malik [219] and show the role of graph Laplacian in providing a set of feature vectors for clustering purposes.

Let $G = (V, E)$ be a weighted graph with V being the set of vertices and E being the set of connections between them. W is the affinity matrix associated with G and defined as before, hence, it is symmetric, connected, and $W_{ij} \geq 0$ for all i, j. The ultimate goal of N-Cuts is

to find an optimal cut to partition W into K pre-defined clusters, such that vertices within a cluster are more similar to each other than those in different clusters.

For subdividing W into two disjoint clusters A and B, where $A \cup B = V$, the normalised-cut cost function is defined as

$$Ncut(A, B) = cut(A, B)\Big(\frac{1}{vol(A)} + \frac{1}{vol(B)}\Big), \tag{5.9}$$

where $cut(A, B) = \sum_{i \in A, j \in B} W_{ij}$ represents the sum of the weights on edges that connect vertices in A to vertices in B [284] and $vol(\cdot)$ represents the total connections from within-cluster vertices to all vertices in the graph. *Ncut* normalises the cost as a fraction of the sum of the edge weights connecting vertices in a cluster to every other vertex in the graph. It simultaneously minimises the inter-cluster similarity and maximises the within-cluster similarity, and hence, achieves a balanced partitioning.

While the minimisation problem is NP-hard [219], a near global-optimal solution can be approximated by allowing relaxation of the indicator functions to real values and then solving the relaxed problem [26].

Recall that

$$\mathbf{x}^T L \mathbf{x} = \sum_{i,j} (x_i - x_j)^2 w_{ij}. \tag{5.10}$$

Let $a = vol(A)$ and $b = vol(B)$. Put

$$x_i = \begin{cases} \frac{1}{vol(A)}, & \text{if } V_i \in A \\ -\frac{1}{vol(B)}, & \text{if } V_i \in B \end{cases}. \tag{5.11}$$

We now have

$$\mathbf{x}^T L\mathbf{x} = \sum_{i,j}(x_i - x_j)^2 w_{ij} = \sum_{V_i \in A, V_j \in B} \Big(\frac{1}{a} + \frac{1}{b}\Big)^2 cut(A,B) \tag{5.12}$$

and

$$\begin{aligned} \mathbf{x}^T D\mathbf{x} = \sum_i x_i^2 d_{ii} &= \sum_{V_i \in A} \frac{1}{a^2} d_{ii} + \sum_{V_i \in B} \frac{1}{b^2} d_{ii} \\ &= \frac{1}{a^2} vol(A) + \frac{1}{b^2} vol(B) = \frac{1}{a} + \frac{1}{b}, \end{aligned} \tag{5.13}$$

where D is the degree matrix as defined before. Thus,

$$\frac{\mathbf{x}^T L\mathbf{x}}{\mathbf{x}^T D\mathbf{x}} = cut(A,B)\Big(\frac{1}{a} + \frac{1}{b}\Big) = Ncut(A,B). \tag{5.14}$$

The relaxed problem is minimising Eq. 5.14, subject to $\mathbf{x}^T D\mathbf{1} = \mathbf{0}$ where $\mathbf{1}$ is a column vector of ones. Given that D is invertible, G has no isolated vertices (which is ensured while constructing the adjacency matrix) and substituting $\mathbf{y} = D^{1/2}\mathbf{x}$ we can rewrite Eq. 5.14 as

$$\frac{\mathbf{x}^T L\mathbf{x}}{\mathbf{x}^T D\mathbf{x}} = \frac{\mathbf{y}^T D^{-1/2} L D^{-1/2}\mathbf{y}}{\mathbf{y}^T\mathbf{y}} = \frac{\mathbf{y}^T \mathcal{L}\mathbf{y}}{\mathbf{y}^T D\mathbf{y}}, \tag{5.15}$$

where $\mathbf{x}$ and $\mathbf{y}$ are perpendicular to $D^{1/2}\mathbf{1}$. Here the matrix $\mathcal{L} = D^{-1/2} L D^{-1/2}$ is the *normalised* graph Laplacian, a symmetric and positive semidefinite matrix, which serves in the approximation of the minimisation of $Ncut$. That is, the solution to the minimisation problem is given by the eigenvector associated with the smallest non-zero eigenvalue of $\mathcal{L}$.

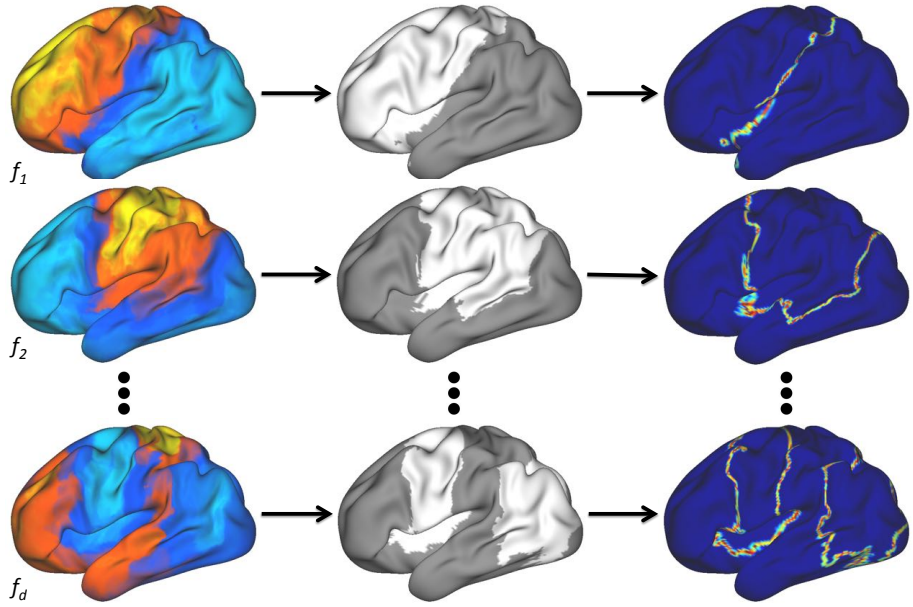
Similarly to the multi-dimensional embedding problem discussed above, this solution can be easily generalised to obtain a K-way partitioning of the graph where $K > 2$. That is, by using the set of eigenvectors corresponding to the K smallest non-zero eigenvalues of $\mathcal{L}$, we can obtain a near global-optimal solution to the multi-way normalised-cut problem and feed these eigenvectors to a general purpose clustering algorithm, e.g. k-means [271] or use discretisation to transform the real-valued representation to hard clusters [219].

Eigenvector Discretisation

As discussed in the preceding section, the locality-preserving process of dimensionality reduction not only provides an optimal embedding for non-linear manifold learning, but also imposes a soft clustering of the data [26]. Therefore, the parcellation problem can be cast as a graph partitioning problem and one would attempt to subdivide the connectivity graph with spectral clustering, e.g. using the normalised cuts criterion [285]. In particular, each of the smallest eigenvectors corresponds to a real-valued solution that optimally sub-partitions the graph. These partitions can be approximated by transforming the real valued eigenvectors into discrete forms (i.e. discretisation), ideally by dividing them into two parts with respect to a splitting point [219]. This can further be generalised towards a multi-way partitioning with a recursive or simultaneous discretisation of the smallest eigenvectors [219], and thus, can be used to obtain a parcellation [54].

However, by definition, our affinity matrix does not impose any spatial constraints, hence such spectral methods cannot guarantee spatial contiguity within the parcels. Instead, we propose a more effective way of deriving parcellations from discrete eigenvectors, i.e. we use them to identify cortical boundaries where connectivity patterns show distinct changes, and later show that this method can produce more reliable parcellations compared to spatially constrained spectral clustering.

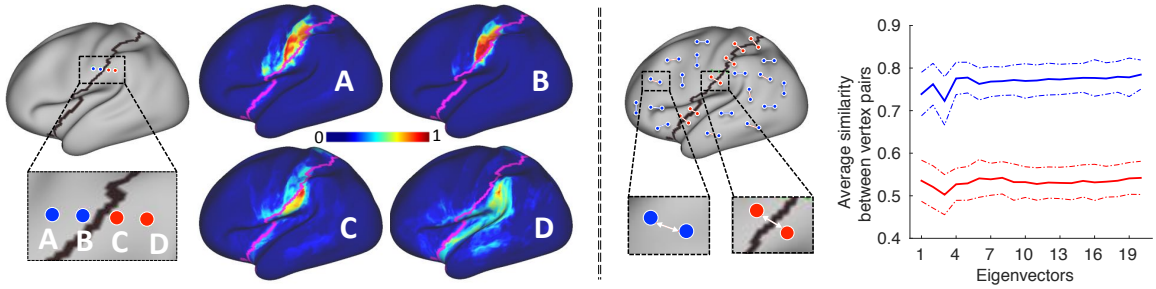
To this end, we first subdivide each real-valued eigenvector into two regions using k-means as illustrated in Fig. 5.3. The edges between these subregions potentially provide good separation points towards obtaining a parcellation, as the vertices within the same subregions tend to have



(a) (b) (c)

Figure 5.3: (a) Eigenvectors overlaid onto the lateral left hemisphere. (b) Discretised eigenvectors obtained via k-means. (c) Boundary estimation from discrete eigenvectors.

similar connectivity properties, whilst the points closer to the boundary attribute to the cortical areas where the connectivity is in transition. For example, Fig. 5.4(a) shows that connectivity profiles of different vertices may exhibit similar or varying patterns, depending on their relative location to an edge. In order to show that this tendency holds across the whole cortex, we randomly selected vertices from one subregion adjacent to the edge and paired them with their closest neighbours residing in the other subregion. Keeping the distance between the vertices in pairs approximately the same, we selected new pairs of vertices, but this time from within the same subregions. We then measured the average correlation between the paired vertices' connectivity profiles in each set and repeated this for all eigenvectors and subjects. Fig. 5.4(b) shows that, the similarity between the connectivity profiles of vertices drops by at least 20% if they reside on different sides of a boundary.



(a) (b)

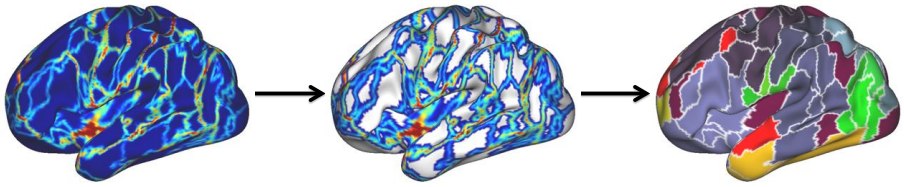
Figure 5.4: (a) Connectivity profiles of vertices from different sides of a boundary. (b) *Left:* Illustration of the vertex selection procedure. *Right:* Average similarity (correlation) between paired vertices for each eigenvector. Dotted lines show the variability across subjects as means of standard deviations.

Boundary Map Generation and Cortical Parcellation

To locate the connectivity transition points and construct a boundary map, we first transfer the discrete eigenvectors back to the native high-dimensional space (project them onto the cortical surface) as shown in Fig. 5.3(b). We then compute the first spatial derivative of each eigenvector across the cortical surface [Fig. 5.3(c)] and combine them into a boundary map as illustrated in Fig. 5.5(a). This map constitutes a more robust substitution for the boundary maps based on gradients directly calculated from the spatial correlations [93], since it can adjust for possible spurious gradients.

In order to obtain the final parcellations from the boundary map, we use a marker-controlled watershed algorithm [267], a powerful image processing technique used for the segmentation of salient objects in images [10, 93]. The marker-controlled watershed technique typically consists of a marker detection stage and a marking function, where the former identifies markers (or seeds) that correspond to estimated locations of structures of interest in the image. The marking function governs the watershed transformation in which markers grow until they touch each other or reach watershed ridge lines.

In our case, we use a watershed implementation adapted for cortical meshes [93]. To this end, we define a set of markers on the boundary map where each marker corresponds to an estimated parcel position as shown in Fig. 5.5(b). The marker definition is typically performed by choosing



(a) (b) (c)

Figure 5.5: (a) A boundary map is obtained by combining gradients of the discrete eigenvectors. (b) Local minima detection: White regions correspond to markers that initiate the watershed process. (c) Output of watershed segmentation is a set of spatially contiguous, non-overlapping cortical parcels subdividing the whole cerebral cortex.

a threshold on the boundary map. We set this threshold to the 25th percentile of the boundary map intensities, since in many empirically tested cases, this effectively revealed approximate parcel locations to be used as ideal markers for a watershed transformation. Driven by the boundary map, the marking function grows these markers until high-intensity cortical vertices are reached or two regions touch each other in the flooding process of the watershed. As a post-processing step, very small regions (comprising 25 or fewer vertices) are merged with their largest neighbours. The final cortical parcellation is shown in Fig. 5.5(c).

5.3 Experiments

5.3.1 Quantitative and Qualitative Evaluation Measures

In the absence of a gold standard parcellation for comparison, we adapted two well-known cluster validity measures, homogeneity [13, 93] and Silhouette analysis [54, 66], to assess how well a parcellation reflects the underlying connectivity. Both methods are previously explained in Chapter 3 and summarised here. Homogeneity measures a parcellation's ability to cluster vertices with similar connectivity by calculating average cross-correlations within each parcel after applying Fisher's r-to-z transformation to resulting correlation coefficients. Silhouette analysis combines within-parcel similarity with inter-parcel separation and measures how vertices within a parcel are similar to each other, compared to the vertices within the nearest parcels [245, 66].

The goodness-of-fit is estimated with respect to the structural connectivity data from which the parcellations have been derived. However, performing the evaluation based on the same connectivity data initially used to obtain the parcellations may lead to a circular analysis [66]. To reduce the impact of any bias towards model selection, we also evaluate parcellations by measuring their extent to reflect the underlying connectivity estimated from resting-state fMRI. Although dMRI-based cortical parcellations do not necessarily reflect the functional organisation of the cortex derived from rs-fMRI, this type of evaluation provides an external data source for validation [66].

Standard silhouette coefficient represents how well a vertex lies in its cluster [208] and is typically computed for each vertex individually. However, when the model generation and validation are based on different data sources, both within-cluster and out-of-cluster similarity of a vertex are likely to be very low. As a result, Silhouette coefficients tend to be close to zero, indicating neither bad nor good clustering regardless of the clustering configuration. In such a case, considering all vertices within a parcel collectively and computing a parcel-wise Silhouette coefficient may provide a more robust means to assess the quality of clustering [54]. To this end, we make use of the modified Silhouette analysis proposed in [54], which normalises the average similarity within a parcel by the average out-of-parcel similarity obtained from the parcels in its neighbourhood.

We not only evaluate the parcellations from a clustering perceptive, but also assess the agreement of the connectivity-driven parcellations with the cyto- and neuro-anatomical architecture of the cortex, identified with Brodmann's areas and myelin maps, respectively. Visual examination of parcellation boundaries with Brodmann's areas is accompanied by Dice-based overlapping maps to provide a quantitative basis for evaluation. We achieve this by considering the Brodmann atlas as an alternative parcellation of the cerebral cortex and measuring the amount of overlap between the proposed parcellations and Brodmann's areas using Dice coefficients. To this end, we first match each parcel with a Brodmann area, if their overlap ratio is ≥ 0.5 (i.e. the Brodmann area comprises at least half of the other parcel) and then calculate the Dice coefficient between the matching pairs. It is worth noting that several parcels can be matched to the same cytoarchitectonic region and therefore merged into a larger parcel, which

could potentially yield a bias towards small, evenly sized regions having high Dice scores after merging.

In addition, we assess the inter-subject variability/consistency across parcellations by quantitatively identifying the cortical areas that are repeatedly assigned to the same/different parcels across subjects. This is achieved by computing a consistency map for all subjects/resolutions using the Dice-based matching algorithm, as defined in Chapter 3. This method calculates a consistency score for each vertex within a range of $[0, 1]$. Cortical regions with a high inter-subject consistency (values close to 1) are assumed to be parcellated in a similar way (e.g. having parcels of similar size/shape) for most of the subjects. On the other hand, consistency scores close to 0 indicate a high variability across subjects in terms of subdividing a particular cortical area (i.e. vertices being assigned to parcels with highly different size/shape).

It is worth noting that, one should interpret such maps with great care and always account for possible biases towards the underlying modality and/or processing techniques before drawing any conclusions. For example, one known issue associated with tractography algorithms is the propensity of streamlines terminating within gyri instead of sulci [260]. This gyral bias may influence the parcellation boundaries obtained from dMRI and consequently yield an inherent alignment with cortical folding [188]. To understand the impact of the gyral bias on the proposed parcellations, we make use of the average 'sulc' map, computed across all subjects in the dataset. A sulc map provides information about the depth of vertices within sulci and allows to visualise the macro-anatomy of the folded cortex for inflated meshes [73]. In other words, it indicates how deep and how high the brain folds are with respect to the distance of cortical vertices to the gyri and sulci. In our experiments, the average sulc map is obtained from the sulcal depth surface files, which are generated by the HCP structural preprocessing pipelines and made publicly available for each subject.

5.3.2 Comparison Methods

We compare our parcellations to the ones obtained by several other clustering methods, including spectral clustering with normalised cuts [54] (*N-Cuts*) applied to a spatially constrained

version of the affinity matrix computed for the proposed method; k-means clustering (*K-Means*) applied to the low-dimensional embedding; Ward's hierarchical clustering applied to the low-dimensional embedding (*HC-Low*) and driven by the connectivity profiles in the high-dimensional space (*HC-High*); a connectivity-driven, multi-scale spectral clustering technique based on dMRI [188] (*M-Scale*); random parcellations by Poisson disk sampling (*Random*); geometric parcellations, i.e. k-means clustering of the vertex coordinates [244] (*Geometric*).

M-Scale parcellations are obtained from the spatially-constrained affinity matrices using the publicly available code at `https://github.com/parisots/SpectralParcellation`. All other methods are adapted as described in Chapter 3, unless otherwise noted. Apart from *K-Means*, which is exclusively driven by the connectivity information, all methods are spatially constrained and do not yield disjoint parcels. *M-Scale* and *HC-High* are based on an initial connectivity-based over-parcellation of the cortex to compensate for the noise, and thus, to obtain higher accuracy (1000, 2000 and 3000 regions for *M-Scale*; 3000 regions for *HC-High*). *Random* and *Geometric* parcellations do not account for any connectivity information, therefore provide a baseline for the assessment [244].

5.4 Results

5.4.1 Parameter Selection

As there is no known optimal number of parcels, we evaluate the proposed method at different scales, determined by the number of eigenvectors incorporated into the boundary map. We present results for $d = 10$, 15, and 20 eigenvectors per hemisphere, which on average, yield parcellations with around 180, 230, and 280 regions for each subject, respectively. For a fair comparison, other methods are tuned to use the same number of parcels as inferred by our models.

The number of parcels obtained by the proposed method as a function of the number of eigenvectors is given in Fig. 5.6. As shown in the figure, fewer eigenvectors lead to very coarse

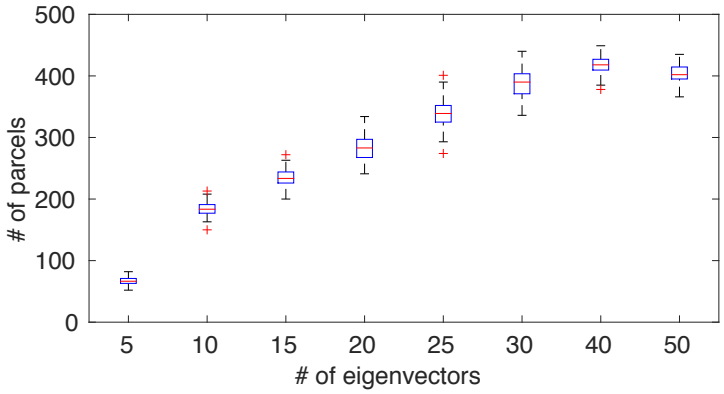


Figure 5.6: The number of parcels generated with respect to different number of eigenvectors used in the spectral embedding. Box plots show the variability across subjects.

parcellations, which may potentially under-represent the underlying connectivity. Increasing d beyond 20 eigenvectors, although at first introduce more parcels, gradually leads to lower-resolution parcellations with unusual borders. This is due to the fact that incorporating too many eigenvectors to the embedding yield very dense boundary maps and consequently reduce the number of markers that drive the watershed transform.

Another external parameter is the number of nearest neighbours k that is used to determine the connectedness of the affinity matrix. Intuitively, k should be large enough to effectively capture the local connectivity patterns across the cerebral cortex. Our experiments with values of k within a range of $[50, 200]$ revealed that $k = 100$ tends to provide locality-preserving affinity matrices for all subjects with most connections having high correlation coefficients ($r > 0.5$). If k is set to a very low value, the affinity matrix is likely to capture stable features regarding the spatial geometry of the cortical mesh. This may consequently result in regularly shaped parcels with similar size, and hence, may not be able to accurately represent the underlying structural connectivity. Picking a very large k leads to edges with negative correlation coefficients, indicating that connections tend to become weaker with increasing k. It also comes with an additional computational bottleneck, as the size of the matrix dramatically increases when more edges are formed between vertices.

5.4.2 Quantitative Assessment of the Parcellation Quality

We present the cluster validity results, including homogeneity values and Silhouette coefficients, based on structural and functional connectivity data in Fig. 5.7 and Fig. 5.8, respectively. Statistical significance of the results between the winner and the runner-up method is tested using a Wilcoxon sign-rank test.

Fig. 5.7 shows that our method surpasses other approaches in terms of homogeneity and Silhouette coefficients. The results are statistically significant across all resolutions at $p < 0.05$. It is followed by two clustering techniques, *K-Means* and *HC-Low*, both are driven by the low-dimensional embedding. This may indicate that, Laplacian eigenmaps can successfully reveal the intrinsic geometry of the underlying connectivity, and hence, provides a more robust set of features towards parcellating the cortical surface. In addition, the way we utilize discrete eigenvectors for boundary mapping allows the delineation of more distinct parcels compared to the traditional clustering approaches directly applied to the spectral coordinates. Considering the results obtained by *HC-High*, we can infer that non-linear dimensionality reduction can identify local connectivity patterns which may not be directly detected in the high dimensional space.

On the other hand, *M-Scale* and *N-Cuts* can provide accurate parcellations only to some extent and are even surpassed by *Random* for most of the resolutions. This may be linked to the fact that spatially-constrained spectral techniques are more likely to capture the structure of the cortical mesh rather than the connectivity patterns [244], since only the spatially-adjacent vertices are taken into consideration while constructing the connectivity matrix. *Geometric* unsurprisingly yields the least favourite performance regardless of the resolution, potentially indicating that spatial information alone may not be sufficient to reveal the connectivity-based organisation of the cerebral cortex.

The difference in performance between our approach and the others becomes more prominent with the resting-state functional connectivity results as indicated in Fig. 5.8. Both homogeneity and Silhouette analysis show a tendency in favour of the proposed method. Despite the fact

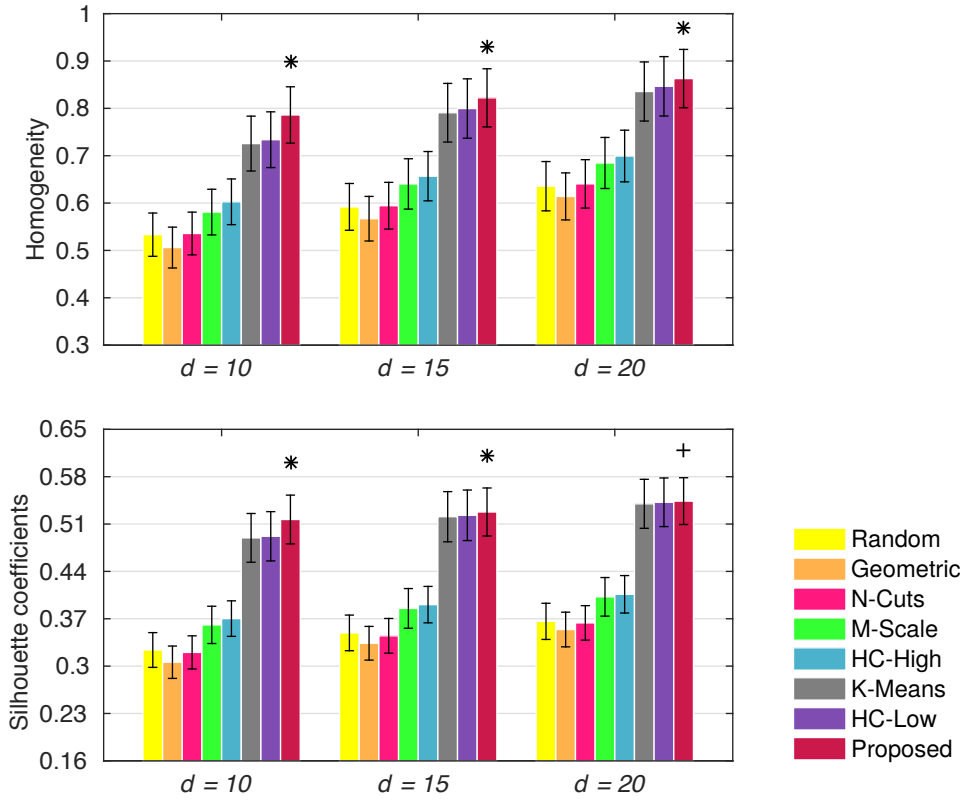


Figure 5.7: Homogeneity values and Silhouette coefficients based on structural connectivity estimated from dMRI. Error bars show the variability across subjects. Stars (*) and cross (+) indicate statistical significance between the winner and the runner-up with $p < 0.01$ and $p < 0.05$, respectively.

that parcellations derived from dMRI do not necessarily represent the functional connectivity estimated with rs-fMRI, these trends indicate that the proposed method better reflect the underlying function with respect to the other approaches. *HC-Low* is the only other method that consistently shows a higher performance than *Random*. Other parcellations, although yield homogeneous regions to some degree, generally fail to separate vertices with different signals from each other, as indicated by Silhouette coefficients.

5.4.3 Visual and Inter-Modality Assessment

Parcellations of the lateral cortex of the left hemisphere derived from one subject for varying resolutions are given in Fig. 5.9 for visual inspection. Parcellation resolutions are determined

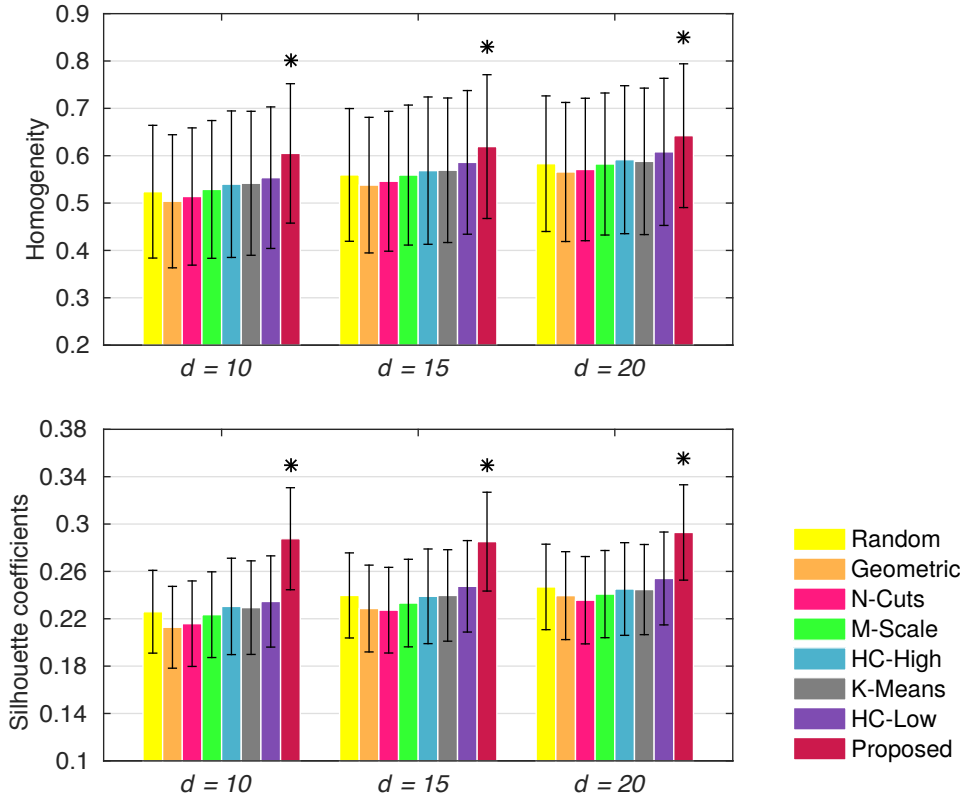


Figure 5.8: Homogeneity values and Silhouette coefficients based on functional connectivity estimated from resting-state fMRI. Error bars show the variability across subjects. Stars (*) indicate statistical significance between the winner and the runner-up with $p < 0.01$.

by the three sets of eigenvectors ($d = 10, 15, 20$) used in the proposed framework. The consistency/variability of the proposed parcellations across subjects/resolutions is shown with a consistency map in Fig. 5.10(a). A complimentary sulc map is also provided alongside the consistency map to show the average cortical-folding organisation across subjects, in which bright regions represent gyral crowns and dark regions represent buried cortex (the darker the shading the deeper the sulcus) [261]. This map allows to better understand the degree of gyral bias towards the proposed parcellations (see Discussion).

Myelin maps and Brodmann's areas projected onto the cortical surfaces of randomly selected subjects along with the proposed parcellations are given in Figs. 5.11 and 5.12, respectively. The white borders in the figures show the parcellation boundaries obtained by the proposed method at the highest computed resolution (using $d = 20$ eigenvectors). The Brodmann parcellations contain labels for the primary somato-sensory cortex (BA[3,1,2]), the primary motor cortex

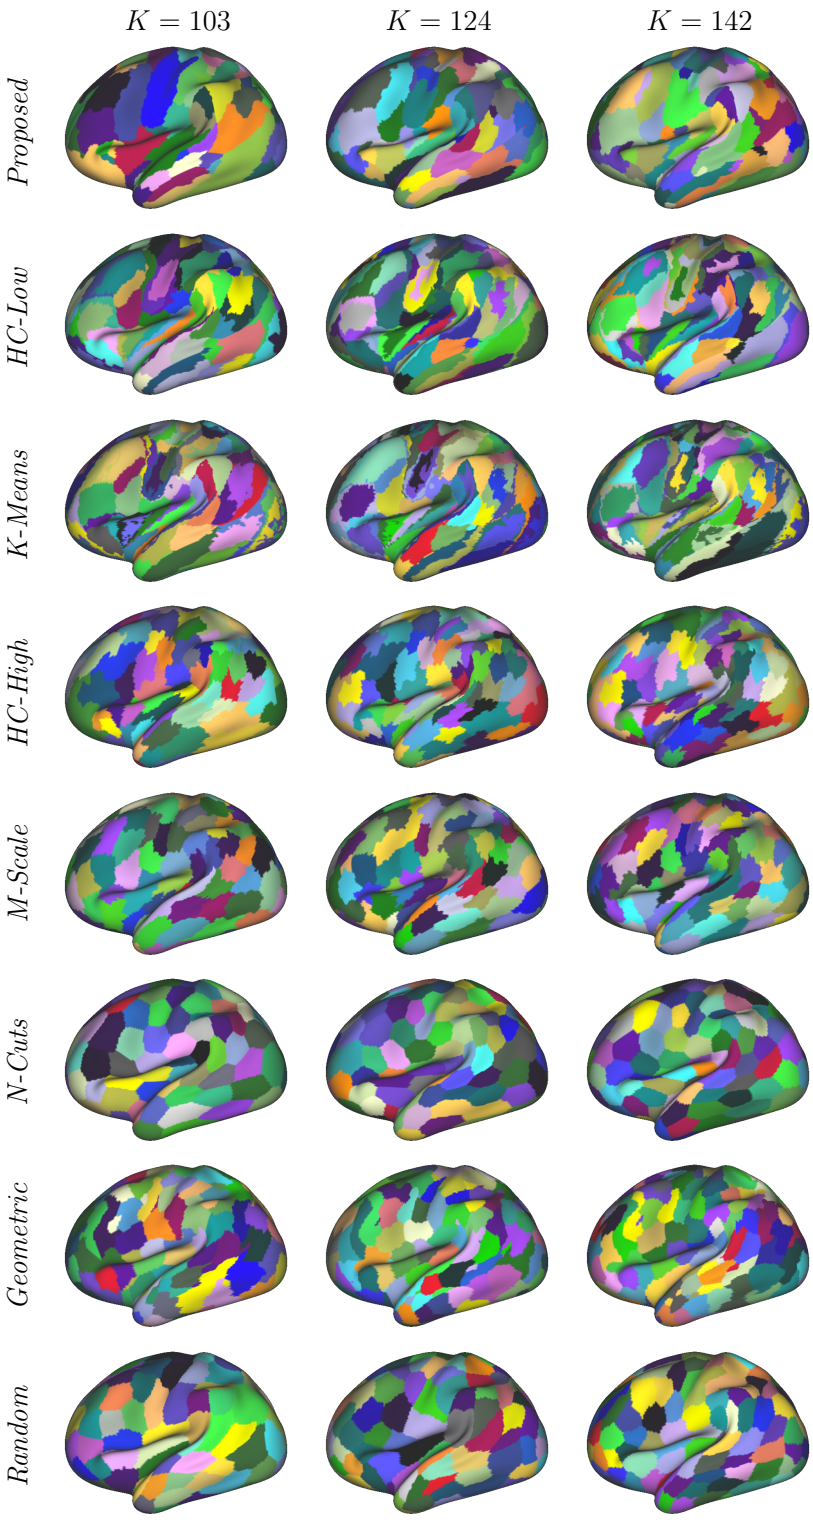
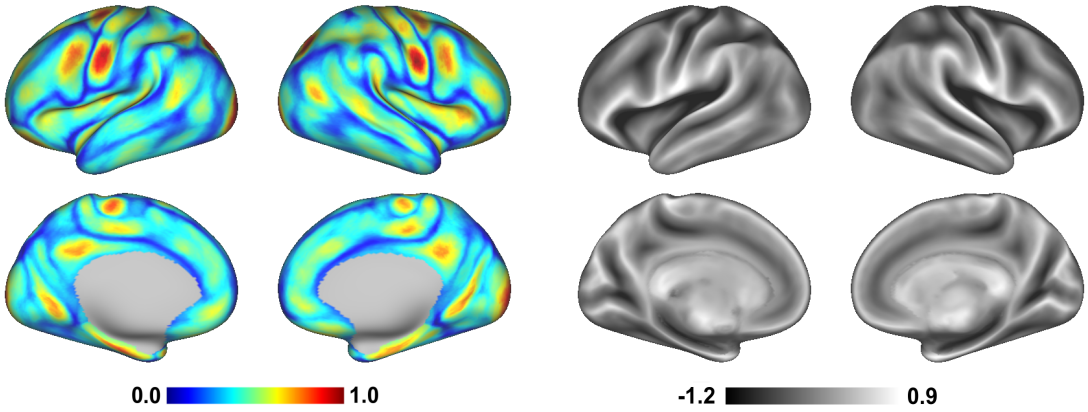


Figure 5.9: Parcellations of the lateral cortex of the left hemisphere derived from one subject. K is determined by the number of eigenvectors used in the proposed framework.



(a) (b)

Figure 5.10: (a) Inter-subject consistency map obtained from all subjects/resolutions. Values are normalised to $[0, 1]$ for better visualisation. Hotter colours indicate a higher consistency (better match) across subjects. (b) Average sulc map obtained from the complimentary sulc images provided by HCP for all subjects. The bright regions represent gyral crowns and dark regions represent buried cortex (the darker the shading the deeper the sulcus).

(BA4), the pre-motor cortex (BA6), Broca's area (BA[44,45]), the visual cortex (BA17 and MT), and the perirhinal cortex (BA[35,36]) as shown in Fig. 5.13(a).

Visual assessment of the parcellation borders reveals some alignment with highly-myelinated cortical regions and several Brodmann areas, especially around the motor area and the visual cortex. Although we only present parcellations of several subjects, it appears that a similar agreement is observed across many subjects as indicated by the inter-subject consistency map in Fig. 5.10. The degree of the alignment with respect to different Brodmann areas is further shown in Fig. 5.13. Average Dice scores obtained from all computed resolutions/subjects are overlaid on each Brodmann area in Fig. 5.13(b) and the box plots showing the variability across subjects are presented in Fig. 5.13(c). The highest overlap is observed within the primary somato-sensory cortex (BA[3,1,2]), pre-motor cortex (BA6) and primary visual cortex (BA17).

5.5 Discussion

In this chapter, we introduced a new connectivity-driven parcellation approach based on dMRI. The proposed method encapsulates the local connectivity characteristics with manifold learn-

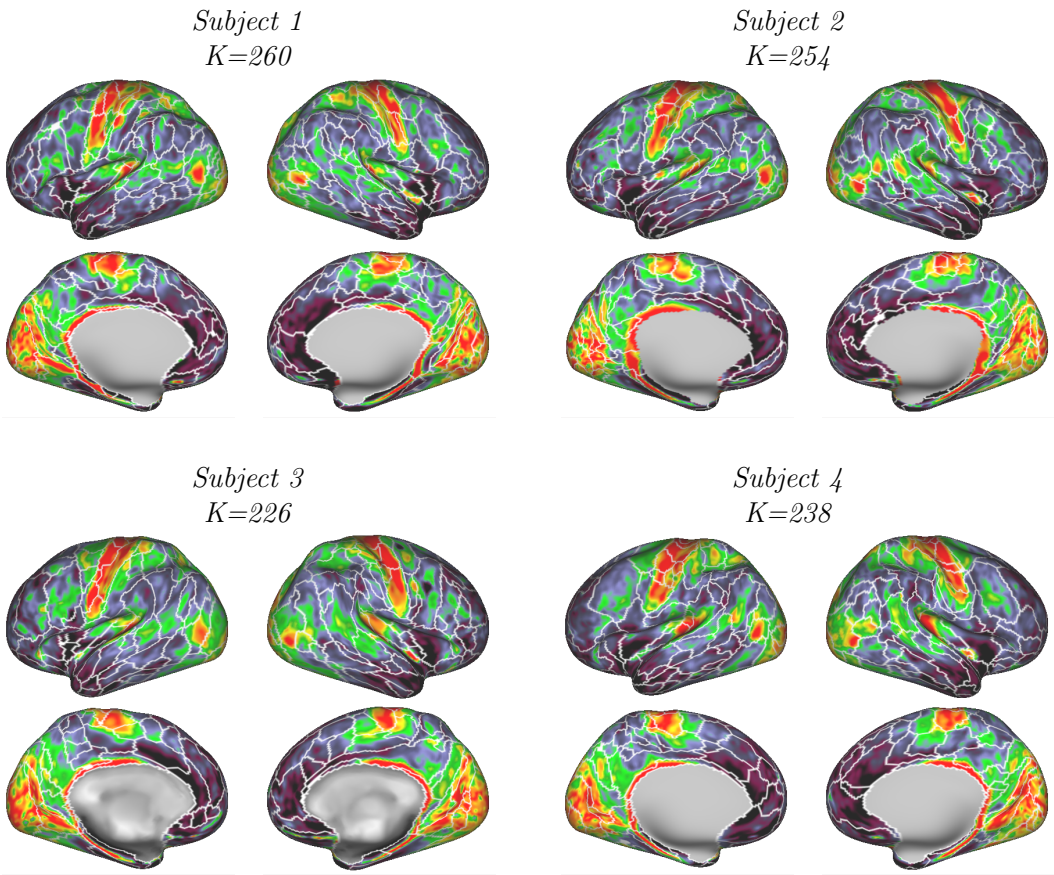


Figure 5.11: Parcellations of four different subjects obtained via the proposed approach are overlaid onto subject-specific myelin maps. The white borders show the parcellation boundaries delineated at the highest computed resolution (using $d = 20$ eigenvectors).

ing and uses the low-dimensional embedding to identify locations where connectivity patterns change. Particularly, these transition locations are interpreted as an abstraction of the parcellation boundaries and consequently allow to delineate distinct parcels at different scales. In the absence of a connectivity-driven cortical atlas, we showed that the proposed parcellations provide a more accurate representation of the underlying connectivity compared to different connectivity-based, anatomical, and random parcellations with respect to well-known clustering quality measures. Our additional experiments with resting-state fMRI revealed the potential of the proposed method to reflect the functional organisation of the brain, at least more faithfully than the other tested parcellations.

The visual and quantitative comparisons with the myelin maps and cytoarchitectonic regions showed that the proposed parcellations reflect the myelo- and cyto-architecture of the cerebral

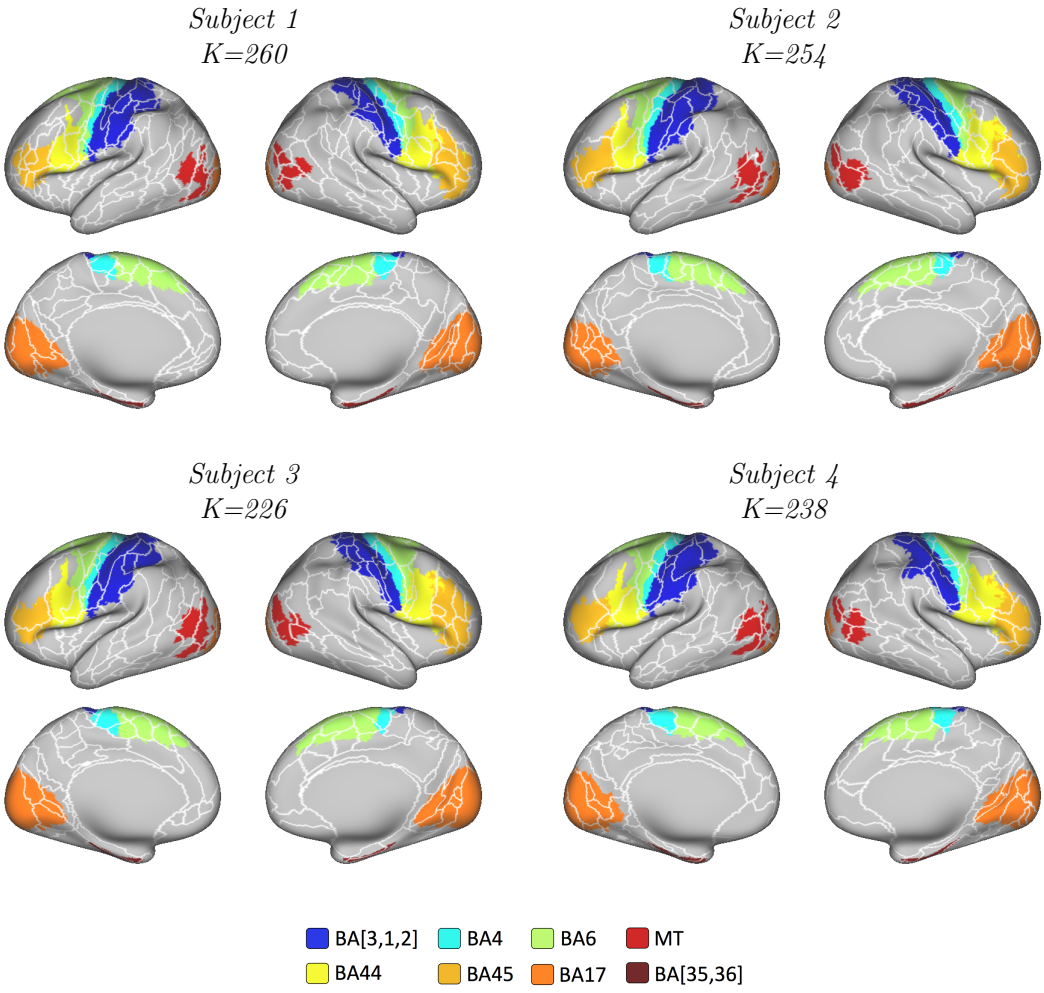


Figure 5.12: Parcellations of four different subjects obtained via the proposed approach are overlaid onto subject-specific cytoarchitectonic regions labelled with the Brodmann atlas. The white borders show the parcellation boundaries delineated at the highest computed resolution (using $d = 20$ eigenvectors).

cortex to some extent, especially within the motor area. Although, this trend appears to be consistent across parcellations with respect to the consistency map [Fig. 5.10(a)], such inter-modality comparisons should always be interpreted carefully. First of all, it should be noted that connectivity-driven parcellations obtained from dMRI do not necessarily align with the cortical organisation of the brain estimated via resting-state fMRI (or other neuro-anatomical properties), as each modality provides a different source of information that does not need to agree with the other [172]. In addition, cortical parcellations obtained from different data sources may be plagued by modality-specific noise and biases, e.g. physiological noise and

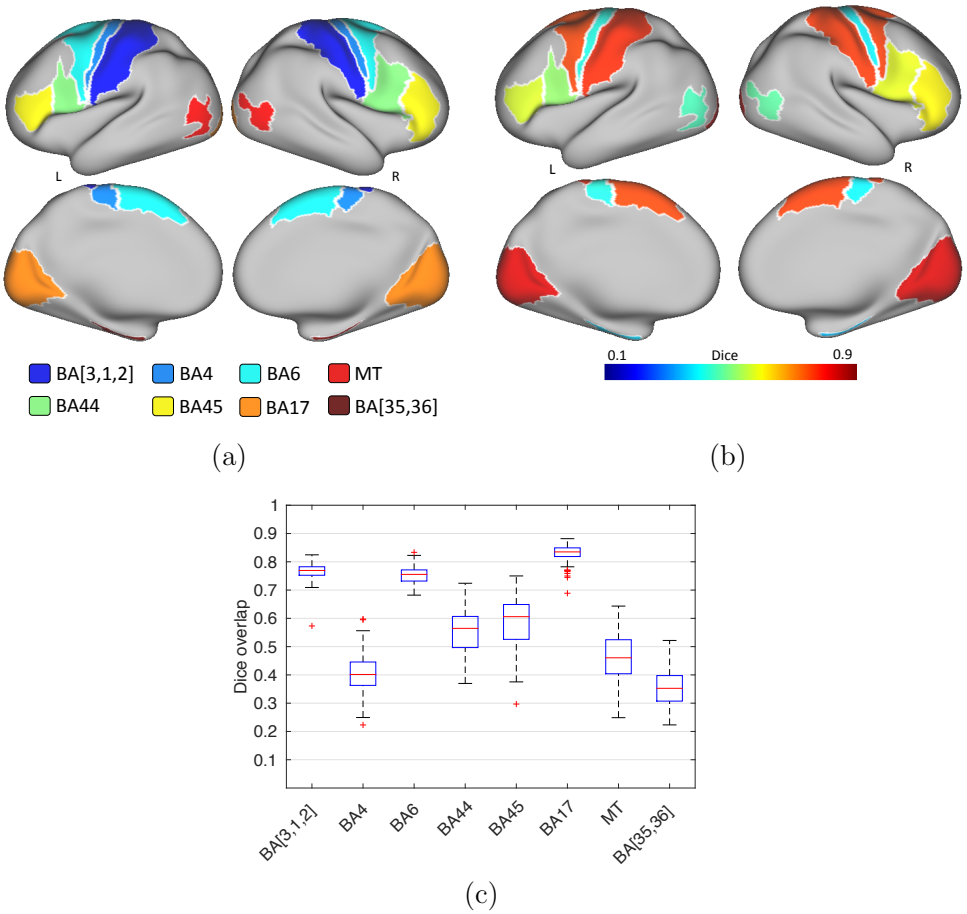


Figure 5.13: (a) Different cytoarchitectural areas as delineated from the Brodmann atlas. (b) Dice overlaps averaged across all computed resolutions (i.e. for $d = 10, 15, 20$) and subjects are displayed on each Brodmann area. Hotter colours indicate a higher overlap. (c) Box plots show the inter-subject variability for each Brodmann area.

head motion may yield spurious functional correlations, while tractography errors may bias structural connectivity [55, 66]). However, some degree of agreement with certain parts of the cerebral cortex that are known to have distinct neuro-biological features, or function, would potentially express the reliability of the proposed parcellation scheme, at least for the cortical areas with which high alignment is observed.

In addition, it is important to consider the possible limitations/biases introduced by the data processing pipelines or the evaluation techniques used to measure the agreement between different modalities. For example, the Brodmann areas are mapped onto each subject's surface using a cortical folding driven registration algorithm. As a result, a better alignment can be

achieved in the motor and visual cortex, where the folding patterns are more consistent across subjects [188]. On a similar note, the overlap-based measure used to assess the degree of agreement between parcellations and different Brodmann areas is biased by the size and shape of the parcels. Evenly sized/shaped parcels are easier to match with their target parcels, while differences in Dice scores will be much more striking when comparing small parcels over big ones. For example, despite the fact that the parcellation boundaries usually align well with the primary motor cortex (BA4) as shown in Fig. 5.12, the quantitative results given in Fig. 5.13(c) indicates less favourable overlapping scores for BA4 compared to the other motor-related Brodmann areas (BA[3,1,2] and BA6), which typically span a larger proportion of the motor cortex.

Another critical consideration point is the gyral bias inherent to probabilistic tractography, which is known to affect the estimated structural connectivity [260] and to influence the parcellation boundaries, such that they naturally align well with cortical folding. The sulc map in Fig. 5.10(b) can provide a clearer intuition regarding the impact of the gyral bias on the connectivity-driven parcellations, as it shows how the cortical sheet is folded. When compared to the consistency map (Fig 5.10), it can be clearly seen that the parcellation boundaries have a strong alignment with the gyral crowns (brighter areas). In other words, parcellation boundaries follow the cortical folding patterns, as the underlying streamlines that are used to estimate structural connectivity typically terminate in gyri rather than sulci [260].

In general, several other limitations should also be taken into account when interpreting parcellations derived from dMRI or comparing them with other modalities. First, dMRI provides a very indirect measurement of white matter connectivity (i.e. diffusion of water molecules in the brain). In addition, the processing methods (e.g. tractography) can suppress existing structural connections, and thus, alleviate the reliability of the connectome analysis. For example, while log transformation helps reduce tractography's bias towards short-range connections, it does not necessarily guarantee an accurate estimation of the long range connections [85, 190]. Other limitations include the dominance of large fiber bundles, impaired detection of crossing and kissing fibers, as well as difficulty to determine the origin or termination of the tracts [260, 172]. Despite all these drawbacks, dMRI constitutes a very important aspect of brain mapping, as it remains the best non-invasive way of measuring the physical (structural) connectivity within

the brain. While the approximated tracts may not be entirely accurate, the similarity between the estimated connectivity patterns to characterise the anatomical organisation of the brain could still be valid [126]. Consequently, the proposed parcellations hold a great potential for brain mapping studies.

On the other hand, when parcellations are considered from a neuro-biological point of view, using only a single modality might be rather limited in order to reveal the complex structure of the cerebral cortex, which consists of a mosaic of multiple properties nested at different levels of detail [86]. In addition, considering the limitations inherent to different modalities and processing techniques, the integration of multiple modalities to the parcellation generation task may potentially provide more accurate and robust cortical segregation of the cerebral cortex, as shown in the recently proposed multi-modal cortical parcellation [86]. A prospective future work, therefore, would be to approach the parcellation problem from a multi-modal point of view and develop a flexible parcellation framework that would allow combining information from different modalities.

While the proposed parcellations can successfully represent the structural organisation of the cerebral cortex on a single subject basis, the high variability in parcel size, shape, and number do not allow to construct spatial correspondence across subjects, which is generally a default prerequisite for group-wise connectome analysis. Nevertheless, the proposed approach can be potentially expanded for the generation of a parcellation that can reliably represent the shared characteristics within a population. For example, one can derive an average feature set from the subject-level spectral embeddings by first bringing them to a common space [132, 131] and then applying boundary mapping. Such an approach would provide more robust parcellations, as less distinct connectivity patterns are likely to be averaged out and not propagated to the parcellation stage.

However, it should be noted that averaging (on top of feature reduction and alignment) may also lead to loss of important information regarding connectivity features. As a result, averaging, although being highly popular, may not be the most reliable solution to the parcellation problem at the group level. Next chapter investigates different ways of obtaining group-wise parcel-

lations within the context of functional connectivity and proposes a new algorithm based on spectral decomposition for a whole-brain group-wise parcellation of the cerebral cortex, which can effectively reveal the fundamental properties of a population while preserving individual subject characteristics.

Chapter 6

Joint Spectral Decomposition of the Cerebral Cortex

This chapter is based on:

S. Arslan, S. Parisot, and D. Rueckert. *Joint Spectral Decomposition for the Parcellation of the Cerebral Cortex Using Resting-State fMRI.* Information Processing in Medical Imaging (IPMI) Lecture Notes in Computer Science, vol. 9123, pp. 85-97, 2015.

Abstract

This chapter describes a robust cerebral cortex parcellation method based on spectral graph theory and resting-state fMRI based connectivity that generates reliable groupwise parcellations across multiple subjects. Our method combines within- and inter-subject connectivity patterns to construct a multi-layer graph, which effectively captures the fundamental properties of the whole group as well as preserves individual subject characteristics. Spectral decomposition of this joint graph is used to obtain whole-brain parcellations of the cerebral cortex. Using rs-fMRI data collected from 100 healthy subjects, we show that our proposed algorithm computes more reproducible parcellations across different groups of subjects and at varying levels of resolution

compared to two other state-of-the-art spectral methods. We also report that our group-wise parcellations are functionally more consistent, thus, can potentially be used to represent a population in network analysis.

6.1 Introduction

The human brain is assembled into interacting regions that coordinate the nervous system. Identification of these regions is critical for a better understanding of the functional organisation of the brain and to reveal the interactions of underlying subsystems [233] that are responsible for different functions, including but not limited to, reasoning, memory, visual processing, and movement coordination. Functional connectivity studies have identified several subsystems, each of which spans across different cortical areas and associated with a different functional ability [212]. This has further advanced the analysis of the functional architecture of the brain by constructing graphical models of the connections within individual subsystems and their interactions with each other at different levels of detail [245, 202].

In whole-brain connectivity analysis, a critical stage is the parcellation of the cerebral cortex in order to compute a set of network nodes that can effectively represent cortical regions with uniform connectivity patterns. It is of great importance for a parcellation framework to be able to cluster vertices with similar functional connectivity together, thus, the average signal can reliably represent each part of the cortical region. It is also highly critical to generate a reliable group-wise representation that reflects the common functional characteristics of the population, yet is tolerant to changes in the functional organisation at the individual subject level that may emerge due to functional and anatomical variability across subjects. In this chapter, our motivation is to compute parcellations consisting of functionally homogeneous and spatially contiguous regions which can be used as the network nodes for a whole-brain connectivity analysis. To this end, we present a robust cerebral cortex parcellation framework based on spectral graph theory and resting-state fMRI connectivity (RSFC).

Parcellating the cerebral cortex based on RSFC can potentially identify functional organisa-

tion of the brain without relying on the anatomical features or an external stimulus/cognitive process [227]. RSFC-based cortical parcellation literature mainly consists of methods that subdivide the cerebral cortex into varying number of regions according to the requirements of the subsequent network analysis and topological network features across the cerebral cortex [59]. Some techniques [21, 255, 245, 202] parcellate the cortex at a very coarse level (less than hundred subregions), with the aim of identifying the functional communities, or resting-state networks, spanning across the cortex or some fractions of it. In this case, these parcellations cannot be reliably used for network node identification, as each region itself constitutes a gross, complex functional system, possibly exhibiting non-uniform functional patterns [226].

Other methods [54, 218, 93] typically generate a few hundred clusters without losing the ability of capturing the functional organisation of the cortex. This is generally achieved by (1) computing an average dataset from the RSFC data of each individual subject and submitting this to a parcellation algorithm or (2) first deriving individual subject parcellations and then generating a group-wise parcellation with a second-level algorithm running atop the subject-level parcellations. The group-wise parcellations generated by the former approach are not likely to reflect the individual RSFC characteristics as the inter-subject variability is usually over-smoothed during the averaging process. While the latter method allows the propagation of subject-level RSFC patterns into the group level, it still suffers from the adaptability of group representation to individual single subjects, as it is very unlikely that a group-wise parcellation obtained in such a setting would match with the subject-level parcellations. Besides, this method might result in slower computation as the parcellation should be computed for all subjects individually before generating a group-wise representation.

In this chapter, we propose an alternative approach for obtaining group-wise parcellations: We make use of spectral graph decomposition techniques and represent the population in a multi-layer graph which effectively captures the fundamental properties of the whole population while preserving individual subject characteristics. We show that the parcellations obtained in this setting are (a) more reproducible across different groups of subjects and (b) better reflect functional and topological features shared by multiple subjects in the group compared to the parcellations obtained by the other two approaches. Using rs-fMRI data collected from 100

healthy subjects, we show that our proposed algorithm computes highly reproducible parcellations across different subsets of the same population and at varying levels of detail with an average Dice score of 0.85, achieving up to 11% better group-to-group reproducibility. We also report that our group-wise parcellations are functionally more consistent, thus, can be reliably used to represent the population in network analysis. These aspects of the proposed method differentiate it from the previous parcellation algorithms and constitute our main contributions in this chapter. Finally, our framework can be used to generate parcellations with different number of subregions, allowing users to conduct a network analysis at different levels of detail.

In the remainder of this chapter, we first summarise the methodology (Section 6.2), starting with the dataset and following with a detailed explanation of the joint spectral decomposition method. In Section 6.3, we describe the experimental setup and provide details of the evaluation measures and comparison methods used to validate the proposed approach. In Section 6.4 we provide visual and quantitative results to show the effectiveness of the proposed method over other methods. Finally, in Section 6.5, we conclude the chapter by discussing the impact/limitations of the proposed method and provide some insight towards future research directions.

6.2 Methodology

6.2.1 Data

We conducted our experiments on Dataset 1, which contains 100 subjects (54 female, 46 male healthy adults, ages 22-35). The details of the dataset are provided in Chapter 2 and briefly summarised here. The data for each subject was acquired in two sessions that were held on different days and divided into four runs of approximately 15 minutes each. Data was preprocessed and denoised by the HCP minimal preprocessing pipelines [87]. As an additional step, 15-minute scans of each subject were temporally concatenated, after being normalised to zero-mean and unit-variance, yielding 60-minute rs-fMRI datasets for all subjects.

6.2.2 Joint Spectral Decomposition

We propose a general framework based on spectral decomposition to identify whole-cortex parcellations that can effectively capture the connectivity across multiple subjects. At the subject level, the cerebral cortex is represented as an affinity matrix, in which the within-subject connectivity is encoded as the edge weights between the interacting cortical units. Each individual affinity matrix is transformed to the spectral domain via an eigenspace decomposition and the corresponding eigenvectors are used to identify connectional similarities between subjects. Within- and inter-subject connectivity is then combined into a multi-layer graph, which is capable of representing the fundamental properties of the underlying functional organisation of the population. Similar to the subject-level graph decomposition, this multi-layer graph can then be decomposed into its eigenvectors, creating a feature matrix in the spectral domain (e.g. a spectral embedding) that can be fed into a clustering algorithm for producing the final parcellations. A visual summary of the approach is given in Fig. 6.1.

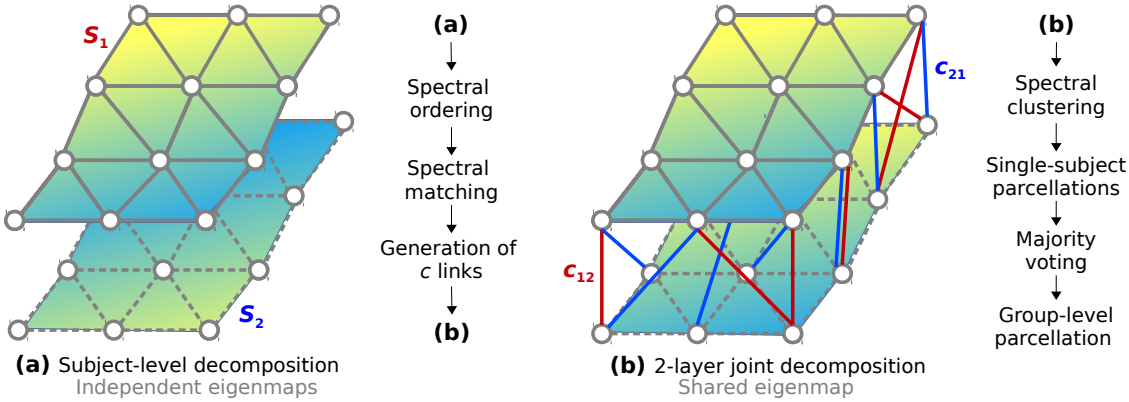


Figure 6.1: Visual representation of the parcellation pipeline with an emphasis on (a) single-subject and (b) joint spectral decomposition, illustrated on the patches cropped from the cortical surfaces S_1 and S_2. The red and blue edges correspond to the mappings c_{12} and c_{21}, obtained by matching the closest vertices in S_1 and S_2, respectively.

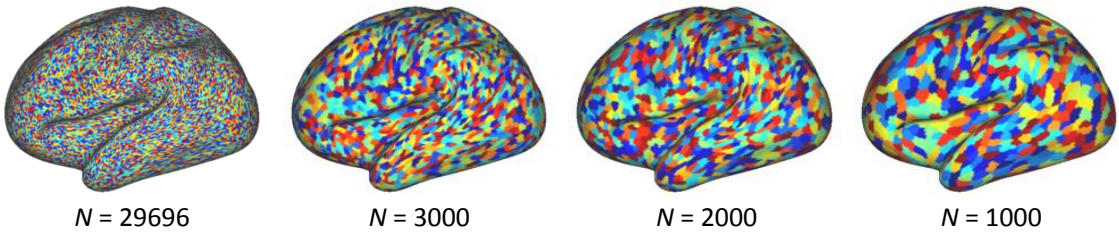
Sparse Affinity Matrix

The cerebral cortex is represented as a smooth, triangulated mesh with no topological defects, in which each vertex is associated with a timeseries that represents the fMRI signal. Typically, an

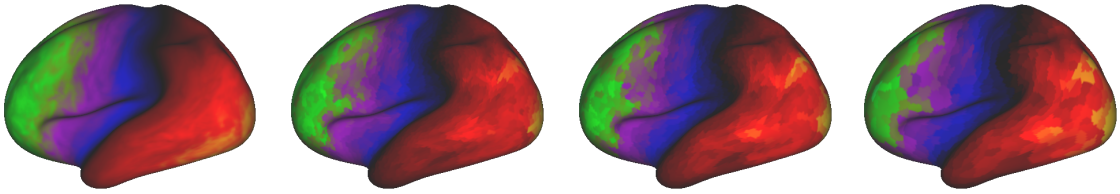
affinity matrix can be defined to capture the functional similarity across the cortex by modelling the cortical vertices and their associations as a weighted adjacency graph. However, mapping each cortical vertex in the mesh to a node in this graph can become computationally prohibitive due to high dimensionality of the data at the vertex level. For this reason, we pre-cluster cortical vertices into *supervertices* with the supervertex generation algorithm introduced in Chapter 4. In order to maintain the spatial correspondence across subjects, we define the supervertices on the average Conte69 cortical mesh provided by the HCP [259] and use the same set of supervertices for all subjects in the dataset. Supervertex parcellations obtained for different resolutions are shown in Fig. 6.2(a). Each supervertex is associated with the average timeseries of its constituent vertices, which further helps alleviate the impact of low SNR inherent in fMRI. For all supervertex pairs (v_i, v_j) an edge in the affinity matrix $W \in \mathbb{R}^{N \times N}$ is defined as follows.

$$W_{ij} = \begin{cases} r(t_i, t_j) & \text{if } i \text{ and } j \text{ are neighbours} \\ 0 & \text{otherwise} \end{cases} \tag{6.1}$$

where N denotes the number of nodes in the network (i.e. supervertices) and $r(t_i, t_j)$ is the Pearson's correlation coefficient between the timeseries t_i and t_j of v_i and v_j, respectively. Edges with negative correlations are set to zero in order to ensure that the affinity matrix is positive-semidefinite (i.e. all $W_{ij} \geq 0$) for each subject. Here, the benefit of introducing a spatial constraint and defining an edge only between the adjacent vertices in the mesh is two-fold: (1) It ensures that the resulting parcellations are spatially contiguous after applying spectral clustering and (2) it yields a sparse affinity matrix, and hence, avoids computational overheat during spectral decomposition.



(a)



(b)

Figure 6.2: (a) Cortical meshes visualised at different spatial resolutions. The first column shows the original mesh, in which each vertex is associated with a node in the affinity matrix W. It is followed by three supervertex parcellations at decreasing resolutions, denoted by N. (b) Eigenvectors associated with the first non-zero eigenvalues are superimposed onto the cerebral cortex. Spectral embedding are obtained by decomposing the corresponding $N \times N$ affinity matrices.

Spectral Embedding

Whereas parcellating the cortical surface into supervertices provides a means of *spatial* dimensionality reduction, the spectral decomposition of the affinity matrix yields a low-dimensional embedding, and hence, reduces the *temporal* resolution without losing the ability to represent the functional connectivity patterns encoded in the affinity matrix. A spectral embedding is typically obtained by using the first eigenvectors obtained through the spectral decomposition of the graph Laplacian as explained in Chapter 5. Given the affinity matrix W, the graph Laplacian can be defined as $L = D^{-1/2}(D - W)D^{-1/2}$, where D is a diagonal matrix with each entry $D_{ii} = \sum_j W_{ij}$ representing the degree of v_i. L is a diagonalisable matrix which can be factorised as $L = U\Lambda U^{-1}$, where $U = (u_0, u_1, ..., u_{N-1})$ is the eigensystem, with u_i representing each eigenvector and Λ is a diagonal matrix that contains the eigenvalues, represented as $\Lambda_{ii} = \lambda_i$. The first eigenvectors obtained after decomposing a randomly selected subject with respect to different spatial resolutions are shown in Fig. 6.2(b). This figure indicates that the intrinsic connectivity at the subject level can be captured by the spectral embedding, even if

the spatial resolution of the affinity matrix is immensely reduced by the use of supervertices.

Eigenvectors are powerful tools in terms of encapsulating valuable information extracted from the decomposed matrix in a lower dimensionality. In particular, sorting the eigenvalues as $0 = \lambda_0 \leq \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_{N-1}$ and organising the corresponding eigenvectors accordingly, we can use the set of first d eigenvectors after omitting the eigenvector corresponding to eigenvalue 0 to define a spectral feature matrix, i.e. $F = (u_1, u_2, \cdots, u_d)$, which is capable of representing the most important characteristics of the decomposed matrix. Thus, each node in the adjacency graph can be represented by its corresponding row in F, without losing any critical information, i.e. $v_i \mapsto (u_1(i), \ldots, u_d(i))$.

Spectral Matching

The idea of spectral matching is computing a mapping between two embeddings by comparing their eigenvectors [149]. Our observations on the cortical surfaces in the spectral domain revealed that eigenvectors show similar characteristics across subjects. This attribute can be exploited to obtain a joint spectral embedding that reflects functional features shared by the subjects in the group, while preserving individual subject characteristics.

However, spectral coordinates are not directly comparable. There might exist a sign ambiguity when computing eigenvectors[1], requiring sign checks of each eigenvector in the spectral embeddings. Similarly, the ordering of the eigenvectors may change [108] if the eigenvectors correspond to the same eigenvalue, as the solver may not necessarily compute the eigenvectors in the same order in both embeddings [148]. As a result, the same cortical information represented with the eigenvector u_i in F_1 can be decoded in the eigenvector u_j in F_2, without the necessity of being in the same order or having the same sign. Therefore, an additional correction must be carried out in order to find the corresponding eigenvectors on both cortical surfaces. We overcome this problem by aligning two spectral embeddings using Procrustes analysis [34, 133, 131]. This technique computes a linear transformation between a reference (e.g. F_1) and a target matrix (e.g. F_2), so that the points in the target matrix best match

[1] Given the general eigenvector problem, if $Ax = \lambda x$, then $A(-x) = \lambda(-x)$.

the points in the reference matrix with respect to a goodness-of-fit criterion, e.g. the sum of squared distances between two matrices. This typically handles the sign ambiguity and ensures that both embeddings have the same ordering.

After spectral ordering, the spectral matching problem can be solved by computing pairs of closest vertices on cortex S_1 and cortex S_2 with respect to their spectral feature matrices F_1 and F_2 as illustrated in Fig. 6.1. For points (supervertices) x on S_1 and y on S_2, the mappings $c_{12} : x_i \mapsto y_{c_{12}(i)}$ and $c_{21} : y_i \mapsto x_{c_{21}(i)}$ for $i = 1, \cdots, N$ can be established by running a k-nearest neighbour search with respect to F_1 and F_2. These mappings provide a set of inter-subject *links* connecting cortical surfaces to each other based on the similarity of their functional connectivity captured in the low dimensional space.

Definition of a Multi-Layer Graph Across Subjects

The use of spectral matching to find the mappings between pairs of cortical surfaces can be extended to generate a multi-layer graph for representing the whole group of subjects. A critical part in such a setting is the definition of edge weights that constitute the links between cortical surfaces. Using the correlations of rs-fMRI timeseries for this purpose is not sensible, since the mapped supervertices do not belong to the same subject. Instead, we use the correlations of the connectivity fingerprints to compute the connection strength of inter-subject links. A connectivity fingerprint is a feature vector that shows how a cortical subunit (e.g. a vertex, supervertex, or parcel) is functionally connected to the rest of the cerebral cortex [49]. It is defined by correlating the rs-fMRI timeseries associated with a supervertex with the rest of the supervertices and applying Fisher's r-to-z transformation to the resulting correlation coefficients. Connectivity fingerprints effectively reflect the functional organisation of the cerebral cortex and intuitively, we expect the fingerprints of linked supervertices to be similar as they are mapped to each other based on their functional connectivity patterns.

After defining inter-subject links for all subject pairs, we combine them with the individual sparse affinity matrices within a multi-layer graph $\mathbb{W} = (W^{(i,j)} \mid \forall\ i, j \in [1, M])$, in which the diagonal entities $W^{(i,i)} = W$ corresponds to single-subject affinity matrices reflecting the

within-subject connectivity and the off-diagonal entities $W^{(i,j)}$ $(i \neq j)$ denote the set of links c_{ij} and c_{ji} between cortical surfaces S_i and S_j, representing the *inter-subject* connectivity. $\mathbb{W}$ is an M-layer graph, with the size of $(N \times M) \times (N \times M)$ where M is the number of subjects in the group and N is the number of supervertices. A small patch taken from a 2-layer graph is illustrated as an example in Fig. 6.1(b).

Generation of Group-wise and Subject-Level Parcellations

The joint spectral decomposition of $\mathbb{W}$ provides a shared spectral embedding $\mathbb{F}$, representing every subject in the group. We compute the joint-embedding by following the steps described in *Spectral Embedding* (Section 6.2.2), since $\mathbb{W}$ holds the same properties as the subject-level sparse affinity matrices, i.e. it is symmetric and positive-semidefinite. Due to the fact that each eigenvector in the joint embedding characterises the shared connectivity patterns in the population, it can be used in a spectral clustering setting in order to compute parcellations across subjects. At this stage, applying k-means to the low-dimensional embedding can be the typical strategy for its simplicity and applicability [11], however, k-means can be rather unstable due to its dependency on a random initialization. As a result, the parcellations generated by k-means may not be reproducible on a run-to-run basis. To overcome this, we replaced k-means with an alternative technique that discretises the spectral embedding by learning a rotation of the eigenvectors [285] as suggested in [54, 244].

The output of spectral clustering is a label vector $\mathbb{L}$ of length $N \times M$ that assigns a parcel to each supervertex on all cortical surfaces used to define $\mathbb{W}$. By dividing $\mathbb{L}$ into sequential sub-vectors of length N, we can obtain a parcellation for each subject and identify brain regions that are consistent across subjects. A simple majority voting across the individual subject parcellations can then be used to generate the group parcellation.

6.3 Experiments

6.3.1 Evaluation

We evaluate our parcellation framework with respect to three criteria: (a) reproducibility, (b) fidelity of the parcellations to the underlying connectivity data, and (c) functional consistency between subject-level and group-wise parcellations.

Reproducibility is a widely-accepted technique for evaluating the robustness of a parcellation algorithm with respect to the extent of alignment in boundaries between different parcellations [54, 32, 218]. We use two well-known reproducibility measures, Dice coefficient and adjusted Rand index, to evaluate group-to-group reproducibility. Dice coefficient measures the overlap between two parcellations [62], while adjusted Rand index [103] evaluates the agreement of two parcellations without the necessity of matching parcels *a priori*, and hence, can be more effective when parcellations with different numbers of clusters are compared [163]. The details of these techniques are previously given in Chapter 4.

In order to assess how well a parcellation fits the underlying data, we make use of two cluster validity techniques, homogeneity and Silhouette analysis, both of which are previously introduced in Chapter 4 and only summarised here. Homogeneity measures a parcellation's ability to group vertices with similar connectivity into same clusters and is calculated as the average similarity between every pair of vertices assigned to a parcel. A global homogeneity value is obtained by averaging the homogeneity values across all parcels [54]. Silhouette analysis measures how well vertices fit in their assigned parcels. Different from homogeneity, not only it evaluates the compactness of parcels, but also their degree of separation from each other. It is achieved by combining within-parcel dissimilarity defined as the average distance to all other vertices in the same parcel and inter-parcel dissimilarity obtained from those assigned to other parcels. A global score is obtained by averaging the Silhouette coefficients across all vertices.

Another critical performance measure for group-wise parcellations is their ability to represent individual subjects in terms of functional consistency. We expect the variability in connectivity

networks to be consistent and minimal in order to reliably use the group-wise representation in place of each subject in the group. To this end, we evaluate functional consistency by computing the sum of absolute differences between functional connectivity networks, when we replace a subject-level parcellation with its corresponding group-wise parcellation without changing the underlying RSFC data. A connectivity network is simply defined by cross correlating the average timeseries associated with each parcel. In order to compare two networks, we first match the single- and group-level parcellations using the Dice-based similarity method and exclude the non-matching parcels from the comparison in order to allow an objective comparison of all methods. This procedure is illustrated in Fig. 6.3.

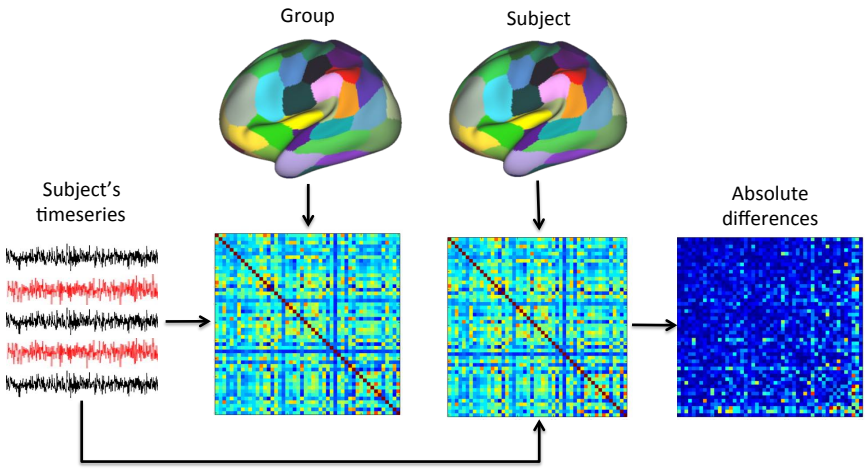


Figure 6.3: Illustration of how to compute the sum of absolute differences (SAD) between a group-wise and subject-level parcellation.

6.3.2 Comparison Methods

We compare our framework with two other parcellation strategies, both are driven by spectral clustering with normalised cuts (NCUTS) [54]. The first strategy consists in performing clustering for each subject individually and applying a second level clustering to subject-level parcellations in order to obtain group-wise parcellations (i.e. two-level approach). The second approach consists of computing an average matrix from the population and submitting this matrix to a clustering algorithm (i.e. group-average approach). We chose NCUTS as the clustering algorithm for the comparison methods, as our method also relies on NCUTS to compute

parcellations from the multi-layer graph. Therefore, it provides a consistent means for assessing the performance of the proposed framework (joint spectral decomposition) with respect to two alternative strategies (two-level and group-average).

Two-level approach. This approach is similar to majority voting, in the sense that vertices assigned to the same region across subject-level parcellations are clustered together. As a result, group-wise parcellations obtained via this method can capture the shared characteristics of the population as approximated by the individual parcellations. To this end, a graphical model of the 'parcel stability' is computed across all individual parcellations [54, 255]. We obtain subject-level parcellations by applying NCUTS to individual affinity matrices as defined in *Sparse Affinity Matrix* (Section 6.2.2). We then construct an adjacency matrix (i.e. stability graph), in which an edge between vertices v_i and v_j is weighted by the number of times both vertices are assigned to the same parcel across all individual subject parcellations. Since the initial parcellations are spatially contiguous, only vertices sharing the same cluster membership can be connected in the adjacency matrix, and hence, the spatial integrity of the second level parcellation is also guaranteed. Following the steps described in *Spectral Embedding* (Section 6.2.2), a low dimensional representation of the stability graph is computed and submitted to NCUTS. An illustration is provided in Fig. 6.4 that explains the construction of a stability graph with 4 toy parcellations. This method will be referred to as *Two-Level* throughout the rest of the chapter.

Group-average approach. This technique aims to capture shared patterns between individuals within a population, by computing a group average representation of connectivity. This is achieved by averaging the individual affinity matrices (after applying Fisher's r-to-z transformation) and then submitting the average matrix to NCUTS. This method will be referred to as *Group-Avr* throughout the rest of the chapter.

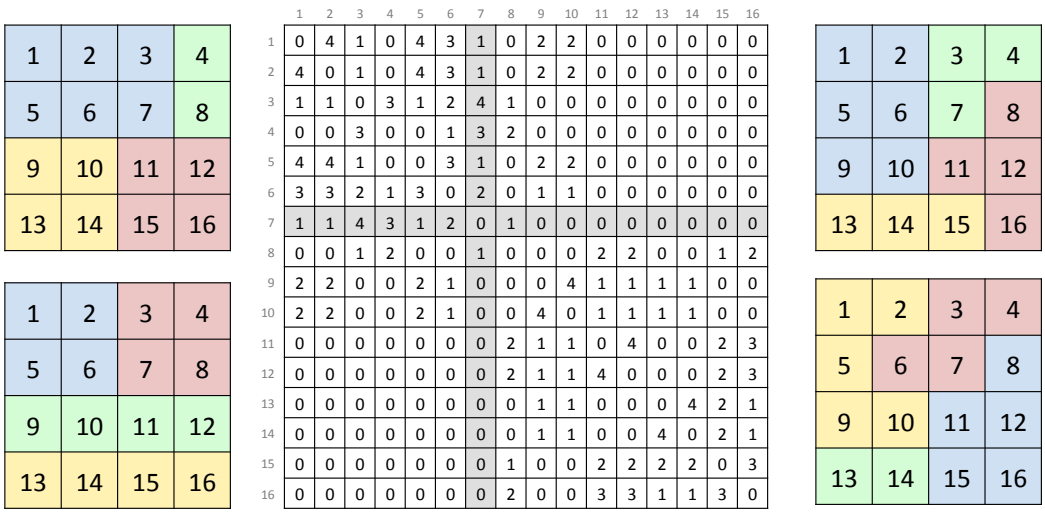


Figure 6.4: Illustration of how to compute a symmetric 16×16 adjacency matrix A from four toy parcellations (4×4 matrices) where each colour represents a different label/parcel. For example, vertex v_7 (with its corresponding row and column highlighted in A) is assigned to the same parcel as v_3, v_4, and v_6, for 4, 3, and 2 times, respectively, giving $A_{7,3} = 4$, $A_{7,4} = 3$ and $A_{7,6} = 2$. $A_{7,1} = A_{7,2} = A_{7,5} = A_{7,8} = 1$, since it shares the same label with v_1, v_2, v_5, and v_8 in just one parcellation, while the rest of the entries in row 7 of the adjacency matrix are 0, since there does not exist a shared label between the other vertices and v_7 in any of the toy parcellations.

6.3.3 Experimental Setup

In our experiments, we subdivide the dataset into two equally-sized, mutually exclusive groups (Subgroup 1 and Subgroup 2) by random selection and compute group-level parcellations for each subgroup (of 50 subjects) by running the algorithms separately on the left and right hemispheres. For each method, we generate parcellations containing $K = 100$ to 500 regions (i.e. 50 to 250 per hemisphere, in increments of 50). This process was repeated for 100 times, each time forming new subgroups of 50 subjects and generating two new group-wise parcellations for the same population. We assess reproducibility by comparing parcellations obtained from Subgroup 1 and Subgroup 2 in each repetition. For the homogeneity and Silhouette analysis, we only evaluate the parcellations obtained from Subgroup 1, but with respect to the data from Subgroup 2 so as to reduce the possible bias towards parcellation generation methods.

6.3.4 Parameter Selection

Other than the number of parcels K, the proposed algorithm has three external model parameters: (1) the number of eigenvectors retained in the spectral feature matrix for spectral matching, d, (2) the number of supervertices used to pre-cluster the cortical vertices, N, and (3) the number of links that connect each supervertex to other cortical surfaces, k.

The eigenvectors corresponding to the smallest eigenvalues represent the most important characteristics of the input dataset. However, there is still no consensus on how to select the optimal number of eigenvectors in the low-dimensional embedding. We retained the first 50 eigenvectors for the spectral matching process, since our experiments showed that further increasing d did not change the mappings between subjects, thus had no impact on the final parcellations.

The number of supervertices (N), and the number of links that interconnect subjects (k), have a more critical role in the proposed framework as they might directly affect the computed parcellations. Decreasing the spatial dimensionality through an initial clustering stage to some extent is primarily beneficial to alleviate the impact of noise at the vertex level. A very coarse initial parcellation, however, would lead to over-smoothing of the subject-level connectivity, and may consequently generate a bias towards capturing features regarding the spatial geometry of the multi-layer graph, rather than the underlying function. Similarly, k determines the number of links that map one subject to another, which are consequently used to define the inter-subject connectivity in the multi-layer graph. Therefore, k can be considered as a weighting parameter that controls the influence of the inter-subject connectivity over the intra-subject connectivity. The more number of links are constructed between subjects, the more the multi-layer graph is dominated by inter-subject variability, which may consequently lead to less reproducible group-wise parcellations. In our experiments, we selected these parameters as $N = 2000$ and $k = 10$. We further analyse the impact of the selection of these parameters to parcellation performance in the following section.

6.4 Results

We present the group-wise reproducibility results obtained by the proposed algorithm as well as the comparison methods in Fig. 6.5. Both Dice coefficients and adjusted Rand indices indicate that our joint spectral decomposition approach is able to obtain more reproducible parcellations at each level of resolution, with at least an average Dice score of 0.82. *Group-Avr* provides the least favourable results compared to the other two approaches. This might indicate that methods that propagate information from subject-level parcellations in order to generate group-wise representations yield more reproducible results, as they already provide a means of spatial smoothing. There is a general decreasing trend in all methods with the increasing parcellation resolution, which can be attributed to the fact that, as K gets larger, the functional variability across subjects gets more prominent, thus, reducing the similarity between the common characteristics within different groups and leading to less reproducible group parcellations.

In Fig. 6.6, we present the group-wise parcellations obtained by our approach from one group of 50 randomly selected subjects with different number of parcels and visualize the reproducibility of each group-level parcel across individual subject parcellations. Cross-subject reproducibility is measured by Dice similarity. In order to compute how reproducible a region is across different subjects, we follow the following procedure: For each group-level parcel we find its match in a subject-level parcellation and record their Dice score. We repeat the same process for all 50 subjects and average the Dice scores to obtain the reproducibility measure for that group-level parcel. Hot colours indicate a higher reproducibility across single-subject parcellations. Due to functional and structural variability between different individuals and the varying levels of SNR in the rs-fMRI data, it is not possible to obtain high Dice scores for each part of the cerebral cortex [93]. In general, the similarity between the group-wise parcellation and the individual subject parcellations drops with increasing resolution, as the impact of inter-subject variability gets stronger at higher level of granularity. Nevertheless, our approach is robust enough to achieve an average Dice score of at least 0.7 for each group-wise parcel.

Goodness-of-fit of the parcellation approaches in terms of homogeneity and Silhouette coeffi-

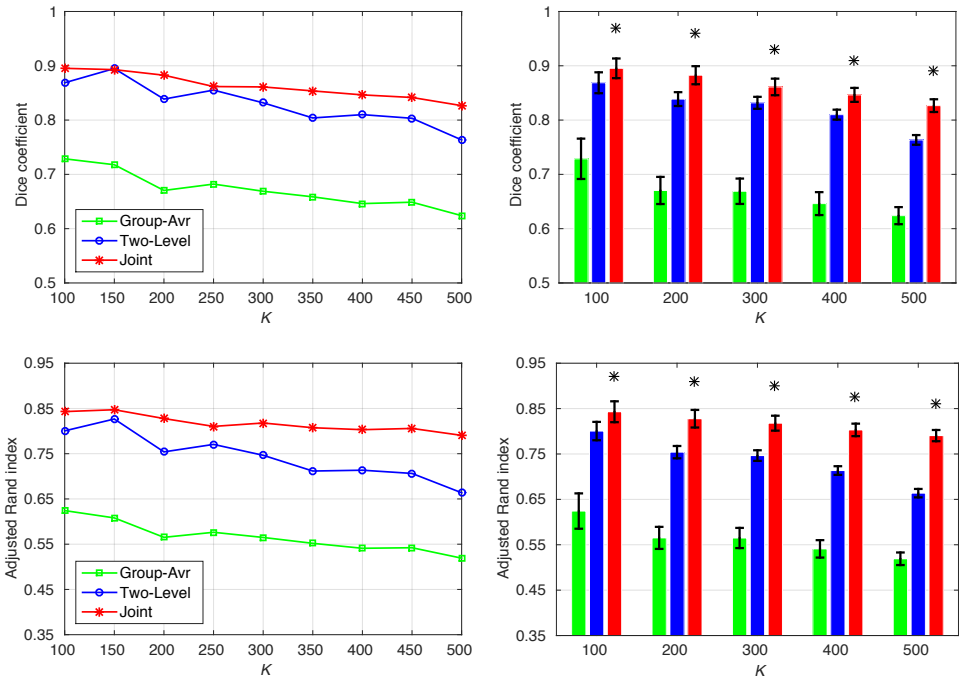


Figure 6.5: Group-to-group reproducibility results for different number of parcels, measured by Dice similarity (*top*) and adjusted Rand index (*bottom*). Reproducibility is computed between two equally-sized, mutually exclusive groups of 50 subjects, randomly selected from the dataset. This is repeated for 100 times, each time forming two subgroups and measuring reproducibility between the parcellations generated for each subgroup. Whiskers in the bar graphs indicate the variability across repetitions as measured by standard deviation. Stars (*) show statistical significance between the winner and the runner-up based on Wilcoxon signed rank test with $p < 0.01$.

cients is reported in Fig. 6.7. Although, homogeneity results show a relatively similar performance for all of the approaches, the proposed method consistently generate the most homogeneous parcellations for almost all resolutions, apart from the highest level (i.e. $K = 500$) where *Two-Level* appears to perform slightly better. Similar trends are observed with Silhouette coefficients, in which the proposed approach shows a more favourable performance for most of the resolutions, particularly for $K = 100$ to 300, but is outperformed by the other approaches at the highest level of granularity. As a general trend in both measures, all methods show a performance increasing with the number of parcels computed. This might be linked to the fact that both measurements are dependent on the size of the parcels (e.g. smaller parcels yield better results) [93, 54].

In Fig. 6.8, we present the average sum of absolute differences (SADs) between the functional

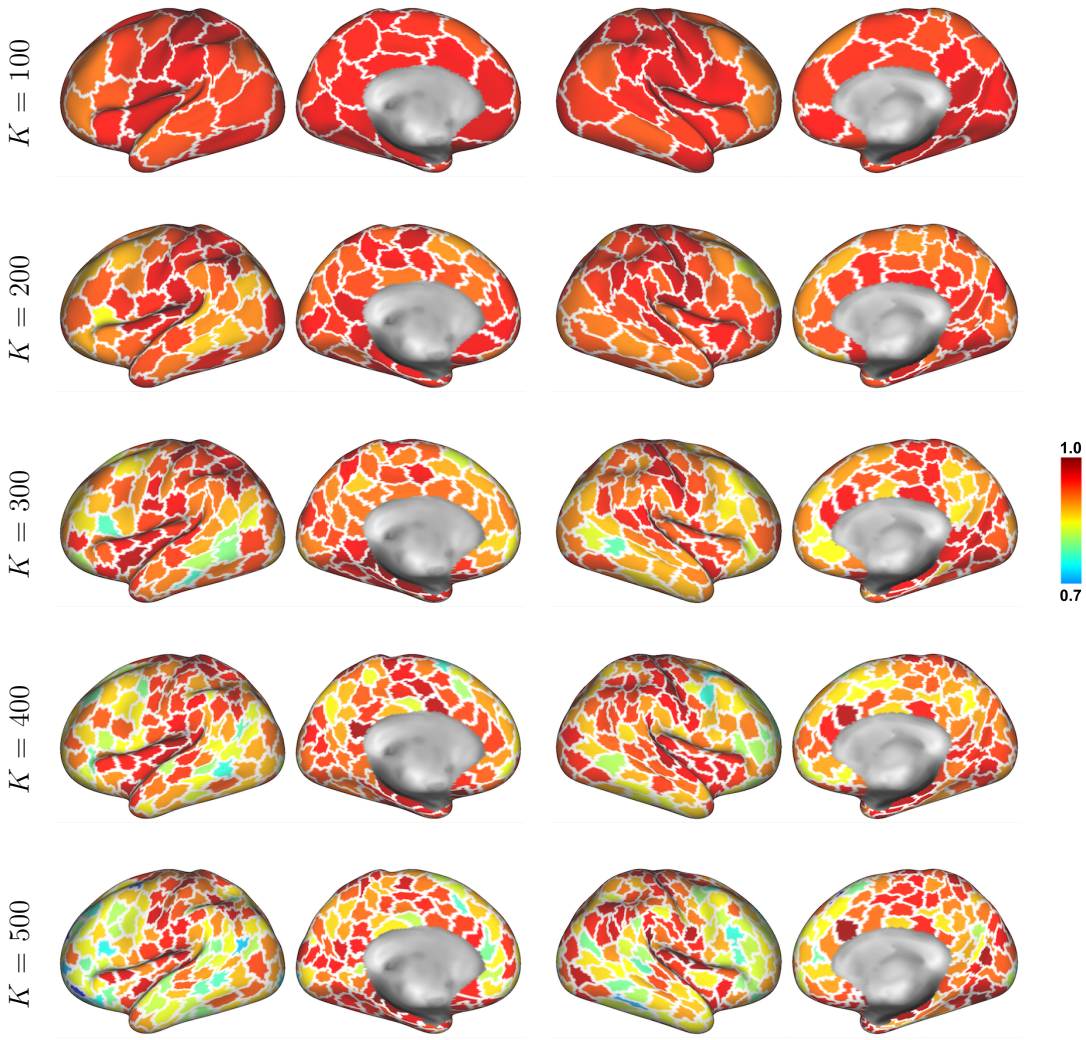


Figure 6.6: Group-wise whole-brain parcellations obtained by the joint spectral decomposition method for $K = 100$, 200, 300, 400, and 500 parcels. The colouring indicates the average reproducibility score (Dice coefficients) of each parcel across subject-level parcellations.

connectivity networks obtained by the individual parcellations and their group-wise representations. We excluded *Group-Avr* from this experiment, since it does not provide individual subject parcellations as part of its pipelines. The results show a clear tendency in favour of the proposed method at all resolutions, indicating that the group-wise parcellations obtained by *Joint* are functionally more consistent, and thus, can potentially represent the population with a higher reliability in network analysis. As expected, SADs show a steadily-rising trend for both approaches, due to increasing variability between the group- and subject-level parcellations with respect to network dimensionality, which directly depends on the parcellation

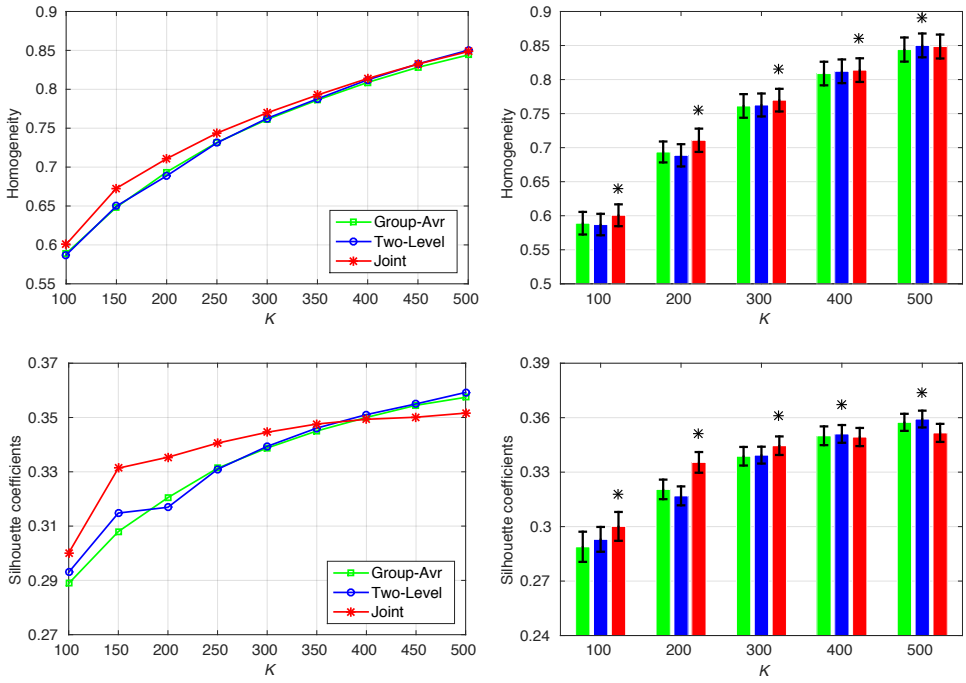


Figure 6.7: Homogeneity (*top*) and Silhouette analysis (*bottom*) results for different number of parcels. Both measures are computed for 100 different group-wise parcellations, each of which is obtained from a set of 50 randomly selected subjects. Whiskers in the bar graphs indicate the variability across runs as measured by standard deviation. Stars (*) show statistical significance between the winner and the runner-up based on Wilcoxon signed rank test with $p < 0.01$.

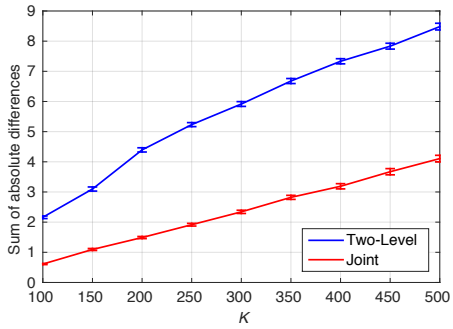


Figure 6.8: Functional consistency results of the proposed method (*Joint*) and *Two-Level* for different number of parcels. The sum of absolute differences (SADs) are computed between the connectivity networks obtained by a group-wise parcellation and its corresponding subject-level parcellations, and repeated for 100 different group parcellations, each of which is obtained from a set of 50 randomly selected subjects. The results shown here are the average SADs across all runs, normalised by the number of parcels at each level of granularity. Whiskers in the bar graphs indicate the variability across runs as measured by standard deviation. Lower SADs indicate better performance.

resolution.

6.4.1 Parameter Analysis

We analyse the impact of the external model parameters on the parcellation performance. We run our algorithm with different values for each parameter, while fixing the value of the other parameter and observe the changes in four different performance measures, namely, group-to-group reproducibility, homogeneity, Silhouette coefficients, and sum of absolute differences, over a spectrum of parcellation resolutions. This analysis is conducted based on the experimental setting described above, but repeated for a fewer number of runs (i.e. 10) and only on the left hemisphere. We consider any combination of the following parameter values $k = \{1, 10, 50\}$ and $N = \{1000, 2000, 3000\}$ per hemisphere. We present the results obtained for each parameter in Fig. 6.9.

The first parameter is the number of inter-subject links k, which is used to define the mappings in spectral matching and stands for as a weight parameter between the intra- and inter-subject connectivity represented in the multi-layer graph. Increasing the value of this parameter leads to an inclination towards inter-subject connectivity in the graphical model, yielding less reproducible group-wise parcellations and higher sum of absolute differences between subject and group parcellations. However, lowering k does not always reduce the group-to-group reproducibility and having around 10 links per supervertex appears to produce more robust parcellations. Considering the fidelity of the parcellations to the underlying data, increasing k helps obtain more faithful representations of the functional connectivity within the population as indicated by the homogeneity values and Silhouette coefficients.

The second parameter is the number of supervertices, N, used for spatial dimensionality reduction of intra-subject connectivity matrices to alleviate the impact of noise at the vertex level. Using too coarse parcellations lowers reproducibility, which might be attributed to the fact that functional connectivity at the subject level is over-smoothed, and hence, the parcellation framework relies more on the inter-subject connectivity for the generation of group-wise parcellations. A similar trend can also be observed in the SAD results, where the dissimilarity between subject and group-wise networks increases more rapidly with decreasing supervertex resolutions. The homogeneity and Silhouette analysis results are less significantly affected by

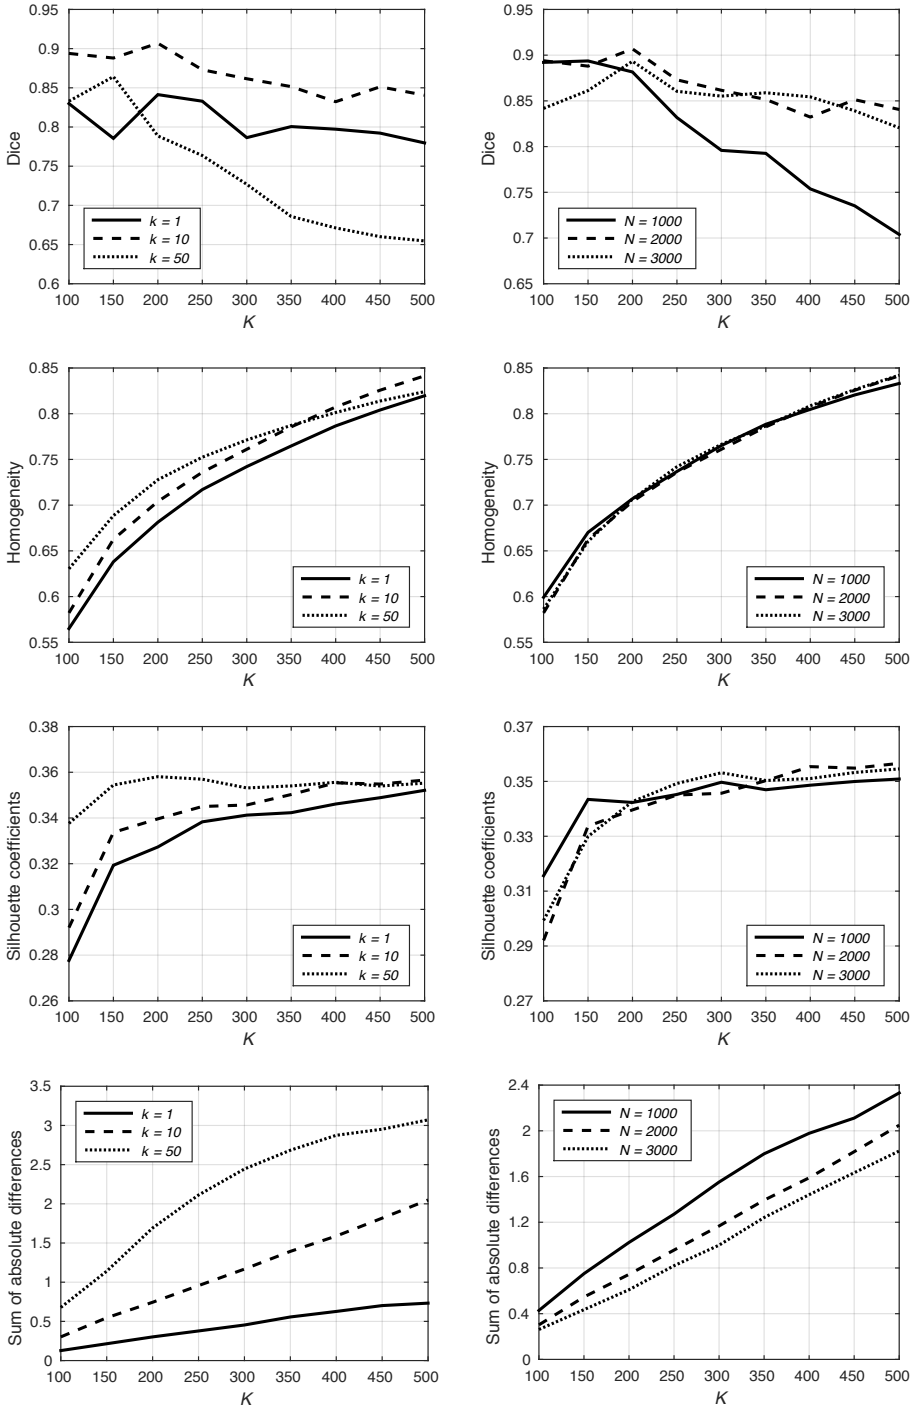


Figure 6.9: Evolution of Dice-based group-to-group reproducibility, homogeneity, Silhouette coefficients, and sum of absolute differences as a function of (*left*) the number of inter-subject links k and (*right*) the number of supervertices N, with respect to varying number of parcels K.

the number of supervertices, which might also indicate that the fidelity of the model is more dependent on the number of links between the subjects.

It is worth noting that, computational performance also remains as a critical consideration point regarding the selection of k and N. Together with the number of subjects in the population, both parameters have a direct impact on the size of the multi-layer graph, and hence, can dramatically increase the computational space/time requirements for joint spectral decomposition. In Table 6.1, we present the computational performance measures of the decomposition process for different combinations of k and N. Table 6.1(a) shows the amount of memory allocated for each multi-layer graph based on a different parameter setup, while Table 6.1(b) presents the computation time during the execution of their spectral decomposition. Each multi-layer graph is defined as a sparse matrix from a set of 50 subjects and the spectral decomposition is performed using the *eigs* function in Matlab (ver R2015a), on a 32-Core CPU server with 256GB RAM, operated by 64bit Ubuntu.

Table 6.1: (a) Size of memory and (b) computation time for the spectral decomposition of multi-layer graphs obtained with varying k and N (mb - mega bytes, gb - gigabytes, min - minutes).

	N = 1000	N = 2000	N = 3000
k = 1	47 mb	393 mb	1.78 gb
k = 10	103 mb	816 mb	3.58 gb
k = 50	163 mb	1.24 gb	5.44 gb

(a)

	N = 1000	N = 2000	N = 3000
k = 1	1.35 min	1.62 min	5.09 min
k = 10	3.21 min	3.54 min	10.71 min
k = 50	5.18 min	6.92 min	21.93 min

(b)

6.5 Discussion

In this chapter, we presented a spectral graph decomposition approach to parcellate the cerebral cortex with respect to functional connectivity estimated from resting-state fMRI. Our experiments demonstrated that the proposed algorithm can produce robust parcellations with higher reproducibility and can better reflect functional and topological features shared by multiple subjects compared to other parcellation methods. The performance of our framework can be

attributed to the proposed multi-layer graphical model, which combines functional connectivity captured at the subject level with the general functional tendency of the population.

Group-wise parcellations obtained by our approach can be used to define network nodes in connectome analysis, as they represent functionally uniform regions across the cerebral cortex. In general, such parcellations can be useful for comparing different groups of healthy subjects, for example, in order to identify how connectivity changes through ageing or from one gender to the other. This could further enable generation of a functional connectivity atlas, assuming that a very large cohort is used to minimise the inter-subject variability and noise-related artefacts. In addition, analysing connectivity patterns in a group of subjects with a brain disorder could help derive biomarkers in order to better understand disease-related differences in the brain connectivity.

Spectral clustering can effectively capture the shared connectional features across the cerebral cortex, however due to hard spatial constraints imposed to the connectivity data, parcellations can be biased by the geometry of the cortical mesh [244]. Since the normalised-cuts optimisation simultaneously maximises the inter-cluster dissimilarity and within-cluster similarity, parcellations obtained via spectral clustering tend to consist of uniformly sized regions. This constitutes a well-known limitation of spectral techniques, as parcels of uniform size cannot effectively adapt to cortical regions that are naturally irregular in shape and size [66]. One solution to overcome this could be to perform parcellation at a higher resolution, i.e. over-parcellate the cortex, and look for consistently identified boundaries across different scales. However, it should be noted that, spectral clustering algorithms do not inherently provide parcellations that are hierarchically consistent, hence, parcellations obtained at different resolutions do not necessarily share similar boundaries.

We developed our approach from functional connectivity captured in the form of spontaneous resting-state fMRI correlations, however, the proposed joint decomposition framework can easily be generalized to operate on other types of connectivity data, such as structural connectivity estimated from dMRI. Although, the relationship between structural and functional connectivity is an active research area [101, 256, 172], a comparison of parcellations derived from

different sources of connectivity is under-explored, and can provide an interesting future work towards locating cortical areas that are consistently assigned to the same regions across different parcellations/resolutions. In addition to comparing dMRI and rs-fMRI driven parcellations, combining them could also be an interesting future direction towards identifying the cortical organisation of the brain more accurately than by using just a single modality.

In the absence of a gold standard parcellation, we evaluated the accuracy of parcellations with respect to widely-accepted clustering validity techniques, such as reproducibility, parcel homogeneity, Silhouette analysis, and sum of absolute differences. While these measures are particularly important to show the robustness of the parcellations and their ability to fit the connectivity in the population, they do not directly provide a means of assessing the reliability of parcellations from a network analysis point of view. We also limited our experiments only to cover the spectral parcellation techniques, so as to demonstrate the advantage of the proposed joint decomposition framework over the two other alternative approaches. The following chapter provides a more comprehensive evaluation of the state of the art by incorporating other connectivity-driven, anatomical, and random brain parcellations from the literature. We also expand the scope of evaluation by including simple network analysis tasks to our experiments and further assess the agreement of the parcellations with complementary information, such as task fMRI and the myelo-/cyto-architecture of the cerebral cortex.

Chapter 7

A Systematic Comparison of Parcellation Methods

This chapter is based on:

S. Arslan, S. I. Ktena, A. Makropoulos, E. C. Robinson, D. Rueckert, and S. Parisot, *Human Brain Mapping: A Systematic Comparison of Parcellation Methods for the Human Cerebral Cortex*, NeuroImage, 2017. (*In Press*)

Abstract

This chapter provides a systematic comparison of connectivity-driven, anatomical, and random parcellation methods proposed in the thriving field of brain parcellation. Using functional MRI data and a plethora of quantitative evaluation techniques investigated in the literature, we evaluate 24 group-wise parcellation methods at different resolutions. We assess the accuracy of parcellations from four different aspects: (1) reproducibility across different groups, (2) fidelity to the underlying connectivity data, (3) agreement with fMRI task activation, myelin maps, and cytoarchitectural areas, and (4) network analysis. This extensive evaluation of different parcellations generated at the group level highlights the strengths and shortcomings of the various

methods and aims to provide a guideline for the choice of parcellation technique and resolution according to the task at hand. The results obtained in this study suggest that there is no optimal method able to address all the challenges faced in this endeavour simultaneously.

7.1 Introduction

In this chapter, we propose a systematic comparison of existing parcellation methods using publicly available resources and evaluation measures that are widely used in the literature through a structured experimental pipeline. We focus on rs-fMRI, as the majority of data-driven parcellation methods we are investigating have been developed and tested using this modality. We aim to provide some insight into the reliability of parcellations in terms of reflecting the underlying mechanisms of cognitive function, as well as, revealing their potential impact on network analysis.

There have been several recent studies that provide an overview of the current state of the parcellation literature [66, 244, 59]. In particular, Eickhoff et al. [66] review many design choices that may emerge in the process of parcellating the brain, including but not limited to definition of ROIs, choice of clustering method, selection of the number of clusters, and the evaluation methods, as well as discuss possible limitations and pitfalls of connectivity-driven approaches. Thirion et al. [244] specifically focus on the selection of a clustering algorithm for brain parcellation and provide a detailed analysis of three widely-used clustering algorithms with respect to their appropriateness for parcellation. de Reus and van den Heuvel [59] approach the problem from a different perspective and address the impact of parcellation model/resolution on network analysis. However, to the best of our knowledge, this is the first large-scale systematic comparison of the state-of-the-art parcellation methods that encompasses all these different aspects in a unified experimental setting.

The main contributions of our study are the following: (1) We evaluate 24 group-level methods using publicly available datasets provided by the Human Connectome Project [261]. (2) Our experiments consist of quantitative assessments of parcellations at both subject and group levels

and for different resolutions. (3) We evaluate parcellations not only from a data clustering point of view but also with regards to network analysis and multi-modal consistency. Our evaluation includes reproducibility (e.g. Dice coefficient and adjusted Rand index), cluster validity analysis (e.g. Silhouette coefficient and parcel homogeneity) and multi-modal comparisons with task fMRI activation, myelin and cytoarchitectural maps. In addition, we devise simple network-based tasks (such as gender classification and individual subject prediction) to evaluate the potential impact of parcellations on capturing representative features at different scales.

The remainder of this paper is organised as follows: Section 7.2 summarises the procedures pursued during the generation and evaluation of parcellations. Experimental results are presented in Section 7.3. In Section 7.4, we discuss the reliability and applicability of parcellations for network analysis and summarise the impact/limitations of this study.

7.2 Materials and Methods

A summary of the processing pipelines is given in Fig. 7.1. A brief description of the parcellation methods is provided in Tables 7.1 and 7.2, respectively.

7.2.1 Data

This study is carried out using data from the publicly available Human Connectome Project (HCP) database [261]. All connectivity-driven parcellations are derived from Dataset 1, which consists of 100 subjects (54 female, 46 male adults, aged 22-35). For evaluation purposes, we use Dataset 2 comprising randomly selected 50 male and 50 female adults of age 22-35. The evaluation is performed on Dataset 2 so as to reduce the possible bias towards parcellations computed from Dataset 1 with respect to the provided ones. The details of both datasets are previously given in Chapter 2 and briefly summarised here.

We use rs-fMRI as our primary data modality for the generation and evaluation of parcellations. This is because most methods selected for this study were developed for rs-fMRI driven parcel-

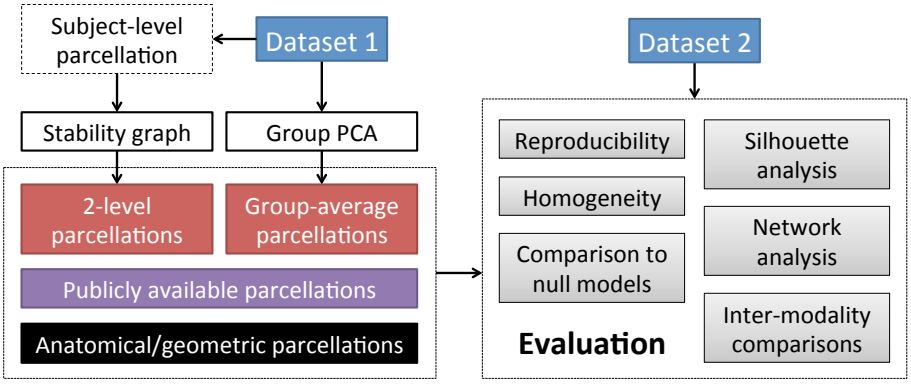


Figure 7.1: Visual outline of parcellation generation and evaluation steps.

lation, and rs-fMRI allows test-retest measurements across acquisitions, subjects, and groups. The rs-fMRI scans for each subject were conducted in two sessions, consisting of a total of four runs of approximately 15 minutes each. All subjects were preprocessed by the HCP structural and functional minimal preprocessing pipelines [87]. The output of these pipelines for each subject is a standard set of cortical vertices with inter-subject correspondence. Following these preprocessing steps, each timeseries was temporally normalised to zero-mean and unit-variance.

7.2.2 Parcellation Methods

In order to provide a comprehensive evaluation of the state of the art on surface-based brain parcellation at the group level, we gathered 24 parcellation methods from the literature. The methods included in this study satisfy at least one of the following criteria:

1. An implementation is publicly available.
2. Pre-computed parcellations are publicly available. Both surface-based and volumetric parcellations are considered.
3. The method can easily be re-implemented.

All the computed and publicly available methods considered in this study are respectively presented in Tables 7.1 and 7.2, along with their associated names, and briefly explained below.

Name	Reference	Resolution	Description
JOINT	Arslan et al. (2015a); codes available from www.doc.ic.ac.uk/~sa1013/codes.html	Varying	The parcellation method proposed in Chapter 6. An initial finer parcellation of 2000 supervertices per hemisphere is used for spatial feature reduction in order to compensate for the computational cost.
2-LEVEL	Group-level parcellations from subject-level *Arslan, Blumensath, Bellec, Ward, K-Means*, and *N-Cuts* parcellations (See Table 3.1 for details).	Varying	2-level approach is applied to the subject-level parcellations of various methods to obtain group-level parcellations. Parcellations computed in this setting are referred as the method's name followed by "2" (e.g. Ward-2).
Ward-AVR	Ward (1963); in-house implementation featuring scikit-learn's AgglomerativeClustering function.	Varying	The group average matrix is fed into the Ward's agglomerative hierarchical clustering algorithm using the same setting as the subject-level *Ward* parcellations derived in in Chapter 4. Only adjacent clusters are joined into a higher level in order to ensure the spatial contiguity.
K-Means-AVR	K-means clustering as described in Thirion et al. (2014); in-house implementation featuring scikit-learn's KMeans and PCA functions.	Varying	The group average matrix is fed into *k*-means clustering after being concatenated with the average spatial coordinates to improve spatial contiguity of parcellations.
N-Cuts-AVR	Craddock et al. (2012); in-house implementation of spectral decomposition featuring discretisation (Yu and Malik 2003).	Varying	A temporal correlation matrix is derived from the group average matrix and transformed into a spatially constrained affinity matrix. Spectral clustering with normalised cuts is used as in the same setting as the subject-level *N-Cuts* parcellations derived in in Chapter 4.
GRASP	Honnorat et al. (2015); codes available from cbica.upenn.edu/sbia/software/grasp/index.html	Varying	A method based on Markov Random Field (MRF) that can subdivide the cortex into spatially contiguous parcels by using shape priors. The group average matrix is parcellated into 10000 initial clusters by running the method in the hierarchical clustering mode. Final parcellations are derived from this low-dimensional matrix.
GRAMPA	Parisot et al. (2016); codes available from www.doc.ic.ac.uk/~sparisot	Varying	An MRF model that iteratively updates parcel centres and parcel assignments based on modality specific costs. The parcellation is computed using the group average matrix.
Geometric	Geometric parcellations as described in Thirion et al. (2014); in-house implementation featuring scikit-learn's KMeans function.	Varying	*K*-means clustering is applied to the average spatial vertex coordinates. No connectivity information is accounted for.

Table 7.1: Computed parcellation methods.

Computed Parcellations

The two more popular ways of computing a data-driven group-wise parcellation [54] are (1) performing parcellation for each subject individually and applying a second level clustering algorithm to subject-level parcellations (i.e. *2-level approach*), and (2) computing a representative feature matrix from the population, for instance by concatenating BOLD timeseries across subjects, and submitting this combined matrix to a parcellation method (i.e. *group-average approach*).

The *2-level* approach is explained in the previous chapter in more detail and summarised here. It is based on the idea of identifying shared features of the population as approximated by the individual parcellations. For this purpose, an adjacency matrix is constructed from subject-level parcellations, in which an edge between vertices v_i and v_j is weighted by the number of times both vertices are assigned to the same parcel across all individual subject parcellations.

This matrix is then subdivided into different number of regions using spectral clustering with normalised cuts [255, 54]. We use this approach to generate a group-level parcellation from the individual subject parcellations obtained by (1) the proposed hierarchical clustering method introduced in Chapter 4, which will be denoted as *Arslan* in the remainder of this chapter, and (2) other connectivity-driven subject-level methods such as *K-Means*, *Ward*, *N-Cuts*, *Blumensath*, and *Bellec*, all of which are previously explained in Chapter 4 and summarised in Table 4.1. The only difference is in the implementation of *K-Means*, in which we combined the PCA components with spatial vertex coordinates to improve spatial contiguity of parcellations.

The *group-average approach* aims to capture shared patterns between individuals within a population by computing a group average representation of connectivity. This may be achieved by concatenating the timeseries of each subject and applying PCA for dimensionality reduction before parcellation [244, 225]. However, using the full-concatenated timeseries with traditional PCA quickly becomes computationally prohibitive when the population's size increases. To overcome this, we follow the methodology employed by the HCP for the generation of group average matrices. We use FSL's incremental group PCA [225], a technique developed for computing 'pseudo timeseries' that can (to good approximation) estimate the real PCA output applied to the original combined dataset, while relying on a limited amount of memory. We apply this technique to generate group-level pseudo timeseries from both Dataset 1 and Dataset 2. Group-level parcellations are computed from each of these datasets using our in-house implementations of clustering techniques (*K-Means*, *Ward* and *N-Cuts*) as well as connectivity-driven methods for which implementations are available [102, 190].

Alternative to these two most commonly-used approaches, we also incorporate parcellations obtained by the joint spectral decomposition technique which is introduced in Chapter 6 and will be denoted as *JOINT* hereafter. The only difference in the parameter setting is the number of subjects used to construct the multi-layer graph (50 vs. 100 subjects).

Finally to obtain a baseline for the connectivity-driven methods, 'geometric parcellations' are derived using k-means clustering of the average spatial coordinates of all cortical vertices as described in [244] and Chapter 4.

Name	Reference	Resolution	Description
Gordon	Gordon et al. (2016); parcellation available from www.nil.wustl.edu/labs/petersen/Resources.html	333 (161 L, 172 R)	A surface-based parcellation computed from the average gradients of resting-state functional connectivity networks. Provided parcellation is iteratively dilated to cover the entire cortical surface.
Power	Power et al. (2011), parcellation available from balsa.wustl.edu/study/show/WG33	130 (65 L, 65 R)	Resting-state communities originally identified in volume space are projected onto the cortical surface and made publicly available by Van Essen et al. (2016). Connected components within each parcel are relabeled to ensure spatial contiguity.
Yeo	Yeo et al. (2011); parcellation available from balsa.wustl.edu/study/show/WG33	96 (49 L, 47 R)	17-cluster resting-state networks originally derived in volume space from average resting-state functional connectivity data using a GMM-based clustering algorithm are projected onto the cortical surface and made publicly available by Van Essen et al. (2016). Connected components in each parcel are relabeled to ensure spatial contiguity.
ICA	Group-ICA parcellations available from db.humanconnectome.org/data/projects/HCP_500	Varying	Group-average parcellations by means of group-ICA (Beckmann et al. 2004) are obtained at several different dimensionalities (25, 50, 100, 200, 300), using a group-PCA output (Smith et al. 2014) from the HCP S500 subjects. Connected components within each parcel are relabeled.
Baldassano	Baldassano et al. (2015); parcellation available from www.princeton.edu/~chrisb/code.html	171 (84 L, 87 R)	A multi-purpose clustering algorithm based on nonparametric Bayesian modeling is applied to dense connectome derived from the HCP S500 group PCA output (Smith et al. 2014) in order to compute a surface-based parcellation.
Glasser	Glasser et al. (2016); parcellation available from balsa.wustl.edu/study/show/RVVG	360 (180 L, 180 R)	A cortical parcellation generated from multi-modal images of 210 adults from the HCP, using a semi-automated approach. Cortical regions are delineated with respect to function, connectivity, cortical architecture, and topography, as well as, expert knowledge and meta-analysis results from the literature.
Fan	Fan et al. (2016); parcellation available from atlas.brainnetome.org	210 (105 L, 105 R)	A volumetric brain parcellation is obtained using both anatomical landmarks and connectivity-driven information. Anatomical regions defined by Desikan et al. (2006) are parcellated into subregions using functional and structural connectivity data from 40 adults provided by the HCP. Cortical parcels are projected from volume to surface.
Shen	Shen et al. (2013); parcellation available from www.nitrc.org/frs/?group_id=51	200 (102 L, 98 R)	A spectral clustering approach is used to compute a volumetric groupwise parcellation based on an optimization process that guarantees functional homogeneity within each parcel and ensures that computed parcels are consistent across subjects. Volumetric parcels from the provided 1 mm sampled 268-parcel atlas are projected to cortical surface.
AAL	Tzourio-Mazoyer et al. (2002); available from www.gin.cnrs.fr/AAL2_files/aal2_for_SPM12.tar.gz	82 (41 L, 41 R)	A popular volumetric atlas that is manually delineated with respect to anatomical landmarks, in particular, by following the sulci course in the brain. Cortical parcels are projected from volume to surface.
Destrieux	Fischl et al. (2004); parcellations available from db.humanconnectome.org	150 (75 L, 75 R)	Majority voting used across subject-level *Destrieux* parcellations to obtain a group-level parcellation.
Desikan	Desikan et al. (2006); parcellations available from db.humanconnectome.org	70 (35 L, 35 R)	Majority voting used across subject-level *Desikan* parcellations to obtain a group-level parcellation.

Table 7.2: Pre-computed, publicly available parcellation methods.

Publicly available parcellations

Several pre-computed, publicly available group-level parcellations are included in this study [93, 245, 202, 15, 68, 218, 225, 86]. Details on the method and the resolution of the parcellations are provided in Table 7.2. In particular, it should be noted that the parcellations provided by [15] (*Baldassano*) and the *ICA* parcellations [21, 225] are computed from a much larger HCP cohort (group average of 500 subjects) which can comprise our evaluation dataset. This may introduce a bias in the evaluation of both methods.

The methods proposed by [245] (*Yeo*) and [202] (*Power*), as well as the *ICA* parcellations [21, 225] were originally developed for identifying communities or resting-state networks, hence do not naturally provide spatially contiguous parcellations. Since this can affect the evaluation measures, we overcome this by relabelling connected components within each parcel. We then remove very small parcels and slightly dilate the remaining ones to adjust for vertices lost. *K-Means* (both 2-level, and group-average versions) and another connectivity-driven approach, *GRAMPA*, can also provide spatially disjoint parcels. In our experiments, we do not apply any post-processing to the parcellations derived by these methods, as we aim to obtain roughly the same number of regions for all computed parcellations for the sake of consistency. Nonetheless, we perform additional experiments to analyse the impact of relabelling connected components for these methods and discuss how their performance changes compared to the original parcellations.

The multi-modal parcellation of the human cerebral cortex [86] (*Glasser*) is computed through expert manual annotation of imaging data from several modalities, including function, connectivity and cortical architecture.

We also incorporate anatomical atlases in our study, including the volumetric Automated Anatomical Labelling (*AAL*) brain atlas [254] and two surface-based parcellations provided by the HCP [74, 61], namely *Destrieux* and *Desikan*. We obtain a group-wise representation of the *Destrieux* and *Desikan* parcellations using majority voting across their subject-level versions.

Several parcellations are only available in volume space [254, 218, 68]. We use volume-to-surface and surface-to-surface sampling techniques to project volumetric parcels onto the HCP average cortical atlas (Conte69) [259]. *AAL* [254] and the volumetric parcellation by [218] (*Shen*) are projected onto the cortical surface generated from the Colin27 brain [100] using FreeSurfer [72], which is then registered to the Conte69 standard space using multimodal surface matching [205]. Our last volumetric parcellation, *Fan* [68], is provided in the HCP volumetric space, and is therefore directly projected onto the HCP's standard surface. Finally, all volumetric parcellations are post-processed and each parcel is slightly dilated to fill holes that may have emerged dur-

Evaluation technique	Description	Quantitative measurements	Previously used in literature
Reproducibility	Assesses the similarity between two sets of parcellations obtained from different groups (group-to-group reproducibility).	Dice coefficient; Adjusted Rand index	Craddock et al. (2012); Blumensath et al. (2013); Shen et al. (2013); Thirion et al. (2014); Honnorat et al. (2015); Arslan et al. (2015a); Arslan et al. (2015b); Parisot et al. (2016a)
Cluster validity analysis	Evaluates the quality of parcellations from a clustering point-of-view by measuring the faithfulness of the parcellation to the underlying data source.	Homogeneity; Homogeneity relative to null models; Silhouette coefficient	Yeo et al. (2011); Craddock et al. (2012); Arslan et al. (2015a); Arslan et al. (2015b); Parisot et al. (2016); Gordon et al. (2016); Arslan et al. (2016)
Agreement with cytoarchitecture	Assesses the overlap with known cyto-architectonic areas as delineated by the Brodmann atlas.	Dice coefficient	Blumensath et al. (2013); Arslan et al. (2016); Parisot et al. (2016)
Goodness-of-fit to task activation	Evaluates how well the parcellations agree with the task activation maps.	Bayesian information criterion	Thirion et al. (2014); Parisot et al. (2016);
Alignment with myelination	Assesses the agreement between the parcellations and highly myelinated cortical areas, identified by a coarse myelin-driven parcellation.	Dice coefficient	Parisot et al. (2017)
Individual subject prediction	Evaluates the impact of parcellations on a simple individual identification task with respect to connectivity profiles.	Prediction rate	Finn et al. (2015)
Network-based classification	Evaluates the ability of parcellations to capture population differences with a simple gender classification task on functional connectivity networks.	Classification accuracy	Vergun et al. (2013); Satterthwaite et al. (2014)

Table 7.3: Parcellation evaluation techniques.

ing projection. Unfortunately, volume-to-surface resampling is not a straightforward process, and hence, it is impossible to retain all volume-based parcels after projection. However, we ensure that the parcellation boundaries and relative positions of parcels to each other remain as faithful to the original atlas as possible.

7.2.3 Parcellation Evaluation Techniques

We gather here the most commonly used evaluation techniques from the literature to evaluate parcellations with respect to varying resolutions. These techniques can be separated into four categories with regards to the parcellation aspects they assess: (1) reproducibility, (2) clustering validity measures, e.g. homogeneity and Silhouette analysis, (3) multi-modal comparisons with cytoarchitecture, task fMRI activation, and myelination, and (4) network analysis. A summary of the evaluation techniques is given in Table 7.3.

Reproducibility

We use reproducibility techniques to measure the degree of similarity between two parcellations obtained from different groups. A robust parcellation method should yield very similar parcellations for different groups of subjects, assuming that the groups comprise a relatively large number of subjects with similar characteristics (e.g. healthy adults from same age range). To this end, we compare the parcellations obtained from Dataset 1 with the ones derived from Dataset 2 (i.e. group-to-group reproducibility). Unfortunately, we are limited to performing the reproducibility analysis only for the computed parcellations, as only one parcellation/atlas is publicly available from each external source.

For measuring reproducibility on a group-to-group basis we adapt the two widely-used measures, Dice coefficient [62] and adjusted Rand index (ARI) [103], both of which have been previously used for reproducibility analysis in Chapters 4 and 6. Dice coefficient relies on a parcel-to-parcel comparison by measuring the overlap between pre-matched regions while ARI provides a more direct evaluation of parcellation configurations and can be more reliable when parcellations with different resolutions are compared [163]. Similarly to the subject-level analysis in Chapter 4, we also use the modified version of Dice coefficient [32], that runs a merging stage to account for the over-parcellated regions from one parcellation to the next. All measures provide reproducibility scores within a range of $[0, 1]$, where 1 implies a perfect match (identical parcellations) and 0 implies that the parcellations do not agree on any of the labels. Please refer to Chapter 4 for a more detailed explanation of each technique.

Cluster Validity Measures

This second category of validation measures aims to evaluate parcellations from a clustering point of view. Among many tools targeted at evaluating clustering solutions, we focus on two measures that are extensively being used for brain parcellations, namely homogeneity and Silhouette coefficients. In addition, we adopt the evaluation technique proposed in [93] that compares parcellations to a set of 'null models' obtained by randomly relabelling the parcellation

without altering the relative parcel locations with respect to each other.

***Homogeneity*:** This technique measures the degree of similarity between vertices that are assigned to the same parcel and is extensively used in the literature for the evaluation of parcellations [54, 218, 93, 13, 188, 102]. Parcellating the cerebral cortex into highly homogeneous regions might be particularly important for network analysis, as the network nodes are typically represented by the average signal (e.g. BOLD timeseries) within each parcel [218, 93]. In this study, we utilise the same homogeneity technique presented in Chapter 4 for the subject-level parcellations. That is, we first compute the similarity between every pair of vertices (represented by their connectivity fingerprints) within a parcel using Pearson's correlation coefficient. The average of pair-wise correlations are then used to measure the homogeneity of each parcel. A global homogeneity value is obtained by averaging the homogeneity scores across all parcels. Different from the subject-level analysis in Chapter 4, we use a weighted arithmetic mean, where each parcel's contribution to the global homogeneity is proportional to its size. This is to avoid possible unfairness that can emerge from the parcel size distributions of the tested parcellations. Smaller parcels tend to have a higher homogeneity than large ones, such that, a parcellation mostly composed of many small parcels and a few large regions may likely perform better than one with similarly sized parcels.

We obtain homogeneity results from the average connectivity fingerprints of all subjects, as well as on a per subject basis by using each subject's connectivity fingerprints and then averaging across subjects.

***Comparison to null models*:** While computing homogeneity by means of a weighted mean reduces the bias towards small parcels, homogeneity values remain dependent on the resolution of parcellations so that fair comparison between different resolutions is not possible. An alternative is proposed in [93] which consists of comparing a parcellation with the so-called 'null models' of the same resolution.

In order to obtain such null models, we perform the procedure illustrated in Fig. 7.2. For each

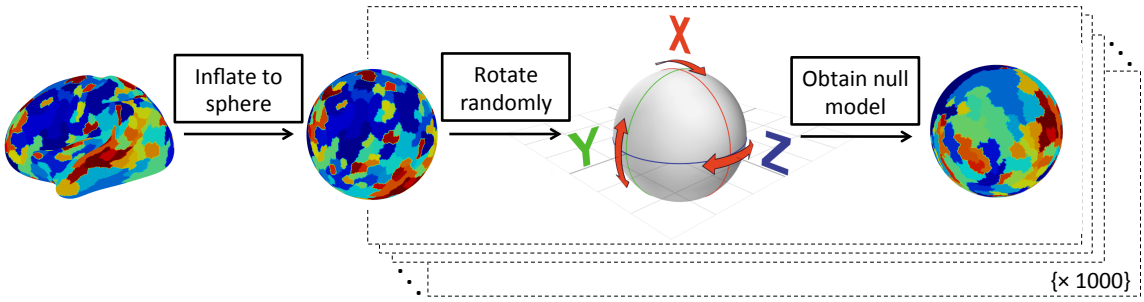


Figure 7.2: Illustration of the generation of null models.

hemisphere, we project the parcellation onto a standard spherical surface provided by the HCP and randomly rotate each point in this sphere around the x, y, and z axes. This process moves each parcel to a new location on the cortical surface without altering their relative positions. We then measure the homogeneity of the rotated parcellation and repeat the same process for 1000 different null models. Parcels that move to the medial wall, where no connectivity information is available, are discarded from computations. The advantage of this approach is that it reduces the observed biases with respect to parcel shape and size, as the parcellations are compared to their rotated versions, which have the same resolutions and similar parcel shapes.

In order to quantitatively evaluate parcellations with respect to their null models, we (1) count the number of rotated parcellations with lower homogeneity scores than the original parcellation and (2) compute the difference between the homogeneity of the original parcellation and the mean homogeneity score of null models, scaled by their standard deviation (i.e. z scores relative to null models) [93].

***Silhouette analysis*:** Another popular clustering validity tool that can be used to quantify parcellation reliability is the Silhouette coefficient (SC) [208]. It is typically used as an indicator of how well vertices fit in their assigned cluster [245, 54]. This is achieved by measuring the similarity of a vertex to the other vertices within the same cluster, with respect to its dissimilarity to the out-of-cluster vertices. As a result, SC not only evaluates the compactness (homogeneity) of a parcellation, but also its ability to separate regions with distinct features. For our group-level Silhouette analysis, we rely on the same SC definition as given in Chapter 4.

A global score is obtained for each parcellation by averaging the Silhouette coefficients across all vertices. SC can range within $[-1, +1]$, in which a value close to 1 indicates that the vertex is clustered with a high degree of confidence.

Similarly to the homogeneity analysis, Silhouette coefficients are obtained from the average connectivity fingerprints of all subjects. However, we also perform the same analysis on a per subject basis by using each subject's connectivity fingerprints and then averaging across subjects.

Comparisons with Other Modalities

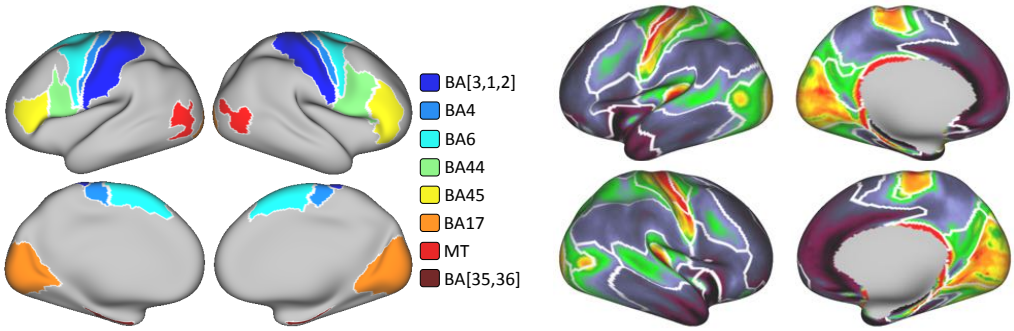
The previously proposed measurements assess the accuracy of parcellations from a clustering point of view. However, when defining regions of interest for neuro-anatomical purposes, the consistency of these areas with well-defined neuro-biological features also constitutes a critical aspect of evaluation. To this end, we expand our comparisons to those with other modalities. We test the parcellation quality by evaluating their agreement with task activation maps and myelination patterns, and their overlap with well-known cortical regions delineated from cytoarchitectonic features.

It should be noted that *Glasser* is the only parcellation that is simultaneously driven by task fMRI, myelination, and neuro-anatomist input, and as a result, will develop an inevitable positive bias towards these modalities. Therefore, the performance of this approach with respect to the inter-modality comparisons should be interpreted by taking this into account.

***Goodness-of-fit to task activation*:** Bayesian information criterion (BIC) is proposed in [244] as a means of quantifying the agreement of parcellations with task activation. Each vertex is associated with a task activation map (or the concatenated task activation maps of all subjects as in our case, since we consider group-wise parcellations). The BIC criterion measures the goodness-of-fit of a probabilistic model of the concatenated task activation maps by penalising the negative log likelihood by the complexity of the model, determined by the number of parcels.

***Overlap with cytoarchitectonic areas*:** We measure the agreement of our parcellations with the Brodmann areas [38]. Although functional connectivity obtained from BOLD time-series does not necessarily reflect the cytoarchitecture of the cerebral cortex [280], agreement (to a certain extent) with some known cytoarchitectonic areas could indicate a parcellation's ability to reflect the underlying cortical segregation [93]. Our standpoint to include comparisons with the cytoarchitecture is to show the extent of such agreement with at least certain areas, such as the motor and visual cortex, for which several parcellation techniques report a noticeable alignment [32, 280, 93]. To this end, we use the Brodmann atlas provided by the HCP, which contain labels for the primary somatosensory cortex (BA 3, 1, and 2), the primary motor cortex (BA 4), the premotor cortex (BA 6), Broca's area (BA 44, 45), the visual cortex (BA 17 and MT), and the perirhinal cortex (BA 35, 36) as shown in Fig. 7.3(a).

Quantitative comparisons are performed with the joined Dice coefficient approach explained in Chapter 5. In brief, overlapping parcels are matched with the Brodmann areas if their overlap ratio is ≥ 0.5. We then compute the Dice coefficient between the matching pairs. It is important to note that several parcels can be matched to the same area and therefore merged into a larger parcel.



(a) (b)

Figure 7.3: Cyto- and myelo-architecture of the cerebral cortex as defined respectively by (a) Brodmann's areas and (b) a coarse-resolution myelin-driven parcellation.

***Agreement with structured myelination patterns*:** Strong similarities have been observed between myelin maps and resting-state fMRI gradients [88]. We should therefore expect the boundaries of rs-fMRI driven parcellations to align with myelination patterns. To evaluate

this, we compute a coarse-resolution myelin-driven parcellation (25 parcels) using the method described in [190] from the average myelin map of all subjects. This method simply regroups vertices with similar myelin values and, as shown in Fig. 7.3(b), effectively delineates the major changes in myelination across the cortex.

We compare the parcellations obtained by different methods to these coarse parcellations using the joined Dice coefficient approach described above. To this end, we only consider the highly myelinated cortical areas, i.e., cortical areas with a mean myelin value below a certain threshold are discarded.

Network Analysis

Parcellations can significantly reduce the dimensionality of the dense human connectome without eliminating valuable information about the interactions between different brain regions and the mechanisms through which these interconnections give rise to complex cognitive processes. It has been common practice in recent neuroscience studies to explore several neurological [247, 77] and neuro-developmental disorders [107, 145, 60, 78] from a network perspective. These disorders have often been linked to a disruption or abnormal integration of spatially distributed brain regions that would normally be part of a single large-scale network, leading to their characterisation as disconnection syndromes [45]. Apart from the clinical value of network analysis, efforts to explore potential correlations between connectivity patterns and certain phenotypes like fluid intelligence [221], or to predict an individual's biological age [203, 206, 185] have been made. Therefore, a parcellation method can also be evaluated in terms of its ability to capture the inter-individual variability and to reveal patterns that explain observed cognitive performance.

Once the parcellation has been generated, a network representation can easily be obtained by mapping each network node to a parcel. The edge weights in functional networks usually represent the statistical dependency between the brain regions underlying the connected nodes. In our analysis of functional networks, we use temporal correlation of the representative timeseries as an estimate of the connection strength between two brain parcels.

We explore different ways in which the underlying parcellation can affect network analysis through (1) a network-based classification task and (2) individual subject identification from connectivity profiles.

***Network-based classification*:** Several studies suggest differences in both structural and functional connectivity between genders [91]. More specifically, in terms of functional connectivity derived from rs-fMRI data, which is the focus of the current survey, significant differences in the topological organisation of functional networks have been found between males and females [246]. For this reason we choose a gender classification task to evaluate the impact of the parcellation on network-based classification tasks. We employ linear Support Vector Machines (SVMs) [43], a well-established classifier from the machine learning literature and a 10-fold cross-validation procedure to get an estimate of each method's performance. Previous studies [206, 265, 213] have used SVM as a machine learning classifier, which is designed to make predictions based on high-dimensional data, to investigate sex differences in functional connectivity.

Given a set of p-dimensional feature vectors, SVM aims to identify a $(p-1)$-dimensional hyperplane that represents the largest separation or margin between the feature vectors of the two classes. The hyperplane is chosen in a way that the distance from the nearest data points of each class is maximised. The weights assigned to the normalised features to obtain a low-dimensional representation of the feature vectors can, additionally, be used to describe how heavily weighted the connectivity feature is within the multivariate model [213].

Since node correspondences are ensured with group-wise parcellations, an embedding of each subject's connectivity matrix can be employed to get a general vector representation [262], rendering the use of the aforementioned classifier straightforward. This approach is often referred to as 'bag of edges' [55] and has been widely used when the underlying parcellation is the same among all subjects. Connectivity networks are generated using the same set of nodes for all subjects in Dataset 2. The nodes correspond to non-overlapping regions specified by the provided or data-driven parcellations obtained from Dataset 1. In order to explore the

performance of different parcellation methods in capturing population differences, we show the results of a gender prediction task with r-to-z transformed full correlation networks. Before the classification step, dimensionality reduction through PCA is performed [198] and the components explaining 100% of the variance [204] in the training data are preserved for both training and testing.

***Individual subject identification*:** Most of the human neuroimaging studies typically focus on identifying patterns of brain connectivity that are shared across individuals. However, it has been recognised that brain structure [42] and function [168] may exhibit a great amount of individual variability. Despite many shared characteristics discovered across individuals, evidence suggests that the human connectome possesses connectional traits that are unique to each subject [168, 16]. A recent study by Finn et al. [71] has shown that connectivity networks obtained from rs-fMRI can be used as 'fingerprints' to identify individuals with a high accuracy. Results suggest that the intrinsic connectivity captured during rest (or task) is (1) robust and reliable across different acquisitions, and (2) distinct enough to successfully distinguish one individual from another. Finn et al. [71] explored several factors that may affect the identification accuracy, including the impact of the parcellation used to obtain the connectivity networks, however, this has been covered rather limitedly. Here, we adapt these experiments to the focus of the current study and explore the extent to which the underlying parcellation and its resolution affect a similar subject identification task.

For each subject, two r-to-z connectivity networks are generated from two different acquisitions (rs-fMRI-1 and rs-fMRI-2) using one of the group-level parcellations and Pearson's correlation to estimate the connection strength of the edges between each pair of nodes. rs-fMRI-1 networks are denoted as the 'target set', whereas networks obtained from the rs-fMRI-2 timeseries are referred to as the 'database'. The identification is performed on a per subject basis as illustrated in Fig. 7.4. One matrix is selected from the target set and compared with all the other matrices in the database to find the matrix with the maximum similarity. If the target and the predicted matrix belong to the same subject, that iteration is assigned a score of 1, and 0 otherwise (i.e. actual and predicted matrices correspond to different subjects). The similarity is measured after

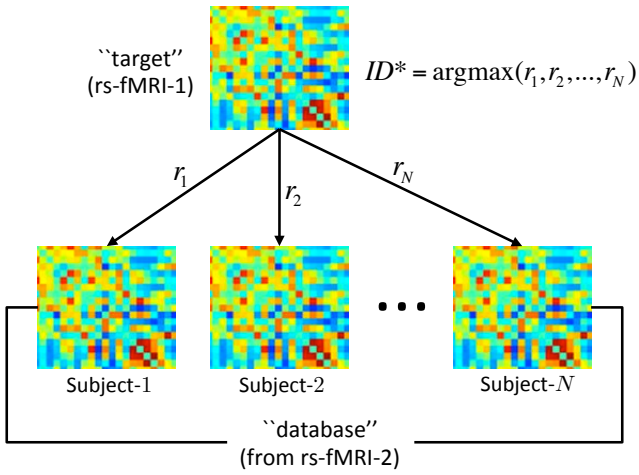


Figure 7.4: Individual subject identification procedure. Given a 'target' matrix, we compute the correlations between this matrix and all the 'database' connectivity matrices. The predicted identity $ID*$ corresponds to the subject with the highest correlation coefficient (r).

embedding all edges in a one-dimensional vector [262]. At the end of this process, the prediction (success) rate is calculated as the ratio of correctly identified subjects to the population size. The identification task is repeated for all group-level parcellations and available resolutions to assess their impact on the success rate.

7.3 Results

7.3.1 Experimental Setup

All computed parcellations are generated from Dataset 1, in which each subject is represented by two 30-min rs-fMRI scans (rs-fMRI 1 and rs-fMRI 2) that were obtained by concatenating the timeseries of two 15-min scans acquired on the same day [87]. The 2-level parcellations are generated from the individual subject parcellations obtained from the rs-fMRI 1 set. This set is also submitted to MIGP to obtain the group-PCA matrix, which is subsequently used to compute parcellations using the group-average approach. rs-fMRI 2 set is only used for the individual subject identification task and are not involved in the generation or evaluation of parcellations otherwise.

Dataset 2 is primarily used to evaluate the group-wise parcellations (publicly available and computed ones from Dataset 1). A second set of group-level parcellations is also generated using Dataset 2 in order to assess reproducibility across different groups. It is worth noting that, this second set is solely used to assess group-to-group reproducibility and excluded from any other stage of the analysis pipelines.

Most of the pre-computed/publicly available parcellations comprise a fixed number of regions, while the methods for which an implementation is available can be explored at different resolutions, allowing us to assess the sensitivity of quantitative measures with respect to the number of parcels. For these methods, we generate parcellations containing between 50 and 500 regions (i.e. 25 to 250 per hemisphere), in increments of 50.

Finally, results are reported using the following naming scheme: parcellations obtained using the 2-level approach will be referred to via their associated method name followed by '2' (e.g. *Ward-2*), whereas parcellations derived from the group-average approach will be accompanied by the 'AVR' suffix (e.g. *Ward-AVR*).

7.3.2 Reproducibility Results

The reproducibility values (Fig. 7.5) are only reported for methods that allow the derivation of multiple parcellations. As expected, the spectral techniques have the best reproducibility, with *N-Cuts* and *JOINT* leading the others. In general, more favourable results are achieved by 2-level parcellations. This may be attributed to the fact that these parcellations are obtained from a set of individual parcellations that already provide a means of spatial smoothing. Furthermore, the parcellations are computed using normalised cuts, a technique known to increase the reproducibility of parcellations [54, 32]. Among the parcellations derived from the average matrix, *Ward-AVR* shows the least favourable performance. MRF-based methods (i.e. *GRASP* and *GRAMPA*) and *K-Means-AVR* also have a relatively poor performance. While joining over-parcellated regions generally increases reproducibility for the group-average approaches, it has a lesser impact on the 2-level parcellations as most of them only show a marginal improvement.

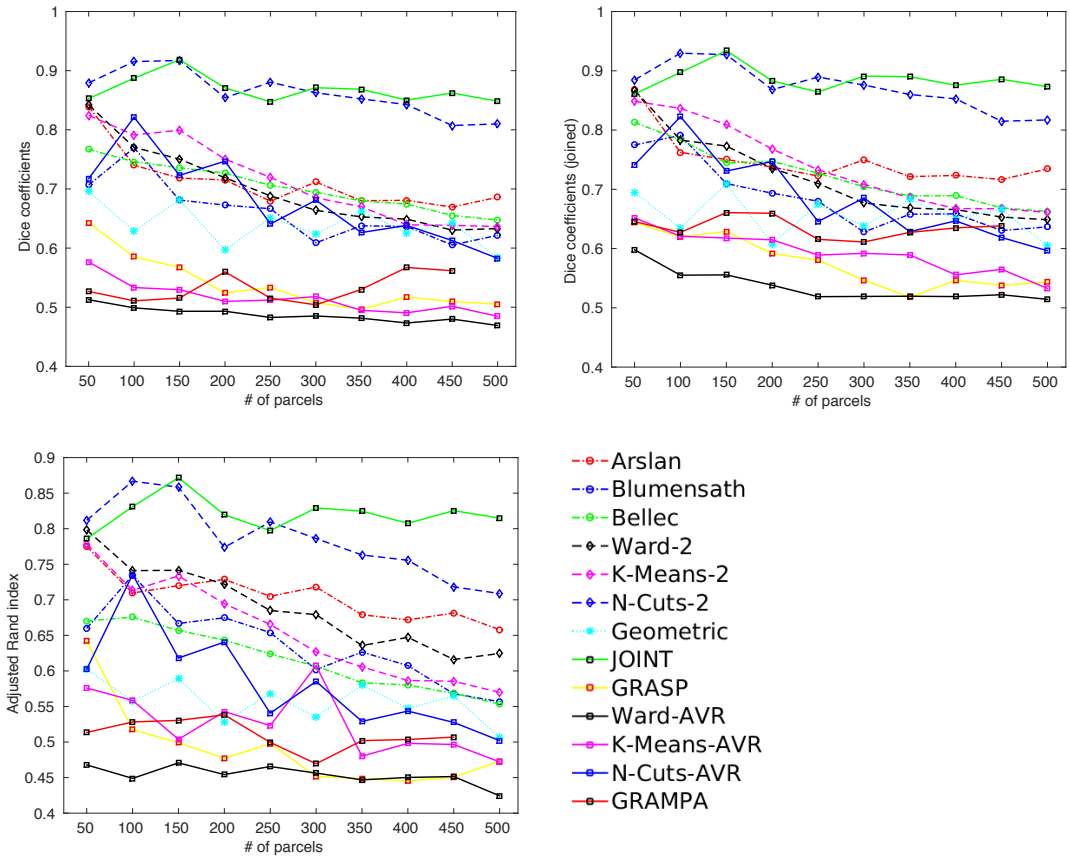
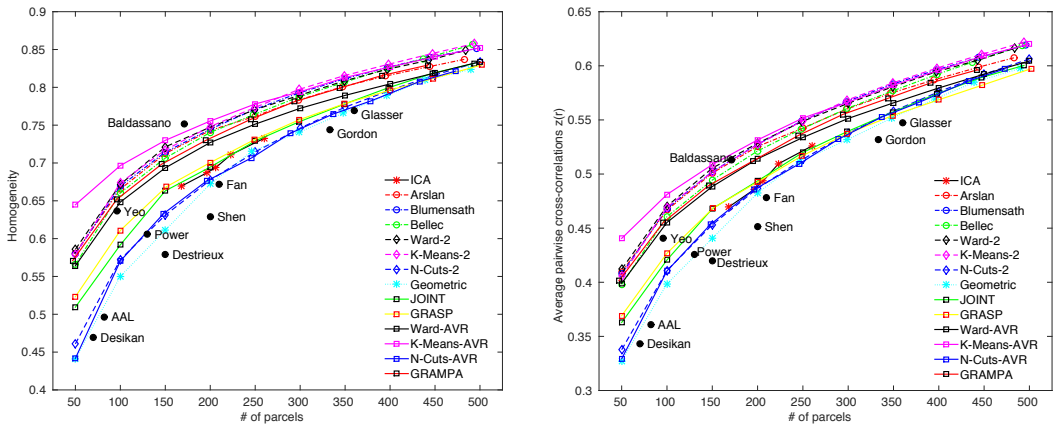


Figure 7.5: Group-level reproducibility results. Reproducibility values for each method are obtained using Dice coefficient (top, left), Dice coefficient after joining over-parcellated regions (top, right), and adjusted Rand index (bottom).

7.3.3 Cluster Validity Results

Clustering validity results in terms of parcellation homogeneity are summarised in three figures. First of all, homogeneity values obtained by each method/resolution are given in Fig. 7.6. The homogeneity of each method for a set of selected resolutions together with the homogeneity of their respective null parcellations are presented in Fig. 7.7. The difference between the homogeneity of the computed parcellations and the homogeneity distribution of null models measured as z scores is shown in Fig. 7.8.

Homogeneity results in Fig. 7.6 show a relatively poor performance for most of the provided parcellations. The methods that generate the most reproducible parcellations (e.g. spectral methods *JOINT*, *N-Cuts-2*, and *N-Cuts-AVR*) as well as *Geometric* also obtain poor homogene-



(a) (b)

Figure 7.6: Group-level homogeneity results obtained (a) from the average connectivity fingerprints of all subjects and (b) on a per subject basis by using each subject's connectivity fingerprints and then averaging across subjects. Whereas lines show the homogeneity values for all computed resolutions, black dots correspond to the homogeneity scores obtained from the publicly available parcellations with fixed resolutions.

ity values. In general, other connectivity-driven computed parcellations tend to generate highly homogeneous parcellations with the group-average and 2-level methods obtaining very similar results. Among them, *K-Means-AVR* especially excels at lower resolutions, but is outperformed by *Baldassano*, one of the publicly available parcellations based on functional connectivity when similar resolutions are considered. It should be noted, though, that *Baldassano* is obtained from a larger HCP cohort (500 subjects) which may contain our evaluation set and positively bias homogeneity results.

As shown in Fig. 7.7 and 7.8, we observe similar performance trends for most of the computed parcellations when compared to null models. Anatomical parcellations (*AAL*, *Destrieux*, and *Desikan*), and some of the provided parcellations (*Fan*, *Gordon*, and *Shen*), regardless of their respective resolutions perform similar to or worse than their null models. Among the publicly available parcellations, *Baldassano* is on par with *K-Means-AVR*, while *Yeo*, *Power*, and *ICA* also yield good results.

Group-level Silhouette coefficients (Fig. 7.9) mostly follow the tendency observed in homogeneity. *K-Means-AVR* outperforms the other approaches at all resolutions. It is followed by

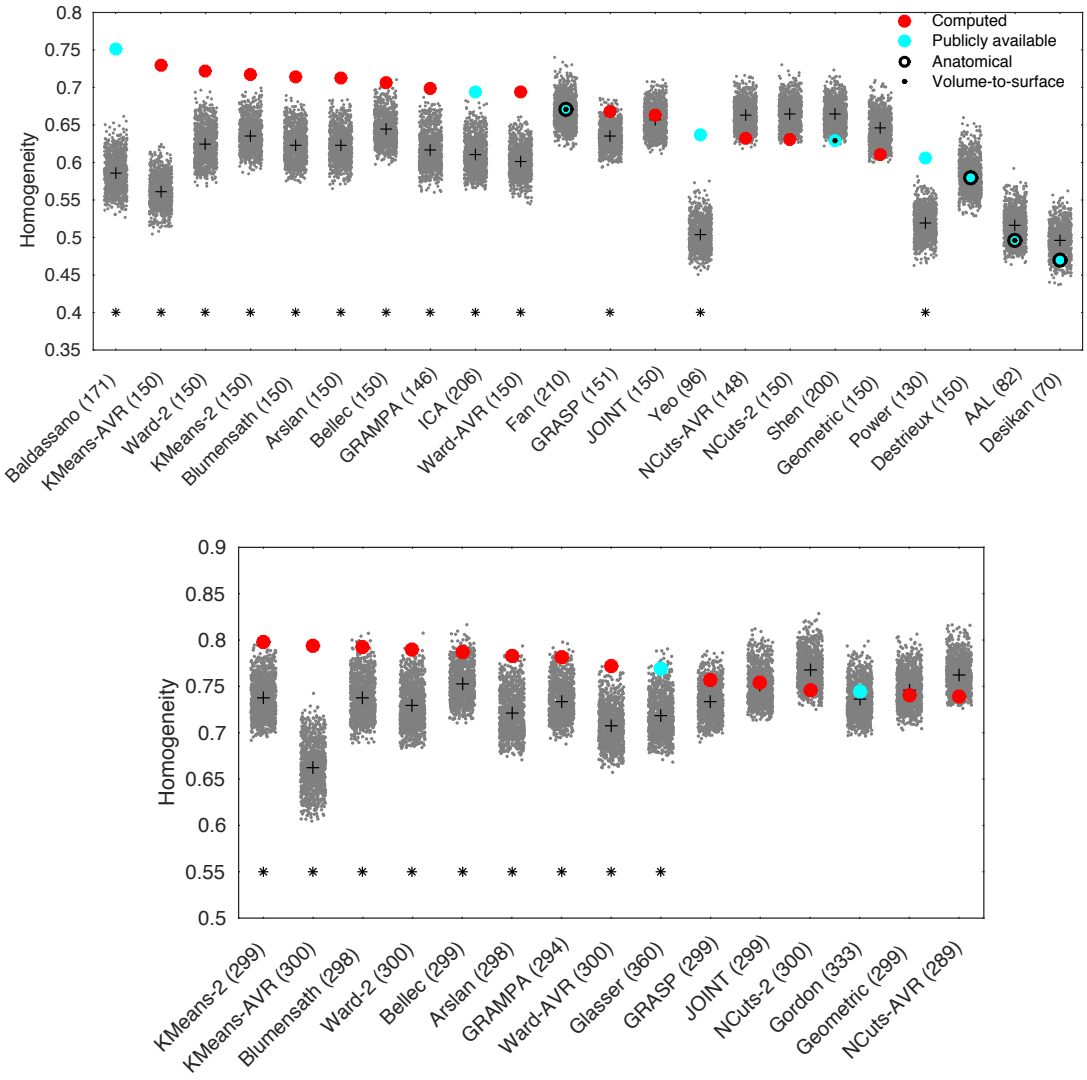


Figure 7.7: Homogeneity of each parcellation (red and purple dots) and their respective 1000 null models (gray dots). Null models yield different homogeneity scores due to variation across parcel size and location. Cross (+) indicates the average homogeneity obtained by each set of null parcellations. Asterix (∗) indicates that the computed homogeneity is higher than at least 950 of its null parcellations (i.e. $p < 0.05$). *Top:* Results of publicly available parcellations with relatively low resolutions (comprising around or fewer than 200 regions) and the computed parcellations with 150 parcels. *Bottom:* Results of publicly available parcellations with higher resolutions (e.g. comprising around or greater than 300 regions), with the computed parcellations having a fixed resolution of 300 parcels. The exact number of parcels for each method is indicated aside the method name in parentheses.

another group-average technique, *GRAMPA*, which shows a good performance at low levels of granularity. All 2-level approaches, apart from *N-Cuts-2*, perform equivalently well and produce more distinct parcels than most of the group-average methods. In contrast to the homogeneity

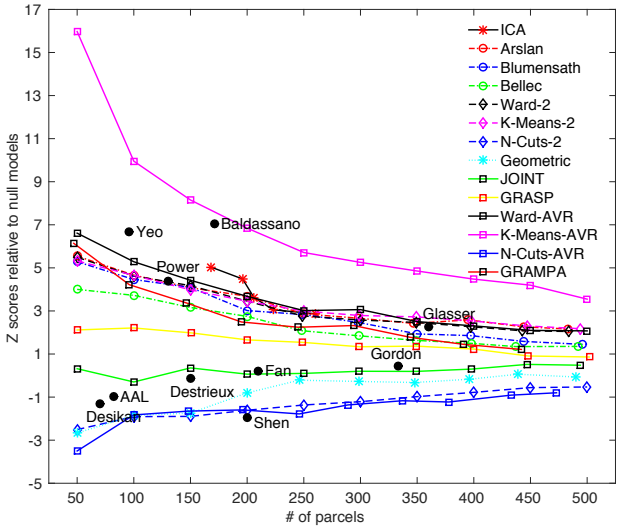
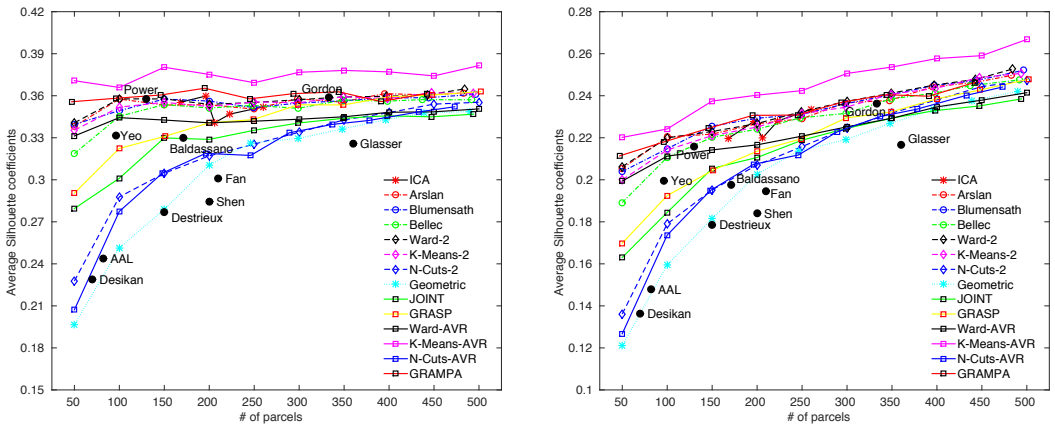


Figure 7.8: Difference between the actual homogeneity and the homogeneity distribution of null models. Lines show the z scores relative to null models for all computed resolutions, while black dots correspond to the z scores obtained from the publicly available parcellations with fixed resolutions.

results, *Gordon* and *Power* are the top-performing provided parcellations. Interestingly, despite producing homogeneous parcellations, *Baldassano*, *Yeo*, and *ICA* show an average performance in terms of Silhouette coefficients. This shows that generating homogeneous parcellations does not necessarily guarantee a good separation between parcels. Overall, spectral techniques perform poorly but still surpass the anatomical and geometric parcellations.

Impact of relabelling connected components in disjoint parcellations. Considering all computed parcellation methods, two k-means variants, *K-Means-AVR*, and *K-Means-2* as well as *GRAMPA* can generate spatially disjoint parcellations. In particular, *K-Means-AVR* yields many discontinuous parcels, which significantly increases the total number of parcels after the relabelling process as shown in 7.10. This unsurprisingly yields more homogeneous regions, as homogeneity depends on the resolution and likely to increase when the cortex is parcellated into more subregions (i.e. homogeneous regions still stay homogeneous when subdivided). On the other hand, as we alter the clustering configuration unnaturally by forcing parcels to split, fidelity to the underlying data is negatively affected, yielding lower Silhouette coefficients and z



(a) (b)

Figure 7.9: Group-level Silhouette analysis results obtained (a) from the average connectivity fingerprints of all subjects and (b) on a per subject basis by using each subject's connectivity fingerprints and then averaging across subjects. Lines show the Silhouette coefficients (SC) for all computed resolutions, while black dots correspond to the SC obtained from the publicly available parcellations with fixed resolutions.

scores relative to null models. However, this change in the spatial structure of the parcellations in general appears to yield a positive impact on reproducibility, as indicated by the joined Dice coefficients and adjusted Rand indices. The other two methods, *GRAMPA* and *K-Means-2*, only produce few parcels that are discontinuous, thus relabelling does not lead to a significant change in their evaluation measures.

7.3.4 Multi-Modal Comparisons

The agreement with concatenated single-subject task activation maps is reported in Fig. 7.11. In general, all provided parcellations yield relatively poor BIC values compared to the computed parcellations with similar resolutions. The 2-level approaches tend to yield better results than their group-average (AVR) counterparts, in particular at higher resolutions, with *K-Means-2* showing the best performance for most resolutions. This could be linked to the fact that the parcellations are derived from the subject level, where the individual task activation is also estimated from. The only provided methods that show a competitive performance are *Yeo* and *Baldassano*, while *GRASP* yields the worst results amongst the computed parcellations. *Glasser*

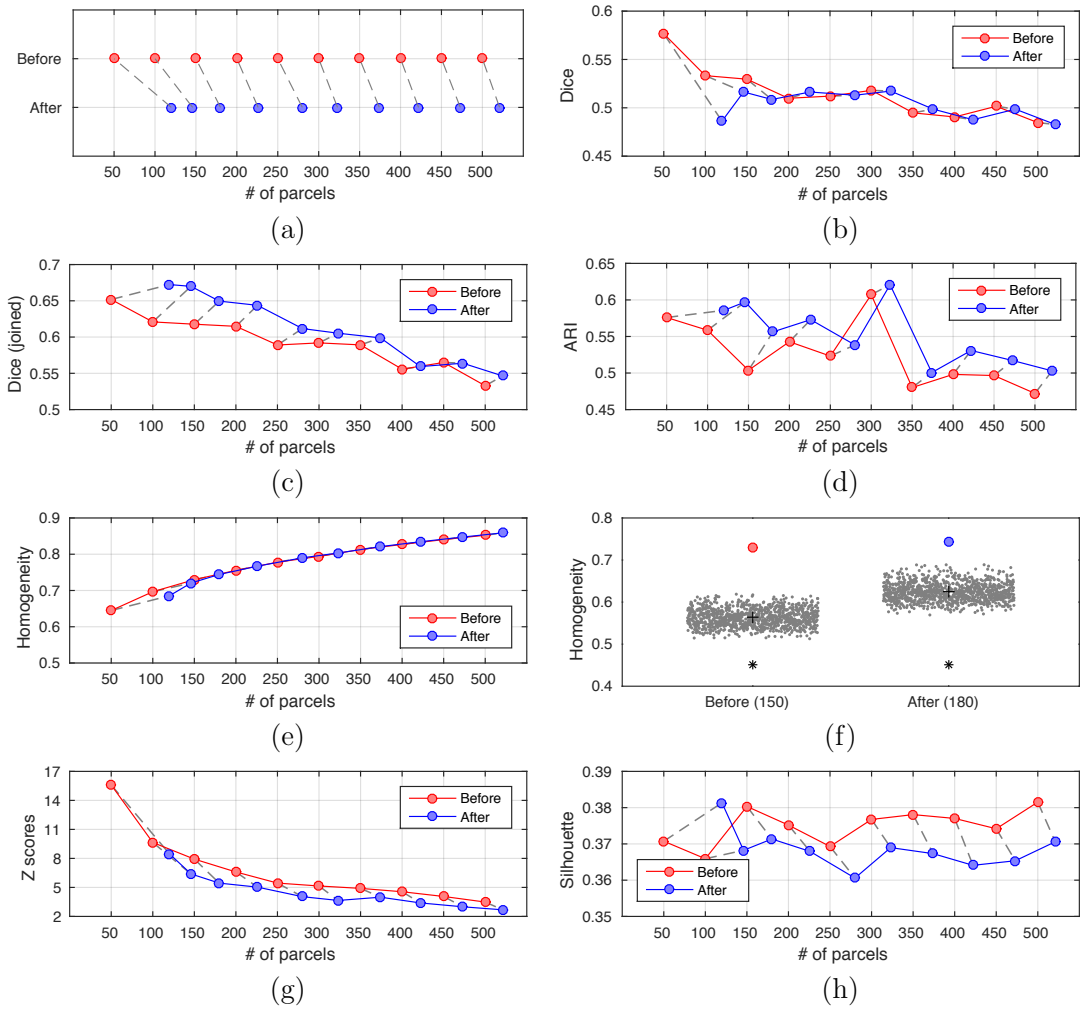


Figure 7.10: Quantitative evaluation measures obtained from *K-Means-AVR* parcellations, *before* and *after* disjoint parcels are split into spatially contiguous regions. Points representing the original and relabelled parcellations (shown in red and blue, respectively) are matched with dashed lines for ease of comparison. The blue points correspond to the number of parcels acquired at each resolution after splitting, and therefore, are plotted further to the right with respect to the red points, which align with the resolutions of the original parcellations (50 to 500, in increments of 50) along the x axis. (a) The number of parcels before and after the split process. (b-d) Group-to-group reproducibility obtained via Dice similarity, joined Dice similarity, and adjusted Rand index. (e-h) Clustering accuracy measured via parcel homogeneity, comparison to null models (only for one resolution), z scores relative to null models, and Silhouette analysis.

has a poor performance despite being driven by task average data. This can be attributed to the fact that it is generated from a different dataset which does not necessarily reflect the single subject task data in our evaluation set. Furthermore, this parcellation is obtained using a different registration method (based on cortical folding, myelination and functional

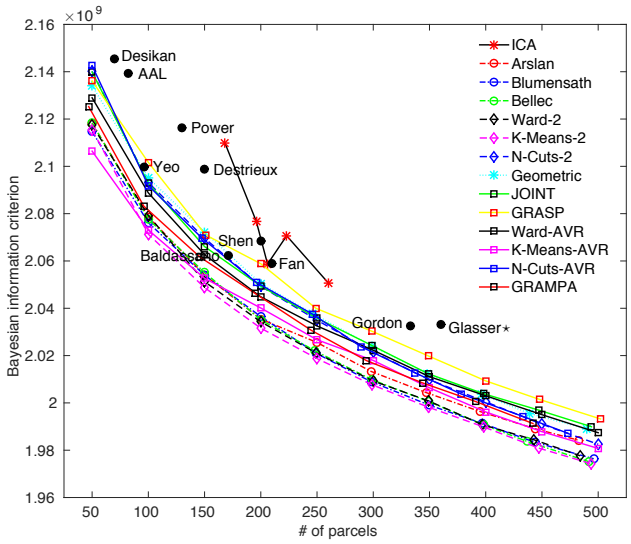


Figure 7.11: Group-level Bayesian information criterion (BIC) results for measuring the agreement with task activation. Lines show the BIC values for all computed resolutions, while black dots correspond to the BIC obtained from the publicly available parcellations with fixed resolutions. Lower BIC indicates higher agreement with the task activation. ⋆: It should be noted that *Glasser* is derived from group average task activation maps, which can influence this evaluation.

connectivity) than our evaluation set, which can also negatively influence the results.

The overlap between the groupwise parcellations and the Brodmann areas (BA) for all resolutions is given in Fig. 7.12. Many methods show a high degree of agreement with the primary somato-sensory cortex (BA[3,1,2]), pre-motor cortex (BA6), and primary visual cortex (BA17). Relatively low measures are obtained for the rest of the Brodmann areas, especially for the perinatal cortex (BA[35,36]). Overall, *Glasser* shows the best performance and yields the highest overlap for most areas. Similarly, other provided parcellations *Fan* and *Gordon*, as well as the anatomical parcellations show a relatively high performance. *Yeo*, *Power*, and *ICA* yield the lowest overlap measures, and in contrast to the general tendency in the group, do not align well with BA[3,1,2]. Interestingly, *K-Means-AVR* produces the poorest results amongst the computed parcellations for all resolutions. This can be linked to the fact that *K-Means-AVR* parcels are not necessarily spatially contiguous and may be spread across the cortex. In particular, the 2-level approaches perform better than the group-average ones, with *Bellec* leading them at almost all levels of granularity.

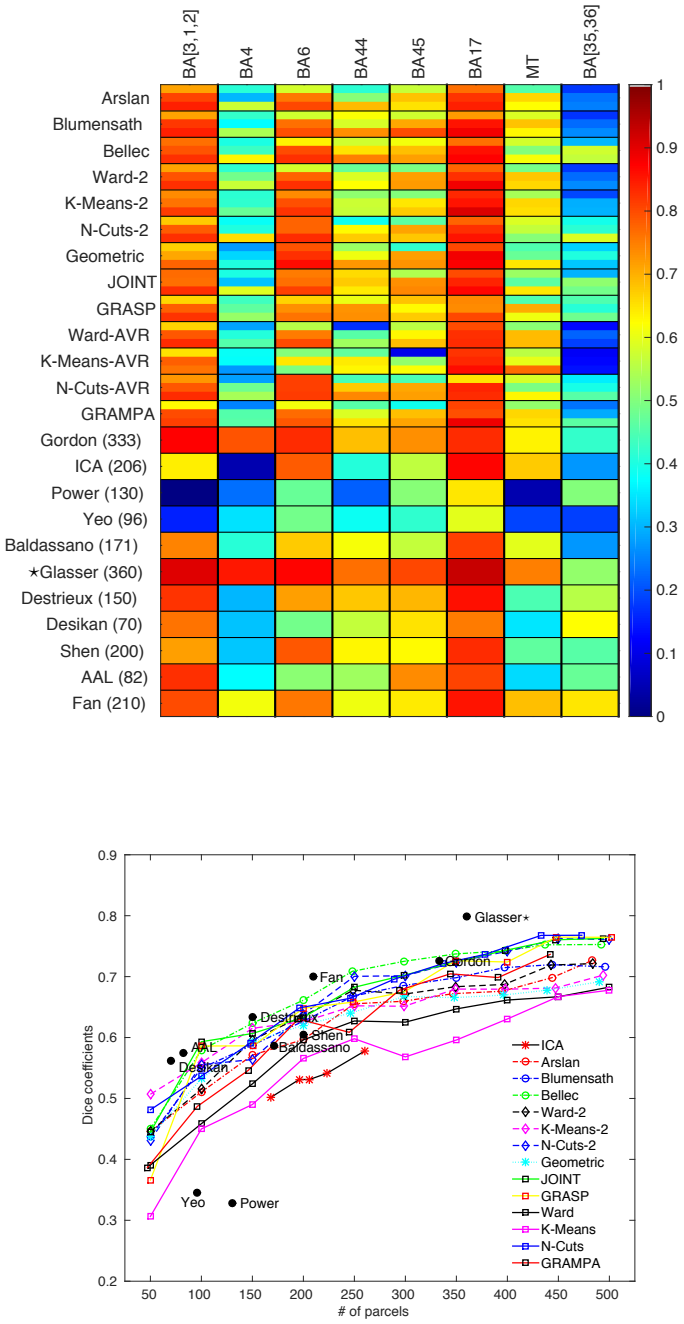


Figure 7.12: *Top*: Agreement of all group-level parcellations with several Brodmann areas. For the computed parcellations (top 13 rows), each cell shows Dice coefficients for $K = 100$, 200, and 300 regions, respectively from top to bottom. For the other parcellations, resolutions are indicated aside their names in parentheses. *Bottom*: Average Dice coefficients for each method/resolution.⋆: It should be noted that *Glasser* uses expert knowledge and priors from the neuro-anatomical literature for the delineation of parcellation borders, which can influence this evaluation.

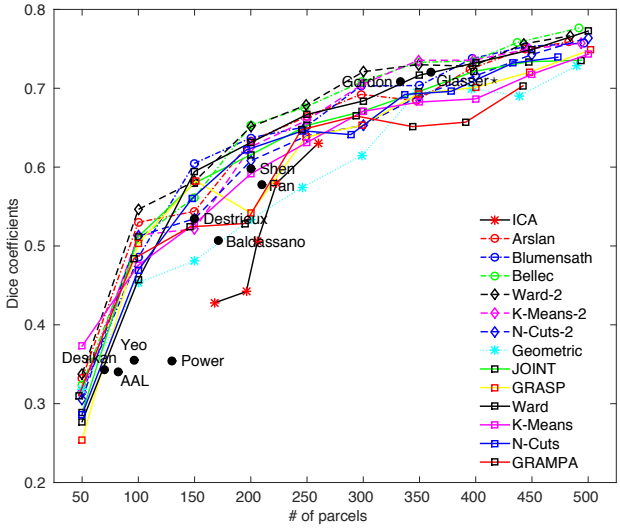


Figure 7.13: Dice-based overlap measures of all group-level parcellations with highly myelinated cortical areas, derived from a coarse parcellation of the average myelination map. ⋆: It should be noted that *Glasser* is derived from myelin maps and is therefore expected to have a good performance here.

Average overlap scores with myelin based parcellations are given in Fig. 7.13. In general, the 2-level approaches show similar performance and outperform the group-average methods for most resolutions. *Bellec*, *Ward-2* and *K-Means-2* have the highest agreement among the computed parcellations, while *GRAMPA* and *Geometric* yield relatively poor measures. *Glasser* and *Gordon* show the best performance amongst provided parcellations and outperform most of the other approaches when similar resolutions are considered. This is to be expected for *Glasser* since it is derived from myelin maps. Other provided parcellations generally yield relatively low measures.

7.3.5 Network Analysis

The results obtained for gender classification are illustrated in Figure 7.14. Although there is no single winner across all different resolutions, anatomical parcellations are generally outperformed by several data-driven methods with similar number of parcels. Overall, results obtained with SVM are not very consistent across resolutions, since there is no obvious up-

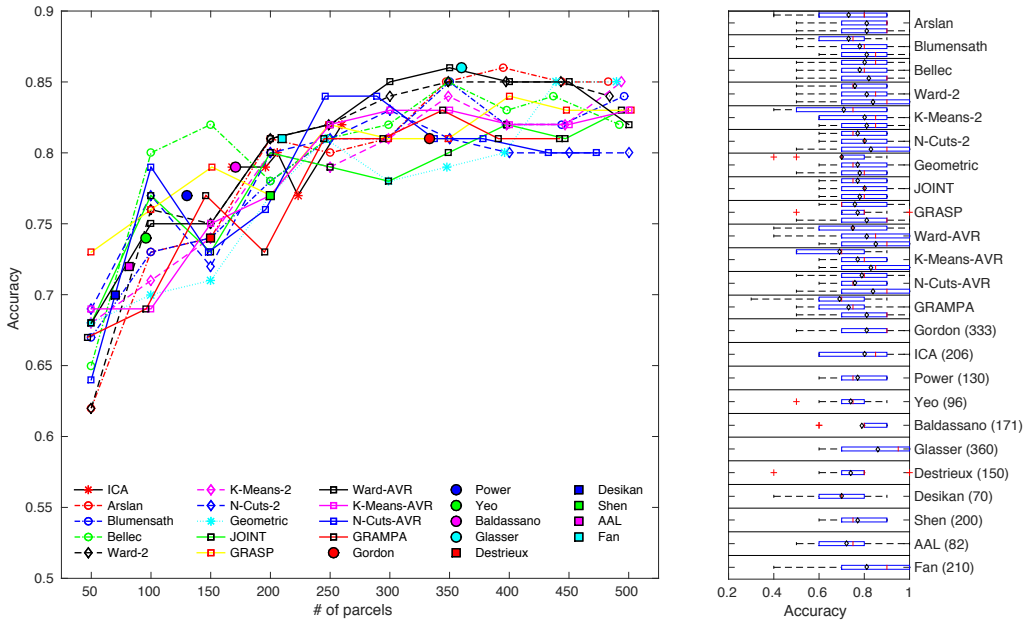


Figure 7.14: Gender classification results. *Left:* Average Gender classification accuracy on 100 subjects with SVM. *Right:* Variation across results is shown with respect to 10-fold cross-validation.

ward/downward trend with increasing resolution. In fact, most methods demonstrate a similar average performance, being able to classify males and females with above 60% accuracy for granularities below 150 parcels and above 70% for higher resolutions. More specifically, *Geometric* tends to perform poorly compared to the rest of the methods, both at lower and higher resolutions. The highest SVM classification accuracy (86%) is achieved with *Ward-AVR* and *Glasser* at the scale of 350 and 360 parcels, respectively. Moreover, we can observe that increasing the resolution of the parcellation in data-driven approaches beyond a certain value (350 parcels) does not necessarily provide additional information about population differences. However, lower resolutions lead to lower classification scores, perhaps due to the fact that functional information valuable for the discrimination between the two classes fades by averaging the signal in larger parcels. Interestingly, *N-Cuts-AVR*, *Bellec* and *Arslan* perform quite well for several resolutions, while *GRASP* yields the top accuracy among all methods for 50 parcels across the cortex. It is also worth mentioning that the parcellations provided by *Yeo*, *Shen* and *Gordon* have below average performance, while *Fan* and *Glasser* have good performance

compared to parcellations with similar resolutions.

Individual subject identification results presented in Fig. 7.15 show a clearer trend than gender classification: connectivity networks derived from relatively high-resolution parcellations yield a higher success rate. In particular, *K-Means-AVR* appear to obtain the highest accuracy at most of the granularity levels, whereas for lower resolutions, *Yeo*, *Power* and *ICA* outperform the other approaches. The rest of the methods have similar performances with *GRASP*, anatomical and spectral parcellations generally having the lowest success rate for most resolutions. The scale of the parcellation appears to have a much stronger impact on the identification accuracy than the parcellation scheme itself. This is likely due to the fact that the variability across individuals cannot be as effectively captured by low-resolution parcellations.

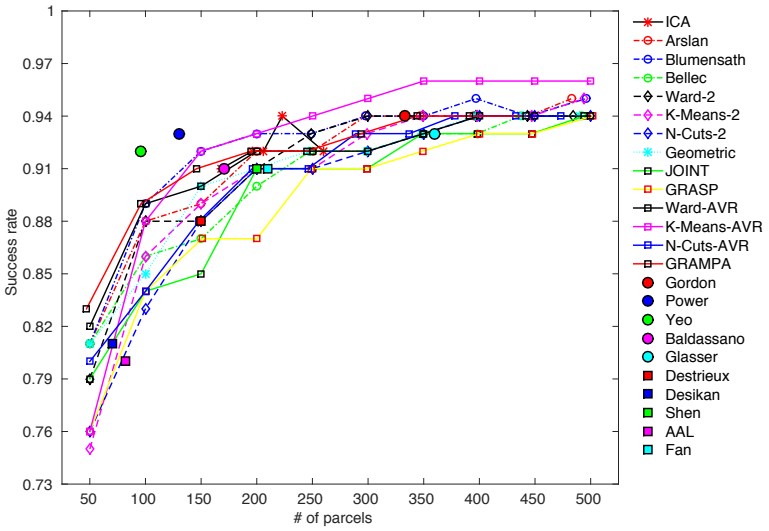


Figure 7.15: Individual subject prediction results.

7.4 Discussion

In this chapter, we proposed a large-scale comparison of existing group-level parcellation methods using state-of-the-art evaluation measures and publicly available data provided by HCP. The generation and evaluation of the parcellations is based on resting-state functional connectivity, which is thought to express the network behaviour underlying high level cognitive processes.

We considered several criteria simultaneously to evaluate the quality of the parcellations, such as reproducibility, parcel homogeneity, and Silhouette analysis. These measurements assessed the performance from a cluster quality point of view. The neuro-biological interpretation of the obtained parcels is also investigated by comparing parcel boundaries with well-defined neuro-biological properties, such as cytoarchitecture and myelination, as well as task activation. In addition, we devised two simple network analysis task, i.e. gender classification and subject identification, in order to measure the impact of the underlying parcellation on network analysis.

Our experiments show that there is no clear trend in favour of a specific method - or type of method - regarding all evaluation metrics considered. For instance, k-means clustering appears to be largely leading in terms of clustering quality. It, however, shows a poor performance regarding reproducibility and agreement with other modalities. In addition, while cortical delineation intrinsically requires a relatively large number of parcels, this does not appear to be a requirement for effective network analysis. This may suggest that different types of parcellations are to be investigated depending on the task at hand (e.g. one should use different methods when considering network analysis or cortical delineation).

We observe that connectivity-driven parcellations have a much better agreement with the underlying rs-fMRI connectivity compared to anatomical and random parcellations as expected. The benefit of using connectivity to parcellate the brain is not as clear regarding the delineation of cortical areas (agreement with other modalities and established brain delineations) and subsequent network analysis. In particular, anatomical parcellations appear to yield equivalent or better results with respect to cytoarchitecture. A general conclusion regarding network analysis would be to use any parcellation available, or learn it from the data, as it seems to have a limited impact. However, while this may be true for simple analysis of healthy subjects, it would have to be investigated further in the context of largely different brains (such as large age gaps or diseases).

Connectivity estimated from resting state fMRI and its impact on parcellations

Resting state fMRI is the most commonly used state-of-the-art technique to map whole-brain functional connectivity, with its high spatial resolution favouring its application over alternative electro-physiological recordings, like EEG and MEG. Its effectiveness to map the function of the brain has been consistently shown across a wide range of studies [56, 212, 255, 202]. However, the true biological interpretation of the BOLD signals is still unknown [66], and its low temporal resolution (commonly at the order of seconds) is a limiting factor for the observation of high-frequency patterns. The structure of functional connectivity can also be a shortcoming in itself. Several sources of noise can influence BOLD signals, including head movement, cardiac and respiratory pulsations. On top of this, one known issue with functional connectivity is the possible indirect connections mediated by third-party regions [226, 66]. As a result, it can be difficult to separate the actual signal from the noise, (or the direct connections from the indirect) and to assess the reliability of functional connectivity derived from fMRI and their interpretation. In particular, the noise in the data and the different hypotheses and processing decisions taken for a given parcellation technique could explain why different connectivity driven parcellation methods perform better or worse with respect to certain evaluation measures. For example many techniques alter the structure of the connectivity network by applying thresholding to exclude negative and weak correlations from the connectivity matrices with the aim of obtaining more robust parcellations [255, 202, 54, 11].

Evaluation of parcellations from a clustering point of view

When parcellations are evaluated, both implicit constraints inherent to the method and explicit constraints imposed to the data should be taken into consideration, as they yield inevitable biases towards the computed parcellations [32]. It is, therefore, highly critical to evaluate clustering accuracy from different perspectives.

K-means, hierarchical clustering, and spectral clustering (as well as their variants) are frequently used to obtain connectivity-driven parcellations, ultimately serving the task of brain

mapping [66]. Their impact on the parcellation configuration as well as their limitations and advantages over each other have been extensively reviewed in [244, 66] and further discussed in Chapter 3. In general, our results align with the previous literature (as well as with those derived from subject-level parcellations in Chapter 4) regarding the performance of these clustering algorithms. For example, k-means generally provides the best performing groupings of the data, but suffers from low reproducibility due to the fact that it does not inherently rely on hard spatial constraints. On the contrary, spectral techniques are usually dominated by spatial constraints and consequently yield highly reproducible parcellations, but at the cost of inaccurate alignment with the brain's underlying functional organisation. Hierarchical clustering offers the advantage of generating reproducible parcellations to a certain extent, while still capturing the functional connectivity with high fidelity.

Several parcellations computed on a different dataset yield relatively good cluster quality results. One can infer from this observation that similar characteristics shared by healthy adults can be robustly detected across different datasets as long as the analysis is performed on a large cohort (for example *ICA* and *Baldassano* are originally obtained from a group of 500 subjects where this number increases to 1000 for *Yeo*). It should be also noted that, *ICA* and *Baldassano* can also comprise some subjects from our test dataset as they are computed from a larger HCP cohort. This may constitute an important factor promoting a more favourable performance for these two methods compared to the others.

Predictably, anatomical parcellations yield the lowest performance in terms of clustering quality. However, they allow a more intuitive neuro-biological interpretation which can make network analysis more insightful. On top of that, our network-based experiments show that a better quality clustering does not necessarily benefit network analysis. One limitation is their relatively low resolution which is typically addressed by partitioning each parcel into subunits without altering the anatomically delineated boundaries. This can be achieved randomly [96, 101] or using functional connectivity [194, 68]. This approach is adapted by *Fan*, but appears to provide a limited improvement compared to anatomical parcellations.

Agreement of parcellations with other neuro-biological properties of the cortex

The anatomical parcellations based on cortical folding, i.e. *Desikan* and *Destrieux*, as well as the anatomo-functional atlas based on the Desikan parcels (i.e. *Fan*) interestingly show a high degree of agreement with the cytoarchitecture of the cerebral cortex. Although these results may reflect a better alignment between anatomy and cytoarchitectural atlases than with rs-fMRI, this might also be linked to registration errors as the Brodmann maps are registered to each individual subject based on cortical folding. While we can expect a good overlap in the motor and visual cortex, where the folding patterns are more consistent across subjects, stronger misalignments could occur in other regions.

Similar observations can be made for connectivity-driven parcellations, in which case a higher degree of alignment is found within the motor and visual cortex. Despite the fact that functional connectivity obtained from BOLD timeseries is not necessarily expected to reflect the cytoarchitecture of the cerebral cortex, these results agree with several rs-fMRI based studies that report similar findings regarding these regions [32, 280, 93]. On the other hand, a more consistent agreement can be expected between the connectivity-driven parcellations and highly myelinated areas, as the gradients in rs-fMRI-driven connectivity have been observed to align well with the myelination patterns [88].

One should also take into consideration the reliability of the evaluation techniques used to compare the different modalities. For example, overlap-based measures, such as the Dice coefficient, are biased by the size of the parcels. Evenly sized/shaped parcels are easier to match with their target parcels, while differences in Dice scores will be much more striking when comparing small parcels over big ones. This bias can lead to more favourable results for some of the parcellations, such as *Geometric*, *N-Cuts*, and *Random*, all of which comprise more uniformly shaped/sized parcels than the rest of the approaches. Although such quantitative measures can provide a means of comparing different methods, the quality of a parcellation with respect to cytoarchitecture or myelin content should also be visually assessed before drawing any conclusion.

Similarly, the Bayesian information criterion has a bias towards more complex models, i.e. parcellations with higher resolution are always favoured [244]. It should be also noted that there may exist redundant and contradictory information in the different tasks/contrasts which could bias the results. On top of that, the SNR in the task activation maps is low, therefore, it is likely that the results might be compromised by noise. Finally, our experiments have compared group-level parcellations to single subject level task activation maps. While the objective is to evaluate whether these group parcellations provide a good representation of the population, one could also consider comparing to group average task activation maps. This would alleviate single subject noise and could yield better results, in particular for provided parcellations. For example, the *Glasser* parcellation is expected to have a much better performance with respect to group level task maps on which it is derived.

Additionally, this multi-modal parcellation (*Glasser*) can give a clearer intuition on the behaviour of inter-modality comparisons. This method does not only rely on resting-state functional connectivity, but also embodies information from task activation, myelin content, and the cortical architecture. It yields very good overlap with the Brodmann areas and myelin content, especially on some parts of the cortex (e.g. motor cortex, highly-myelinated areas), indicating that the overlap measures used for multi-modal comparisons do provide accurate information.

Impact of parcellations on network analysis

When network analysis is concerned, our experiments based on two different tasks, gender classification and individual subject prediction, yield different results. The former does not favour any particular method to subdivide the brain into regions that would better reflect population differences. On the other hand, when individual subject prediction task is taken into consideration, we observe a dominance of connectivity-driven approaches, with a rather robust trend of increasing performance with respect to parcellation resolution. As the difference between the two sets of results can be linked to the specific network analysis tasks at hand, their ability to assess the parcellation quality should be interpreted separately.

Although classification analysis has previously been applied in studies of functional connectiv-

ity to predict demographic measures including gender [213, 206] and age [265], our experiments suggest that the classification score alone is not a valuable tool for the evaluation of parcellation quality. Instead, the number of features selected (edges in the connectivity matrix) to achieve the same classification performance might be a better means of evaluation. Provided that a larger number of subjects is available, a good parcellation should give a sparse selection of features and a more interpretable result. The results do not favour any parcellation, either anatomy, cytoarchitecture, or data driven, to better reflect population differences. In fact, anatomical parcellations appear to perform as well as data-driven approaches when the similar resolutions are considered. This could be attributed to the specific task at hand, since anatomical and, more specifically, cerebral volume differences have been reported between males and females that significantly influence the volume of white and gray matter [141]. Therefore, volume/anatomy-related differences and sex-related differences are hard to disentangle under the current experimental setting, despite the fact that all subjects have been registered to the same anatomical space.

When individual subject identification is concerned, connectivity-driven approaches appear to show a consistently more favourable performance compared to the other parcellation schemes. However, not particularly the parcellation scheme itself, but its resolution boosts the performance of the prediction task, as any parcellation with more than 200 regions can accurately predict at least 90 out 100 subjects. This may be linked to the fact that, parcellations with fewer regions may not be able to capture relevant features to distinguish one subject from another. Similar findings are also reported in [71], where higher accuracy is obtained by using a more fine-grained parcellation. It is also worth noting that, more robust trends are observed with the prediction task than gender classification. While one can infer that the classification task is simply more sensitive than the prediction task, this could also be attributed to the fact that capturing characteristic features on a single-subject basis is much easier than identifying population differences from a whole group of subjects.

Limitations

In this paper, we considered both surface-based and volumetric parcellations. Whilst efforts are made to be fair to all methods, several important methodological choices have been made, which may have an impact on the evaluation and possibly promote some parcellations over the others. In particular, parcellations are not the products of the same processing pipeline. Most of the publicly available parcellations have been generated under different assumptions, from different sets of subjects with varying cohort size and after being subject to a series of processing steps. Additional processing was applied to certain methods to make parcellations comparable on a more standard basis. Parcellations that do not naturally provide spatially contiguous cortical areas (e.g. *Yeo*, *Power*, *ICA*) were relabelled while methods that do not cover the entire cortical surface (e.g. *Gordon*) were dilated. Similarly, we used the group-average *Glasser* parcellation in our experiments, despite the fact that this method also provides individual parcellations tailored to each subject. In particular, the performance of parcellations sampled from a volumetric space should be interpreted carefully due to the complicated transformation steps. Nevertheless, we believe these parcellations are an essential aspect of our evaluation. Please see Figs. 7.16 and 7.17, for figures showing group-wise parcellations used in this study, respectively. All the parcellations and evaluation code are made publicly available via the web page: `http://biomedia.doc.ic.ac.uk/brain-parcellation-survey`, in case one may need access to these parcellations for their own analysis on a different dataset.

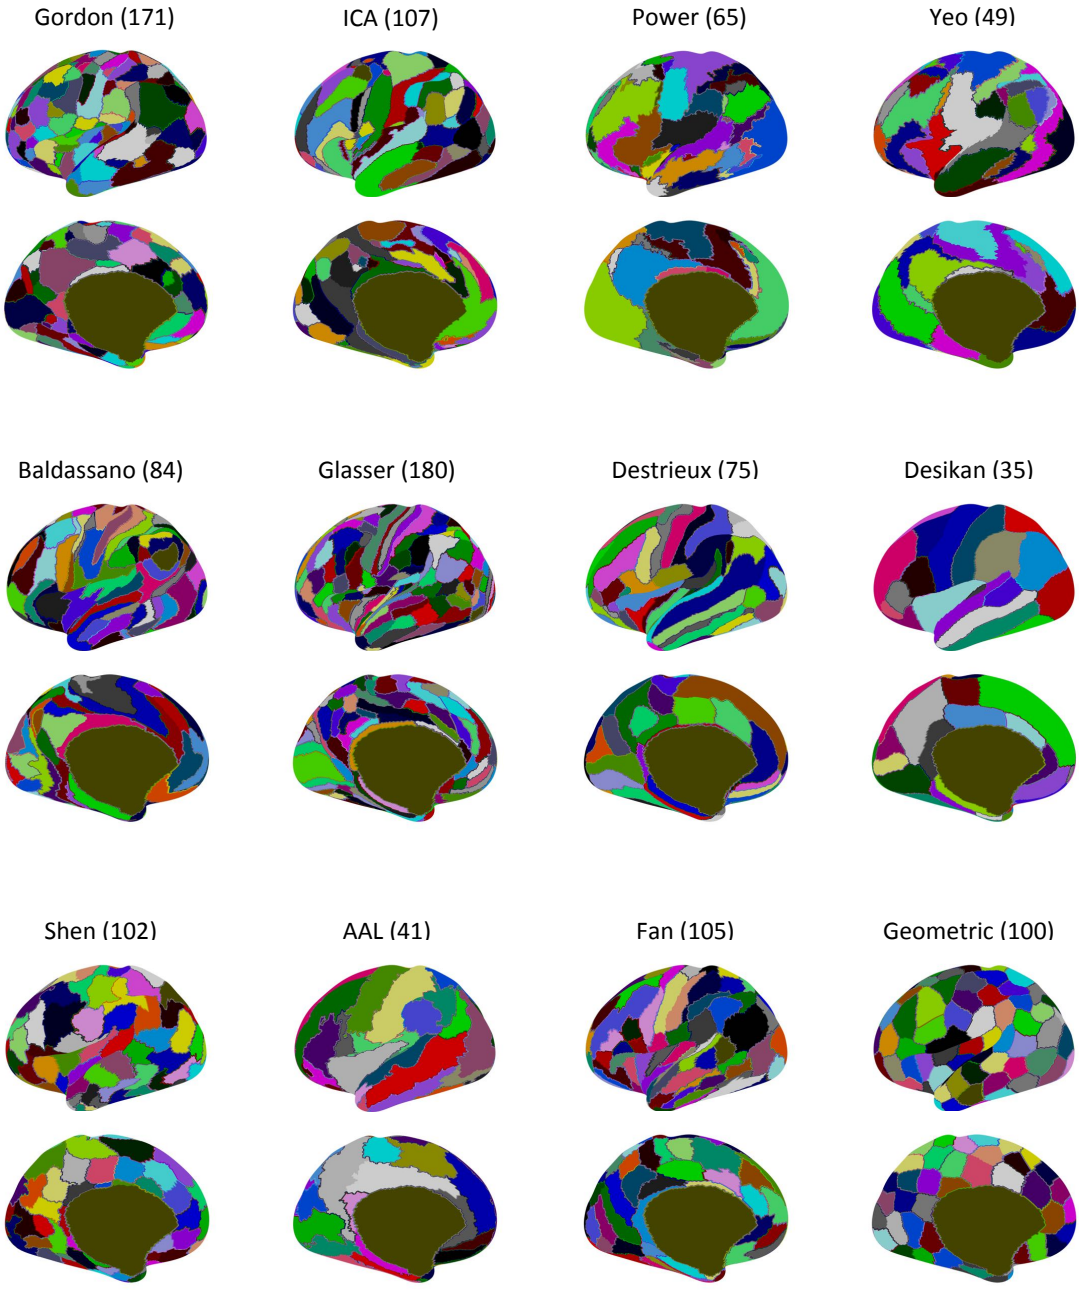


Figure 7.16: The first half of the parcellations considered in this study as projected on the left hemisphere. The number of parcels is shown along with the name of the method above each parcellation. Computed parcellations are only shown for a fixed resolution of 100 regions.

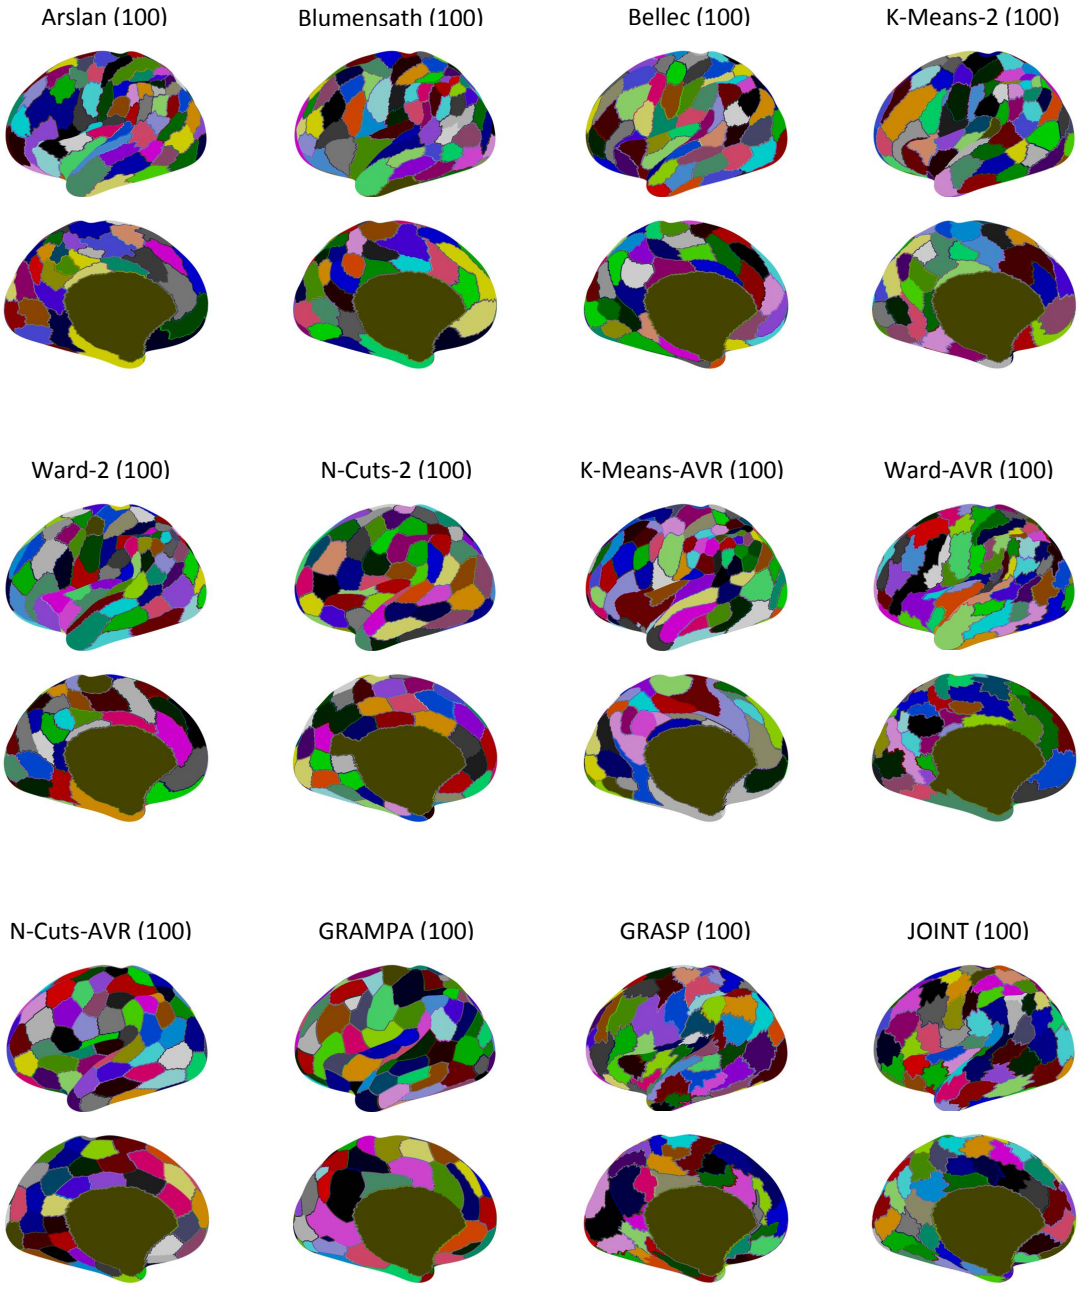


Figure 7.17: The second half of the parcellations considered in this study as projected on the left hemisphere. The number of parcels is shown along with the name of the method above each parcellation. Computed parcellations are only shown for a fixed resolution of 100 regions.

Chapter 8

Conclusions

This thesis presents a number of different connectivity-driven parcellation methods to subdivide the human cerebral cortex into non-overlapping, spatially contiguous, and homogeneous regions that can be used for studying the functional and structural organisation of the human brain. Our technical contributions offer new solutions to different aspects of the parcellation problem. As features for the proposed methods, we rely on the connectivity information obtained from rs-fMRI and dMRI, the most-widely used neuroimaging techniques for mapping the human brain. We propose multi-scale parcellations both for identifying the organisation of individual brains and representing shared patterns of connectivity within a relatively large cohort of healthy subjects.

In Chapter 4, we introduce a new parcellation method for segregating the cerebral cortex on a single subject basis. The proposed methodology relies on two different clustering approaches combined into a two-layer framework, in which the first layer stands as a spatial feature reduction stage and provides a relatively high-level representation of the cortex. Running atop this initial parcellation, the second layer computes a spectrum of parcellations, nested in a hierarchical fashion, yielding subject-level parcellations at different levels of granularity. We show that the proposed parcellations subdivide the cerebral cortex into regions with distinct patterns of connectivity and with higher reproducibility and homogeneity compared to another two-layer approach. We further discuss the advantages and limitations of our method over sev-

eral other data-driven clustering techniques, and show that the functional organisation of the individual human brain can be more accurately modelled by connectivity-driven parcellations with respect to random or anatomical ones.

In Chapter 5, we again address the subject-level parcellation problem, but from the perspective of a different modality, i.e. structural connectivity acquired via dMRI. We reformulate the parcellation problem as a feature reduction problem and utilise a non-linear manifold learning technique to identify connectivity characteristics of an individual subject that may not be directly visible in high dimensional space. Assuming that sharp changes in connectivity patterns residing in this low-dimensional embedding correspond to parcellation boundaries, we use edge detection techniques to subdivide the cerebral cortex into spatially contiguous and non-overlapping regions. Our experiments show that the proposed method is capable of providing distinct parcels with uniform connectivity profiles at varying resolutions with respect to several other parcellation techniques. We further show that our parcellations can potentially represent both the structural and functional organisation of the brain and hence provide a more reliable set of network nodes for connectome analysis.

In Chapter 6, we propose an alternative methodology to group-wise parcellation approaches currently used in the literature. We suggest combining within- and between-subject connectivity patterns into a multi-layer graphical model in order to better capture the fundamental connectional characteristics of the whole population, while still preserving connectivity profiles unique to each subject. Using spectral decomposition techniques, we partition this joint graph into cortical parcellations at varying levels of granularity. Experimenting with resting-state functional connectivity, we show that the proposed framework is able to delineate more reproducible parcellations across different groups of subjects and can more accurately model the underlying data with respect to two popular ways of obtaining group-wise parcellations. We also report that individual subjects can be more reliably represented by our group-wise parcellations compared to the other approaches.

In Chapter 7, we provide a large-scale systematic comparison of connectivity-driven, anatomical, and random parcellation methods proposed for brain mapping purposes. We rely on rs-

fMRI and widely-used quantitative evaluation techniques to evaluate 24 group-wise parcellation methods at different resolutions. We assess the parcellation accuracy from different aspects, including reproducibility, fidelity to the underlying data, similarity with task activation and myelo-/cyto-architecture of the cortex, as well as performance in network analysis. Taking the experimental results into consideration, we extensively review the strengths and limitations of the various methods over each other, with the aim of providing a guideline for choosing the most suitable parcellation for future connectome studies. We conclude that there is not a single 'best' parcellation that can simultaneously address all the challenges and yield high performance measures for all the tested cases.

8.1 Future Work

In this thesis, we developed both individual and group-wise parcellation methods -as well as evaluated many more- with the aim of providing some insight into different levels of representing the brain's functional and structural organisation.

Analysing individual parcellations in collaboration with group-wise models can potentially enhance our capacity in understanding inter-subject variability. While we explored the variations between individuals with respect to subject-level parcellations, it could be highly interesting to locate cortical regions that are most consistent and/or least similar across subjects by comparing individual parcellations to a group representation. Similarly, a subject-to-group analysis from a network theoretic point of view could provide a reliable means to reveal the cortical areas with similar/different connectivity patterns. In this context, understanding the source of variability across subjects constitutes an additional challenge. While differences in connectivity can be linked to a possible impairment of cognitive abilities or brain functioning, these could also be associated with genetic variations, between-subject topological differences, or simply caused by imaging/processing artefacts or low SNR in the data [132, 63, 92]. Given the proposed methods at multiple levels, studying this variability could be an interesting challenge to tackle as a follow-up research project.

The results presented as part of this thesis may indicate parcellation techniques and/or resolutions that are more appropriate for the problem under investigation. For example, one can prefer a more homogeneous parcellation to define network nodes in connectome analysis, as they are more likely to represent regions with uniform patterns of connectivity and reduce information loss. Similarly, a more reproducible parcellation can be selected for comparing different groups of subjects, for instance, to identify how connectivity changes longitudinally. While our experiments only covered healthy subjects, one can use a data-driven parcellation as reference for analysing connectivity patterns in a cohort with a specific brain disorder (after learning a parcellation for this new population), assuming the organisation of the brain is not too altered, for example, due to existence of a tumour. This could help derive biomarkers in order to better understand disease-related differences in brain connectivity, and thus, constitutes one of our primary future directions.

Another point worth noting is the fact that, not only the parcellation scheme itself, but its resolution might have an impact on network analysis, depending on the task at hand [9]. A recent study suggests that increasing the parcellation resolution yields more reliable biomarkers for studying brain disorders [3]. Similarly, using more ROIs for network analysis is reported to improve the performance of age prediction tasks [143]. This might be linked to the fact that parcellations with fewer ROIs may not be able to capture structural patterns of interest from the underlying data due to their resolutions. In this case, connectivity-driven parcellations provide a greater flexibility to study the impact of resolution on connectome analysis, as they allow to subdivide the cortex at different resolutions, as opposed to pre-computed parcellations/atlases with fixed resolutions. As a matter of fact, parcellation resolution should also be taken under consideration in future work that aims to identify biomarkers for studying brain disorders.

The parcellations we presented in this thesis can be used to model the functional/structural organisation of the brain and/or derive distinct features for network analysis. However, additional information might be required to enhance the information provided by rs-fMRI or dMRI alone. Recent evidence suggests that a single modality is too limited to reveal the complex structure of the cerebral cortex, which consists of a mosaic of multiple properties nested at different levels of detail [189, 86, 66]. From a neuro-biological point of view, the integration of

other modalities to the parcellation generation task may provide more accurate and robust cortical segregation of the cerebral cortex, as shown in the recently proposed multi-modal cortical parcellation [86]. A prospective future work therefore would be to use a similar technique and expand the current evaluation pipeline towards parcellations obtained from different modalities and their combinations. Last but not least, features obtained from multiple modalities may yield more robust biomarkers for the prediction of neuro-cognitive disorders. In this context, a recent study has shown that multimodal integration of anatomical features (such as cortical thickness) and functional connectivity improves brain-based age prediction [143].

Deep learning has recently gained a lot attention due to its notable performance in classification and pattern analysis tasks [199]. These tasks are also critical for neuroimaging and connectomics research, and thus, making deep models attractive to the researchers working in the related fields. For example, deep learning is actively being used for exploring neurological disorders and shown to be an effective tool for identifying brain-based biomarkers (see [266] for a review), studying disease-related differences in populations [191], and functional network modelling [128]. These works are mainly made possible by the advancements in *geometric deep learning*, which allows to generalise structured deep models to non-Euclidean domains, such as graphs [39]. Since brain connectivity can be represented as a graph of interacting units (possibly derived from a parcellation), one of our future works is to use convolutional neural networks (CNNs) through geometric deep learning in order to explore within- and inter-subject connectivity. On the other hand, while CNNs that are supervised by anatomical parcellations are already being used for brain segmentation problems [140, 58], learning a connectivity-based parcellation directly from the data is still not possible due to lack of a ground truth parcellation. Yet, unsupervised deep learning models, such as auto-encoders can be used to learn a latent representation of the underlying data, which in turn, may be used for parcellation purposes. For example such a non-linear representation can potentially be considered to substitute the Laplacian eigenmaps used in Chapter 5. Auto-encoders in conjunction with CNNs can be used to reconstruct the input brain signals (e.g. timeseries) in an unsupervised setting and feature maps obtained throughout such a deep model may reveal the intrinsic connectivity patterns and potentially be used to derive parcellations.

List of Publications

- **S. Arslan**, S. I. Ktena, A. Makropoulos, E. C. Robinson, D. Rueckert, S. Parisot, *Human Brain Mapping: A Systematic Comparison of Parcellation Methods for the Human Cerebral Cortex*, NeuroImage, 2017, *In Press*.

- S. I. Ktena, **S. Arslan**, S. Parisot, D. Rueckert, *Exploring Heritability of Functional Brain Networks with Inexact Graph Matching*, International Symposium on Biomedical Imaging (ISBI), 2017.

- S. Parisot, B. Glocker, S. I. Ktena, **S. Arslan**, M. D. Schirmer, D. Rueckert, *A Flexible Graphical Model for Multi-modal Parcellation of the Cortex*, NeuroImage, 2017, NeuroImage, vol. 162, pp. 226-248, 2017.

- **S. Arslan**, S. Parisot, D. Rueckert, *Boundary Mapping through Manifold Learning for Connectivity-Based Cortical Parcellation*. International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI), vol. 9900 of LNCS. Springer, pp. 115-122, 2016. -Received a travel award-

- S. Parisot, **S Arslan**, J. Passerat-Palmbach, W. M. Wells, D. Rueckert, *Group-wise Parcellation of the Cortex through Multi-scale Spectral Clustering*, NeuroImage, vol. 136, pp. 68-83, 2016. -Featured on issue cover-

- **S. Arslan** and D. Rueckert, *Multi-level Parcellation of the Cerebral Cortex Using Resting-State fMRI*, International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI), vol. 9351 of LNCS. Springer, pp. 47-54, 2015.

- **S. Arslan**, S. Parisot, D. Rueckert, *Joint Spectral Decomposition for the Parcellation of the Cerebral Cortex using Resting-State fMRI*, Information Processing in Medical Imaging (IPMI) Lecture Notes in Computer Science, vol. 9123, pp. 85-97, 2015. -Oral presentation-

- S. Parisot, **S. Arslan**, J. Passerat-Palmbach, W. M. Wells III, D. Rueckert, *Tractography-Driven Groupwise Multi-Scale Parcellation of the Cortex*, Information Processing in Medical Imaging (IPMI) Lecture Notes in Computer Science, vol. 9123, pp. 600-612, 2015.

- **S. Arslan**, S. Parisot, D. Rueckert, *Comparing Connectivity-based Groupwise Parcellations Generated from Resting-state fMRI and DTI Data: Preliminary results*, Symposium on Big Data Initiatives for Connectomics Research, London, 2015.

- S. I. Ktena, **S. Arslan**, D. Rueckert, *Gender Classification and Manifold Learning on Functional Brain Networks*," Symposium on Big Data Initiatives for Connectomics Research, London, 2015.

- **S. Arslan**, S. Parisot, D. Rueckert, *Supervertex Clustering of the Cerebral Cortex Using Resting-state fMRI*," Organization for Human Brain Mapping (OHBM), Honolulu, 2015.

- **S. Arslan**, S. Parisot, D. Rueckert, *How to Represent Subregions in a Parcellated Brain for fMRI Analysis?*," Organization for Human Brain Mapping (OHBM), Honolulu, 2015.

Bibliography

[1] The Developing Human Connectome Project. `http://www.developingconnectome.org/`, 2013.

[2] A. Abraham, E. Dohmatob, B. Thirion, D. Samaras, and G. Varoquaux. Extracting brain regions from rest fmri with total-variation constrained dictionary learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 607–615. Springer, 2013.

[3] A. Abraham, M. Milham, A. Di Martino, R. C. Craddock, D. Samaras, B. Thirion, and G. Varoquaux. Deriving reproducible biomarkers from multi-site resting-state data: An autism-based example. *NeuroImage*, 2016.

[4] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Susstrunk. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE T Pattern Anal*, 34(11):2274–2282, 2012.

[5] E. Amaro and G. J. Barker. Study design in fMRI: Basic principles. *Brain and cognition*, 60(3):220–232, 2006.

[6] J. Andersson, J. Xu, E. Yacoub, E. Auerbach, S. Moeller, and K. Ugurbil. A comprehensive Gaussian process framework for correcting distortions and movements in diffusion images. In *Proceedings of the 20th Annual Meeting of ISMRM*, page 2426, 2012.

[7] J. L. Andersson, S. Skare, and J. Ashburner. How to correct susceptibility distortions in spin-echo echo-planar images: Application to diffusion tensor imaging. *Neuroimage*, 20(2):870–888, 2003.

[8] A. Anwander, M. Tittgemeyer, D. Y. von Cramon, A. D. Friederici, and T. R. Knösche. Connectivity-based parcellation of Broca's area. *Cerebral Cortex*, 17(4):816–825, 2007.

[9] S. Arslan, S. I. Ktena, A. Makropoulos, E. C. Robinson, D. Rueckert, and S. Parisot. Human brain mapping: A systematic comparison of parcellation methods for the human cerebral cortex. *NeuroImage*, 2017.

[10] S. Arslan, E. Ozyurek, and C. Gunduz-Demir. A color and shape based algorithm for segmentation of white blood cells in peripheral blood and bone marrow images. *Cytometry Part A*, 85(6):480–490, 2014.

[11] S. Arslan, S. Parisot, and D. Rueckert. Joint spectral decomposition for the parcellation of the human cerebral cortex using resting-state fMRI. In N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, editors, *IPMI 2015 LNCS*, pages 85–97, 2015.

[12] S. Arslan, S. Parisot, and D. Rueckert. Boundary mapping through manifold learning for connectivity-based cortical parcellation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 115–122. Springer International Publishing, 2016.

[13] S. Arslan and D. Rueckert. Multi-level parcellation of the cerebral cortex using resting-state fMRI. In S. Ourselin, D. Alexander, C.-F. Westin, and M. Cardoso, editors, *MICCAI 2015 LNCS*, pages 47–54, 2015.

[14] F. A. Azevedo, L. R. Carvalho, L. T. Grinberg, J. M. Farfel, R. E. Ferretti, R. E. Leite, W. J. Filho, R. Lent, and S. Herculano-Houzel. Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *Journal of Comparative Neurology*, 513(5):532–541, 2009.

[15] C. Baldassano, D. M. Beck, and L. Fei-Fei. Parcellating connectivity in spatial maps. *PeerJ*, 3:e784, 2015.

[16] D. M. Barch, G. C. Burgess, M. P. Harms, S. E. Petersen, B. L. Schlaggar, M. Corbetta, M. F. Glasser, S. Curtiss, S. Dixit, C. Feldt, D. Nolan, E. Bryant, T. Hartley, O. Footer,

J. M. Bjork, R. Poldrack, S. Smith, H. Johansen-Berg, A. Z. Snyder, and D. C. V. Essen. Function in the human connectome: Task-fMRI and individual differences in behavior. *NeuroImage*, 80:169–189, 2013.

[17] S. T. Barnard, A. Pothen, and H. Simon. A spectral algorithm for envelope reduction of sparse matrices. *Numerical Linear Algebra with Afpplications*, 2(4):317–334, 1995.

[18] P. J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44(4):625–632, 2000.

[19] D. S. Bassett, E. Bullmore, B. A. Verchinski, V. S. Mattay, D. R. Weinberger, and A. Meyer-Lindenberg. Hierarchical organization of human cortical networks in health and schizophrenia. *J Neurosci*, 28(37):9239–9248, 2008.

[20] M. F. Bear, B. W. Connors, and M. A. Paradiso. *Neuroscience: Exploring the brain.* Wolters Kluwer, 2016.

[21] C. F. Beckmann and S. M. Smith. Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Trans Med Imag*, 23(2):137–152, 2004.

[22] M. Beckmann, H. Johansen-Berg, and M. F. Rushworth. Connectivity-based parcellation of human cingulate cortex and its relation to functional specialization. *J Neurosci*, 29(4):1175–1190, 2009.

[23] T. Behrens, H. J. Berg, S. Jbabdi, M. Rushworth, and M. Woolrich. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage*, 34(1):144–155, 2007.

[24] T. E. Behrens and S. Jbabdi. MR diffusion tractography. *Diffusion MRI, from Quantitative Measurement to in vivo Neuroanatomy*, pages 333–352, 2009.

[25] T. E. Behrens, M. Woolrich, M. Jenkinson, H. Johansen-Berg, R. Nunes, S. Clare, P. Matthews, J. M. Brady, and S. M. Smith. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic Resonance in Medicine*, 50(5):1077–1088, 2003.

[26] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput*, 15(6):1373–1396, 2003.

[27] P. Bellec, V. Perlbarg, S. Jbabdi, M. Pelegrini-Issac, J.-L. Anton, J. Doyon, and H. Benali. Identification of large-scale networks in the brain using fMRI. *NeuroImage*, 29(4):1231–1243, 2006.

[28] P. Bellec, P. Rosa-Neto, O. C. Lyttelton, H. Benali, and A. C. Evans. Multi-level bootstrap analysis of stable clusters in resting-state fMRI. *NeuroImage*, 51(3):1126–1139, 2010.

[29] J. C. Bezdek. *Pattern recognition with fuzzy objective function algorithms*. Springer Science & Business Media, 2013.

[30] B. Biswal, F. Zerrin Yetkin, V. M. Haughton, and J. S. Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnet Reson Med*, 34(4):537–541, 1995.

[31] B. B. Biswal, M. Mennes, X.-N. Zuo, S. Gohel, C. Kelly, S. M. Smith, C. F. Beckmann, J. S. Adelstein, R. L. Buckner, S. Colcombe, A.-M. Dogonowski, M. Ernst, D. Fair, M. Hampson, M. J. Hoptman, J. S. Hyde, V. J. Kiviniemi, R. Ktter, S.-J. Li, C.-P. Lin, M. J. Lowe, C. Mackay, D. J. Madden, K. H. Madsen, D. S. Margulies, H. S. Mayberg, K. McMahon, C. S. Monk, S. H. Mostofsky, B. J. Nagel, J. J. Pekar, S. J. Peltier, S. E. Petersen, V. Riedl, S. A. R. B. Rombouts, B. Rypma, B. L. Schlaggar, S. Schmidt, R. D. Seidler, G. J. Siegle, C. Sorg, G.-J. Teng, J. Veijola, A. Villringer, M. Walter, L. Wang, X.-C. Weng, S. Whitfield-Gabrieli, P. Williamson, C. Windischberger, Y.-F. Zang, H.-Y. Zhang, F. X. Castellanos, and M. P. Milham. Toward discovery science of human brain function. *P Natl Acad Sci*, 107(10):4734–4739, 2010.

[32] T. Blumensath, S. Jbabdi, M. F. Glasser, D. C. Van Essen, K. Ugurbil, T. E. Behrens, and S. M. Smith. Spatially constrained hierarchical parcellation of the brain with resting-state fMRI. *NeuroImage*, 76:313–324, Aug 2013.

[33] J. W. Bohland, H. Bokil, C. B. Allen, and P. P. Mitra. The brain atlas concordance problem: Quantitative comparison of anatomical parcellations. *PLoS One*, 4(9):e7200, 2009.

[34] F. L. Bookstein. Two shape metrics for biomedical outline data: Bending energy, procrustes distance, and the biometrical modeling of shape phenomena. In *Shape Modeling and Applications, 1997. Proceedings., 1997 International Conference on*, pages 110–120. IEEE, 1997.

[35] J. R. Booth, L. Wood, D. Lu, J. C. Houk, and T. Bitan. The role of the basal ganglia and cerebellum in language processing. *Brain Research*, 1133:136–144, 2007.

[36] I. Borg and P. J. Groenen. *Modern multidimensional scaling: Theory and applications*. Springer Science & Business Media, 2005.

[37] M. Bota, O. Sporns, and L. W. Swanson. Architecture of the cerebral cortical association connectome underlying cognition. *Proceedings of the National Academy of Sciences*, 112(16):E2093–E2101, 2015.

[38] K. Brodmann. Vergleichende lokalisationslehre der gro hirnrinde. 1909.

[39] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: Going beyond euclidean data. *IEEE Signal Proc Mag*, 34(4):18–42, 2017.

[40] R. L. Buckner and J. L. Vincent. Unrest at rest: Default activity and spontaneous network correlations. *NeuroImage*, 37(4):1091–1096, 2007.

[41] E. Bullmore and O. Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186–198, 2009.

[42] U. Burgel, K. Amunts, L. Hoemke, H. Mohlberg, J. M. Gilsbach, and K. Zilles. White matter fiber tracts of the human brain: Three-dimensional mapping at microscopic resolution, topography and intersubject variability. *NeuroImage*, 29(4):1092–1105, 2006.

[43] C. J. Burges. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, 2(2):121–167, 1998.

[44] R. B. Buxton, K. Uludağ, D. J. Dubowitz, and T. T. Liu. Modeling the hemodynamic response to brain activation. *NeuroImage*, 23:S220–S233, 2004.

[45] M. Catani and D. H. Ffytche. The rises and falls of disconnection syndromes. *Brain*, 128(10):2224–2239, 2005.

[46] R. C. Chan, D. Shum, T. Toulopoulou, and E. Y. Chen. Assessment of executive functions: Review of instruments and identification of critical issues. *Archives of Clinical Neuropsychology*, 23(2):201 – 216, 2008.

[47] P. S. Churchland. *Neurophilosophy: Toward a unified science of the mind-brain*. MIT Press, 1989.

[48] L. L. Cloutman and L. Ralph. Connectivity-based structural and functional parcellation of the human cortex using diffusion imaging and tractography. *Frontiers in neuroanatomy*, 6:34, 2012.

[49] A. L. Cohen, D. A. Fair, N. U. 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.

[50] R. R. Coifman and S. Lafon. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30, 2006.

[51] D. M. Cole, S. M. Smith, and C. F. Beckmann. Advances and pitfalls in the analysis and interpretation of resting-state fMRI data. *Frontiers in Systems Neuroscience*, 4(8), 2010.

[52] D. L. Collins, P. Neelin, T. M. Peters, and A. C. Evans. Automatic 3D intersubject registration of MR volumetric data in standardized Talairach space. *Journal of Computer Assisted Tomography*, 18(2):192–205, 1994.

[53] D. Cordes, V. Haughton, J. D. Carew, K. Arfanakis, and K. Maravilla. Hierarchical clustering to measure connectivity in fMRI resting-state data. *Magn Reson Im*, 20(4):305–317, 2002.

[54] R. C. Craddock, G. A. James, P. E. Holtzheimer, X. P. Hu, and H. S. Mayberg. A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Hum Brain Mapp*, 33(8):1914–1928, 2012.

[55] R. C. Craddock, S. Jbabdi, C.-G. Yan, J. T. Vogelstein, F. X. Castellanos, A. Di Martino, C. Kelly, K. Heberlein, S. Colcombe, and M. P. Milham. Imaging human connectomes at the macroscale. *Nature Methods*, 10(6):524–539, 2013.

[56] 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. *P Natl Acad Sci*, 103(37):13848–13853, 2006.

[57] C. Davatzikos, K. Ruparel, Y. Fan, D. Shen, M. Acharyya, J. Loughead, R. Gur, and D. Langleben. Classifying spatial patterns of brain activity with machine learning methods: Application to lie detection. *NeuroImage*, 28(3):663 – 668, 2005.

[58] A. de Brebisson and G. Montana. Deep neural networks for anatomical brain segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 20–28, 2015.

[59] M. A. de Reus and M. P. van den Heuvel. The parcellation-based connectome: Limitations and extensions. *NeuroImage*, 80:397–404, 2013. Mapping the Connectome.

[60] E. L. Dennis, N. Jahanshad, J. D. Rudie, J. A. Brown, K. Johnson, K. L. McMahon, G. I. de Zubicaray, G. Montgomery, N. G. Martin, M. J. Wright, S. Y. Bookheimer, M. Dapretto, A. W. Toga, and P. M. Thompson. Altered structural brain connectivity in healthy carriers of the autism risk gene, CNTNAP2. *Brain Connectivity*, 1(6):447–459, 2011.

[61] R. S. Desikan, F. Sgonne, 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.

[62] L. R. Dice. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302, 1945.

[63] J. Dubois and R. Adolphs. Building a science of individual differences from fMRI. *Trends in Cognitive Sciences*, 20(6):425–443, 2016.

[64] S. Eickhoff, N. B. Walters, A. Schleicher, J. Kril, G. F. Egan, K. Zilles, J. D. Watson, and K. Amunts. High-resolution MRI reflects myeloarchitecture and cytoarchitecture of human cerebral cortex. *Hum Brain Mapp*, 24(3):206–215, 2005.

[65] S. B. Eickhoff, D. Bzdok, A. R. Laird, C. Roski, S. Caspers, K. Zilles, and P. T. Fox. Co-activation patterns distinguish cortical modules, their connectivity and functional differentiation. *NeuroImage*, 57(3):938–949, 2011.

[66] S. B. Eickhoff, B. Thirion, G. Varoquaux, and D. Bzdok. Connectivity-based parcellation: Critique and implications. *Hum Brain Mapp*, 36(12):4771–4792, 2015.

[67] T. Elbert and B. Rockstroh. Reorganization of human cerebral cortex: the range of changes following use and injury. *The Neuroscientist*, 10(2):129–141, 2004.

[68] L. Fan, H. Li, J. Zhuo, Y. Zhang, J. Wang, L. Chen, Z. Yang, C. Chu, S. Xie, A. R. Laird, P. T. Fox, S. B. Eickhoff, C. Yu, and T. Jiang. The human brainnetome atlas: A new brain atlas based on connectional architecture. *Cerebral Cortex*, 26(8):3508, 2016.

[69] Q. Fang and D. Boas. Tetrahedral mesh generation from volumetric binary and grayscale images. In *Proc. I S Biomed Imaging ISBI '09*, pages 1142–1145, 2009.

[70] D. J. Felleman and D. C. Van Essen. Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1(1):1–47, 1991.

[71] E. S. Finn, X. Shen, D. Scheinost, M. D. Rosenberg, J. Huang, M. M. Chun, X. Papademetris, and R. T. Constable. Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity. *Nature Neuroscience*, 18:16641671, 2015.

[72] B. Fischl. Freesurfer. *NeuroImage*, 62(2):774–781, 2012.

[73] 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, 1999.

[74] B. Fischl, A. van der Kouwe, C. Destrieux, E. Halgren, F. Sgonne, 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. *Cerebral Cortex*, 14(1):11–22, 2004.

[75] G. Flandin, F. Kherif, X. Pennec, D. Riviere, N. Ayache, and J. B. Poline. Parcellation of brain images with anatomical and functional constraints for fMRI data analysis. In *Biomedical Imaging, 2002. Proceedings. 2002 IEEE International Symposium on*, pages 907–910, 2002.

[76] FlickR. Carolyn's brain, by ClintJCL, 2006. Accessed on January 21, 2017.

[77] A. Fornito, A. Zalesky, and M. Breakspear. The connectomics of brain disorders. *Nature Reviews Neuroscience*, 16(3):159–172, 2015.

[78] A. Fornito, A. Zalesky, C. Pantelis, and E. Bullmore. Schizophrenia, neuroimaging and connectomics. *NeuroImage*, 62(4):2296–2314, 2012.

[79] P. Fransson. Spontaneous low-frequency BOLD signal fluctuations: An fMRI investigation of the resting-state default mode of brain function hypothesis. *Human Brain Mapp*, 26(1):15–29, 2005.

[80] J. Friedman, T. Hastie, and R. Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.

[81] K. Friston. Statistical parametric mapping and other analysis of functional imaging data. In *Brain Mapp Methods*, pages 363–385. Academic Press, 1996.

[82] K. Friston. Beyond phrenology: What can neuroimaging tell us about distributed circuitry? *Annual Review of Neuroscience*, 25(1):221–250, 2002.

[83] K. Friston, C. Frith, P. Liddle, and R. Frackowiak. Functional connectivity: the principal-component analysis of large (PET) data sets. *Journal of Cerebral Blood Flow & Metabolism*, 13(1):5–14, 1993.

[84] C. Garel, E. Chantrel, H. Brisse, M. Elmaleh, D. Luton, J.-F. Oury, G. Sebag, and M. Hassan. Fetal cerebral cortex: Normal gestational landmarks identified using prenatal MR imaging. *American Journal of Neuroradiology*, 22(1):184–189, 2001.

[85] G. Girard, K. Whittingstall, R. Deriche, and M. Descoteaux. Towards quantitative connectivity analysis: Reducing tractography biases. *NeuroImage*, 98:266–278, 2014.

[86] M. F. Glasser, T. S. Coalson, E. C. Robinson, C. D. Hacker, J. Harwell, E. Yacoub, K. Ugurbil, J. Andersson, C. F. Beckmann, M. Jenkinson, S. M. Smith, and D. C. Van Essen. A multi-modal parcellation of human cerebral cortex. *Nature*, 2016.

[87] M. F. Glasser, S. N. Sotiropoulos, J. A. Wilson, T. S. Coalson, B. Fischl, J. L. Andersson, J. Xu, S. Jbabdi, M. Webster, J. R. Polimeni, D. C. Van Essen, and M. Jenkinson. The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, 80:105–124, 2013.

[88] M. F. Glasser and D. C. Van Essen. Mapping human cortical areas in vivo based on myelin content as revealed by T1- and T2-weighted MRI. *J Neurosci*, 31(32):11597–11616, 2011.

[89] P. Golland, D. Lashkari, and A. Venkataraman. Spatial patterns and functional profiles for discovering structure in fMRI data. In *Signals, Systems and Computers, 2008 42nd Asilomar Conference on*, pages 1402–1409. IEEE, 2008.

[90] Y. Golland, P. Golland, S. Bentin, and R. Malach. Data-driven clustering reveals a fundamental subdivision of the human cortex into two global systems. *Neuropsychologia*, 46(2):540–553, 2008.

[91] G. Gong, Y. He, and A. C. Evans. Brain connectivity gender makes a difference. *The Neuroscientist*, 17(5):575–591, 2011.

[92] E. M. Gordon, T. O. Laumann, B. Adeyemo, A. W. Gilmore, S. M. Nelson, N. U. Dosenbach, and S. E. Petersen. Individual-specific features of brain systems identified with resting state functional correlations. *NeuroImage*, 2016.

[93] E. M. Gordon, T. O. Laumann, B. Adeyemo, J. F. Huckins, W. M. Kelley, and S. E. Petersen. Generation and evaluation of a cortical area parcellation from resting-state correlations. *Cerebral Cortex*, 26(1):288–303, 2016.

[94] C. Goutte, P. Toft, E. Rostrup, F. Nielsen, and L. K. Hansen. On clustering fMRI time series. *NeuroImage*, 9(3):298–310, 1999.

[95] R. Guillemaud and M. Brady. Estimating the bias field of mr images. *IEEE Transactions on Medical imaging*, 16(3):238–251, 1997.

[96] 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.

[97] S. Haller and A. J. Bartsch. Pitfalls in fMRI. *European Radiology*, 19(11):2689–2706, 2009.

[98] R. Heller, D. Stanley, D. Yekutieli, N. Rubin, and Y. Benjamini. Cluster-based analysis of fMRI data. *NeuroImage*, 33(2):599–608, 2006.

[99] M. Hernández, G. D. Guerrero, J. M. Cecilia, J. M. García, A. Inuggi, S. Jbabdi, T. E. Behrens, and S. N. Sotiropoulos. Accelerating fibre orientation estimation from diffusion weighted magnetic resonance imaging using GPUs. *PLoS One*, 8(4):e61892, 2013.

[100] C. J. Holmes, R. Hoge, L. Collins, R. Woods, A. W. Toga, and A. C. Evans. Enhancement of MR images using registration for signal averaging. *J Comput Assist Tomo*, 22(2):324–333, 1998.

[101] C. Honey, O. Sporns, L. Cammoun, X. Gigandet, J.-P. Thiran, R. Meuli, and P. Hagmann. Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences*, 106(6):2035–2040, 2009.

[102] N. Honnorat, H. Eavani, T. Satterthwaite, R. Gur, R. Gur, and C. Davatzikos. GraSP: Geodesic graph-based segmentation with shape priors for the functional parcellation of the cortex. *NeuroImage*, 106:207–221, 2015.

[103] L. Hubert and P. Arabie. Comparing partitions. *Journal of Classification*, 2(1):193–218, 1985.

[104] S. A. Huettel, A. W. Song, and G. McCarthy. *Functional magnetic resonance imaging*, volume 1. Sinauer Associates Sunderland, 2004.

[105] Human Connectome Project. WU-Minn HCP 900 subjects data release: Reference manual. `http://www.humanconnectome.org/documentation/S900/HCP_S900_Release_Reference_Manual.pdf`. Document created on December 8, 2015. Accessed on February 2, 2017.

[106] C. Hutton, A. Bork, O. Josephs, R. Deichmann, J. Ashburner, and R. Turner. Image distortion correction in fMRI: A quantitative evaluation. *NeuroImage*, 16(1):217–240, 2002.

[107] M. Jafri, G. Pearlson, M. Stevens, and V. Calhoun. A method for functional network connectivity among spatially independent resting-state components in schizophrenia. *NeuroImage*, 39(4):1666–1681, 2008.

[108] V. Jain and H. Zhang. Robust 3D shape correspondence in the spectral domain. In *IEEE International Conference on Shape Modeling and Applications 2006 (SMI'06)*, pages 19–19. IEEE, 2006.

[109] R. J. Janssen, P. Jylänki, and M. A. van Gerven. Lets not waste time: Using temporal information in clustered activity estimation with spatial adjacency restrictions (CAESAR) for parcellating fmri data. *PloS ONE*, 11(12):e0164703, 2016.

[110] S. Jbabdi, S. N. Sotiropoulos, A. M. Savio, M. Graña, and T. E. Behrens. Model-based analysis of multishell diffusion MR data for tractography: How to get over fitting problems. *Magnetic Resonance in Medicine*, 68(6):1846–1855, 2012.

[111] 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, 2009.

[112] R. Jenatton, A. Gramfort, V. Michel, G. Obozinski, F. Bach, and B. Thirion. Multi-scale mining of fMRI data with hierarchical structured sparsity. In *Pattern Recognition in NeuroImaging (PRNI), 2011 International Workshop on*, pages 69–72, 2011.

[113] B. Jeurissen, A. Leemans, J.-D. Tournier, D. K. Jones, and J. Sijbers. Investigating the prevalence of complex fiber configurations in white matter tissue with diffusion magnetic resonance imaging. *Hum Brain Mapp*, 34(11):2747–2766, 2013.

[114] H. Johansen-Berg, T. Behrens, M. Robson, I. Drobnjak, M. Rushworth, J. Brady, S. Smith, D. Higham, and P. Matthews. Changes in connectivity profiles define functionally distinct regions in human medial frontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 101(36):13335–13340, 2004.

[115] S. C. Johnson. Hierarchical clustering schemes. *Psychometrika*, 32(3):241–254, 1967.

[116] D. K. Jones and M. Cercignani. Twenty-five pitfalls in the analysis of diffusion MRI data. *NMR in Biomedicine*, 23(7):803–820, 2010.

[117] D. K. Jones, A. Simmons, S. C. Williams, and M. A. Horsfield. Non-invasive assessment of axonal fiber connectivity in the human brain via diffusion tensor MRI. *Magnetic Resonance in Medicine*, 42(1):37–41, 1999.

[118] T. Kahnt, L. J. Chang, S. Q. Park, J. Heinzle, and J.-D. Haynes. Connectivity-based parcellation of the human orbitofrontal cortex. *Journal of Neuroscience*, 32(18):6240–6250, 2012.

[119] E. R. Kandel, J. H. Schwartz, T. M. Jessell, and S. Mack, editors. *Principles of neural science*. McGraw-Hill Medical, New York, Chicago, San Francisco, 2013.

[120] H. Kennedy, D. C. Van Essen, and Y. Christen. *Micro-, meso-and macro-connectomics of the brain*. Springer, 2016.

[121] 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.

[122] V. Kiviniemi, J.-H. Kantola, J. Jauhiainen, A. Hyvärinen, and O. Tervonen. Independent component analysis of nondeterministic fMRI signal sources. *NeuroImage*, 19(2):253–260, 2003.

[123] V. Kiviniemi, T. Starck, J. Remes, X. Long, J. Nikkinen, M. Haapea, J. Veijola, I. Moilanen, M. Isohanni, Y.-F. Zang, and O. Tervonen. Functional segmentation of the brain cortex using high model order group PICA. *Hum Brain Mapp*, 30(12):3865–3886, 2009.

[124] 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):204–211, 2007.

[125] S. B. Klein and B. M. Thorne. *Biological psychology*. Macmillan, 2006.

[126] T. Knsche and M. Tittgemeyer. The role of long-range connectivity for the characterization of the functionalanatomical organization of the cortex. *Frontiers in Systems Neuroscience*, 5:58, 2011.

[127] B. Kolb and I. Q. Whishaw. *Fundamentals of human neuropsychology*. Macmillan, 2009.

[128] S. I. Ktena, S. Parisot, E. Ferrante, M. Rajchl, M. Lee, B. Glocker, and D. Rueckert. Distance metric learning using graph convolutional networks: Application to functional brain networks. In M. Descoteaux, L. Maier-Hein, A. Franz, P. Jannin, D. L. Collins, and S. Duchesne, editors, *MICCAI 2017 LNCS*, pages 469–477, 2017.

[129] A. R. Laird, P. M. Fox, S. B. Eickhoff, J. A. Turner, K. L. Ray, D. R. McKay, D. C. Glahn, C. F. Beckmann, S. M. Smith, and P. T. Fox. Behavioral interpretations of intrinsic connectivity networks. *Journal of Cognitive Neuroscience*, 23(12):4022–4037, 2011.

[130] J. L. Lancaster, M. G. Woldorff, L. M. Parsons, M. Liotti, C. S. Freitas, L. Rainey, P. V. Kochunov, D. Nickerson, S. A. Mikiten, and P. T. Fox. Automated Talairach atlas labels for functional brain mapping. *Hum Brain Mapp*, 10(3):120–131, 2000.

[131] G. Langs, P. Golland, and S. S. Ghosh. Predicting activation across individuals with resting-state functional connectivity based multi-atlas label fusion. In N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, editors, *MICCAI 2015. LNCS*, volume 9350, pages 313–320. Springer, Heidelberg, 2015.

[132] G. Langs, A. Sweet, D. Lashkari, Y. Tie, L. Rigolo, A. J. Golby, and P. Golland. Decoupling function and anatomy in atlases of functional connectivity patterns: Language mapping in tumor patients. *NeuroImage*, 103:462–475, 2014.

[133] G. Langs, Y. Tie, L. Rigolo, A. Golby, and P. Golland. Functional geometry alignment and localization of brain areas. In *Advances in Neural Information Processing Systems*, pages 1225–1233, 2010.

[134] G. Langs, D. Wang, P. Golland, S. Mueller, R. Pan, M. R. Sabuncu, W. Sun, K. Li, and H. Liu. Identifying shared brain networks in individuals by decoupling functional and anatomical variability. *Cerebral Cortex*, 26(10):4004–4014, 2016.

[135] D. Lashkari, E. Vul, N. Kanwisher, and P. Golland. Discovering structure in the space of fMRI selectivity profiles. *NeuroImage*, 50(3):1085–1098, 2010.

[136] T. O. Laumann, E. M. Gordon, B. Adeyemo, A. Z. Snyder, S. J. Joo, M. Y. Chen, A. W. Gilmore, K. B. McDermott, N. U. Dosenbach, B. L. Schlaggar, J. A. Mumford, R. A. Poldrack, and S. E. Petersen. Functional system and areal organization of a highly sampled individual human brain. *Neuron*, 87(3):657–670, 2015.

[137] M. Lazar and A. L. Alexander. An error analysis of white matter tractography methods: Synthetic diffusion tensor field simulations. *NeuroImage*, 20(2):1140–1153, 2003.

[138] D. Le Bihan, C. Poupon, A. Amadon, and F. Lethimonnier. Artifacts and pitfalls in diffusion MRI. *Journal of Magnetic Resonance Imaging*, 24(3):478–488, 2006.

[139] M. H. Lee, C. D. Hacker, A. Z. Snyder, M. Corbetta, D. Zhang, E. C. Leuthardt, and J. S. Shimony. Clustering of resting state networks. *PLoS ONE*, 7(7):e40370, 2012.

[140] N. Lee, A. F. Laine, and A. Klein. Towards a deep learning approach to brain parcellation. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pages 321–324. IEEE, 2011.

[141] C. M. Leonard, S. Towler, S. Welcome, L. K. Halderman, R. Otto, M. A. Eckert, and C. Chiarello. Size matters: Cerebral volume influences sex differences in neuroanatomy. *Cerebral Cortex*, 18(12):2920–2931, 2008.

[142] D. Levine and P. D. Barnes. Cortical maturation in normal and abnormal fetuses as assessed with prenatal MR imaging. *Radiology*, 210(3):751–758, 1999.

[143] F. Liem, G. Varoquaux, J. Kynast, F. Beyer, S. K. Masouleh, J. M. Huntenburg, L. Lampe, M. Rahim, A. Abraham, R. C. Craddock, S. Riedel-Heller, T. Luck, M. Loeffler, M. L. Schroeter, A. V. Witte, A. Villringer, and D. S. Margulies. Predicting brain-age from multimodal imaging data captures cognitive impairment. *NeuroImage*, 148:179 – 188, 2017.

[144] M. A. Lindquist. The statistical analysis of fMRI data. *Statistical Science*, pages 439–464, 2008.

[145] Y. Liu, M. Liang, Y. Zhou, Y. He, Y. Hao, M. Song, C. Yu, H. Liu, Z. Liu, and T. Jiang. Disrupted small-world networks in schizophrenia. *Brain*, 131(4):945–961, 2008.

[146] C.-Y. Lo, P.-N. Wang, K.-H. Chou, J. Wang, Y. He, and C.-P. Lin. Diffusion tensor tractography reveals abnormal topological organization in structural cortical networks in Alzheimer's disease. *Journal of Neuroscience*, 30(50):16876–16885, 2010.

[147] N. K. Logothetis and B. A. Wandell. Interpreting the BOLD signal. *Annu Rev Physiol*, 66:735–769, 2004.

[148] H. Lombaert, L. Grady, J. R. Polimeni, and F. Cheriet. Fast brain matching with spectral correspondence. In *Biennial International Conference on Information Processing in Medical Imaging*, pages 660–673. Springer, 2011.

[149] H. Lombaert, J. Sporring, and K. Siddiqi. Diffeomorphic spectral matching of cortical surfaces. In J. Gee, S. Joshi, K. Pohl, W. Wells, and L. Zllei, editors, *IPMI 2013*, volume 7917 of *LNCS*, pages 376–389. Springer Heidelberg, 2013.

[150] M. Lowe, B. Mock, and J. Sorenson. Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. *NeuroImage*, 7(2):119–132, 1998.

[151] Y. Lu, T. Jiang, and Y. Zang. Region growing method for the analysis of functional MRI data. *NeuroImage*, 20(1):455–465, 2003.

[152] M. D. Luca, C. Beckmann, N. D. Stefano, P. Matthews, and S. Smith. fMRI resting state networks define distinct modes of long-distance interactions in the human brain. *NeuroImage*, 29(4):1359 – 1367, 2006.

[153] J. H. Lui, D. V. Hansen, and A. R. Kriegstein. Development and evolution of the human neocortex. *Cell*, 146(1):18–36, 2011.

[154] H. Lutkepohl. Handbook of matrices. *Computational Statistics and Data Analysis*, 2(25):243, 1997.

[155] D. Mantini, M. G. Perrucci, C. Del Gratta, G. L. Romani, and M. Corbetta. Electrophysiological signatures of resting state networks in the human brain. *Proc Natl Acad Sci USA*, 104(32):13170–13175, 2007.

[156] D. S. Marcus, M. P. Harms, A. Z. Snyder, M. Jenkinson, J. A. Wilson, M. F. Glasser, D. M. Barch, K. A. Archie, G. C. Burgess, M. Ramaratnam, M. Hodge, W. Horton, R. Herrick, T. Olsen, M. McKay, M. House, M. Hileman, E. Reid, J. Harwell, T. Coalson, J. Schindler, J. S. Elam, S. W. Curtiss, and D. C. V. Essen. Human Connectome Project informatics: Quality control, database services, and data visualization. *NeuroImage*, 80(0):202–219, 2013.

[157] D. S. Margulies, A. 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.

[158] G. Marrelec, A. Krainik, H. Duffau, M. Pélégrini-Issac, S. Lehéricy, J. Doyon, and H. Benali. Partial correlation for functional brain interactivity investigation in functional MRI. *NeuroImage*, 32(1):228–237, 2006.

[159] J. C. Mazziotta, A. W. Toga, A. Evans, P. Fox, and J. Lancaster. A probabilistic atlas of the human brain: Theory and rationale for its development. The International Consortium for Brain Mapping (ICBM). *NeuroImage*, 2(2):89–101, 1995.

[160] J. C. Mazziotta, A. W. Toga, A. C. Evans, P. T. Fox, and J. L. Lancaster. Digital brain atlases. *Trends in Neurosciences*, 18(5):210–211, 1995.

[161] 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.

[162] V. Michel, A. Gramfort, G. Varoquaux, E. Eger, C. Keribin, and B. Thirion. A supervised clustering approach for fMRI-based inference of brain states. *Pattern Recog*, 45(6):2041–2049, 2012.

[163] G. W. Milligan and M. C. Cooper. A study of the comparability of external criteria for hierarchical cluster analysis. *Multivariate Behavioral Research*, 21(4):441–458, 1986.

[164] D. Moreno-Dominguez, A. Anwander, and T. R. Knösche. A hierarchical method for whole-brain connectivity-based parcellation. *Hum Brain Mapp*, 35(10):5000–5025, 2014.

[165] S. Mori. *Introduction to diffusion tensor imaging*. Elsevier, 2007.

[166] S. Mori, B. J. Crain, V. P. Chacko, and P. Van Zijl. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology*, 45(2):265–269, 1999.

[167] S. Mori and P. van Zijl. Fiber tracking: Principles and strategies–a technical review. *NMR in Biomedicine*, 15(7-8):468–480, 2002.

[168] S. Mueller, D. Wang, M. D. Fox, B. T. Yeo, J. Sepulcre, M. R. Sabuncu, R. Shafee, J. Lu, and H. Liu. Individual variability in functional connectivity architecture of the human brain. *Neuron*, 77(3):586–595, 2013.

[169] J. A. Mumford, S. Horvath, M. C. Oldham, P. Langfelder, D. H. Geschwind, and R. A. Poldrack. Detecting network modules in fMRI time series: A weighted network analysis approach. *NeuroImage*, 52(4):1465–1476, 2010.

[170] L. Nanetti, L. Cerliani, V. Gazzola, R. Renken, and C. Keysers. Group analyses of connectivity-based cortical parcellation using repeated k-means clustering. *NeuroImage*, 47(4):1666–1677, 2009.

[171] S. M. Nelson, A. L. Cohen, J. D. Power, G. S. Wig, F. M. Miezin, M. E. Wheeler, K. Velanova, D. I. Donaldson, J. S. Phillips, B. L. Schlaggar, and S. E. Petersen. A parcellation scheme for human left lateral parietal cortex. *Neuron*, 67(1):156–170, 2010.

[172] B. Ng, G. Varoquaux, J. B. Poline, and B. Thirion. Implications of inconsistencies between fMRI and dMRI on multimodal connectivity estimation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 652–659. Springer, 2013.

[173] L. Nie, P. M. Matthews, and Y. Guo. Inferring individual-level variations in the functional parcellation of the cerebral cortex. *IEEE Transactions on Biomedical Engineering*, 63(12):2505–2517, 2016.

[174] A. Nieto-Castanon, S. S. Ghosh, J. A. Tourville, and F. H. Guenther. Region of interest based analysis of functional imaging data. *NeuroImage*, 19(4):1303–1316, 2003.

[175] R. Nieuwenhuys. The myeloarchitectonic studies on the human cerebral cortex of the Vogt-Vogt school, and their significance for the interpretation of functional neuroimaging data. In *Microstructural Parcellation of the Human Cerebral Cortex*, pages 55–125. Springer, 2013.

[176] R. Nieuwenhuys, J. Voogd, and C. Van Huijzen. *The human central nervous system: A synopsis and atlas*. Springer Science & Business Media, 2007.

[177] Y. Nir, U. Hasson, I. Levy, Y. Yeshurun, and R. Malach. Widespread functional connectivity and fMRI fluctuations in human visual cortex in the absence of visual stimulation. *NeuroImage*, 30(4):1313 – 1324, 2006.

[178] J. Nolte. *The human brain: An introduction to its functional anatomy*. Mosby, 2002.

[179] K. A. Norman, S. M. Polyn, G. J. Detre, and J. V. Haxby. Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences*, 10(9):424–430, Sep 2006.

[180] S. Ogawa, T.-M. Lee, A. R. Kay, and D. W. Tank. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences*, 87(24):9868–9872, 1990.

[181] S. Ogawa, D. W. Tank, R. Menon, J. M. Ellermann, S. G. Kim, H. Merkle, and K. Ugurbil. Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging. *Proceedings of the National Academy of Sciences*, 89(13):5951–5955, 1992.

[182] S. W. Oh, J. A. Harris, L. Ng, B. Winslow, N. Cain, S. Mihalas, Q. Wang, C. Lau, L. Kuan, A. M. Henry, M. T. Mortrud, B. Ouellette, T. N. Nguyen, S. A. Sorensen, C. R. Slaughterbeck, W. Wakeman, Y. Li, D. Feng, A. Ho, E. Nicholas, K. E. Hirokawa, P. Bohn, K. M. Joines, H. Peng, M. J. Hawrylycz, J. W. Phillips, J. G. Hohmann, P. Wohnoutka, C. R. Gerfen, C. Koch, A. Bernard, C. Dang, A. R. Jones, and H. Zeng. A mesoscale connectome of the mouse brain. *Nature*, 508(7495):207–214, 2014.

[183] H. P. Op de Beeck, J. Haushofer, and N. G. Kanwisher. Interpreting fMRI data: Maps, modules and dimensions. *Nature Reviews Neuroscience*, 9(2):123–135, Feb 2008.

[184] OpenStax. Brodmann's areas of the cerebral cortex. `https://cnx.org/contents/FPtK1zmh@8.25:94lv8wHH@4/The-Central-Nervous-System`, 2016. Version 8.25 from the Textbook OpenStax Anatomy and Physiology. Published on May 18, 2016. Accessed on February 7, 2017. Used under Creative Commons Attribution 4.0 International License.

[185] A. Pandit, E. Robinson, P. Aljabar, G. Ball, I. Gousias, Z. Wang, J. Hajnal, D. Rueckert, S. Counsell, G. Montana, and A. Edwards. Whole-brain mapping of structural connectivity in infants reveals altered connection strength associated with growth and preterm birth. *Cerebral Cortex*, 24(9):2324–2333, 2014.

[186] N. G. Papadakis, K. M. Martin, J. D. Pickard, L. D. Hall, T. A. Carpenter, and C. L.-H. Huang. Gradient preemphasis calibration in diffusion-weighted echo-planar imaging. *Magnetic Resonance in Medicine*, 44(4):616–624, 2000.

[187] S. Parisot, S. Arslan, J. Passerat-Palmbach, W. M. Wells, and D. Rueckert. Tractography-driven groupwise multi-scale parcellation of the cortex. In S. Ourselin, D. Alexander, C.-F. Westin, and M. Cardoso, editors, *IPMI 2015 LNCS*, pages 600–612, 2015.

[188] S. Parisot, S. Arslan, J. Passerat-Palmbach, W. M. Wells, and D. Rueckert. Group-wise parcellation of the cortex through multi-scale spectral clustering. *NeuroImage*, 2016.

[189] S. Parisot, B. Glocker, S. I. Ktena, S. Arslan, M. D. Schirmer, and D. Rueckert. A flexible graphical model for multi-modal parcellation of the cortex. *NeuroImage*, 162:226–248, 2017.

[190] S. Parisot, B. Glocker, M. D. Schirmer, and D. Rueckert. Grampa: Graph-based multi-modal parcellation of the cortex using fusion moves. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 148–156. Springer International Publishing, 2016.

[191] S. Parisot, S. I. Ktena, E. Ferrante, M. Lee, R. G. Moreno, B. Glocker, and D. Rueckert. Spectral graph convolutions for population-based disease prediction. In M. Descoteaux, L. Maier-Hein, A. Franz, P. Jannin, D. L. Collins, and S. Duchesne, editors, *MICCAI 2017 LNCS*, pages 177–185, 2017.

[192] S. Parisot, M. Rajchl, J. Passerat-Palmbach, and D. Rueckert. A continuous flow-maximisation approach to connectivity-driven cortical parcellation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 165–172. Springer, 2015.

[193] R. E. Passingham, K. E. Stephan, and R. Kötter. The anatomical basis of functional localization in the cortex. *Nature Reviews Neuroscience*, 3(8):606–616, 2002.

[194] R. S. Patel, D. Borsook, and L. Becerra. Modulation of resting state functional connectivity of the brain by naloxone infusion. *Brain Imaging Behav*, 2(1):11–20, 2008.

[195] J. Pearce. Brodmann's cortical maps. *Journal of neurology, neurosurgery, and psychiatry*, 76(2):259, 2005.

[196] J. Peng, P. Wang, N. Zhou, and J. Zhu. Partial correlation estimation by joint sparse regression models. *Journal of the American Statistical Association*, 104(486):735–746, 2009.

[197] W. D. Penny, K. J. Friston, J. T. Ashburner, S. J. Kiebel, and T. E. Nichols. *Statistical parametric mapping: The analysis of functional brain images*. Academic Press, 2011.

[198] F. Pereira, T. Mitchell, and M. Botvinick. Machine learning classifiers and fMRI: A tutorial overview. *NeuroImage*, 45(1):S199–S209, 2009.

[199] S. M. Plis, D. R. Hjelm, R. Salakhutdinov, E. A. Allen, H. J. Bockholt, J. D. Long, H. J. Johnson, J. S. Paulsen, J. A. Turner, and V. D. Calhoun. Deep learning for neuroimaging: A validation study. *Frontiers in neuroscience*, 8, 2014.

[200] R. A. Poldrack, J. A. Mumford, and T. E. Nichols. *Handbook of functional MRI data analysis*. Cambridge University Press, 2011.

[201] M. Poustchi-Amin, S. A. Mirowitz, J. J. Brown, R. C. McKinstry, and T. Li. Principles and applications of echo-planar imaging: A review for the general radiologist 1. *Radiographics*, 21(3):767–779, 2001.

[202] J. D. Power, A. L. Cohen, S. M. Nelson, G. S. Wig, K. A. Barnes, J. A. Church, A. C. Vogel, T. O. Laumann, F. M. Miezin, B. L. Schlaggar, and S. E. Petersen. Functional network organization of the human brain. *Neuron*, 72(4):665–678, 2011.

[203] A. Qiu, A. Lee, M. Tan, and M. K. Chung. Manifold learning on brain functional networks in aging. *Medical Image Analysis*, 20(1):52–60, 2015.

[204] E. C. Robinson, A. Hammers, A. Ericsson, A. D. Edwards, and D. Rueckert. Identifying population differences in whole-brain structural networks: A machine learning approach. *NeuroImage*, 50(3):910–919, 2010.

[205] E. C. Robinson, S. Jbabdi, M. F. Glasser, J. Andersson, G. C. Burgess, M. P. Harms, S. M. Smith, D. C. Van Essen, and M. Jenkinson. MSM: A new flexible framework for multimodal surface matching. *NeuroImage*, 100:414–426, 2014.

[206] E. C. Robinson, M. Valstar, A. Hammers, A. Ericsson, A. D. Edwards, and D. Rueckert. Multivariate statistical analysis of whole brain structural networks obtained using probabilistic tractography. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 486–493. Springer, 2008.

[207] 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. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 935–942. Springer, 2009.

[208] P. J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53 – 65, 1987.

[209] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500):2323–2326, 2000.

[210] S. Ryali, T. Chen, K. Supekar, and V. Menon. A parcellation scheme based on von Mises-Fisher distributions and Markov random fields for segmenting brain regions using resting-state fMRI. *NeuroImage*, 65:83–96, 2013.

[211] G. Salimi-Khorshidi, G. Douaud, C. F. Beckmann, M. F. Glasser, L. Griffanti, and S. M. Smith. Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical fusion of classifiers. *NeuroImage*, 90:449–468, 2014.

[212] R. Salvador, J. Suckling, M. R. Coleman, J. D. Pickard, D. Menon, and E. Bullmore. Neurophysiological architecture of functional magnetic resonance images of human brain. *Cerebral Cortex*, 15(9):1332–1342, 2005.

[213] T. D. Satterthwaite, D. H. Wolf, D. R. Roalf, K. Ruparel, G. Erus, S. Vandekar, E. D. Gennatas, M. A. Elliott, A. Smith, H. Hakonarson, R. Verma, C. Davatzikos, R. E. Gur, and R. C. Gur. Linked sex differences in cognition and functional connectivity in youth. *Cereb Cortex*, 25(9):2383–2394, 2015.

[214] M. D. Schirmer. *Developing brain connectivity: Effects of parcellation scale on network analysis in neonates*. PhD thesis, King's College London, 2015.

[215] Z. Shehzad, A. M. C. Kelly, P. T. Reiss, D. G. Gee, K. Gotimer, L. Q. Uddin, S. H. Lee, D. S. Margulies, A. K. Roy, B. B. Biswal, E. Petkova, F. X. Castellanos, and M. P. Milham. The resting brain: Unconstrained yet reliable. *Cereb Cortex*, 19(10):2209–2229, 2009.

[216] X. Shen and F. G. Meyer. Nonlinear dimension reduction and activation detection for fMRI dataset. In *Computer Vision and Pattern Recognition Workshop, 2006. CVPRW'06. Conference on*, pages 90–90. IEEE, 2006.

[217] X. Shen, X. Papademetris, and R. Constable. Graph-theory based parcellation of functional subunits in the brain from resting-state fMRI data. *NeuroImage*, 50(3):1027–1035, 2010.

[218] X. Shen, F. Tokoglu, X. Papademetris, and R. T. Constable. Groupwise whole-brain parcellation from resting-state fMRI data for network node identification. *NeuroImage*, 82:403–415, 2013.

[219] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans Pattern Anal Mach Intell*, 22(8):888–905, 2000.

[220] R. Sladky, K. J. Friston, J. Tröstl, R. Cunnington, E. Moser, and C. Windischberger. Slice-timing effects and their correction in functional mri. *NeuroImage*, 58(2):588–594, 2011.

[221] S. Smith. Linking cognition to brain connectivity. *Nature Neuroscience*, 19(1):7–9, 2016.

[222] S. M. Smith. Overview of fMRI analysis. *The British Journal of Radiology*, 77:S167–S175, 2004.

[223] S. M. Smith, C. F. Beckmann, J. Andersson, E. J. Auerbach, J. Bijsterbosch, G. Douaud, E. Duff, D. A. Feinberg, L. Griffanti, M. P. Harms, M. Kelly, T. Laumann, K. L. Miller, S. Moeller, S. Petersen, J. Power, G. Salimi-Khorshidi, A. Z. Snyder, A. T. Vu, M. W. Woolrich, J. Xu, E. Yacoub, K. Uurbil, D. C. V. Essen, and M. F. Glasser. Resting-state fMRI in the Human Connectome Project. *NeuroImage*, 80:144 – 168, 2013.

[224] 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. *Proceedings of the National Academy of Sciences*, 106(31):13040–13045, 2009.

[225] S. M. Smith, A. Hyvärinen, G. Varoquaux, K. L. Miller, and C. F. Beckmann. Group-PCA for very large fMRI datasets. *NeuroImage*, 101:738–749, 2014.

[226] 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*, 54(2):875–891, 2011.

[227] S. M. Smith, D. Vidaurre, C. F. Beckmann, M. F. Glasser, M. Jenkinson, K. L. Miller, T. E. Nichols, E. C. Robinson, G. Salimi-Khorshidi, M. W. Woolrich, D. M. Barch, K. Uurbil, and D. C. V. Essen. Functional connectomics from resting-state fMRI. *Trends Cogn Sci*, 17(12):666–682, 2013.

[228] J. Soares, P. Marques, V. Alves, and N. Sousa. A hitchhiker's guide to diffusion tensor imaging. *Frontiers in neuroscience*, 7:31, 2013.

[229] R. R. Sokal. A statistical method for evaluating systematic relationships. *Univ Kans Sci Bull*, 38:1409–1438, 1958.

[230] S. N. Sotiropoulos, S. Jbabdi, J. Xu, J. L. Andersson, S. Moeller, E. J. Auerbach, M. F. Glasser, M. Hernandez, G. Sapiro, M. Jenkinson, D. A. Feinberg, E. Yacoub, C. Lenglet,

D. C. V. Essen, K. Ugurbil, and T. E. Behrens. Advances in diffusion {MRI} acquisition and processing in the human connectome project. *NeuroImage*, 80:125 – 143, 2013.

[231] O. Sporns. The human connectome: A complex network. *Ann N Y Acad Sci*, 1224(1):109–125, 2011.

[232] O. Sporns. Structure and function of complex brain networks. *Dialogues Clin Neurosci*, 15(3):247–262, 2013.

[233] O. Sporns, G. Tononi, and R. Ktter. The Human Connectome: A structural description of the human brain. *PLoS Comput Biol*, 1(4):e42, 2005.

[234] C. Stam, B. Jones, G. Nolte, M. Breakspear, and P. Scheltens. Small-world networks and functional connectivity in Alzheimer's disease. *Cerebral Cortex*, 17(1):92–99, 2007.

[235] E. O. Stejskal and J. E. Tanner. Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. *Journal of Chemical Physics*, 42(1):288–292, 1965.

[236] F. L. Stevens, R. A. Hurley, K. H. Taber, R. A. Hurley, L. A. Hayman, and K. H. Taber. Anterior cingulate cortex: Unique role in cognition and emotion. *The Journal of Neuropsychiatry and Clinical Neurosciences*, 23(2):121–125, 2011.

[237] F. T. Sun, L. M. Miller, and M. D'Esposito. Measuring interregional functional connectivity using coherence and partial coherence analyses of fMRI data. *NeuroImage*, 21(2):647–658, 2004.

[238] K. Supekar, V. Menon, D. Rubin, M. Musen, and M. D. Greicius. Network analysis of intrinsic functional brain connectivity in Alzheimer's disease. *PLoS Comput Biol*, pages 1–11, 2008.

[239] J. Talairach and P. Tournoux. *Co-planar Stereotaxic Atlas of the Human Brain: 3-dimensional Proportional System: An Approach to Cerebral Imaging*. Thieme Publishers Series. Thieme, 1988.

[240] I. Tavor, O. P. Jones, R. B. Mars, S. M. Smith, T. E. Behrens, and S. Jbabdi. Task-free MRI predicts individual differences in brain activity during task performance. *Science*, 352(6282):216–220, 2016.

[241] J. B. Tenenbaum, V. De Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.

[242] B. Thirion and O. Faugeras. Nonlinear dimension reduction of fmri data: The Laplacian embedding approach. In *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on*, pages 372–375. IEEE, 2004.

[243] B. Thirion, G. Flandin, P. Pinel, A. Roche, P. Ciuciu, and J. B. Poline. Dealing with the shortcomings of spatial normalization: multi-subject parcellation of fMRI datasets. *Hum Brain Mapp*, 27(8):678–693, 2006.

[244] B. Thirion, G. Varoquaux, E. Dohmatob, and J. B. Poline. Which fMRI clustering gives good brain parcellations? *Front Neurosci*, 8:167, 2014.

[245] B. T. Thomas Yeo, F. M. Krienen, J. Sepulcre, M. R. Sabuncu, D. Lashkari, M. Hollinshead, J. L. Roffman, J. W. Smoller, L. Zöllei, J. R. Polimeni, B. Fischl, H. Liu, and R. L. Buckner. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J Neurophysiol*, 106(3):1125–1165, 2011.

[246] L. Tian, J. Wang, C. Yan, and Y. He. Hemisphere-and gender-related differences in small-world brain networks: A resting-state functional MRI study. *NeuroImage*, 54(1):191–202, 2011.

[247] B. Tijms, C. Möller, H. Vrenken, A. Wink, W. de Haan, W. van der Flier, C. Stam, P. Scheltens, and F. Barkhof. Single-subject grey matter graphs in Alzheimer's disease. *PloS ONE*, 8(3):e58921, 2013.

[248] A. W. Toga. *Brain mapping: An encyclopedic reference*. Academic Press, 2015.

[249] A. W. Toga, K. A. Clark, P. M. Thompson, D. W. Shattuck, and J. D. Van Horn. Mapping the human connectome. *Neurosurgery*, 71(1):1–5, 2012.

[250] 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):10259–10269, 2007.

[251] D. S. Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52(6):1358–1372, 2004.

[252] D. S. Tuch, T. G. Reese, M. R. Wiegell, N. Makris, J. W. Belliveau, and V. J. Wedeen. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magnetic resonance in medicine*, 48(4):577–582, 2002.

[253] A. Tucholka, B. Thirion, M. Perrot, P. Pinel, J.-F. Mangin, and J.-B. Poline. Probabilistic anatomo-functional parcellation of the cortex: How many regions? In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 399–406. Springer, 2008.

[254] N. Tzourio-Mazoyer, B. Landeau, D. Papathanassiou, F. Crivello, O. Etard, N. Delcroix, B. Mazoyer, and M. Joliot. Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage*, 15(1):273–289, 2002.

[255] M. van den Heuvel, R. Mandl, and H. Hulshoff Pol. Normalized cut group clustering of resting-state fMRI data. *PLoS ONE*, 3(4):e2001, 2008.

[256] M. P. Van Den Heuvel, R. C. Mandl, R. S. Kahn, H. Pol, and E. Hilleke. Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain. *Hum Brain Mapp*, 30(10):3127–3141, 2009.

[257] M. P. Van Den Heuvel and H. E. H. Pol. Exploring the brain network: A review on resting-state fMRI functional connectivity. *European Neuropsychopharmacology*, 20(8):519–534, 2010.

[258] K. R. Van Dijk, T. Hedden, A. Venkataraman, K. C. Evans, S. W. Lazar, and R. L. Buckner. Intrinsic functional connectivity as a tool for human connectomics: Theory, properties, and optimization. *Journal of Neurophysiology*, 103(1):297–321, 2010.

[259] D. C. Van Essen, M. F. Glasser, D. L. Dierker, J. Harwell, and T. Coalson. Parcellations and hemispheric asymmetries of human cerebral cortex analyzed on surface-based atlases. *Cerebral Cortex*, 22(10):2241–2262, 2012.

[260] D. C. Van Essen, S. Jbabdi, S. N. Sotiropoulos, C. Chen, K. Dikranian, T. Coalson, J. Harwell, T. E. Behrens, and M. F. Glasser. Mapping connections in humans and nonhuman primates: aspirations and challenges for diffusion imaging. *Diffusion MRI, 2nd edition (eds. Johansen-Berg, H. & Behrens, TEJ)*, pages 337–358, 2013.

[261] D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub, and K. Ugurbil. The WU-Minn Human Connectome Project: An overview. *NeuroImage*, 80:62 – 79, 2013.

[262] G. Varoquaux and R. C. Craddock. Learning and comparing functional connectomes across subjects. *NeuroImage*, 80:405–415, 2013.

[263] G. Varoquaux, A. Gramfort, F. Pedregosa, V. Michel, and B. Thirion. Multi-subject dictionary learning to segment an atlas of brain spontaneous activity. *IPMI*, 6801:562–573, 2011.

[264] G. Varoquaux, S. Sadaghiani, P. Pinel, A. Kleinschmidt, J.-B. Poline, and B. Thirion. A group model for stable multi-subject ICA on fMRI datasets. *NeuroImage*, 51(1):288–299, 2010.

[265] S. Vergun, A. S. Deshpande, T. B. Meier, J. Song, D. L. Tudorascu, V. A. Nair, V. Singh, B. B. Biswal, M. E. Meyerand, R. M. Birn, and V. Prabhakaran. Characterizing functional connectivity differences in aging adults using machine learning on resting state fMRI data. *Frontiers in Computational Neuroscience*, 7:38, 2013.

[266] S. Vieira, W. H. Pinaya, and A. Mechelli. Using deep learning to investigate the neuroimaging correlates of psychiatric and neurological disorders: Methods and applications. *Neuroscience & Biobehavioral Reviews*, 2017.

[267] L. Vincent and P. Soille. Watersheds in digital spaces: An efficient algorithm based on immersion simulations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(6):583–598, 1991.

[268] N. X. Vinh, J. Epps, and J. Bailey. Information theoretic measures for clusterings comparison: Is a correction for chance necessary? In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 1073–1080. ACM, 2009.

[269] C. Vogt and O. Vogt. *Allgemeine ergebnisse unserer hirnforschung*, volume 21. JA Barth, 1919.

[270] K. von Economo and G. Koskinas. *Die Cytoarchitektonik der Hirnrinde des erwachsenen Menschen*. Springer, Berlin, 1925.

[271] U. Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.

[272] D. Wang, R. L. Buckner, M. D. Fox, D. J. Holt, A. J. Holmes, S. Stoecklein, G. Langs, R. Pan, T. Qian, K. Li, J. T. Baker, S. M. Stufflebeam, K. Wang, X. Wang, B. Hong, and H. Liu. Parcellating cortical functional networks in individuals. *Nat Neurosci*, 18(12):1853–1860, 2015.

[273] D. Wang and H. Liu. Functional connectivity architecture of the human brain not all the same. *The Neuroscientist*, 20(5):432–438, 2014.

[274] J. Wang, L. Wang, Y. Zang, H. Yang, H. Tang, Q. Gong, Z. Chen, C. Zhu, and Y. He. Parcellation-dependent small-world brain functional networks: a resting-state fMRI study. *Hum Brain Mapp*, 30(5):1511–1523, 2009.

[275] L. Wang, C. Zhu, Y. He, Y. Zang, Q. Cao, H. Zhang, Q. Zhong, and Y. Wang. Altered small-world brain functional networks in children with attention-deficit/hyperactivity disorder. *Hum Brain Mapp*, 30(2):638–649, 2009.

[276] J. H. Ward. Hierarchical grouping to optimize an objective function. *J Amer Statist Assoc*, 58(301):236–244, 1963.

[277] V. J. Wedeen, P. Hagmann, W.-Y. I. Tseng, T. G. Reese, and R. M. Weisskoff. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magnetic Resonance in Medicine*, 54(6):1377–1386, 2005.

[278] V. J. Wedeen, R. Wang, J. D. Schmahmann, T. Benner, W. Tseng, G. Dai, D. Pandya, P. Hagmann, H. D'Arceuil, and A. J. de Crespigny. Diffusion spectrum magnetic resonance imaging (DSI) tractography of crossing fibers. *NeuroImage*, 41(4):1267–1277, 2008.

[279] G. S. Wig, T. O. Laumann, A. L. Cohen, J. D. Power, S. M. Nelson, M. F. Glasser, F. M. Miezin, A. Z. Snyder, B. L. Schlaggar, and S. E. Petersen. Parcellating an individual subject's cortical and subcortical brain structures using snowball sampling of resting-state correlations. *Cereb Cortex*, pages 2036–2054, 2013.

[280] G. S. Wig, T. O. Laumann, and S. E. Petersen. An approach for parcellating human cortical areas using resting-state correlations. *NeuroImage*, 93:276–291, 2014.

[281] Wikimedia Commons. The human brain and skull. `https://commons.wikimedia.org/wiki/File:Skull_and_brain_normal_human.svg`, 2006. Accessed on January 21, 2017. Used under Creative Commons Attribution 2.5 License 2006 / cropped and removed background.

[282] Wikimedia Commons. The tensor field in a diffusion MRI image. `https://en.wikipedia.org/wiki/File:DiffusionMRI_glyphs.png`, 2012. Accessed on February 4, 2017. Used under Creative Commons Attribution-ShareAlike 3.0 Licence / cropped and removed background.

[283] Wikimedia Commons. Magnetic resonance imaging (MRI). `https://commons.wikimedia.org/wiki/File:Blausen_0654_MRI_01.png`, 2013. Accessed on February 4, 2017. Used under Creative Commons Attribution 3.0 Unported Licence / cropped and removed background.

[284] Z. Wu and R. Leahy. An optimal graph theoretic approach to data clustering: Theory and its application to image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(11):1101–1113, 1993.

[285] S. X. Yu and J. Shi. Multiclass spectral clustering. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pages 313–319. IEEE, 2003.

[286] K. Zilles and K. Amunts. Receptor mapping: Architecture of the human cerebral cortex. *Curr Opin Neurol*, 22(4):331–339, 2009.

[287] K. Zilles, K. Amunts, and K. Brodmann. Centenary of Brodmann's map–conception and fate. *Nat Rev Neurosci*, 11(2):139–145, 2010.

[288] K. Zilles and N. Palomero-Gallagher. Cyto-, myelo-, and receptor architectonics of the human parietal cortex. *NeuroImage*, 14(1):S8–S20, 2001.

[289] K. Zilles, N. Palomero-Gallagher, and K. Amunts. Development of cortical folding during evolution and ontogeny. *Trends in Neurosciences*, 36(5):275–284, 2013.

[290] K. Zilles, N. Palomero-Gallagher, and A. Schleicher. Transmitter receptors and functional anatomy of the cerebral cortex. *J Anat*, 205(6):417–432, 2004.

[291] K. Zilles, A. Schleicher, N. Palomero-Gallagher, and K. Amunts. Quantitative analysis of cyto- and receptor architecture of the human brain. In A. W. Toga and J. C. Mazziotta, editors, *Brain Mapping: The Methods*, pages 573–602. Academic Press, 2002.