THE UNIVERSITY *of* EDINBURGH

Structural brain networks from diffusion MRI: methods and application

Colin R. Buchanan

Doctor of Philosophy

Institute for Adaptive and Neural Computation
School of Informatics
University of Edinburgh

2015

Abstract

Structural brain networks can be constructed at a macroscopic scale using diffusion magnetic resonance imaging (dMRI) and whole-brain tractography. Under this approach, grey matter regions, such as Brodmann areas, form the nodes of a network and tractography is used to construct a set of white matter fibre tracts which form the connections. Graph-theoretic measures may then be used to characterise patterns of connectivity.

In this study, we measured the test-retest properties of such networks by varying several factors affecting network construction using ten healthy volunteers who underwent a dMRI protocol at 1.5 T on two separate occasions. High resolution T_1-weighted brains were parcellated into regions-of-interest and network connections were identified using dMRI and two alternative tractography algorithms, two alternative seeding strategies, constraints on anatomical plausibility and three alternative network weightings. Test-retest performance was found to improve when: 1) seeding from white matter, rather than grey; and 2) using probabilistic tractography, rather than deterministic. In terms of network weighting, a measure of streamline density produced better test-retest performance than tract-averaged diffusion anisotropy, although it remains unclear which is most representative of the underlying axonal connections.

These findings were then used to inform network construction for two further cohorts: a case-control analysis of 30 patients with amyotrophic lateral sclerosis (ALS) compared with 30 age-matched healthy controls; and a cross-sectional analysis of 80 healthy volunteers aged 25–64 years. In both cases, networks were constructed using a weighting reflecting tract-averaged fractional anisotropy (FA). A mass-univariate statistical technique called network-based statistics, identified an impaired motor-frontal-subcortical subnetwork (10 nodes and 12 bidirectional connections), consistent with upper motor neuron pathology, in the ALS group compared with the controls. Reduced FA for three of the impaired network connections, which involved fibres of the cortico-spinal tract, were significantly correlated with the rate of disease progression. Cross-sectional analysis of the 80 healthy volunteers was intended to provide supporting evidence for the widely reported age-related decline in white matter integrity. However, no meaningful relationships were found between increasing age and impaired connectivity based on global, lobar and nodal network properties – findings which were confirmed with a conventional voxel-based analysis of the dMRI data.

In conclusion, whilst current acquisition protocols and methods can produce networks capable of characterising the genuine between-subject differences in connectivity, it is challenging to measure subtle white matter changes, for example, due to normal ageing. We conclude that future work should be undertaken to address these concerns.

Lay Summary

The structural organisation of the human brain can be mapped using diffusion magnetic resonance imaging (dMRI) and a computational technique called tractography. Brain organisation can be represented as a network where distinct regions of the cerebral cortex form the network nodes. Tractography can map the connections between these nodes by estimating the collective wiring of millions of axonal connections between nerve cells.

In this study, we measured the reliability of such networks by varying several factors affecting network construction using ten healthy human volunteers who underwent a dMRI protocol on two separate occasions. High resolution scans of the brain were automatically divided into regions-of-interest and network connections were identified using dMRI and two alternative tractography approaches, two alternative tractography seeding strategies, constraints on anatomical plausibility and three alternative network weightings. Network reliability was found to improve when: 1) seeding from the white matter of the brain, rather than the grey matter; and 2) using probabilistic tractography, rather than deterministic. However, in terms of network weightings it remained unclear which provided the most accurate representation of the biological connections within the brain.

These findings were then used to inform network construction for two further groups: an analysis of 30 patients with amyotrophic lateral sclerosis (ALS) compared with 30 healthy controls; and an analysis of 80 healthy volunteers aged 25–64 years. A statistical technique called network-based statistics identified an impaired motor network in the ALS group when compared with the controls. This network was consistent with the understanding of the disease and was localised around brain regions known to be associated with motor control and movement. Additionally, three of the impaired network connections, which involved fibres of the cortico-spinal tract, were significantly correlated with the rate of disease progression. Analysis of the 80 healthy volunteers was intended to provide supporting evidence for the widely reported age-related decline in brain organisation. However, no meaningful relationships were found between increasing age and impaired network connectivity.

In conclusion, whilst these methods can produce networks capable of characterising large differences in connectivity from individual to individual, it remains challenging to measure subtle differences in connectivity, for example, due to normal ageing. We conclude that future work should be undertaken to address these concerns.

Acknowledgements

Thanks to Mark for his insight, guidance and perseverance throughout for which I am extremely grateful. Thanks also to Amos for offering his wisdom and advice. Gratitude is also due to: Cyril for his enthusiasm and hard work; Sharon and Lewis for sharing their findings and expertise; many of the current and former students of the Neuroinformatics and Computational Neuroscience Doctoral Training Centre, especially Krzysztof and Ksenia (Kate), for their interest, help and conversation; and those who I worked alongside at the Western General Hospital, Edinburgh. Special thanks is reserved for my Mum and Dad for their support, and to Pamela for her understanding.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(*Colin R. Buchanan*)

Table of Contents

List of Figures

List of Tables

List of Acronyms and Abbreviations

AAL	Automatic anatomical labelling
ACM	Anatomical connectivity map
AD	Alzheimer's disease
ADC	Apparent diffusion coefficient
ALS	Amyotrophic lateral sclerosis
ALSFRS-R	ALS functional rating scale - revised
CNS	Central nervous system
CSF	Cerebrospinal fluid
CV	Coefficient of variation
dMRI	Diffusion magnetic resonance imaging
DOF	Degrees of freedom
DSI	Diffusion spectrum imaging
DTI	Diffusion tensor imaging
DTK	Diffusion toolkit
DWI	Diffusion-weighted imaging
EPI	Echo-planar imaging
FA	Fractional anisotropy
FACT	Fibre assignment continuous tracking
FDT	FSL diffusion toolkit
FLIRT	FMRIB linear image registration tool
fMRI	Functional magnetic resonance imaging
FNIRT	FMRIB nonlinear image registration tool
FSL	FMRIB software library
FTD	Frontotemporal dementia
GM	Grey matter
HARDI	High-angular-resolution diffusion imaging
HC	Healthy control
HO	Harvard-Oxford (atlas)
IBASPM	Individual brain atlases for statistical parametric mapping
ICC	Intraclass correlation coefficient
IQ	Intelligence quotient
LPBA40	LONI probabilistic brain atlas of 40 subjects
MCI	Mild cognitive impairment
MCMC	Markov Chain Monte Carlo
MD	Mean diffusivity
MRI	Magnetic resonance imaging
MS	Multiple sclerosis
NBS	Network-based statistics

NMR	Nuclear magnetic resonance
ODF	Orientation diffusion function
PDF	probability density function
PICo	Probabilistic index of connectivity
RF	Radio frequency
ROI	Region-of-interest
RSN	Resting-state network
SD	Streamline density
SDL	Streamline density length corrected
SPM	Statistical parametric mapping
TBSS	Tract-based spatial statistics
TE	Echo time
TEND	Tensor deflection
TR	Repetition time
WFU	Wake Forest University (atlas)
WM	White matter
WTAR	Wechsler test of adult reading

Chapter 1

Introduction

Connectomics is the study of connectomes, which map the connections of the brain as a network. The ultimate aim of the connectome approach is, perhaps, a full cellular scale map of the human brain, outlining each and every neuron, axon and synapse. However, the technologies to achieve this resolution in mammals have not yet been realised and for now only gross approximations at a macroscopic scale are possible *in vivo* with techniques, such as magnetic resonance imaging (MRI) and tractography. Nevertheless, these techniques have the potential to map the collective wiring of many trillions of axonal fibres and may provide clues on how cerebral white matter structure correlates with function, cognition and behaviour. Although maps of brain connectivity have previously been demonstrated in human volunteers and are widely used to measure patterns of connectivity, currently there are a lack of studies which assess the reliability of the resulting networks. In this work we explored methods for constructing such networks by imaging healthy human volunteers and assessing various aspects affecting reliability. These findings were then used to inform analyses of structural brain connectivity associated with both a neurodegenerative condition and with normal ageing.

The organisation of this thesis is as follows. Chapter 2 begins with a description of brain structure, the basic principles of MRI and diffusion MRI (dMRI), a description of the connectome and an overview of current techniques used to construct structural networks. This leads us to the shortcomings of current approaches and a motivation for this work. Chapter 3 describes our methodology, including MRI acquisition, segmentation of high-resolution structural MRI data, image registration, dMRI processing, tractography and network construction and analysis. Chapter 4 outlines the analysis and findings from building structural networks and assessing the test-retest reliability. We assessed repeat scans of ten healthy volunteers by varying several factors affecting the construction of networks. We compared two alternative tractography algorithms (deterministic and probabilistic), two seeding approaches (grey and white matter), and three alternative network weightings (streamline density, streamline density with length

correction and a measure of tract-averaged diffusion anisotropy). Chapter 5 describes an analysis of the connectivity between brain regions in a group of 30 amyotrophic lateral sclerosis (ALS) patients when compared with a group of age-matched healthy controls. An established statistical technique called network-based statistics (NBS) was used, without *a priori* selected regions, to identify a subnetwork of reduced connectivity in the ALS group compared with the controls. Chapter 6 describes the findings of a structural network analysis in a group of 80 healthy volunteers undergoing normal ageing, between 25 and 64 years of age. Such a cross-sectional analysis has the potential to find supporting evidence, in terms of brain organisation, for the widely reported age-related decline of white matter integrity and we carried out a number of statistical analyses to assess the relationships between age and brain structure. Finally, Chapter 7 presents a discussion of the findings from the three results chapters. We focus on the relative merits and limitations of using such networks obtained from dMRI and tractography in cohorts of healthy volunteers, ageing and disease. This is followed by a discussion of the current challenges facing connectome approaches and some possible avenues of future work.

To our knowledge, this work was the first to directly compare probabilistic and deterministic tractography and alternative seeding strategies in a test-retest network analysis. Our key findings were: 1) Probabilistic tractography was found to perform better than a deterministic method; 2) Performance was improved when seeding from white matter, rather than grey. 3) Thresholding of network weights must be applied with caution in order to remove spurious connections while retaining genuine patterns of connectivity. Whilst current acquisition protocols and methods can produce networks capable of characterising the genuine between-subject differences in connectivity for gross structural differences in ALS, it is challenging to measure subtle white matter changes, for example, due to normal ageing. Overall, our results highlight concerns with reliably measuring nodal network properties. We conclude that future work should be undertaken to address these concerns.

1.1 Notation and terminology

Note that in this work, we are chiefly concerned with the structural network of the adult human brain imaged through dMRI – other connectome technologies and the nervous systems of other organisms are discussed where appropriate. We primarily use the term *network* to refer to a map of brain organisation, but note that the terms *connectome* and *graph* can be considered equivalent and are used interchangeably on occasion. Except where specified, we use the notation, mean $\pm$ standard deviation, in our results.

Chapter 2

Background

2.1 Overview

The human brain is, by most accounts, the most complex known object, formed from an extremely vast and largely unmapped network of interconnected nerve cells. Neuroscience seeks to understand how neurons within this network give rise to neural processing, perception, learning and memory and ultimately higher level cognitive functions and behaviours. Connectomics is a recent development[1] in neuroscience concerned with mapping the structural and functional connections within the brain at various scales (Hagmann, 2005; Sporns et al., 2005). For instance, the structural connectome may be explored at a macroscopic scale *in vivo* through dMRI and whole-brain tractography. Under this approach, segmented cortical areas, such as Brodmann areas, form the nodes of a network and tractography is used to construct a set of white matter fibre tracts which form the connections. Graph-theoretic measures may then be used to characterise patterns of connectivity. Such techniques have the potential to map the collective wiring of many trillions of axonal fibres and may provide detailed information on how cerebral white matter structure correlates with function, and potentially dysfunction, in health and disease.

The organisation of this chapter is as follows. Firstly, a description of the nervous system and the organisation of the brain is outlined. The basic principles of MRI and dMRI are then outlined and the modalities required for producing a structural connectome are described. An overview of tractography techniques is then presented, followed by a description of the connectome and current techniques used in network construction. Finally, this leads us to the shortcomings of the current approaches and a motivation for this work.

[1] or more precisely, a comprehensive reinterpretation of an old idea (Catani et al., 2013).

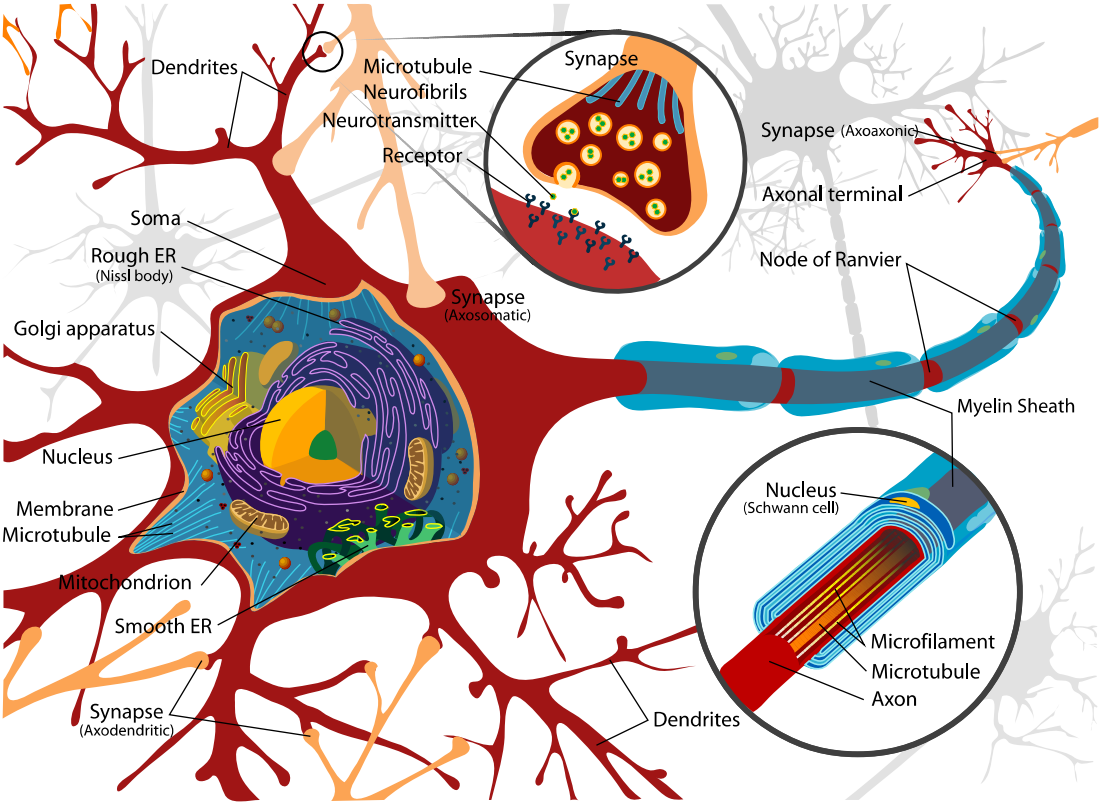

Figure 2.1: Diagram of a typical neuron. (Adapted from Wikimedia Commons, `http://commons.wikimedia.org`.)

2.2 Organisation of the brain

This section outlines some basic concepts of neuroscience and neural connections. For a more comprehensive introduction to the neurosciences refer to Bear et al. (2007) or Purves et al. (2012).

Modern neuroscience is based on the *neuron doctrine*, which has established that the nervous system is formed of discrete interconnected cells (Jones, 1999). *Neurons* (nerve cells) are the electrically excitable cells within the nervous system involved with neural processing and transmission of information. The human brain contains approximately 90 billion neurons and it is estimated that on average each neuron may form *synapses* with up to ten thousand other neurons (Herculano-Houzel, 2009) to create an astonishingly complex computational network. A typical neuron (Figure 2.1) comprises a cell body (soma) and two types of projection called *dendrites* and *axons*. A signal from a 'firing' neuron is caused by the generation of an electrical signal from within the soma. The axon, a cable-like projection originating from the cell body, carries the electrical signal along its length. Axons vary considerably in length, from a fraction of a millimetre to many centimetres. Lengthier axonal projections are typically covered

by layers of *myelin*, a fatty dielectric insulating sheath, which serves to efficiently propagate electrical signals along the axon's length. It is the microscopic structure of this fatty tissue, composed largely of lipids, water and proteins, that permits dMRI to probe axonal organisation. The diameter, length and myelination of axons is known to vary widely by brain region, and a single axon may have up to 100 myelin layers. Dendrites are another type of branched projection originating from a nerve cell that typically receive inputs from other neurons. A *synapse* is the location where an *efferent* neuron connects to the dendritic tree of an *afferent* neuron. At synapses, an electrical potential may trigger the release of *neurotransmitters* as a chemical signal across the synapse.

Although neurons are extremely numerous within the nervous system, another family of cells called neuroglia (glial cells or oligodendrocytes), are even more numerous (Herculano-Houzel, 2009). Various types of glial cells perform functions within the nervous system, such as homeostatic mechanisms, formation of myelin and nourishment of neurons. Glia are not considered to be involved in information processing[2]. Note that the healthy human brain continually undergoes alterations from birth to adulthood to old age. These alterations involve axonal growth, dendritic proliferation, changes in myelination and changes in synaptic plasticity (Bystron et al., 2008).

At a macroscopic scale, neuronal processing is communicated to distant brain regions along white matter pathways comprised of many axons, where coherent bundles of axons with the same trajectory are called *fascicles*. Mapping these pathways is an essential part of understanding brain structure and, ultimately, function. Connections from one brain region to another may be *direct* or *indirect*. For example, the well-studied *central visual pathway* projects from the optic nerve to the primary visual cortex via the optic chiasm and lateral geniculate nucleus. The majority of axonal fibres in the cerebrum connect from one cortical area to another, rather than to subcortical areas. The great majority of fibres connecting the two cerebral *hemispheres* travel through the corpus callosum (Figure 2.2). A small fraction of interhemispheric fibres pass through the anterior commissure, and other routes of communication pass through the hippocampal commissure or, indirectly, via subcortical connections (Funnell et al., 2000).

2.2.1 Neuroanatomy

Neuroscience is often chiefly concerned with the organisation of the *central nervous system* (CNS), formed of the brain and the spinal cord. The CNS is distinct from the peripheral nervous system, which consists of the nerves throughout the rest of the body.

[2]Some researchers have suggested that neuroglia have a role in information processing (Araque et al., 1999), however, this notion has been largely disregarded in contemporary connectomics.

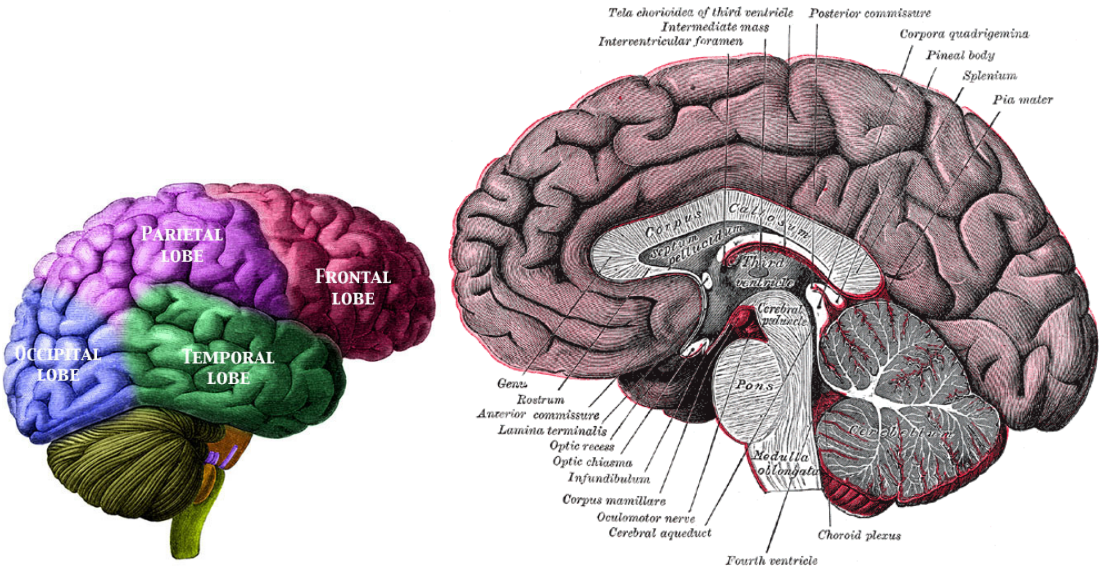

Figure 2.2: Left) Sagittal view of the four major lobes in the cerebrum. Right) Illustration of several anatomical structures visible in a median sagittal section. (Credit: *Manuel de L'anatomiste* published in 1883 by Charles Morel and Mathias Duval, and *Gray's Anatomy of the Human Body*, originally published in 1918.)

Since the earliest days of neuroscience, researchers have sought to classify the structure of the brain. At a macroscopic scale, the simplest distinction is between *grey matter* (comprised mainly of neurons, unmyelinated axons and glial cells) and the inter-connecting *white matter* (myelinated axons and glial cells). The brain is divided into two roughly symmetrical and highly interconnected cerebral hemispheres along the medial longitudinal fissure. The largest white matter structure in the brain, the *corpus callosum*, is formed from a mass of axonal fibres which cross the longitudinal fissure, connecting the cerebral hemispheres, thus facilitating interhemispheric communication (Gazzaniga, 1998).

The grey matter in each cerebral hemisphere is classified into two distinct surfaces, namely the *neocortex* and the *subcortex*. The neocortex (or often simply cortex) is a thin highly folded sheet of neuronal cells on the outer (pial) surface of the brain. Typically, the cerebral cortex has six distinct neuronal layers. A *sulcus* is a characteristic groove in the cortex, and a *gyrus* is a ridge. Each cerebral hemisphere is divided into four distinct lobes (Figure 2.2), namely *occipital*, *parietal*, *temporal* and *frontal*. The subcortical area is an evolutionary more ancient grouping of neuronal units deep within the brain localised around the brain stem. Subcortical structures include the *thalamus*, *amygdala* and *hippocampus*. The *cerebellum* (little brain) is a brain region distinct from the cerebrum and located underneath the cerebral hemispheres. The cerebellum is primarily involved in motor control and although it contains more than three times as many neurons as the neocortex (Herculano-Houzel, 2009), it is rarely considered to have a role in cognitive function and the majority of connectome mapping and cognitive studies

exclude it.

In the early part of the twentieth century cytoarchitecture methods allowed Brodmann to delineate around 47 distinct areas in the cortex using post-mortem Nissl staining. Subsequent research has allowed researchers to further subdivide and refine these areas. Neuroanatomical regions may be divided by various criteria, including cytoarchitecture, anatomy (e.g. gyral and sulcal) and function (i.e. neural units sharing distinguishing functional properties). It is not yet clear whether structurally or functionally defined regions are more appropriate for a connectome mapping task (Hagmann et al., 2010a).

The nervous system is bathed in cerebrospinal fluid (CSF) which is circulated through the ventricular system – a continuous system of channels and reservoirs connecting the ventricles and spinal cord. This fluid serves to provide chemical stability and to cushion the soft tissue of the nervous system. CSF also exhibits well known properties of water molecule diffusion which are measurable with dMRI, specifically high mean diffusivity and very low diffusion anisotropy (see Section 2.4).

2.3 Magnetic resonance imaging

MRI is a non-invasive imaging technique capable of measuring the internal structure of tissue in two or three dimensions. MRI was developed from the principles of nuclear magnetic resonance (NMR) to measure the collective magnetic properties of atomic nuclei. Since the development of the modern MRI scanner in the 1970s, MRI has become an extensively used medical imaging technique, particularly for soft tissues such as muscles, the heart, the brain and tumours. The dMRI modality was developed in the 1980s to measure the diffusion of water molecules in tissue.

This section outlines some basics of MRI and dMRI referred to in this work. For a comprehensive introduction, see McRobbie et al. (2006) and Jones (2010).

2.3.1 Basic principles of magnetic resonance

A proton has basic physical properties including *mass*, *positive electric charge* and *spin*. A proton's spin produces a small but measurable magnetic field. Normally, the magnetic alignment of protons is randomly distributed (Figure 2.3). An MRI scanner uses a superconducting magnet to create a strong external magnetic field, B_0 (measured in Tesla units), to align the magnetisation of protons within tissue. Although protons may actually align with or against B_0, a measurable majority will align with the magnetic field as it is the lowest energy state

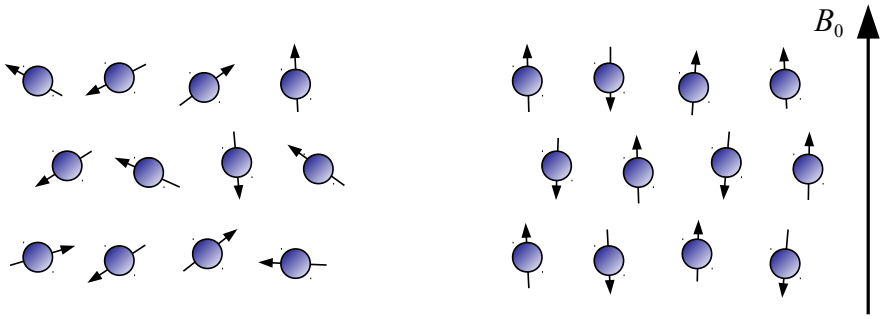

Figure 2.3: Protons with spins randomly distributed (left) and aligned in a strong external magnetic field, B_0 (right).

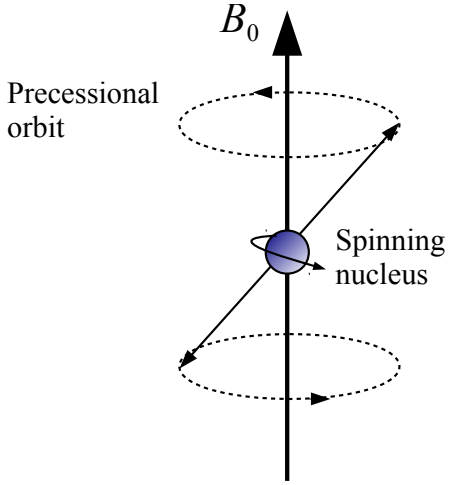

Figure 2.4: Precessional orbit of spinning nucleus around the axis of B_0.

(Figure 2.3). Although the energy of a single proton is very small, it is the large number of protons which produce a measurable signal. Within the body, water has the largest proportion of protons and soft-tissues such as fat and muscle have somewhat fewer.

Spinning protons precess about the axis of the B_0 field (Figure 2.4). The frequency of precession is directly proportional to the strength of the magnetic field as defined by the Larmor equation, $\omega_0 = \gamma B_0$, where ω_0 is the resonance frequency and γ is the gyromagnetic ratio constant measured in Hz/Tesla.

Radio frequency (RF) fields are used to systematically alter the alignment of the magnetic field. If an RF pulse is applied at the resonance frequency some protons will absorb the energy and their spin is altered as a function of the strength and duration of the RF pulse. Following the RF pulse, the absorbed RF energy is retransmitted at the resonance frequency. As the RF energy is retransmitted, the spin alters over a time course T_1, the *spin-lattice relaxation time*, from a high energy state to a low energy sate, thereby realigning with B_0. The recovery over T_1 is dependent on the field strength and the proton density in different tissues.

Following the RF pulse, protons also undergo *spin-spin interaction* as the magnetic field of excited protons interact. These temporary, random interactions between nearby protons cause a cumulative loss of phase, resulting in an exponential signal decay over a time course, T_2, the *spin-spin relaxation time*. Like T_1 relaxation, the signal decay resulting from spin-spin interaction is measurable. The T_2 signal loss is largely determined by the chemical make-up of tissue.

MR acquisition defines the parameters echo time (TE) and repetition time (TR). During MRI acquisition, RF fields are systematically applied to alter the alignment of protons within tissue. This produces a RF signal detectable by the scanner that is recorded to construct an image. By using field gradients in different directions 2D images or 3D volumes can be acquired.

2.3.2 T_1-weighted imaging

T_1-weighted imaging refers to an acquisition which depends on the spin-lattice relaxation time, acquired using either spin-echo or gradient-echo sequences. The T_1-weighted signal is dependent on the field strength and tissue type. The measured signal depends on tissue properties, where water appears dark and denser fatty tissues appear bright. Such imaging is commonly used as a structural reference and to provide contrast between grey and white matter.

2.3.3 T_2-weighted imaging

T_2-weighted imaging refers to an acquisition which depends on the spin-spin relaxation time. The T_2-weighted signal is dependent on the chemical makeup of tissue and to a lesser extent the field strength. Under T_2-weighted imaging fatty tissues appear dark and water or tissues containing fluid appear brighter. Such images are used to distinguish pathology, such as tumours or oedema relating to stroke or brain injury.

2.3.4 Echo-planar imaging

Echo-planar imaging (EPI) is a fast acquisition method developed in the 1990s, which vastly reduces the total acquisition time and motion artefacts by performing continuous signal readout by means of a gradient-echo train (Stehling et al., 1991). EPI uses a pulse sequence in which multiple echoes of different phase steps are acquired using rephasing gradients. EPI is most commonly used for acquiring an entire 2D slice using a sequential 'single-shot' readout in a single TR over a Cartesian grid. However, EPI has limited spatial resolution and single-shot sequences are susceptible to artefacts relating to ghosting, gradient errors, B_0 inhomogeneities and saturation. The inherent geometric distortions from magnetic field inhomogeneities are

due to the rapid sampling of the gradient-echo train. These distortions are partly predictable and correctable, but in practice artefacts cannot be entirely eliminated (Jones et al., 2013).

2.4 Diffusion MRI

dMRI is an acquisition sensitive to the diffusion of water molecules in biological tissue. dMRI has been used clinically for assessing ischaemic damage following stroke (Moseley et al., 1990), but crucially it can also be used to probe the organisation of white matter and provide a quantitative assessment of the brain's white matter microstructure (Le Bihan and Johansen-Berg, 2012).

2.4.1 Diffusion and thermodynamics

Diffusion is a basic property of thermodynamics. Water molecules within a free liquid undergo random motion due to thermal energy, known as *Brownian motion*. The direction of motion is random and molecules will regularly change course as they collide with other molecules. In the absence of barriers water molecules diffuse randomly and equally in all directions and the distribution of displacements will be Gaussian, which is termed *isotropic* diffusion. However, in the presence of microscopic biological barriers, such as cell membranes, the distribution of displacements will tend to depend on the orientation of these barriers, which is termed *anisotropic* diffusion.

Considering diffusion in a single direction, the concentration of fluid is denoted by $C(x,t)$ at location x and time t. The diffusivity of a fluid is denoted by D. If we assume that there are n molecules at a single location $x = 0$ at time $t = 0$, then Diffusion follows,

$$C(x,t) = \frac{n}{\sqrt{4\pi Dt}} \exp\left(-\frac{x^2}{4Dt}\right), \tag{2.1}$$

as described by Einstein (1905). This equation can be extended to three dimensions to account for isotropic and also anisotropic diffusion.

The principle of dMRI is that the properties of biological tissue fundamentally affect the direction of water molecule diffusion at body temperature. dMRI permits quantitive measurement of the gross motion of water molecules *in vivo* along any number of predefined directions. In the brain, the normally random Brownian motion of extracellular water molecules is restricted or hindered by cellular barriers such as cell membranes, microtubules or the myelin of axonal fibres. Typically, white matter tracts are coherent macroscopic structures, formed of microscopically ordered axons and microtubules. The resulting water diffusion is then constrained

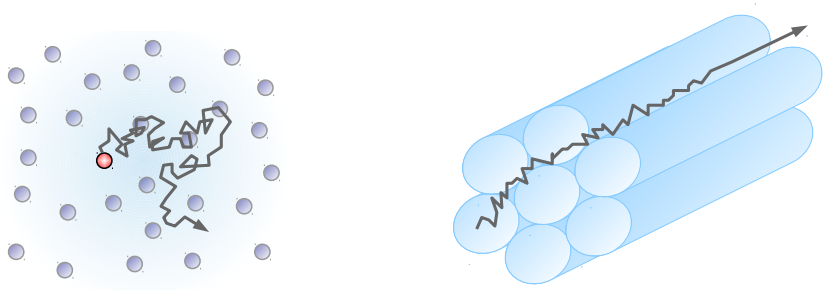

Figure 2.5: Left) A molecule undergoing Brownian motion due to thermal energy. Right) Tubular structures, e.g. axonal fibres, causing a molecule to preferentially diffuse along the direction of the fibre (anisotropic diffusion).

by these barriers, resulting in measurable anisotropic diffusion. The degree to which white matter is ordered will have a direct and measurable impact on the dMRI signal along the observed direction. Due to directional differences in water diffusion across different structures it is possible to estimate the principal orientation of water diffusion which depends on the underlying tissue structure.

2.4.2 Diffusion acquisition

MRI can be sensitised to the diffusion of water molecules in biological tissue by applying additional magnetic field gradients along directions of interest, where the signal corresponds to the magnitude of diffusion. Instead of applying only a homogeneous magnetic field B_0, the homogeneity is varied by a pulsed field gradient. The methods for dMRI acquisition were first developed in the 1980s when Le Bihan introduced the first diffusion-sensitising pulse sequences and the concept of the b-value to combine diffusion-sensitising and imaging gradient RF pulses (Le Bihan et al., 1986). Le Bihan also introduced the *apparent diffusion coefficient* (ADC) to measure diffusion in tissue. The diffusion coefficient is termed *apparent* because the underlying diffusion process is known to be complex in biological tissues and many different mechanisms act together to produce the measure of diffusion. The b-value is used to parameterise the amount of dephasing. Diffusion is not determined from the signal intensity alone, but from the signal loss, when compared to a baseline ($b = 0\ s\ \mathrm{mm}^{-2}$) signal.

A spin-echo pulse sequence can be sensitised to diffusion by using a pair of diffusion-weighted gradients of duration, δ, separated by a 180° refocussing pulse as shown in Figure 2.6. The first gradient pulse offsets the phase of the spins by an amount dependent on their location. The second pulse provides the opposite rephasing in the case that the molecules have not moved. However, water molecules will typically move because of thermal energy, which results in dephasing. The greater the diffusion displacement during the time between the gradient pulses,

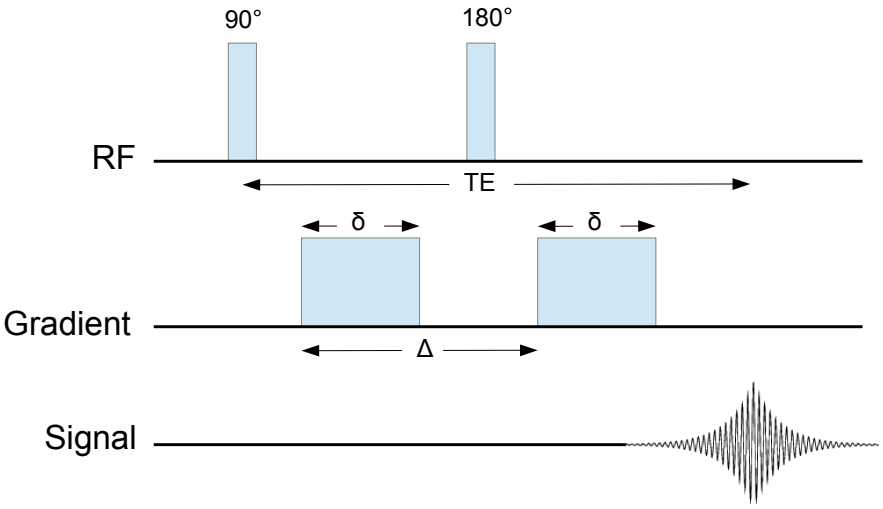

Figure 2.6: Diffusion-weighted spin-echo pulse sequence timing diagram.

Δ, the more attenuated the resulting signal. Diffusion acquisition estimates the log-ratio of the signal intensity with no diffusion sensitisation, A_0, and the signal after the full echo time, $A(b)$. This can be formulated (Basser et al., 1994),

$$\ln\left(\frac{A(\mathbf{b})}{A_0}\right) = -\gamma^2\delta^2\left(\Delta - \frac{\delta}{3}\right)G^2\mathbf{R}^T\mathbf{D}\mathbf{R} = -\sum_{i=1}^{3}\sum_{j=1}^{3} b_{ij}D_{ij}, \tag{2.2}$$

where γ is the gyromagnetic ratio, G is the maximal magnitude of the gradient pulse and $\mathbf{R}$ is a vector describing the direction of the applied gradient. A symmetric matrix $\mathbf{b}$ with elements b_{ij} incorporates the properties of the diffusion gradients. The matrix $\mathbf{D}$ with elements D_{ij} records the effective diffusivity averaged over the time of diffusion acquisition.

Diffusion-weighted imaging (DWI) records a single value reflecting the rate of water diffusion at each voxel. Typically, DWI acquisitions are recorded in the three principal gradient directions in order to estimate the *trace* of average diffusion per voxel. Such acquisitions have proved clinically useful in the imaging of ischaemia or hypoxic tissue, for example due to vascular stroke (Moseley et al., 1990) and for the evaluation of tumours (Koh and Collins, 2007). Early dMRI acquisitions were of lengthy durations and were susceptible to motion artefacts. It was not until the introduction of EPI in the 1990s and later with the development of single-shot EPI, that dMRI became more reliable and practical.

However, dMRI acquisitions have since been developed offering high-resolution probing of microscopic tissue structures by encoding strong bipolar magnetic field gradient ($b >= 1000$ s mm^{-2}) over many predefined directions, typically 12 to 60 or greater. The signal at each voxel corresponds to the diffusion rate at the direction of interest. Increasing the number of directions increases the angular resolution, thereby providing a more accurate sampling of the underlying diffusion process (Jones, 2004). Similarly, repeating angular measurements, for

example repeating 12 directions twice and taking an average, improves the signal-to-noise ratio of the diffusion measurements for each direction. Necessarily, sampling over many directions or repeating measurements requires lengthy acquisition times.

2.4.3 The diffusion tensor model

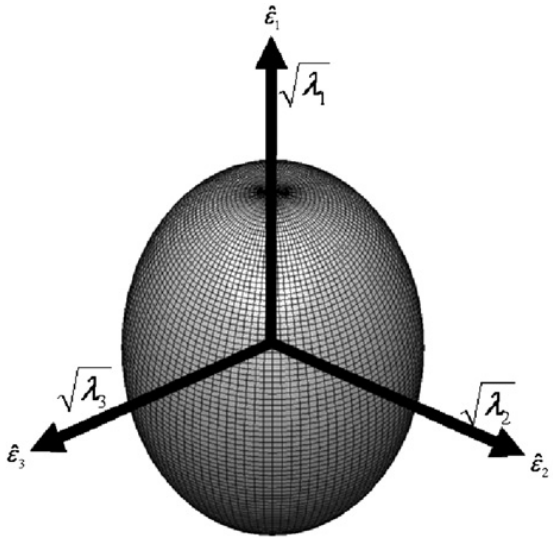

Figure 2.7: An elliptical tensor modelling water diffusion in three directions, with magnitudes λ_1, λ_2 and λ_3.

In the 1990s Michael Moseley and colleagues pioneered the idea of using an elliptical tensor to model diffusion in anisotropic white matter (Moseley et al., 1990). The diffusion tensor models local diffusion by a 3D Gaussian distribution with a covariance matrix proportional to the diffusion tensor (Basser et al., 1994). The diffusion weighted signal, μ_i, is modelled as,

$$\mu_i = S_0 \exp(-b_i \mathbf{r}_i^T \mathbf{D} \mathbf{r}_i), \tag{2.3}$$

along a gradient direction $\mathbf{r}_i$ with b-value b_i for the ith diffusion encoding direction, where S_0 is the signal with no diffusion weighting. The diffusion tensor $\mathbf{D}$ is,

$$\mathbf{D} = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{xy} & D_{yy} & D_{yz} \\ D_{xz} & D_{yz} & D_{zz} \end{bmatrix}. \tag{2.4}$$

These parameters can then be estimated by a least-squares fit to the diffusion data. The tensor model decomposes the diffusion measurements into a set of eigenvectors $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$, the set of *principal directions of diffusion*, and a corresponding set of eigenvalues $(\lambda_1, \lambda_2, \lambda_3)$, the magnitude of these vectors. These parameters define the shape and orientation of the tensor. The model assumes that at each voxel has only a single principal direction of diffusion. The

tensor is most often represented as an ellipsoid (Figure 2.7), where the surface represents the distance to which a molecule will diffuse with equal probability. MRI acquisitions for use with the tensor model are commonly referred to as diffusion tensor imaging (DTI).

2.4.4 Diffusion parameters

The integrity of white matter can be measured by diffusion imaging parameters, such as the *mean diffusivity* (MD) and *fractional anisotropy* (FA), derived from fitting the diffusion tensor to the dMRI signal at each voxel (Le Bihan et al., 2001). These measures are often used in diffusion imaging where they are considered to reflect fibre density and directional coherence.

Axial diffusivity, λ_1, is the diffusivity along the principal axis. The diffusion along the two remaining axis are often averaged to produce a measure of *radial diffusivity*,

$$\lambda_\perp = (\lambda_2 + \lambda_3)/2. \tag{2.5}$$

MD or $\langle D \rangle$ is a scalar value that describes the mean relative degree of diffusion along all directions at a voxel location,

$$\langle D \rangle = (\lambda_1 + \lambda_2 + \lambda_3)/3. \tag{2.6}$$

These measures of diffusivity can be used to assess the restriction of water molecule diffusion that may be due to the underlying density of cell membranes and fibres. FA is a rotationally and translationally invariant scalar value that describes the degree of anisotropy at a voxel location (Basser and Pierpaoli, 1996),

$$\mathrm{FA} = \sqrt{\frac{1}{2}} \frac{\sqrt{(\lambda_1 - \lambda_2)^2 + (\lambda_2 - \lambda_3)^2 + (\lambda_3 - \lambda_1)^2}}{\sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}}. \tag{2.7}$$

An FA value of zero corresponds to isotropic diffusion, i.e. diffusion is equally restricted in all directions, whereas a value of one means that diffusion occurs in only one direction and is fully restricted in all other directions. FA can be close to 0 when measured in ventricular CSF, whereas FA can approach 1 in regions of the corpus callosum where there are many myelinated fibres with directional coherence. Typically, FA in grey matter falls below 0.3 and FA in white matter rarely exceeds 0.7.

Studies have shown distinct architectural differences in white matter organisation across the brain. For example, the corpus callosum is a well-defined white matter structure, which exhibits considerable regional variance in FA along its length, with the thin fibres of the genu having markedly lower FA than the fibres of the splenium or posterior pole, which typically show the greatest diffusion anisotropy in the brain. The greatest FA values are found in the

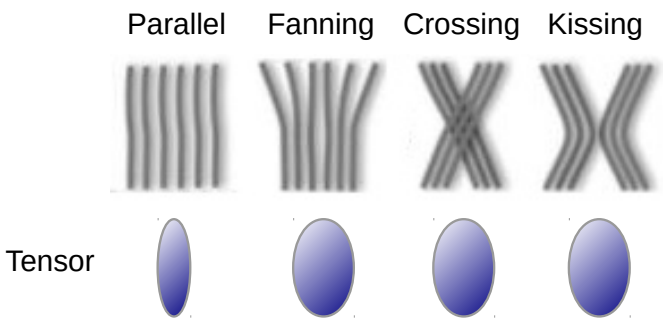

Figure 2.8: Illustration of four intravoxel multi-fibre configurations with the associated tensor models.

white matter regions with the most coherent fibre bundles, such as the splenium (Pfefferbaum et al., 2000). Markedly lower FA values are found in areas consisting of crossing fibres, such as the centrum semiovale and pericallosal regions. Regional differences in FA are not necessarily indicative of white matter abnormality or deterioration, but rather reflect the regional differences in white matter structure. Crucially, axonal structures have been demonstrated to be perturbed by ageing, disease and physical trauma (Tang et al., 1997; Arfanakis et al., 2002). Disruptions of white matter structure measured by dMRI may reflect degradation of myelin, parts of the axonal cytoskeleton or a decrease in axonal density. Furthermore, a seminal study which involved learning of a complex visuomotor skill (juggling) demonstrated that training can induce experienced-based changes in the white matter of the healthy brain (Scholz et al., 2009).

2.4.5 The multi-fibre problem

Research has shown that there are several fibre reconstruction problems with the tensor model (Wiegell et al., 2000; Tuch et al., 2002). The tensor model is rather simplistic in that only a single ellipsoid is modelled, making the assumption that all the axons within a voxel follow the same trajectory. This becomes problematic when estimating the trajectories of multiple fibres within a dMRI voxel, a configuration that has been demonstrated to be common across the brain's inter-connections (Behrens et al., 2007). A number of confounding fibre configurations have been identified, specifically where fibres cross, diverge or turn. Figure 2.8 illustrates four multi-fibre configurations and the associated tensor models. The crossing of multiple fibre trajectories has become known as the *fibre-crossing problem*. The *fibre-fanning problem* describes the case where a bundle of coherent fibres diverges in two or more directions. The *fibre-kissing problem* describes the case where multiple fibres change trajectory within a voxel but do not cross. In practice, all of these fibre problems manifest at some scale throughout the brain and these problems are inherently difficult to resolve at dMRI resolution.

2.4.6 Other dMRI acquisitions

Partly in response to the multi-fibre limitations of the diffusion tensor, more sophisticated dMRI acquisitions have since been developed. For example, high angular resolution diffusion imaging (HARDI; Tuch et al. 2002) and diffusion spectrum imaging (DSI; Wedeen et al. 2005, 2008) have been designed to better model the distribution of diffusion orientations. HARDI uses a high *b*-value diffusion gradient single-shell sampling scheme and a mathematical alternative to the tensor model to model fibre crossing as a function of diffusion gradient orientation. HARDI has been demonstrated to model intra-voxel configurations which cannot be resolved by the tensor model (Tuch et al., 2002). Likewise, DSI has the capacity to image crossing fibres by measuring the 3D spectra of water diffusion (Wedeen et al., 2005). This is achieved by using a large number of gradient encoding directions at various *b*-values. DSI has sufficient angular resolution to map diffusion at each location with a non-Gaussian behaviour, by using orientation diffusion functions (ODFs). This permits resolution of intra-voxel diffusion heterogeneity caused by crossing axonal fibre tracts. DSI has been successfully applied to connectome mapping (Hagmann, 2005) and offers increased sensitivity in estimating complex fibre configurations compared to DTI (Bassett et al., 2010). However, DSI requires lengthy acquisition times and high performance MRI hardware (Hagmann et al., 2010b).

2.5 Tractography

Tractography is the application of image processing algorithms to estimate white matter fibre tracts from dMRI data. The aim is to piece together fibre trajectories based on the continuity of diffusion orientation estimates from voxel to voxel. There are various established techniques for constructing tracts and these can be broadly divided into those using *deterministic* or *probabilistic* approaches. Figure 2.9 illustrates tractography outputs for deterministic and probabilistic methods. Table 2.1 shows a non-exhaustive list of established tractography algorithms.

Early tractography studies were applied to isolate well-studied fibre pathways, for instance the splenium and genu of the corpus callosum. Eventually, as the methods and understanding of white matter pathways has progressed this has enabled whole-brain tract generation.

2.5.1 Deterministic streamline tractography

Deterministic streamline methods were among the earliest established tractography techniques. A *streamline* is constructed from a seed point and then from voxel-to-voxel following only the

Deterministic methods	**Algorithm**
Local path integration	Fibre assignment by continuous tracking (FACT; Mori et al. 1999)
	Conturo et al. (1999)
Diffusion directions	Basser et al. (2000)
Tensor deflection	White matter tractography using tensor deflection (TEND; Lazar et al. 2003)
Proprietary	Diffusion toolkit (DTK; Wang et al. 2007; `http://trackvis.org`)

Probabilistic methods	**Algorithm**
Bayesian	BEDPOSTX/PROBTRACKX, FSL Diffusion Toolkit (FDT; Behrens et al. 2003b)
	Friman et al. (2006)
Monte-Carlo streamline generation	Probabilistic index of connectivity (PICo; Parker et al. 2003)
	Hagmann et al. (2003)

Other methods	**Algorithm**
Dijkstra's algorithm	Iturria-Medina et al. (2007)
Fast marching	Parker et al. (2002)

Table 2.1: Established tractography algorithms.

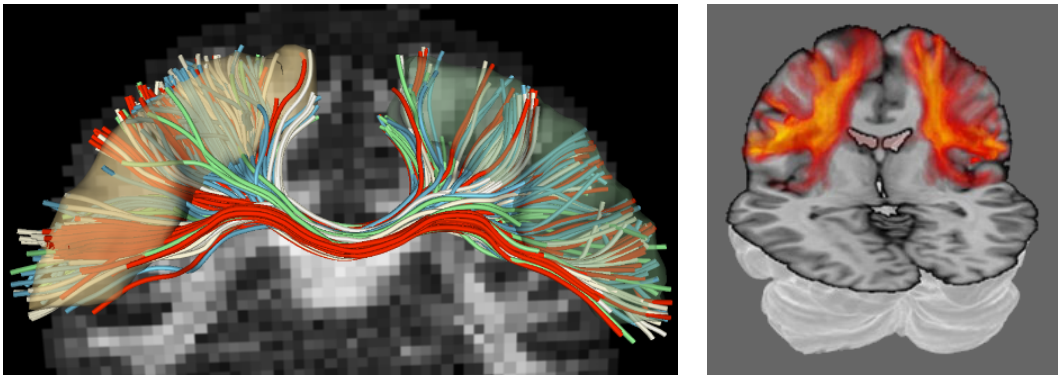

(a) Deterministic (b) Probabilistic

Figure 2.9: Illustration of deterministic and probabilistic tractography: a) a set of streamlines constructed from voxel-to-voxel by deterministic tractography (Credit: Psychiatry Neuroimaging Laboratory, Harvard Medical School); b) Probabilistic tractography outputs showing connectivity distributions, where the colour of each voxel indicates a connection probability (Credit: The Human Connectome Project).

principal eigenvector of the diffusion tensor at each voxel until terminated by some stopping criteria as shown in Figure 2.9(a). Typical stopping criteria include exceeding a streamline angular threshold or falling below an anisotropy threshold. The first established algorithms include fibre assignment by continuous tracking (FACT; Mori et al. 1999; Xue et al. 1999), tracking the direction of fastest diffusion (Conturo et al., 1999) and path integral methods (Basser et al., 2000). While FACT samples a trajectory direction exactly once per voxel, other approaches interpolate the local diffusion data to determine the orientation at a sub-voxel scale (Basser et al., 2000). These algorithms are entirely deterministic and only estimate the best fit of the diffusion tensor.

However, such methods can be strongly affected by image noise and the results can be anatomically ambiguous. These algorithms have been shown to fail in voxels with crossing fibres and other complex fibre configurations and they are strongly affected by MR signal noise (Basser and Pajevic, 2000). To minimise tracking artefacts due to these limitations, modified streamline algorithms were developed, such as use of anatomical priors to regularise tracking (Poupon et al., 2000), and tensor deflection using adaptive stepping (Lazar et al., 2003). However, in order to overcome the inherent limitations of deterministic methods, researchers have adopted algorithms which better model the uncertainty in the estimates of diffusion orientation.

2.5.2 Probabilistic tractography

Probabilistic tractography (Parker et al., 2003; Behrens et al., 2003b, 2007; Friman et al., 2006) is an extension of deterministic streamline tractography. Probabilistic approaches seek to replace the single principal diffusion direction with a distribution of orientations, which capture the uncertainty associated with the diffusion signal at each voxel. During tracking, these algorithms generate a set of streamlines using a Monte Carlo approach, by sampling from a local probability density function (PDF), which describes the uncertainty in orientation at each voxel location. Typically, streamline tracking is repeatedly initiated from a single seed point, which results in many possible end points. Once a full set of streamlines is produced a *connection probability* can be calculated describing the likelihood of connection between any two locations. The probabilistic approach is more robust to noise, allows determination of uncertainty in the reconstructed pathway, and has the potential to construct more reproducible tracts (Behrens et al., 2003b). The probabilistic methods used in this thesis are described more fully in Section 3.6.2.

2.5.3 Other approaches

Various other algorithmic approaches have been applied to tractography, including energy minimisation techniques, such as Dijkstra's algorithm (Iturria-Medina et al., 2007) and fast marching (Parker et al., 2002). Others have applied quantitative tractography to quantify how the shape of an individual tract compares to that of a predefined reference tract (Clayden et al., 2007). Such methods employ unsupervised selection algorithms to reliably segment tracts.

One major limitation of current methods is the spatial scale of the information available. Researchers have proposed super-resolution methods, termed *track density imaging*, which use post-processing methods based on dMRI fibre tracking to estimate structures beyond the resolution of the acquired imaging voxel (Calamante et al., 2010).

2.5.4 Region-of-interest tractography

The earliest tractography studies often involved region-of-interest (ROI) based tractography. Following this approach, an ROI or set of ROIs are selected within white matter and tractography is applied using the chosen seeding strategy and stopping criteria. Typically, an ROI is formed from a contiguous group of voxels in an area of interest. *Seed* ROIs are used as the starting location for tracking and *waypoint* ROIs may also be used which constrain streamlines to pass through one or more waypoint regions (Conturo et al., 1999). Properties of white matter integrity can then be averaged within the resulting tracts and used to assess tract-specific differences between subjects. Early ROI studies often assessed well-studied white matter pathways such as the genu and splenium of the corpus callosum. Manual placement of ROIs has the advantage that it is simple and tractable to perform and the results are often straightforward to interpret. However, by definition, the ROI approach is subjective and limited to *a priori* selected areas and therefore some pertinent white matter tracts may be excluded from analysis by the choice of ROIs.

In comparison, whole-brain tractography approaches using automatic cortical parcellation have the advantage that an objective segmentation can be obtained with little manual effort. However, whole-brain approaches face a challenge to maintain the consistency of segmentations from subject to subject and the segmentations are difficult to validate. In addition, the resulting analysis becomes more complex as the number of regions increases.

2.5.5 Validation

Although tractography is a powerful non-invasive technique it remains largely un-validated against ground-truth studies, primarily because of the huge challenge in observing the true ax-

onal connections, particularly for the human brain. Attempts to validate tractography have been made against post-mortem tracing methods in the Porcine (Dyrby et al., 2007) and Macaque (Hagmann et al., 2008) brains, with good agreement seen in well-studied major white matter tracts. Likewise, synthetic diffusion models have been used to validate various aspects of tractography (Alexander, 2008; Fillard et al., 2011). Hagmann at al. compared connectivity obtained through DSI and deterministic tractography in a Macaque brain against known tracts from the CoCoMac macro-connectivity database[3]. Hagmann's results show 78.9% agreement with known connections but with 6.1% of connections reported as false negatives and 15.0% of connections placed in regions where the presence of pathways was unknown (Hagmann et al., 2008). However, researchers admit that much about the white matter organisation of the human brain remains unknown (Hagmann et al., 2010b).

2.5.6 Subjectivity of tractography

Tractography is subjective in the sense that the researcher is faced with a choice of tractography algorithms, models of multiple fibres, stopping criteria and seed placement. Although these choices can be motivated by the on-going biological validation of the methods, the resulting configuration is often rather arbitrary and these choices inevitably vary from study to study. Although tractography algorithms have been developed to model the underlying white matter structure, the problem remains of how to best perform tract reconstruction in a consistent and reproducible way. This is especially important for case-control and cross-sectional analysis where subjectivity is a confounding factor. For instance, tract reconstruction is strongly affected by the choice of starting location and one seeding strategy can product markedly different streamlines from another strategy (Cheng et al., 2012b; Côté et al., 2013). One advantage of whole-brain tractography over ROI-based tractography is that seeding can be performed objectively and systematically throughout the brain, removing the need for manual placement of seed points. However, this indiscriminate inclusion of seed points means that a large number of seeds may be included which may not be entirely appropriate and the analysis must deal with a greater proportion of spurious streamlines and noise. We revisit seeding strategies for connectome mapping in Section 2.7.2.1 and compare two alternative seeding strategies in Chapter 4.

[3] http://cocomac.g-node.org

2.6 The connectome

The term *connectome* (Hagmann, 2005; Sporns et al., 2005) was supposedly coined independently and simultaneously by Hagmann and Sporns in analogy to the well known genome, the complete mapping of an individual genetic profile that characterises biological function at multiple levels. A connectome is a network which maps the connections of the brain, where some choice of neural units form the nodes of a network and some choice of neural connections between these units form the links of the network. A connectome may involve structural connections and/or functional connections. Refer to Sporns (2011a) for a comprehensive overview of connectome methods and their interpretations. To date, the approach to connectome mapping has chiefly been tackled from two vastly differing scales: a) a cellular scale mapping (bottom up) through techniques such as post-mortem histological tracing and confocal microscopy; and b) a macroscopic scale *in vivo* imaging of the entire cerebrum (top down) through techniques such as dMRI and functional MRI (fMRI). Note that the disparity between these scales is immense, with a typical 1 mm isotropic voxel in cortical grey matter considered to contain approximately 40 thousand neurons. These ongoing developments can be considered something of a response to Crick and Jones reflection on brain connectivity in 1993 which urged the development of radically new techniques to investigate and understand the brain (Crick and Jones, 1993).

Note that a vast amount of pre-connectome research has been concerned with mapping and understanding brain connectivity at various levels, including whole-brain. The distinction between connectome research and earlier studies of connectivity is, of course, rather arbitrary.

2.6.1 The mesoscopic scale

The ultimate aim of the connectome approach is, perhaps, a cellular scale (mesoscopic) map of the human brain, potentially outlining each and every axon and synapse (noting that neuronal connectivity changes with experience and varies from individual to individual). However, technologies to achieve this resolution in mammals are some way off and for now only gross approximations at a macroscopic scale *in vivo* are possible with MRI.

A mesoscopic scale network, in which individual neurons are treated as nodes, are currently only possible for the simplest organisms. Meaningful localisation of circuits in the mammalian brain is challenging because identifying all relevant connections at such a scale is currently intractable. However, on the basis that all organisms share common genetics, neuroscientists have studied simpler model organisms with far fewer neurons. Model organisms involved in connectomics research include the mouse, the fruit fly, and the nematode roundworm.

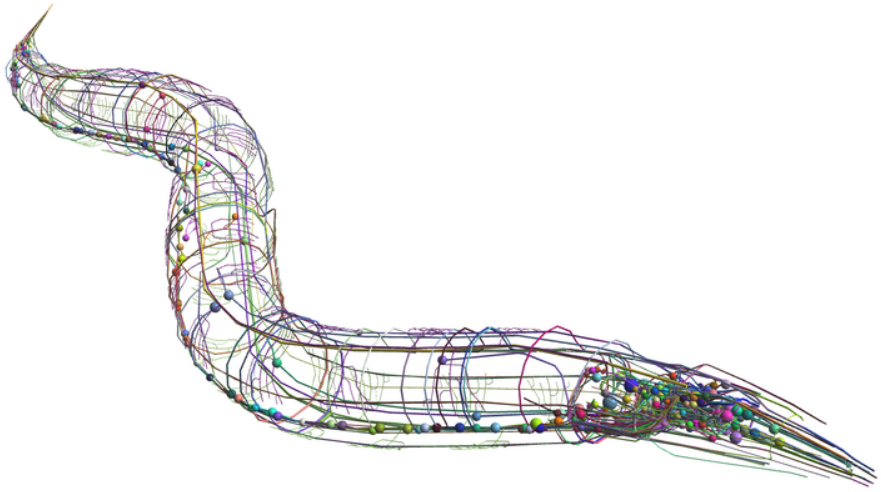

Figure 2.10: 3D model of the nervous system of the nematode round worm, Caenorhabditis elegans. (Credit: The OpenWorm project; Busbice et al. 2013)

For example, the roundworm (*Caenorhabditis elegans*) has a nervous system composed of 302 neurons and roughly 7,500 synapses. This connectome was painstakingly traced from electron micrographs of a single animal by neurobiologists in the 1980s (White et al., 1986). This roundworm is the only organism for which a complete connectome exists. This map (Figure 2.10) has led to a fuller understanding of the functional organisation of this simple nervous system. However, detractors of this approach have also pointed to the lack of neuroscientific benefit arising from this map because of the crucial lack of understanding in how individual neurons behave physiologically[4]. Another criticism of the mesoscopic approach is that neuronal connectivity, in mammals at least, is dependent on experience and it is not clear how mesoscopic connectivity in one individual may relate to connectivity in another individual.

2.6.2 The macroscopic scale

Topological information of the large-scale structural networks in the brain may be obtained by applying tractography algorithms to dMRI data. Since individual analysis of many tracts is impractical at this scale our focus is on the large-scale networks of the cerebral cortex at a systems level, involving the collective wiring of many billions of neurons.

The first mapping approaches combined whole-brain tractography with computational network analysis to quantify the integrity of the entire brain in healthy volunteers (Hagmann et al., 2007; Vaessen et al., 2010; Zalesky and Fornito, 2009). Additionally, advances in computational power and network analysis methods have assisted the characterisation of such large-scale brain networks. Analysis of these networks may determine the normal variation in brain net-

[4] a component also missing from current connectome approaches.

work topology. However, this approach may also prove useful for studying the topological and structural changes that accompany ageing or neurological disease progression in general.

2.6.3 Summary of dMRI connectome studies

The first connectome dMRI studies demonstrated whole-brain network analysis in a small number of healthy volunteers (Hagmann et al., 2007, 2008). These early studies and those that followed demonstrated that whole-brain networks could indeed be constructed from dMRI and the resulting networks showed a remarkable degree of similarity from person to person, but were not identical. These studies have shown that the human brain and the brains of other mammals, have various non-trivial properties in their organisation. For instance, the brain is characterised by highly connected 'hub' nodes, a modular structure and 'small-world' organisation (Hagmann et al., 2008; Honey et al., 2008; Sporns, 2011b; Van Den Heuvel and Sporns, 2011; Yan et al., 2011). These findings are discussed more fully in Section 2.11.

Since the first pioneering publications, many structural network studies have been applied in large cohorts. This includes normal ageing (Gong et al., 2009b; Robinson et al., 2010; Ystad et al., 2011; Wen et al., 2011), gender differences (Gong et al., 2009a), intelligence (Li et al., 2009), Alzheimer's disease (AD) (Lo et al., 2010; Jahanshad et al., 2015), mild cognitive impairment (MCI) (Wee et al., 2011), stroke (Crofts et al., 2010), multiple sclerosis (MS) (Shu et al., 2011), amyotrophic lateral sclerosis (ALS) (Verstraete et al., 2011) and neuropsychiatric disorders (Skudlarski et al., 2010; Zalesky et al., 2011). Broadly, studies can be placed into three categories: a) those demonstrating new methods in network construction or comparing alternative approaches; b) those used to determine network organisation in individuals or populations (cross-sectional); and c) case-control studies, used to identify network differences between two (or more) groups. Table 2.2 provides a list of influential methods based papers and studies which assess aspects of reliability in structural brain networks. Table 2.3 provides a non-exhaustive list of well-cited structural cross-sectional and case-control studies in humans. In each case a summary of the dMRI acquisition, segmentation method and tractography configuration is reported.

However, to date only a small subset of studies have assessed the reliability of the resulting networks (Bassett et al., 2010; Cammoun et al., 2011; Cheng et al., 2012a; Hagmann et al., 2008; Vaessen et al., 2010; Zalesky et al., 2010b), and currently there is a lack of assessment concerning the reproducibility of these approaches.

Notably, the Human Connectome Project[5] (Marcus et al., 2011; Van Essen et al., 2013), is a major ongoing project sponsored by the National Institutes of Health, which aims to build

[5] http://www.humanconnectome.org

a structural and functional map of the healthy human brain at a macroscopic scale, using a multitude of imaging technologies and resolutions. Similarly, the Open Connectome Project[6] (Vogelstein et al., 2010, 2011) has similar goals using human and animal studies but with a largely bottom-up driven approach (from the level of neurons and synapses) and provides open access to its data.

[6] http://openconnecto.me

Study	Cohort (*N*)	DW-MRI	Field Strength (T)	Directions	Parcellation	Nodes	Tractography
Bassett et al. 2010	Test-retest (7)	DTI, DSI	3	30, 258	AAL, HO, LPBA40	90, 110, 54	DTK
Cammoun et al. 2011	Test-retest (5), Healthy (20)	DSI	3	-	FreeSurfer	66, 133, 241, 483, 998	Hagmann et al. 2007
Cheng et al. 2012a	Test-retest (44), Healthy (15)	DTI	3	48	FreeSurfer	68	DTK
Gigandet et al. 2008	Healthy (4)	DSI	3	129	FreeSurfer, as described	66, 998	Hagmann et al. 2007
Hagmann et al. 2007	Healthy (2)	DSI	3	-	as described	500-4000	As described (deterministic)
Hagmann et al. 2008	Test-retest (1), Healthy (5)	DSI	3	-	as described	998	As described (deterministic)
Hagmann et al. 2010b	Juvenile (30)	DTI, Q-ball	3	12, 60	Hagmann et al. 2008	66, 241	Hagmann et al. 2008
Honey et al. 2008	Healthy (5)	DSI	3	129	FreeSurfer	66, 998	Hagmann et al. 2008
Honey et al. 2009	Healthy (5)	DSI	3	129	FreeSurfer, as described	66, 998	Hagmann et al. 2007, 2008
Iturria-Medina et al. 2007	Healthy (5)	DTI	1.5	12	IBASPM	71	as described (probabilistic)
Iturria-Medina et al. 2008	Healthy (20)	-	1.5	12	AAL	90	as described (probabilistic)
Vaessen et al. 2010	Test-retest (6)	DTI	3	6, 15, 32	WFU	111	PICo
Yo et al. 2009	Healthy (1)	HARDI	3	60	Manual segmentation	14	Various
Zalesky and Fornito 2009	Healthy (1)	DTI	1.5	3 0	AAL	70	As described
Zalesky et al. 2010b	Healthy (3)	DTI, HARDI	3	60, 252	AAL and as described	82, 500, 100, 200, 300, 4000	FACT

Table 2.2: A list of human connectome studies using dMRI and tractography primarily reporting on methodology and reliability. Abbreviations used for atlases: Automatic Anatomical Labelling (AAL), Harvard-Oxford (HO), Individual Brain Atlases for Statistical Parametric Mapping (IBASPM), LONI Probabilistic Brain Atlas of 40 subjects (LPBA40), Wake Forest University (WFU).

Study	Cohort (*N*)	DW-MRI	Field Strength (T)	Directions	Parcellation	Nodes	Tractography
Bonilha et al. 2014	Ageing (17)	DTI	3	30	AAL	90	DTK
Chung et al. 2011	Autism (17) : HC (14)	DTI	-	12	as described	1-800	TEND
Crofts et al. 2010	Stroke (9) : HC (18)	DTI	1.5	60	HO	56	FDT
Fischi-Gómez et al. 2014	Infant (52)	DTI	3	30	FreeSurfer	83	Hagmann et al. 2007
Fischer et al. 2014	Ageing (43)	DTI	3	30	HO	111	FACT
Gong et al. 2009b	Healthy (95)	DTI	1.5	30	AAL	78	FDT
Gong et al. 2009a	Healthy (80)	DTI	1.5	6	AAL	78	Mori et al. 1999
Greicius et al. 2009	Healthy (23)	DTI	1.5	12	Manual	4	Basser et al. 2000
Ingalhalikar et al. 2014	Gender (428:521)	DTI	3	64	FreeSurfer	95	FDT
Jahanshad et al. 2011	Healthy (234)	HARDI	4	94	FreeSurfer	70	Aganj et al. 2011
Korgaonkar et al. 2014	Depression (95) : HC (102)	DTI	3	42	FreeSurfer	84	FDT
Kuceyeski et al. 2011	Brain trauma (15) : HC (14)	HARDI	3	55	IBASPM, AAL	116	As described
Li et al. 2009	Healthy (79)	DTI	3	12	AAL	90	FDT
Li et al. 2011	Healthy (94)	DTI	3	60	FreeSurfer	82	FDT
Lo et al. 2010	AD (25) : HC (30)	DTI	1.5	13	AAL	78	FACT
Park et al. 2008	Healthy (1)	DTI	3	45	AAL	73	FDT
Robinson et al. 2008	Healthy (174)	DTI	3	15	Heckemann et al. 2006	83	FDT, Iturria-Medina et al. 2007
Robinson et al. 2010	Healthy (42), Elderly (54)	DTI	3	15	Heckemann et al. 2006	83	FDT
Shu et al. 2009	Blind (17) : HC (17)	DTI	3	12	AAL	90	FACT
Shu et al. 2011	MS (39) : HC (39)	DTI	1.5	6	AAL	90	FACT
Skudlarski et al. 2008	Healthy (41)	DTI	3	12	WFU	26, 5000	FACT
Skudlarski et al. 2010	Schizophrenia (27) : HC (27)	DTI	3	12	WFU	28	FACT
Van Den Heuvel and Sporns 2011	Healthy (21)	DTI	3	30	FreeSurfer	82	FACT
Várkuti et al. 2011	Healthy (23)	DTI	3	12, 64	AAL	116	PICo
Verstraete et al. 2011	ALS (35) : HC (19)	DTI	3	30 ($\times$2)	FreeSurfer	82	FACT
Wee et al. 2011	MCI (10) : HC (17)	DTI	3	25	AAL	90	ExploreDTI (Leemans et al., 2009)
Wen et al. 2011	Elderly (342)	DTI	3	32	FreeSurfer	68	MedInria (`http://med.inria.fr`)
Yan et al. 2011	Healthy (73)	DTI	3	64	AAL	78	FACT
Ystad et al. 2011	Healthy (100)	DTI	1.5	25	FreeSurfer		DTK
Zalesky et al. 2011	Schizophrenia (74) : HC (32)	DTI	1.5	64	AAL	82	FACT

Table 2.3: A non-exhaustive list of cross-sectional and case-control human connectome studies using dMRI and tractography. HC indicates Healthy Controls.

2.6.4 Comparison to other techniques in dMRI

Tract-based spatial statistics (TBSS) is an established voxel-based analysis of white matter tracts using dMRI, which has been described in detail elsewhere (Smith et al., 2006) and applied extensively in analysis of major white matter pathways. Firstly, a nonlinear deformation is used to align the FA map of each subject to a white matter template in standard space. White matter masks are 'skeletonised' (morphologically thinned) in order to obtain the centre-line of the principal white matter pathways, while minimising the impact of registration error and partial volume effects. Hypothesis testing may then be used with case-control or cross-sectional analysis to assess relationships between variables of interest and the voxel-wise FA or MD within white matter. Finally, permutation testing assigns a corrected p-value to each voxel, for which significance can be determined.

To our knowledge, the correspondence between TBSS and structural networks has not been well explored. In Section 5.5.3, we demonstrate a quantitative method to compute a *network-tract overlap* between the two methods, in terms of the white matter regions involved (Buchanan et al., 2014b).

2.7 Methods of network construction

Although there is currently no standard method for constructing dMRI structural networks, previous approaches have typically followed a similar organisation. Firstly, network nodes are formed from segmentation of high resolution 3D T_1-weighted volume scans, often by registration to neuroanatomical atlases (Maldjian et al., 2003; Shattuck et al., 2008; Tzourio-Mazoyer et al., 2002) or surface parcellation based on cortical sulci and gyri (Desikan et al., 2006; Fischl et al., 2004b). The number and choice of nodes requires careful consideration as this affects the resulting measures of connectivity (Zalesky et al., 2010b). Previous approaches have typically divided the cortex into fewer than 100 grey matter nodes, though some researchers have used finer parcellations with thousands of nodes of roughly uniform size, primarily to estimate global network properties (Cammoun et al., 2011; Hagmann et al., 2007, 2008; Zalesky et al., 2010b). Secondly, cross-modal registration (Andersson et al., 2007; Greve and Fischl, 2009; Jenkinson and Smith, 2001; Jenkinson et al., 2002) is typically required to align cortical labels to diffusion space. Thirdly, either deterministic (Basser et al., 2000; Lazar et al., 2003; Mori et al., 1999) or probabilistic (Behrens et al., 2003b, 2007; Parker et al., 2003) tractography is used to construct white matter tracts from dMRI data. Lastly, network connections are computed by quantifying tracts connecting between regions. Network weights typically reflect a count of interconnecting tracts (Hagmann et al., 2008) or some measure of tissue microstruc-

ture, such as diffusion anisotropy, averaged along the length of each tract (Iturria-Medina et al., 2007; Robinson et al., 2010). Connections may then be assessed directly in case-control studies (Zalesky et al., 2010a) or network measures derived from graph-theory (Rubinov and Sporns, 2010) may be used to characterise patterns of connectivity in individuals or across populations. Established network measures include *node degree*, *clustering coefficient* and *characteristic path length* (Watts and Strogatz, 1998), *efficiency* (Latora and Marchiori, 2001), and 'small-world' properties (Humphries and Gurney, 2008). Although the subsequent networks may be weighted or binary, they are typically undirected as tractography cannot distinguish between afferent and efferent connections.

2.7.1 MRI acquisition

For the purpose of building structural networks, typically researchers acquire both a high-resolution T_1-weighted structural volume and dMRI data in the same session. However, some researchers have generated networks from only dMRI volumes (Hagmann et al., 2007). To date, the majority of studies have used conventional dMRI acquisition at angular resolutions ranging from 6 to 64 directions. However, in attempts to resolve the crossing fibre limitations of tensor imaging, a number of recent studies have employed more complex acquisitions, such as DSI (Bassett et al., 2010; Gigandet et al., 2008; Hagmann et al., 2007, 2008; Honey et al., 2008, 2009) and HARDI (Jahanshad et al., 2011; Kuceyeski et al., 2011; Yo et al., 2009; Zalesky et al., 2010a). Acquisitions using EPI typically require some form of eddy-current correction to correct for the known systematic distortions (Jones et al., 2013). Basset et al. compared network properties obtained from both DTI and DSI in six healthy volunteers with results suggesting good agreement between the two acquisitions (Bassett et al., 2010). Zalesky et al. compared DTI and HARDI acquisitions in three volunteers (Zalesky et al., 2010b) at various network scales. They determined that HARDI tractography produced stronger inter-hemispheric connectivity than DTI and yielded more streamlines than DTI, but that some of these streamlines were spurious. Although network properties derived from HARDI and DTI were broadly similar there were significant differences at all the scales investigated. Vaessen et al. compared DTI acquisitions at three different angular resolutions (6, 15 and 32 directions) in six healthy volunteers. They concluded that angular resolution affects the resulting network properties, with the strength of connectivity found to increase with angular resolution. Typically, acquisitions have been made at 3 T field strength, whereas roughly one third of studies we reviewed used 1.5 T acquisition (Table 2.3). Whilst it is commonly accepted that acquisitions at 3 T or greater can improve the signal-to-noise, to our knowledge, the effect of field strength on the resulting networks has not been assessed in a connectome setting. Ultimately, there are many issues and caveats with dMRI acquisitions and preprocessing steps which must

Template-matching atlas	Atlas
AAL	Automatic anatomical labelling (SPM toolbox) (Tzourio-Mazoyer et al., 2002)
HO	Harvard-Oxford atlas (FSL)
IBASPM	Individual Brain Atlases Statistical Parametric Mapping (SPM toolbox) (Alemán-Gómez et al., 2005)
LPBA40	LONI Probabilistic Brain Atlas of 40 subjects (Shattuck et al., 2008)
WFU	Wake Forest University atlas (SPM toolbox) (Maldjian et al., 2003)

Cortical labelling	Atlas
FreeSurfer	Desikan-Killiany (Desikan et al., 2006)
	Destrieux atlas (Destrieux et al., 2010)
	Subcortical segmentation (Fischl et al., 2002, 2004a)

Table 2.4: Established neuroanatomical segmentation approaches.

be considered carefully to obtain an appropriate signal (Jones and Cercignani, 2010).

2.7.1.1 Neuroanatomical segmentation approaches

There is no universally accepted parcellation scheme for obtaining a set of cerebral regions to use as network nodes. Typically, ROIs are formed of a contiguous group of grey matter voxels, although potentially a single seed voxel may constitute a node. Table 2.4 shows a list of established segmentation algorithms/atlases used in dMRI network studies. Several segmentation approaches are based on template registration to fixed anatomical atlases (Maldjian et al., 2003; Shattuck et al., 2008; Tzourio-Mazoyer et al., 2002). The FreeSurfer image analysis suite has become very widely used in network studies and therefore plays an increasingly central role in connectome mapping. The FreeSurfer suite performs grey/white matter tissue estimation, parcellation of the cerebral cortex into units with respect to gyral and sulcal structure (Fischl et al., 2004b; Desikan et al., 2006) and a separate segmentation is used for sub-cortical grey matter structures (Fischl et al., 2002, 2004a). Typically, segmentation methods divide the cerebrum into 50–100 regions, roughly in agreement with Brodmann areas. However, finer grained parcellations have also been developed (Romero-Garcia et al., 2012). Figure 2.11 shows an illustration of two alternative cortical segmentations: the Individual Brain Atlases for Statistical Parametric Mapping (IBASPM; Alemán-Gómez et al. 2005) and the Desikan-Killiany atlas used by FreeSurfer (Desikan et al., 2006). We assess the test-retest performance of two FreeSurfer cortical atlases in Chapter 4.

Segmentation approaches produce different cerebral subdivisions and the volume of ROIs is not usually uniform. For example, structures such as the insula may have a voxel volume several times larger than that of sub-cortical structures such as the amygdala. Accordingly, appropriate normalisation of weighting may be needed (see Section 3.7.3).

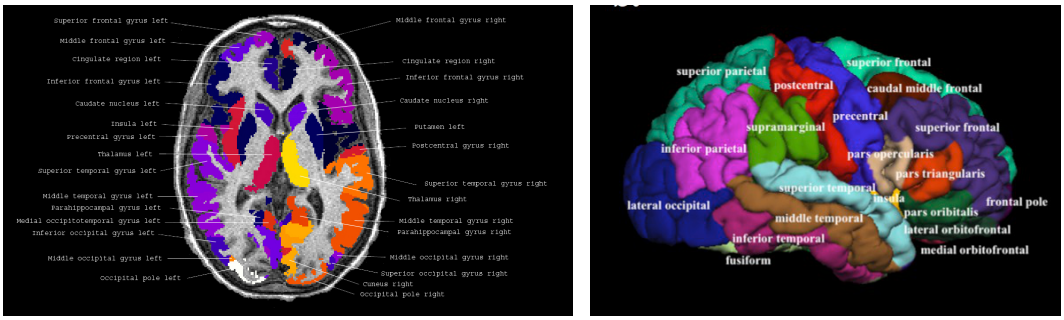

(a) (b)

Figure 2.11: Illustration of two alternative cortical segmentations: a) the IBASPM atlas; b) The FreeSurfer Desikan-Killiany atlas with 34 cortical structures per hemisphere.

However, if cortical correspondence is not a concern, then an alternative approach is to construct networks by dividing the cortex into a number of smaller and arbitrary regions (Bassett et al., 2010; Gigandet et al., 2008; Hagmann et al., 2007, 2008; Honey et al., 2008, 2009). These studies have divided cortical grey matter into roughly 500–4000 nodes of roughly uniform size, with the caveat that identifying spatial correspondence of nodes across subjects is challenging. Local network measures, such as individual node properties cannot meaningfully be compared between subjects. However, this approach has the potential to create a high resolution network, which is useful for estimating global organisational properties.

2.7.2 Tractography algorithms

There is no one established method of performing tractography in connectome studies. To date, researchers have employed a variety of alternative tractography approaches to compute connectivity between brain regions. In general, the deterministic approach estimates the best fit of the diffusion tensor model at each voxel, whereas the probabilistic approach estimates a distribution of possible orientations. There has, however, been a preference for using deterministic methods over probabilistic methods. Deterministic methods have, at the least, the advantage that they are less computationally demanding.

However, the test-retest reliability of these techniques, and how accurately tracts reflect the underlying axonal ground-truths is largely unstudied in terms of false positive and false negative connections. One comparative study compared seven tractography approaches in a connectome context (Yo et al., 2009). Their results suggests that fibre-crossing models identify more connections than a simple tensor model and that probabilistic approaches reveal greater number of connected regions but with lower connectivity values than deterministic methods. However, their study was cross-sectional and does not report test-retest agreement. However, Côté et al.

(2013) found that deterministic tractography produces less invalid tracts which leads to better connectivity results than probabilistic tractography. We compare the performance of deterministic and probabilistic approaches in Chapter 4.

2.7.2.1 Seeding approaches

There are two commonly used tractography seeding approaches, namely white matter (WM) and grey matter (GM) seeding. Under WM-seeding, tracking is initiated from white matter voxels and streamlines are constructed in two collinear directions until terminated by stopping criteria. Under the GM-seeding approach, tracking is instead initiated from grey matter voxels (within an ROI) and streamlines are constructed in a single direction until terminated by stopping criteria. Note that whole-brain seeding can be used to initiate seeding from both grey and white matter. Notably, a sub-class of grey matter seeding had also been employed where seeding is initiated from the white-grey matter interface (Vaessen et al., 2010; Robinson et al., 2010) where myelinated axons are expected to originate/terminate. However, robustly identifying this interface can be challenging. Though WM-seeding is the convention and the most commonly used approach, it is known to introduce a length bias in streamline generation (Hagmann et al., 2007). For instance, given an underlying white matter tract, WM-seeding will produce proportionally more streamlines within a long tract (given that more seed voxels will be placed along its length) than a shorter tract. Researchers have proposed streamline length normalisation to compensate for the accumulated tractography errors which increase with streamline length and to correct for a bias in repeatedly identifying long tracts when seeding from white matter (Hagmann et al., 2008). In contrast, GM-seeding, by only seeding from grey matter voxels is less affected by this length bias, but is affected by a seeding bias due the size of grey matter ROIs (if the number of seed points matches the number of voxels within a region). This too can be corrected by normalising for the number of seed points per region (Hagmann et al., 2008).

Additionally, seed points can be placed either systematically, for example at the centre of every voxel within a seed region, or randomly. Dense seeding methods seek to place multiple seed points within a voxel, which involves sub-voxel placement by interpolation of diffusion measures (Tournier et al., 2007, 2008). There are suggestions that random seed placement may be appropriate for estimating probabilistic connections. Cheng et al. (2012b) showed that seed density has an impact on the resulting structural networks from the local to global level. As the number of seeds were increased, the variance in network properties decreased. However, the authors observed that the network variance is also influenced by other imaging factors, because spurious fibres continue to affect nodal degrees and edge weights. The authors suggest that thresholding of the network weights is necessary to create an appropriate weighted network. Similarly, Côté et al. (2013) reported that multi-seeding has a large impact on tractography

outputs and that the seeding method should be used with care. In Chapter 4 we present a test-retest analysis of WM- and GM-seeding in terms of network reliability and we also assess the performance of the streamline length correction.

2.7.3 Registration

Both *linear* and *nonlinear* deformable transformations have been used to align grey matter regions from T_1 weighted space to diffusion space, or vice versa. Linear transforms can account for *translation*, *rotation*, *scaling* and *shear*, whereas nonlinear techniques (surface splines, multi-quadrics etc.) can account for local deformation or warping but are typically constrained to some degree. Usually a *mutual information* cost function must be used to align intensity gradients between modalities. Different registration algorithms may be suited to different registration targets, but perfect alignment is often not obtainable. Same-subject registration is a challenging open problem in neuroimaging, but also a well researched one. Klein et al. (2009) conducted an evaluation of fourteen nonlinear deformation algorithms applied to brain image registration using 8 different error measures for 80 manually labelled brains. One of the most significant findings of this study is that the relative performances of the registration methods under comparison appear to be little affected by the choice of subject population, labelling protocol, and type of overlap measure.

While registration is not the focus of this study, in Chapter 4 we assessed the test-retest registration agreement between two well-established registration methods, one linear and the other nonlinear. We quantified registration agreement in grey matter, white matter and cortical regions.

2.7.4 Definitions of connectivity

Network connectivity is typically recorded using the graph-theory notation of an *adjacency* (or connectivity) matrix, a, of size $n \times n$, where n is the number of nodes and the entry at position a_{ij} records a measure of association between node i and node j. Inter-regional associations are typically constructed from thousands of streamlines, although many millions can be estimated with probabilistic tractography. Several measures of association between nodes have been used. Most researchers choose to represent the associations in the adjacency matrix by the count of interconnecting streamlines (streamline density) between node i and node j. (Hagmann et al., 2007, 2008; Vaessen et al., 2010; Skudlarski et al., 2008; Zalesky et al., 2010a). However, others record the average tract length from the set of tracts connecting a pair of nodes (Hagmann et al., 2007). Typically, network weightings recording streamline density also undergo length and/or ROI size normalisation (refer to Section 3.7.3). Other associations between nodes are

possible, for example, measures of tissue microstructure, such as diffusion anisotropy, averaged along the length of each tract (Iturria-Medina et al., 2007; Robinson et al., 2010) or estimates of fibre bundle cross-sectional area (Zalesky and Fornito, 2009). Necessarily, these network weightings record different facets of connectivity. It remains unclear which weighting can be calculated most robustly and which best represents the underlying axonal connections.

In most cases, the result of network construction is an undirected positive-weighted graph. For structural networks, connectivity matrices are treated as symmetric because streamline tracking cannot distinguish between afferent and efferent connections. Although self-connections reflect the set of streamlines which start and terminate at the same node, this information is typically discarded from the adjacency matrix before performing network analysis. In Chapter 4 we present a test-retest analysis of three types of network weighting, two based on streamline density and a third based on diffusion anisotropy.

2.7.4.1 Thresholding

Network metrics have often been computed on thresholded and binarised networks, i.e. only connections above a certain strength are considered. Binary approaches discard any weight associated with a connection. The reasons for binarisation are two-fold: 1) thresholding and binarisation may remove noise in the connectivity values obtained; 2) A binary graph representation has the advantage that network metrics are more straightforward to compute. However, now that many network metrics also have weighted (and directed) variants (Rubinov and Sporns, 2010), weighted graphs can be assessed without the need for binarisation. One problem in applying thresholding is that after applying the same threshold to a varied set of adjacency matrices, this will likely result in matrices with different levels of sparsity. While techniques can be applied to match the sparsity across a set of matrices, this alters the weightings differently for each subject (Zalesky et al., 2010a). Approaches which act directly on the connection weighting without thresholding can be used without any bias arising from the setting of arbitrary thresholds, but may be adversely affected by spurious connections.

2.8 Network theory

Network metrics derived from graph theory may then be used to assess brain connectivity between individuals or across populations. Graph theory is the mathematical study of graphs or networks, which are used to define pairwise relationships between objects. Graph theory originated in geometry and was first popularly used in Euler's Königsberg bridge problem in the eighteenth century, a puzzle involving crossing each of seven bridges exactly once in the Prus-

Measure	Description
Node degree	Number of links connected to a node
Node strength	Variant of node degree incorporating connection weights
Network sparsity	Ratio of observed connections to all possible connections (inversely proportional to network degree)
Path length	Shortest path length (distance) in terms of connections, between a pair of nodes
Characteristic path length	Global measure of integration reflecting the average shortest path length between all pairs of nodes
Global efficiency	Global measure of integration inversely related to characteristic path length
Clustering coefficient	measure of how many local clusters exist around a node in terms of its direct neighbours
Transitivity	Normalised variant of clustering coefficient removing the bias due to nodes with low degree
Betweenness centrality	Fraction of shortest paths in the network that pass through a given node
Within module degree z-score	Nodal measure of intra-module connectivity
Participation coefficient	Nodal measure of inter-module connectivity
Small-world measure	Global measure comparing clustering and path length of a given network to an equivalent random network with the same degree distribution

Table 2.5: Summary of network properties

sian city of Königsberg. Network analysis has since been widely applied across many areas of science and for many kinds of both physical and abstract networks. Rubinov and Sporns (2010) have collated a comprehensive overview of network properties and their interpretations in terms of brain connectivity. Established metrics include *node degree*, *clustering coefficient*, *characteristic path length* and *efficiency*. Table 2.5 provides a summary of some commonly used network properties. It is important to note that values of many network measures are greatly influenced by basic network characteristics, such as the number of nodes and links, and the *degree distribution* (Barabasi and Albert, 1999). A *small-world network* is a network which is considerably more clustered than a random network and is simultaneously highly segregated and integrated (Watts and Strogatz, 1998). Mathematical definitions of the network measures used in this work are presented in Section 3.8 along with illustrated examples.

2.9 Network visualisation

A challenge facing connectome researchers is the search for a compact and coherent way to visualise thousands of network connections without losing essential information. Brain networks can been visualised as a connectivity matrix as shown in Figure 2.12(a), a graph with the anatomical location of nodes as shown in Figure 2.12(b), or an abstract graph such as a circle network as shown in Figure 2.12(c). Each representation has relative merits and disadvantages. Adjacency matrices are compact and convey the comparative strength of connectivity well, but are difficult to relate to anatomy. Anatomical networks visualise the anatomy well and can be overlaid against image volumes, but they can be difficult to interpret with networks of thou-

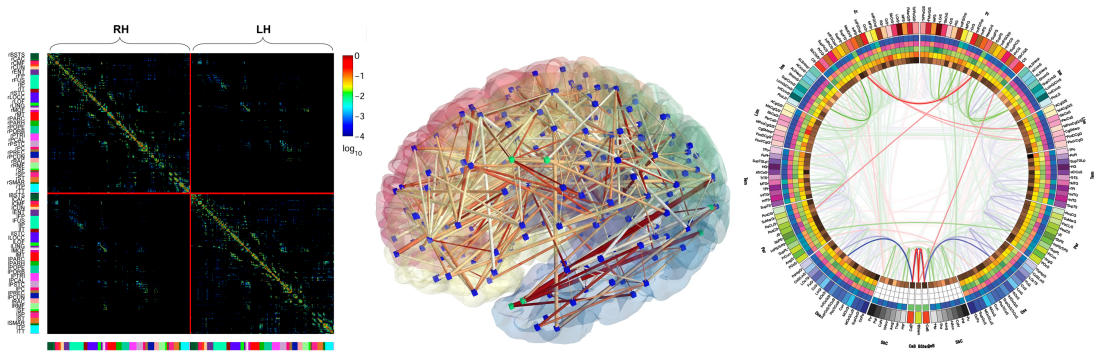

(a) (b) (c)

Figure 2.12: Three alternative network visualisations: a) An adjacency matrix showing the associations between nodes (Hagmann et al., 2008); b) An anatomical graph generated by the Connectome Viewer (Gerhard et al., 2011); c) A circular graph incorporating cortical node properties (Irimia et al., 2012).

sands of connections. Abstract networks have the advantage that nodes can be organised by nodal properties, such that highly interlinked nodes can be viewed adjacently, thus highlighting patterns of connectivity. In most cases, network interpretation requires some simplification in the amount and dimensionality of the network connections in order to display the essential components of brain organisation. It is often not helpful to plot each and every connection.

Neuroinformatics researchers have begun to develop connectome specific file formats and visualisation tools in an effort to create standard methods of storing, exchanging and viewing connectome data. For instance, Gerhard et al. created a "connectome file format", which specifies the required data to define a network (Gerhard et al., 2011). The Connectome Mapping Toolkit[7] is an accompanying software suite which makes use of the connectome file format to visualise networks and adjacency matrices. Additionally, Irimia et al. (2012) developed a circular network representation, termed a *connectogram*, similar to those used to visualise genomic relationships and with claims of a high "data-to-ink ratio". This representation can show network connections alongside cortical node properties, such as volume, area, cortical thickness and curvature, as seen in Figure 2.12(c).

2.10 Network analysis

Analysis of networks can be undertaken at three levels, namely *global*, *nodal* and *edge-wise*. For instance, although metrics can be averaged globally, many metrics can be indexed by node, such as the node degree, or by edge, such as the shortest path length. As a result, a large

[7] http://cmtk.org/viewer/

number of variables may be generated per network. The earliest connectome studies reported global network properties which describe overall organisation, such as small-world organisation. However, as global measures are an approximation of overall connectivity, there has been a subsequent drive towards localised assessment for both nodal properties and individual network connections. Note that global network measures have the advantage of reducing local unreliability by taking an average over all components in the network. On the other hand, global measures may be inherently limited when the network in question is composed of modules (sub-networks) with different organisation. For case-control and cross-sectional analyses, these network properties are often typically assessed by classical statistical approaches, such as ANOVA, ANCOVA or linear regression. However, even for a small network, applying statistical tests on nodal or edge-wise variables must correct for multiple comparisons.

Indeed, whole-brain network analyses, potentially involving thousand of edge-wise comparisons, have received special attention in the form of mass-univariate testing (Zalesky et al., 2010a). Network based statistics (NBS) a non-parametric statistical test used to isolate the components from a $n \times n$ connectivity matrix that differ significantly between two populations, or for regression analysis (Zalesky et al., 2010a). NBS controls the family-wise error rate in the weak sense, where each of the $n(n-1)/2$ edges are tested independently. NBS may offer greater statistical power than conventional procedures for controlling the family-wise error rate, such as the false discovery rate. Hypothesis testing can be performed using the *general linear model* in order to accommodate a number of different statistical models, such as ANOVA and linear regression. The authors demonstrated this method using resting-state fMRI of 12 patients with chronic schizophrenia and 15 healthy volunteers. The method isolated a 'dysconnected subnetwork' in the schizophrenic group which comprised fronto-temporal and occipito-temporal dysconnections. NBS has since been adopted by other researchers and applied in schizophrenia (Zalesky et al., 2010a, 2011), geriatric depression and mild cognitive impairment (Bai et al., 2012) and in an ALS cohort (Verstraete et al., 2011, 2013).

In contrast, rather than using conventional hypothesis testing, some researchers have employed machine learning and dimensionally reduction approaches using cross-validation, for both classification (Robinson et al., 2010) and regression analysis of structural networks (De Boer et al., 2010). These studies have identified discriminatory connections which were used to successfully classify unseen subjects by age and gender.

2.11 Networks of the brain

Many connectome studies have demonstrated that the human brain and the brains of other mammals, have various non-trivial properties in their organisation. For instance, the brain is

characterised by highly connected hub nodes, a modular structure and small-world organisation (Hagmann et al., 2008; Honey et al., 2008; Sporns, 2011b; Van Den Heuvel and Sporns, 2011; Yan et al., 2011). The brain network is highly organised and hierarchical, far beyond the organisation expected of a randomly connected network. Several hub nodes, such as the precuneus and posterior cingulate, participate in many connections and are critical in interhemispheric communication. Other, so called 'provincial hubs' participate in many local connections, but often have a low degree of connectivity outside their local sub-network.

Researchers have reported an exponential edge weight distribution (Hagmann et al., 2007), which shows that there are many low weight connections but very few strong connections within the cortex. Hagmann et al. also reported the tract length distribution, which showed that there are many short connections in the cortex but very few long range connections. The authors suggest these findings indicate that the cortex has developed to optimise total wiring length. Various neuroimaging studies have found the human brain to possess a *scale-free property* (Barabasi and Albert, 1999), particularly in networks of functional connectivity (Van Den Heuvel et al., 2008).

In a seminal study of network properties, Watts and Strogatz, by measuring a high degree of local clustering combined with an average short path length between any pair of nodes, determined that many natural networks have a small-world structure (Watts and Strogatz, 1998). In such networks most nodes are not direct neighbours but can be reached in a small number of steps. A network may be considered small-world if it has both: a) a network clustering coefficient much greater than that of an equivalent randomly connected graph; b) the characteristic path length is comparable to that of random graph with the same number of nodes and edges. Rubinov and Sporns (2010) suggest that the significance of network statistics should be established by comparison with statistics calculated on null-hypothesis networks. Null-hypothesis networks have simple random or ordered topologies but preserve basic characteristics of the original network. The most commonly used null-hypothesis network has a random topology but shares the size, density and binary degree distribution of the original network.

Small-world properties have been reported at various scales in the human brain for both functional and structural networks (Gong et al., 2009a; Hagmann et al., 2007; Iturria-Medina et al., 2008; Vaessen et al., 2010) studies. Researchers have suggested that small-world properties are largely independent of the choice of nodes and spatial resolution. This indicates the presence of highly connected modular sub-networks that are globally interconnected with relatively few long range connections (Hagmann et al., 2010b). Network *centrality* is a measure of the proportion of shortest paths passing through a node. Centrality can be a reliable indicator of a network hub which interconnects sub-networks. Hagmann et al. identified the posterior cingulate and precuneus as hubs nodes in the human cortex (Hagmann et al., 2008). In a widely cited

study, Gong et al. investigated the effect of age and gender on the structural networks obtained from dMRI and tractography in a group of 95 subjects aged from 19 to 85 years (Gong et al., 2009a). They too identified the precuneus and posterior cingulate gyrus as highly connected hub nodes independent of age or gender. The loss or degeneration of such hubs have found to be disruptive to network communication (Honey et al., 2008). However, whether the brain's organisation is different in the ageing or diseased brain, particularly in areas of pathology, remains an open question.

2.12 Relation between brain structure and function

Perhaps, the ultimate purpose of the connectome approach is to relate structural connectivity to brain function. Indeed, functional connectivity must be explained by brain structure in some way. If the human brain can be reliably segmented into a comprehensive structural network, it should be expected that these structural connections will correspond to the functional subdivisions at some level (Hagmann et al., 2010a; Bassett and Gazzaniga, 2011).

Behavioural data provides the basis for relating brain circuits to individual differences in cognition and perception. Although techniques such as electroencephalography and magnetoencephalography can be used to measure functional brain activity, connectome research has primarily focussed on fMRI and resting-state networks (RSNs). fMRI is an MRI acquisition that indirectly measures brain activity, usually by detecting associated changes via the *blood-oxygen-level dependent* contrast – the change in magnetisation between oxygen-rich and oxygen-poor blood. This methods relies on the assumption that cerebral blood flow is coupled with neuronal activation, for which there is good evidence (Logothetis et al., 2001). The spatial resolution of fMRI is roughly equivalent to dMRI and the temporal resolution is in the order of seconds. fMRI has proved useful in the understanding of higher cognitive functions and in the study of neurological conditions, such as Alzheimer's disease, dementia and schizophrenia. Functional connectivity may be determined by short-term temporal correlations in activity between spatially distinct brain regions (van den Heuvel and Hulshoff Pol, 2010). Like structural networks, these interactions can be recorded as pairwise correspondences, which reflect the level of functional communication between network nodes. fMRI experiments have revealed a number of functional networks which are consistently found in healthy subjects and for various functional tasks. RSNs are based on regional interactions that occur when a subject is awake but not performing an explicit task.

Several attempts have been made to reconcile functional connectivity derived from fMRI with structural connectivity derived from dMRI and tractography (Hagmann et al., 2008; Skudlarski et al., 2008; Greicius et al., 2009; Honey et al., 2009; Van Den Heuvel et al., 2009). For

instance, Hagmann et al. (2008) computed both structural networks from DSI and tractography and functional networks from a resting-state fMRI experiment, and then determined that the strengths of the structural connections were highly predictive ($r^2 = 0.62$) of the strengths of functional connections for five human volunteers. Similarly, Greicius et al. (2009) found that resting state functional connectivity reflects structural connectivity (estimated from DTI and tractography) in the default mode network, a network of brain regions including the medial prefrontal cortex, medial temporal lobes and posterior cingulate cortex. These findings show that, in part, structural brain organisation explains functional connectivity, but there are limitations with the spatial and temporal resolution of fMRI. There are also concerns about measurement noise and whether the fMRI signal may be adversely affected by non-neuronal artefacts. Although the relationship between structure and function is far from settled, it is clear that a comprehensive map of the structural connections of the brain is a prerequisite for understanding the basis of functional connectivity.

2.13 Critique of current techniques

Recent test-retest findings suggest that current connectome mapping techniques are adequate for reliably measuring global network measures (Bassett et al., 2010; Cammoun et al., 2011; Hagmann et al., 2008; Vaessen et al., 2010; Zalesky et al., 2010b). However, the reliability of regional network measures are not yet well studied, an aspect we seek to address in this work. Any unreliability in network measures may lead to concerns about the validity of studies based on such measures. Inter-session differences in network measures have been reported previously (Vaessen et al., 2010; Bassett et al., 2010). Some variation may be due to scanner noise and inhomogeneities between sessions and some may be due to systematic variation in processing. Evidently, many intermediate steps are involved in generating structural networks from dMRI data. Given concerns about local unreliability, it may prove informative to assess test-retest agreement at key intermediate stages in network construction, specifically registration, neuroanatomical segmentation and tractography.

The connectome approach is challenging because often a correspondence between network nodes across subjects is required. Achieving neuroanatomical segmentation requires at least: 1) robust anatomical segmentation of the cortex with high test-retest reliability; 2) accurate local registration of ROIs with sub-centimetre spatial error. As the performance of these two requirements degrades the lower the signal-to-noise ratio in the resulting network becomes.

It should be noted that the task of automated cortical labelling is extremely challenging. Error may arise from registration error and further error may arise from systematic variation in the parcellation procedure. Consequently, any subsequent processing based on these segmentations

will be affected by these sources of error. The choice of nodes is hugely important (Hagmann et al., 2007; Zalesky et al., 2010b). We cannot expect to map tracts with a diameter smaller than the imaging resolution. Increasing the number of nodes and making the overall ROI size smaller would increase the mapping resolution but will increase measurement noise. However, using too few large nodes may fail to capture genuine structural connections. Ideally, we wish to establish a network resolution where the contributions from the genuine connections are measurably larger than the noise contribution.

The ability of tractography to estimate accurately the location and orientation of the underlying tracts has not been comprehensively validated in the human brain. Large test-retest variation in tractography parameters may reflect errors in estimating the underlying axonal fibre tracts within a region, with some variation due to both segmentation error and methodological issues in streamline construction. In addition, determining endpoints of a streamline is a well-known weakness of current tractography techniques (Jbabdi and Johansen-Berg, 2011).

Additionally, some commentators have raised fundamental concerns with the purpose of the connectome approach (Jabr, 2012). With contemporary MRI, tissue structure can be resolved at around one millimetre resolution. However, approximately forty thousand neurons may be within 1 mm^3 of cortex (Crick and Jones, 1993), meaning the state-of-the-art resolution estimates of connectivity are far from the neuronal level and must be considered somewhat crude approximations of the underlying structure. It remains difficult to reconcile the disparity between a 90 billion node neuronal network at the mesoscopic scale and a one hundred node network obtained from MRI. It could even be argued that current MRI networks are entirely incompatible with cellular neuroscience.

Arguably, even a mesoscopic connectome does not contain all the necessary information to understand the brain. Firstly, a connectome is merely a snapshot of the brain at a particular time and does not capture the ongoing changes in connectivity or the unique physiology of individual neurons, which may be embodied in various functions. Secondly, it records no information on the glial cells, numerous cells which were originally thought to just help and nourish neurons, which some researchers now believe contribute to information processing (Araque et al., 1999). Finally, determining by any known method how signals pass through a system as complex as the mammalian brain at the cellular and synaptic level, currently seems so intractable that it is unlikely that it can be achieved in the near future. Nevertheless, despite these criticisms we believe that a low resolution macroscopic map derived from MRI will undoubtedly be useful in understanding fundamental brain connectivity and function.

2.14 Motivation

In the following chapters we seek to address some of the concerns relating to brain networks obtained from MRI by evaluating several key steps and representations in the construction of networks. We intend to test the assumptions that current techniques are sufficient to divide grey matter structures into appropriate anatomical nodes and that tractography is capable of reliably and repeatedly measuring the connectivity between these brain regions.

Chapter 3

Methods

3.1 Overview

We developed an automated connectivity mapping pipeline to construct white matter structural networks from dMRI data. Where appropriate we used freely available neuroimaging toolkits to process both T_1-weighted and dMRI data. One aim of this research was to compare a number of alternative techniques and parameter settings at crucial stages in network construction. For example, we assessed linear and nonlinear methods of image registration, alternative cortical atlases, alternative definitions of connectivity and various network weightings. A conceptual view of network construction is presented in Figure 3.1.

For each subject we acquired 3D T_1-weighted volume scans and dMRI data in the same session. Our method involved automated brain extraction from 3D T_1-weighted volumes, grey and white matter segmentation and parcellation of the cortex into 84, 85 or 164 ROIs per subject based on anatomical boundaries (FreeSurfer). A registration step was necessary to align these cortical ROIs to dMRI space, and to a lesser extent correct for systematic EPI distortions associated with the dMRI protocol. Both deterministic and probabilistic tractography was initiated from dMRI voxels using established algorithms with tracts constructed from voxel to voxel until terminated by specific stopping criteria. Structural connections between all brain areas were computed by quantifying tracts found between all ROI pairs in an adjacency matrix. A number of weighting strategies were considered to correct for between-subject differences in brain and ROI size. For each network, we computed various graph-theoretic measures which may be used to characterise patterns of connectivity.

In the remainder of this chapter we discuss the various steps of network construction and analyses used in this work, with particular focus on the tractography techniques, network weightings and network measures.

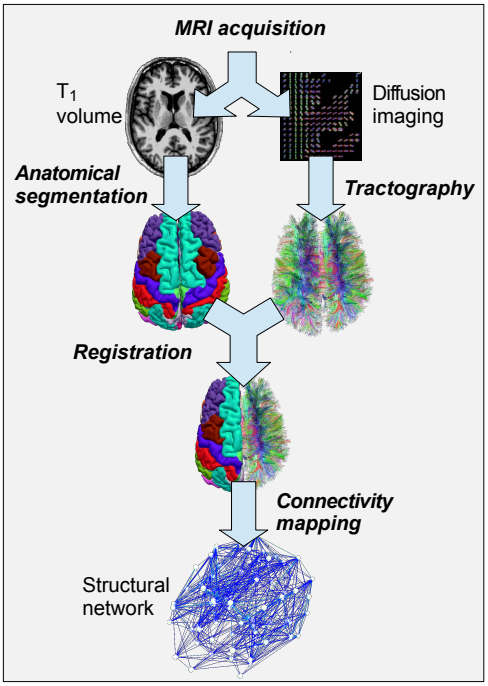

Figure 3.1: Conceptual overview of the connectivity mapping pipeline.

3.2 MRI acquisition

Our work made use of three MRI datasets gathered at the Brain Research Imaging Centre, University of Edinburgh. The first of these was a test-retest MRI dataset involving ten healthy volunteers aged between 50 and 58 years who were scanned twice using an identical protocol 2–3 days apart (Principal Investigator: Dr. Cyril Pernet, University of Edinburgh). The study was approved by the South East Scotland Research Ethics Committee 01. The second study was a case-control ALS dataset involving 30 ALS patients and 30 age- and education-matched healthy controls (Principal Investigator: Prof. Sharon Abrahams, University of Edinburgh). The study was approved by the National Health Service Scotland Research Ethics Committee. The third study was a cross-sectional dataset of 80 healthy, right-handed, volunteers aged between 25 and 65 years (Principal Investigator: Dr. Mark Bastin, University of Edinburgh). The study was approved by the Lothian Research Ethics Committees (REC 05/S1104/45). These three datasets are described in the following chapters. All participants provided informed consent.

All imaging data were acquired using a GE Signa Horizon HDxt 1.5 T clinical scanner (General Electric, Milwaukee, WI, USA) equipped with a self-shielding gradient set (33 mT/m maximum gradient strength) and manufacturer supplied 8-channel phased-array head coil. For the dMRI protocol, single-shot spin-echo echo-planar (EP) diffusion-weighted whole-brain vol-

umes ($b = 1000\ s\ \text{mm}^{-2}$) were acquired with diffusion gradients applied in 64 non-collinear directions, along with seven T_2-weighted ($b = 0\ s\ \text{mm}^{-2}$) volumes (Jones et al., 2002). The repetition and echo times were 16.5 s and 98.3 ms, respectively. Seventy-two contiguous axial 2 mm thick slices were acquired resulting in 2 mm isotropic voxels. In the same session, high resolution 3D T_1-weighted inversion-recovery prepared, fast spoiled gradient-echo volumes were acquired in the coronal plane with 180 contiguous 1.3 mm thick slices resulting in voxel dimensions of $1 \times 1 \times 1.3$ mm. The acquisition took approximately 20 minutes.

3.3 Neuroanatomical segmentation

In order to form a corresponding set of network nodes across subjects it was necessary to divide each T_1-weighted brain into distinct neuroanatomical regions. The choice of nodes is an important consideration for network analysis (Zalesky et al., 2010b). For this purpose, volumetric segmentation and cortical reconstruction was performed with the FreeSurfer image analysis suite[1]. This set of tools performed automated skull stripping and brain extraction (Ségonne et al., 2004), intensity normalisation (Sled et al., 1998), automated Talairach transformation, grey and white matter segmentation, and surface parcellation including complete labelling of cortical sulci and gyri (Desikan et al., 2006; Destrieux et al., 2010; Fischl et al., 2004b). T_1-weighted volumes were conformed to 1 mm isotropic voxels. These processing steps were fully automated using default parameters (FreeSurfer version 5.3.0). Brain extraction and tissue segmentation were visually checked for each subject. FreeSurfer performed segmentation of sub-cortical grey matter structures in a separate processing stream from the cortical segmentation (Fischl et al., 2002, 2004a). Eight subcortical structures were retained per hemisphere, which were the thalamus, hippocampus, amygdala, caudate, putamen, pallidum, accumbens area and ventral diencephalon. In some experiments the brain stem was also used. FreeSurfer morphometric procedures have been demonstrated to show good test-retest reliability across scanner manufacturers and across field strengths (Han et al., 2006; Wonderlick et al., 2009; Reuter et al., 2012), although other studies have shown discrepancies in test-retest reliability (Morey et al., 2010; Gronenschild et al., 2012).

3.3.1 Cortical atlases

Two cortical atlases provided by the FreeSurfer package were used. These were the Desikan-Killiany atlas (Desikan et al., 2006) with 34 cortical structures per hemisphere (Figure 3.2), and the Destrieux atlas (Destrieux et al., 2010) with 74 sulco-gyral structures per hemisphere.

[1] `http://surfer.nmr.mgh.harvard.edu` (Martinos Center for Biomedical Imaging)

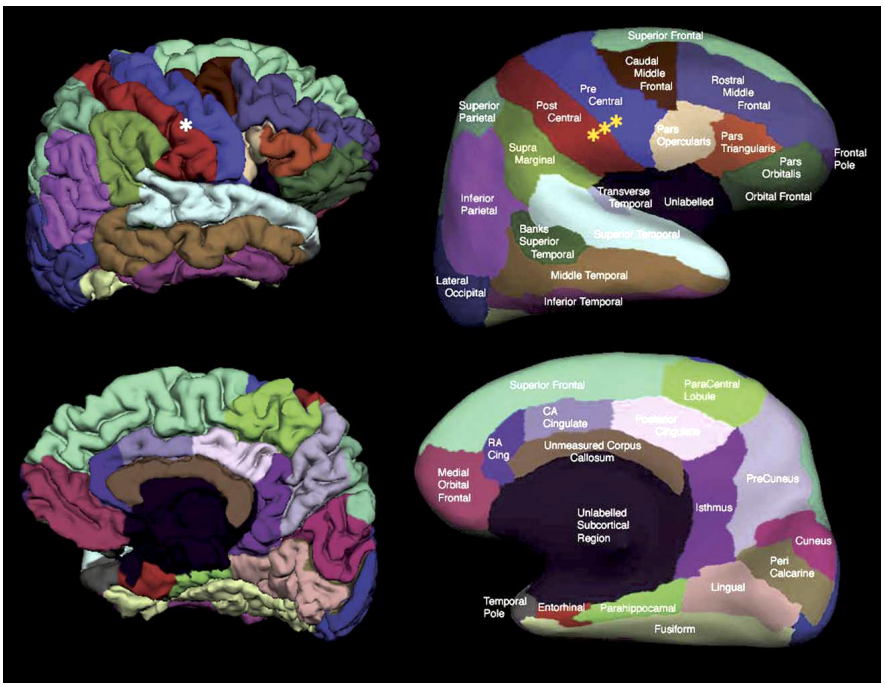

Figure 3.2: Deskian-Killiany cortical atlas: Pial (left) and inflated (right) cortical representations of the regions of interest in one hemisphere. The top row illustrates the lateral view of the hemisphere while the bottom row shows the medial view of the hemisphere. Adapted from (Desikan et al., 2006).

In brief, cortical parcellation is achieved using spatial intensity gradients and by learning associated rules across tissue classes following a supervised probabilistic approach trained with a manually labelled dataset. The Desikan-Killiany atlas is based on a set of 40 manual annotations of MRI scans from three age groups (19–24, 41–57 and 71–86 years of age) with a range of atrophy. The intention in using this varied set of training data is that the supervised learning procedure will generalise to a wide range of ages and be robust to age-related degeneration. Similarly, the Destrieux atlas is based on regions derived from sulcal and gyral boundaries from a manually labelled set of 12 volunteers (18–33 years of age).

We created parcellations using both atlases to allow comparison between node configuration, with the assumption that the Desikan-Killiany atlas may generate more reliable node properties because using fewer and larger regions minimises the effects of image noise, registration artefact and systematic error. Following anatomical segmentation, 84 or 164 grey matter regions were retained per subject. Both configurations included the same 16 sub-cortical structures - only the cortical ROIs differ. For those experiments including the brain stem, 85 or 165 nodes were used per subject. For each subject, the ROIs were labelled and stored in 3D volumes. For each grey matter region, the total ROI volume, in T_1-weighted space, computed by the parcellation procedure was recorded.

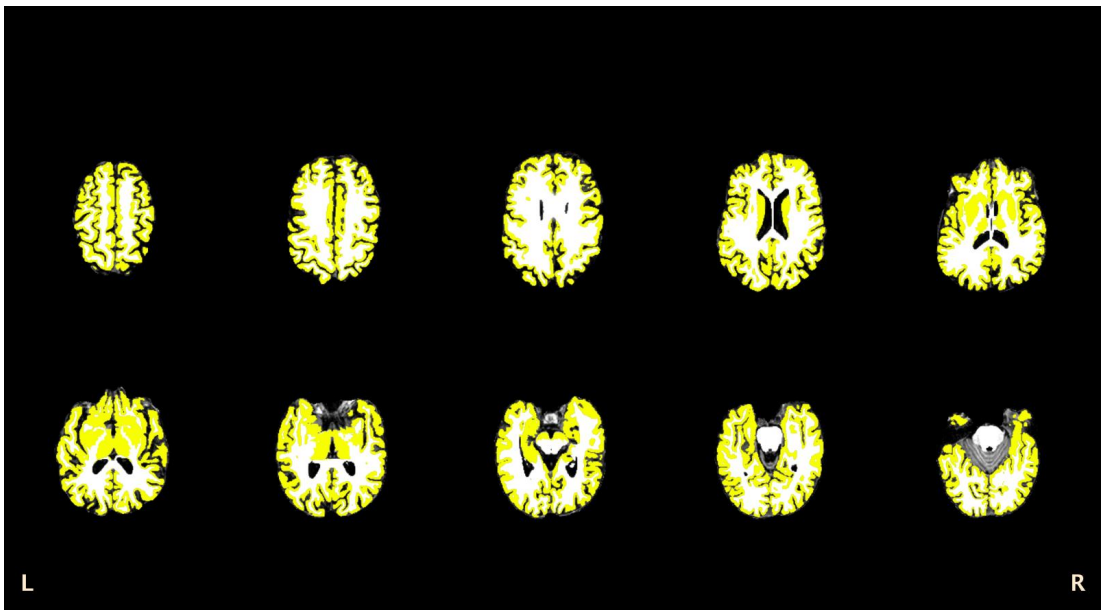

Figure 3.3: Transverse slices showing the grey and white matter masks obtained from the segmentation procedure, overlaid on the T_1-weighted extracted brain (57 year old male).

3.3.2 Mask construction

Additionally, the results of the segmentation procedure were used to construct three binary mask volumes per subject, namely a grey matter mask, a white matter mask and whole-brain tracking mask. The grey and white matter mask were constructed from the results of the tissue segmentation procedure (Fischl et al., 2004b). By definition these tissue classes do not overlap. The cerebellum was removed from both tissue masks as were any regions of CSF identified by the segmentation procedure, including the left and right lateral ventricles, and the third and fourth ventricles. The whole-brain tracking mask was constructed from a union of the grey and white matter masks. These masks were used in subsequent processing to constrain tractography. Note that under this masking approach the spatial limits of tracking are dependent on the accuracy of the segmentation procedure and on the subsequent registration procedure. Figure 3.3 shows the grey and white matter masks for one subject.

3.4 Diffusion processing

Using tools available in the FSL toolkit[2], dMRI (EPI) data were corrected for patient motion and eddy-current induced distortions observed in single-shot EPI volumes. Correction was based on a simple model for eddy-current artefacts that using an affine transform with 12

[2] `http://www.fmrib.ox.ac.uk/fsl/` (FMRIB, Oxford University)

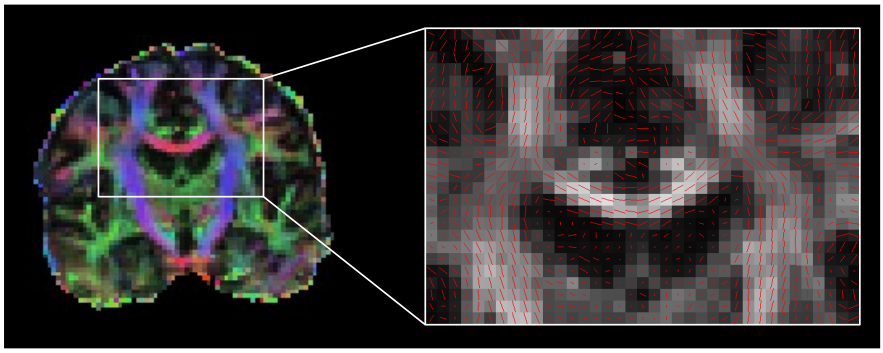

Figure 3.4: Coronal slice of the colour-encoded principal diffusion directions and the corresponding FA values with lines representing the principal orientations and magnitudes of diffusion (57 year old male).

degrees-of-freedom (DOF). EPI volumes were corrected by estimating an affine transform to the first T_2-weighted ($b = 0\ s\ \mathrm{mm}^{-2}$) volume of each subject (Jenkinson and Smith, 2001).

Diffusion tensors were then fitted at each voxel and diffusion imaging parameters, $\langle \mathrm{D} \rangle$ and FA were calculated by Equation 2.6 and 2.7, respectively. Figure 3.4 illustrates the principal orientations and magnitudes of diffusion for one subject. All diffusion parameters and subsequent tractography were computed in native space. Skull stripping and brain extraction was performed (Smith, 2002) on T_2-weighted volumes acquired along with the diffusion-weighted data and applied to the FA volume of each session (Figure 3.5).

3.5 Registration

A registration step was necessary to align each subject's anatomical ROIs from T_1-weighted space to diffusion space. Registration was performed using established algorithms available in FSL. For comparison, we computed both linear and nonlinear transformations. In each case the transform was computed from dMRI space to T_1 weighted space and the inverse transforms computed.

3.5.1 Linear registration

Firstly, we employed a method of computing an affine transform with 12 degrees of freedom (DOF). These 12 parameters define a linear transformation based on *translation*, *rotation*, *scale* and *shear* in three dimensions. A transformation matrix and coordinate vectors, x, y, z, can be multiplied to generate the transformed coordinates x', y', z',

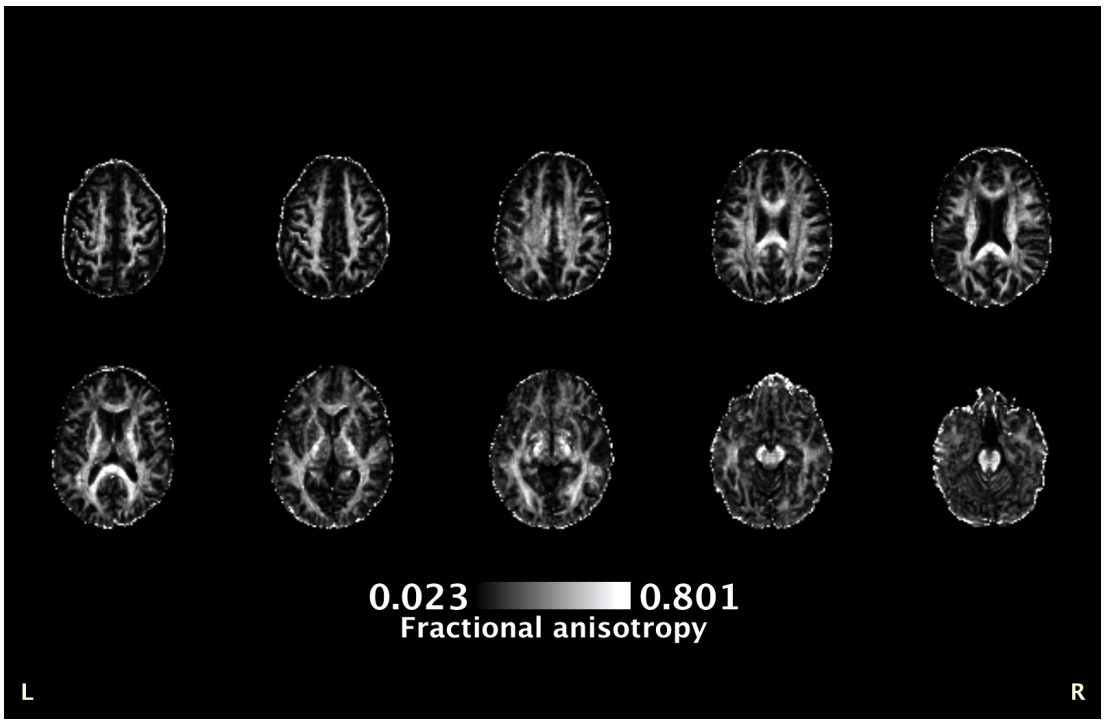

Figure 3.5: Transverse slices of an FA volume showing the estimated diffusion anisotropy at each voxel following skull stripping (51 year old male).

$$\begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{pmatrix} m_{11} & m_{12} & m_{13} & m_{14} \\ m_{21} & m_{22} & m_{23} & m_{24} \\ m_{31} & m_{32} & m_{33} & m_{34} \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}. \tag{3.1}$$

We used an established linear registration algorithm (FLIRT, Jenkinson and Smith 2001; Jenkinson et al. 2002) to estimate an alignment of each brain-extracted *FA* volume to the corresponding FreeSurfer extracted brain. Notably, using the brain extracted FA volume as the registration target was found to improve registration results in comparison to using a T_2-weighted extracted brain (see Section 4.6.2). In principle, same-subject registration only requires 6 DOF (translation, rotation and scale) as a rigid body transform should capture all same-subject transforms. However, in practice 12 DOF was found to improve upon 6 DOF as it allowed further correction of the EP-induced distortions observed in the dMRI volumes. We used a *mutual information* cost function, suitable for cross-modal registration which was minimised to find the optimal set of transformation parameters (Jenkinson and Smith, 2001),

$$H(X,Y) - H(X) - H(Y). \tag{3.2}$$

3.5.2 Nonlinear registration

We also used a nonlinear deformation field based method to refine local alignment (FNIRT, Andersson et al. 2007). The general transformation for nonlinear registration takes the form,

$$\begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \mathbf{M} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} + \begin{pmatrix} d_x(x,y,z) \\ d_y(x,y,z) \\ d_z(x,y,z) \\ 1 \end{pmatrix}, \tag{3.3}$$

where $\mathbf{M}$ is an affine transform matrix as described above and $d_i(x,y,z)$ are the displacement fields, which describe the displacement direction for a sampling point at each coordinate. The algorithm uses splines to determine the deformation field and performs a constrained optimisation of local deformations based on a mutual information cost function using the *FA* parametric map as a registration target. This method requires a sensible initialisation to avoid failed registrations (local minima in the search space), for which we used the affine transformation estimated by the linear registration described above. The inverse warp-field was then computed to obtain the required transforms from T_1-weighted space to diffusion space.

3.5.3 Resampling

Cerebral segmentations and grey matter, white matter and tracking masks were aligned to diffusion space by applying these transforms using nearest neighbour resampling. For visual inspection, each T_1-weighted extracted brain was also aligned to diffusion space using trilinear resampling. These masks were then used to constrain tractography to white matter structures and targeted cerebral voxels.

3.6 Tractography

Whole-brain tractography was performed using two established tractography algorithms: one based on deterministic tensor tractography (Mori et al., 1999), and a probabilistic tractography algorithm modelling two fibre directions at each voxel (Behrens et al., 2003b, 2007). The deterministic approach estimates the best fit of the diffusion tensor model at each voxel, whereas the probabilistic approach estimates a distribution of possible orientations. For both algorithms we incorporated additional streamline termination criteria as described below.

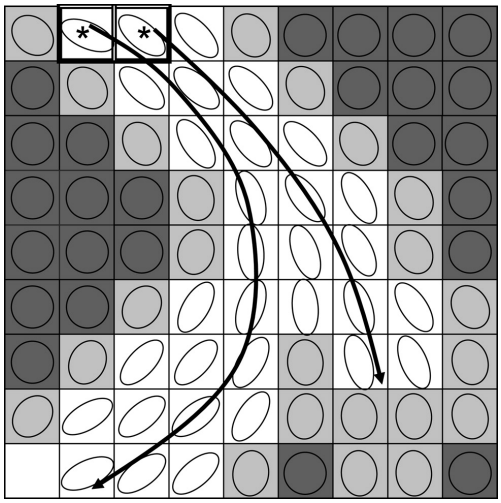

Figure 3.6: An illustration of a deterministic streamline tracking algorithm in 2D. Two streamlines follow the maximal directions of diffusion from a seed point. Adapted from (Mori and Zhang, 2006).

3.6.1 Deterministic tensor tractography

For deterministic tracking we used a modified version of FACT (Mori et al., 1999) with additional streamline termination criteria using the Camino diffusion toolkit[3] (Cook et al., 2006). The tensor model of diffusion (Section 2.4.3) determined the principal orientation and magnitude of diffusion at a voxel location, by the first eigenvector and the associated eigenvalue. Streamlines were simply constructed from voxel-to-voxel using local path integration, such that the principal directions estimated from the diffusion tensor are coherently aligned along fibre tracts of maximal diffusion. Each streamline was constructed from a specified seed voxel in this way, until terminated by the stopping criteria. Figure 3.6 shows an illustration of deterministic streamline tracking in two dimensions.

3.6.2 Probabilistic tractography

Probabilistic tractography made use of the fully model-based method in FSL's diffusion toolkit (FDT), known as Bayesian estimation of diffusion parameters obtained using sampling techniques (BEDPOST; Behrens et al. 2003b). The distributions for tracking were generated with a two-fibre model per voxel (Behrens et al., 2007). This method makes a selection between the single or two-fibre case based on evidence from the diffusion data at each voxel, in order to simplify tracking with multiple fibre orientations. Behrens et al. have previously demonstrated that multi-fibre tractography improves sensitivity in whole-brain fibre tracking but does not dra-

[3] `http://cmic.cs.ucl.ac.uk/camino/` (Microstructure Imaging Group, University College London)

matically change tractography for major white matter pathways with high diffusion anisotropy, compared to a single-fibre model.

This fibre model is a partial volume model, such that the diffusion signal for each fibre orientation is a mixture of two Gaussians, divided into an infinitely anisotropic component and a single isotropic component. The predicted diffusion weighted signal, μ_i was modelled as,

$$\mu_i = S_0 \left((1 - \sum_{j=1}^{N} f_j) \exp(-b_i D) + \sum_{j=1}^{N} f_j \exp(-b_i D r_i^T R_j A R_j^T r_i r_i) \right), \tag{3.4}$$

along a gradient direction r_i with b-value b_i associated with the ith diffusion encoding direction. S_0 is the signal with no diffusion weighting, D is the diffusivity, f_j is the fraction of signal and $R_j A R_j^T$ is the anisotropic diffusion tensor along the jth fibre orientation and $N = 2$ (the maximum number of fibres). Consequently, A is fixed as,

$$A = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \tag{3.5}$$

and R_j rotates A to align with the fibre direction of the voxel, requiring two angle parameters (θ_j, ϕ_j). The noise component is modelled separately as an independently identically distributed Gaussian with a mean of zero and standard deviation of σ. The full model parameter list is $\omega = (S_0, D, f_1, f_2, \theta_1, \theta_2, \phi_1, \phi_2, \sigma, \eta)$, where η is the width of a Beta distribution used in model selection. The probability of observing the data at each voxel Y, given the model M and the set of parameters ω is,

$$P(Y|\omega, M) = \prod_{i=1}^{n} P(Y_i|\omega, M) \tag{3.6}$$

where,

$$P(Y_i|\omega, M) \sim \mathcal{N}(\mu_i, \sigma). \tag{3.7}$$

The authors used a number of non-informative prior distributions on these model parameters, as described in Behrens et al. (2007). The posterior distribution over these parameters is given by Bayes' rule,

$$P(\omega|Y, M) = \frac{P(Y|\omega, M) P(\omega|M)}{P(Y|M)} \tag{3.8}$$

Bayesian estimation was used to fit the parameters of the model to the signal at each voxel. Behrens et al. used a model selection technique called *automatic relevance determination* based on evidence from the diffusion data to select between the single and two-fibre models, in order to avoid poor estimation by fitting an inappropriate fibre model. As the model estimates cannot be solved analytically, estimation was performed using Metropolis-Hastings Markov chain Monte Carlo (MCMC) sampling after marginalising over the model parameters.

Probabilistic tracking (ProbTrack) proceeds by sampling from these distributions repeatedly during streamline tracking. Starting at a specified seed voxel, we used the algorithm to select a random sample (θ_j, ϕ_j) from $P(\theta_j, \phi_j|M)$ and then propagate the streamline in the selected direction by 0.5 mm. Sampling was repeated in this manner at each location until tracking was terminated by the stopping criteria. Interpolation in tracking was used as described in Behrens et al. (2003b) for which a sample is drawn from one of two adjacent voxels in each dimension. We repeated the procedure for generating streamlines 100 times for each seed point. Tracking results in a set of probabilistic streamlines with many possible end points per seed point. This spatial distribution of streamlines permits a connection probability to be calculated, which describes the likelihood of global connection between any two locations.

3.6.3 Seeding

We employed two alternative seeding approaches, namely *WM-seeding* and *GM-seeding*. Under WM-seeding, tracking was initiated from all white matter voxels and streamlines were constructed in two collinear directions until terminated by the stopping criteria. Under the GM-seeding approach, tracking was initiated from all grey matter voxels (within an ROI) and streamlines were constructed in a single direction until terminated by the same stopping criteria.

3.6.4 Termination criteria

The termination criteria was the same for both tractography algorithms and both seeding configurations. Tracking of a streamline was terminated by any of the following constraints:

1. entering a voxel with FA below 0.1,
2. entering an extra-cerebral voxel outside the tracking mask,
3. exceeding a threshold on the in-plane angle from voxel to voxel,
4. exceeding a *distance ratio metric* of 10.

The distance ratio metric (DM) is a simple measure of tortuosity defined as a ratio between the actual path length along a curve, l, and the linear distance between start point and end point, c,

$$\mathrm{DM} = \frac{l}{c}. \tag{3.9}$$

This measure has previously been used to determine the tortuosity of cerebral blood vessels (Bullitt et al., 2003), but in this work was used to identify an implausible accumulation of turns along white matter tracts, which are believed to be efficiently wired (Sporns, 2011a). We

empirically set the threshold at DM > 10, as this was observed to filter the most implausible streamlines but did not impact on the sharp turning 'U-fibres' observed in the corpus callosum. Note that the DM measures an accumulation of turns, but the in-plane angle constraint can only identify a single unacceptable turn along a streamline.

Four thresholds on the in-plane angle curvature were investigated in Section 4.6.5. The value of the anisotropy constraint was set empirically. We determined that an FA threshold > 0.1 can in some cases be too conservative for whole-brain tracking and a threshold of 0.1 was preferable to ensure tracking into grey matter regions. Other network studies have also used such a threshold (Verstraete et al., 2011). Although this threshold is low when compared to conventional tractography, there remains a need for a threshold on the minimum acceptable anisotropy in order to minimise false streamlines in highly isotropic areas of diffusion, for example, in cortical grey matter.

3.6.5 Streamline post-processing

Streamlines may be further filtered depending on the type of macro-connectivity under analysis. We used two waypoint constraints in order to minimise the number of spurious streamlines. For some experiments an entire streamline was discarded if the total path length in white matter was: 1) less than a minimum length; 2) exceeded a maximum length. The motivation for the minimum *waypoint* constraint is to filter short implausible streamlines, which pass from one grey matter region to the another, without passing through white matter, which is possible under GM-seeding. A number of limits on the minimum waypoint length were explored in Chapter 4. The maximum length constraint was set at 20 cm for all experiments and any streamline with a white matter length exceeding this was considered erroneous.

3.6.6 Streamline density

For inspection of tractography data before the connectivity mapping procedure we computed anatomical connectivity maps (ACMs), which record the streamline density per voxel (Embleton et al., 2007). For each voxel in diffusion space, the streamline density is computed as the number of streamlines which pass through that voxel divided by the total number of streamlines seeded (per subject). Such maps (Figure 3.7) have been used to visualise patterns of connectivity and allow quantification of both global and regional differences between tractography algorithms.

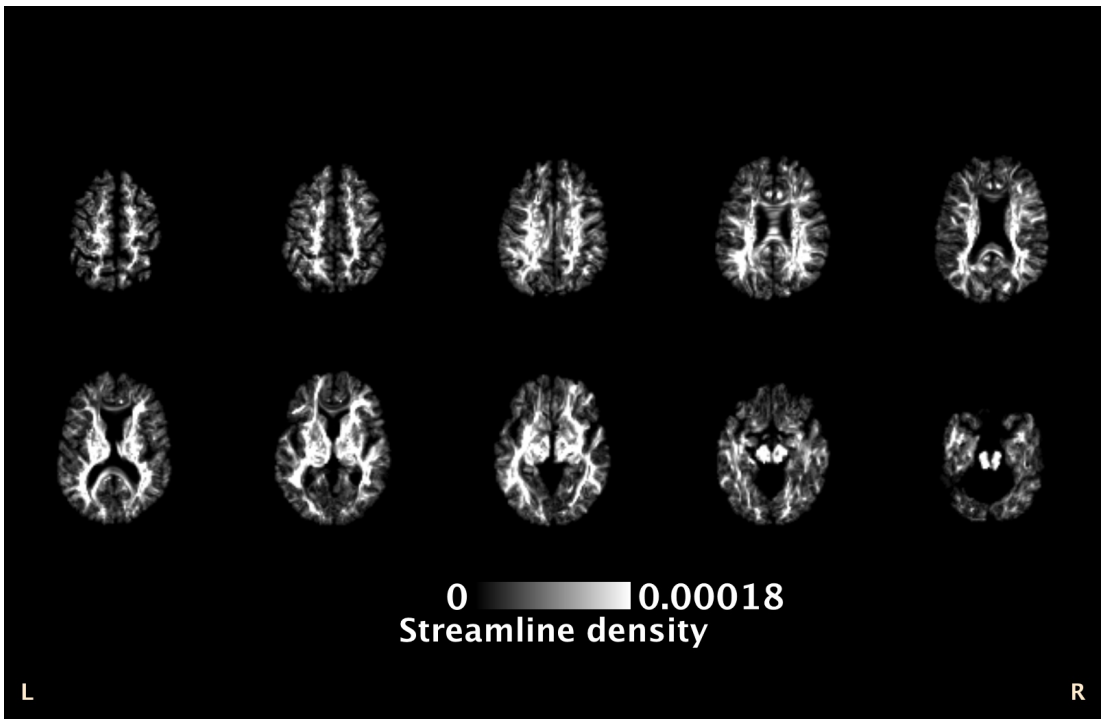

Figure 3.7: Transverse slices indicating the streamline density at each voxel at each voxel (51 year old male). In this case streamlines were computed with 100 iterations of probabilistic tractography using WM-seeding.

3.7 Network construction

Connections between all ROIs (network nodes) may then be computed by quantifying the tracts reconstructed between all possible pairs of nodes. The pairwise relationship between a pair of nodes defines the topological connection.

3.7.1 Streamline connectivity

Network connections were determined by identifying the endpoint of each streamline. Note that the endpoints are strongly affected by the somewhat arbitrary stopping criteria during tracking. Streamline termination is a well-known weakness of current tractography techniques (Jbabdi and Johansen-Berg, 2011); there is no means to prevent some streamlines from either ending prematurely or extending past the true endpoint.

Figure 3.8 illustrates three streamline configurations. Only in the simplest case (**s1**) do we have a single unambiguous connection between two nodes. However, in many cases (**s2** and **s3**), there are a number of alternative ways to consider the connectivity between regions. We implemented four alternative methods for determining the connectivity between grey matter nodes on

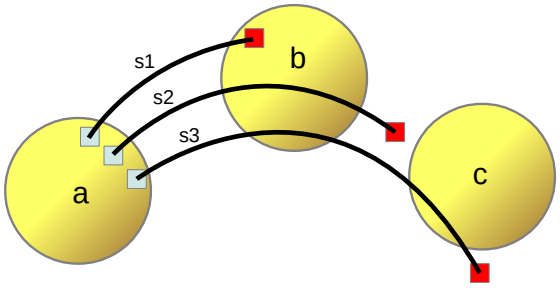

Figure 3.8: Illustration of three example streamlines (**s1**, **s2**, **s3**) and regions (**a**, **b**, **c**): Seed points are highlighted and "absolute" termination points are shown in red. Streamline **s1** unambiguously connects regions **a** and **b**. However, in the case of **s2** and **s3**, there are a number of alternative ways to consider the connectivity between regions.

a streamline, which we termed *absolute*, *first*, *last* and *complete*. The above definitions generate different measures of connectivity in the (not uncommon) case that a streamline apparently passes through multiple grey matter ROIs before it terminates. In each case a streamline is tracked from the seed point. Under *absolute* connectivity the endpoint is simply the ROI (if any) at the position of streamline termination. Under *first* connectivity the endpoint is considered to be the first ROI encountered when tracking from the seed location and any subsequent ROI encountered on the streamline is discounted. Under *last* connectivity the endpoint is considered to be the last ROI encountered when tracking from the seed location and any preceding ROI encountered is discounted. Under "complete" connectivity each ROI falling along the entire length of a streamline is considered to be interconnected (with equal weight). In all cases, streamlines which do not start and end within an ROI are discarded. The four definitions are quantitatively compared in Chapter 4.

3.7.2 The adjacency matrix

For each subject, connections were recorded in a $n \times n$ adjacency matrix a where the entry a_{ij} denotes the measure of connection 'strength' between node i and node j.

3.7.3 Network weighting

In this work, three types of network weighting were considered – two based on streamline density and a third on tract-averaged FA. The first weighting, termed *streamline density* (SD-weighted), records the interconnecting streamline density corrected for ROI size,

$$a_{ij} = \frac{2}{g_i + g_j} |S_{ij}|, \tag{3.10}$$

where S_{ij} is the set of all streamlines found between node i and node j (and $S_{ij} = S_{ji}$), and g_i and g_j are the number of grey matter voxels in nodes i and j, respectively. The second weighting, termed *streamline density with length correction* (SDL-weighted), again records streamline density but with a correction for streamline length (Hagmann et al., 2008),

$$a_{ij} = \frac{2}{g_i + g_j} \sum_{s \in S_{ij}} \frac{1}{l(s)}, \tag{3.11}$$

where $l(s)$ is the length of streamline s between node i and node j. The rationale for the normalisation by g (Eq. 3.10, 3.11) is to correct for between-subject variability in grey matter volume, since the number of possible entry/exit points per region is proportional to grey matter volume (Hagmann et al., 2008). For example, MRI studies have shown differences in brain volume by gender and age (Good et al., 2001; Resnick et al., 2003; Sullivan et al., 2004). Additionally, inter-subject variation in GM volume may affect the number of tracts seeded by tractography. As a result, it may be necessary to normalise weightings in order to allow representative comparison of connectivity between subjects. The rationale for streamline length normalisation (Eq. 3.11) is twofold: 1) to compensate for the accumulated tractography errors which increase with streamline length; and 2) to correct for a bias in repeatedly identifying long tracts when seeding from white matter (Hagmann et al., 2008). *FA-weighted* networks were constructed from the same set of streamlines by recording the mean FA value along interconnecting streamlines. Each entry in the adjacency matrix was computed,

$$a_{ij} = \frac{1}{|S_{ij}|} \sum_{s \in S_{ij}} \frac{\sum_{v \in V_s} \mathrm{FA}(v)}{m_s}, \tag{3.12}$$

where V_s is the set of voxels (of size m_s) found along the streamline s between node i and node j, and FA measures the diffusion anisotropy per voxel.

For each type of weighting, the result is an undirected positive-weighted graph. Each matrix is symmetric by definition as $S_{ij} = S_{ji}$ and tractography cannot distinguish between afferent and efferent connections. Self-connections, a_{ii}, were removed (set to zero). While self-connections reveal something about the grey-white matter interface of each node, these values are typically discarded for network analyses as they have no role in assessing inter-connections[4].

[4] A large number of self-connections may indicate that the granularity of grey matter ROIs is too course and streamlines of short length are misinterpreted as self-loops.

3.8 Network properties

Considering the entire brain as a network of nodes enables use of graph-theoretic measures to assess the connections obtained from tractography. Network properties may then be used to characterise the overall organisation of the brain at a global, nodal or edge-wise level (Rubinov and Sporns, 2010). Where possible, network properties were computed using weighted variants rather than binarising the adjacency matrices.

Before specifying network measures we first define some basic notation (Rubinov and Sporns, 2010). N is the set of all nodes in the network, n is the number of nodes and a is an undirected weighted graph, stored as a $n \times n$ adjacency matrix, where a_{ij} indicates the connection weight between node i and node j. Note that, in this work a graph and an adjacency matrix are equivalent representations of a network. A graph is defined as $G = \{V, E\}$, where V is a set of vertices (nodes) and E is the set of edges (connections) which connect a pair of nodes.

Each weighted adjacency matrix has a corresponding binarised adjacency matrix b where

$$b_{ij} = \begin{cases} 0 & \text{if } a_{ij} = 0, \\ 1 & \text{if } a_{ij} > 0. \end{cases} \tag{3.13}$$

3.8.1 Thresholding

We considered two alternative methods for thresholding of network connections, which may be used to remove weak or spurious connections. A threshold on network weights may be computed,

$$\hat{a}_{ij} = \begin{cases} 0 & \text{if } a_{ij} < \tau, \\ a_{ij} & \text{if } a_{ij} \geq \tau, \end{cases} \tag{3.14}$$

where τ is the network weight at a quantile of interest. Alternatively, a threshold on network connections may be computed by proportion,

$$\hat{a}_{ij} = \begin{cases} 0 & \text{if } \Pi_{ij} < \pi, \\ a_{ij} & \text{if } \Pi_{ij} \geq \pi, \end{cases} \tag{3.15}$$

where the value Π_{ij} is the proportion of subjects for which $a_{ij} > 0$, and π is the threshold on this proportion.

3.8.2 Basic properties

The *degree* of a node i, measures the number of connecting links (see Figure 3.9(a)),

$$k_i = \sum_{j \in N} b_{ij}. \tag{3.16}$$

The *mean node degree* (network density), is a commonly used measure of global connectivity or total "wiring cost",

$$K = \frac{1}{n} \sum_{i \in N} k_i. \tag{3.17}$$

For a weighted graph, the *strength* of a node i is an analogous measure to the node degree incorporating weighting,

$$w_i = \sum_{j \in N} a_{ij}. \tag{3.18}$$

The *mean node strength* (network strength), is calculated

$$W = \frac{1}{n} \sum_{i \in N} w_i. \tag{3.19}$$

The *network sparsity* is calculated,

$$S = \frac{n^2 - n - \sum_{i \in N} k_i}{n^2 - n}, \tag{3.20}$$

reflecting the count of non-zero edges in the observed network, where S is between 0 and 1 and is inversely proportional to the network degree.

3.8.3 The distance matrix

From the adjacency matrix a distance matrix, d, was constructed, recording the *shortest weighted path length* (distance), between node i and node j,

$$d_{ij} = \sum_{a_{uv} \in g_{i \leftrightarrow j}} f(a_{uv}), \tag{3.21}$$

where f is a mapping from weight to length and $g_{i \leftrightarrow j}$ is the shortest weighted path between i and j (Rubinov and Sporns, 2010). The shortest weighted paths can be computed efficiently by Dijkstra's algorithm (Dijkstra, 1959). For instance, we used an inverse mapping from weight to length to reflect that tracts with higher connection probability are 'stronger' links. Note that, although there may exist many paths between a pair of nodes, this measure only considers the shortest, i.e. the path with smallest cumulative weighting (see Figure 3.9(b)).

Many established network measures are derived from the adjacency and distance matrices. The following metrics may be categorised as measures of *integration*, *segregation*, *centrality*, *modularity* or those measuring other properties such as 'small-worldness'.

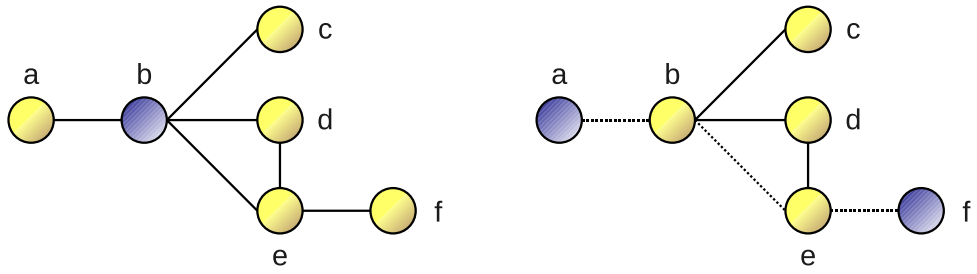

(a) (b)

Figure 3.9: Network example where for simplicity, all edges have a weight of one: a) Node **b** has four direct neighbours and therefore has a node degree of four. Overall, the mean node degree of this network is $\frac{12}{6} = 2$. The clustering coefficient of node **b** is the number of existing connections betweens node **b**'s neighbours (**d-e**) divided by all possible connections (**a-c, a-d, a-e, c-d, c-e, d-e**) between the same set of neighbours, in this case, $\frac{1}{6}$. The mean clustering coefficient of the network is $\frac{1}{3}$; b) The shortest path length between node **a** and node **f** is three, indicated by the dashed line. The characteristic path length (mean shortest path length) of the network is $\frac{52}{36} = \frac{13}{9}$.

3.8.4 Measures of integration

The weighted *characteristic path length* (Watts and Strogatz, 1998) of a network,

$$L = \frac{1}{n} \sum_{i \in N} \frac{\sum_{j \in N, j \neq i} d_{ij}}{n-1}, \tag{3.22}$$

is the average shortest weighted path length between all pairs of nodes in the network and is a commonly used measured of integration. For instance, a small characteristic path length indicates an average short distance between all pairs of nodes.

The *global efficiency* (Latora and Marchiori, 2001) of a network is computed,

$$E = \frac{1}{n} \sum_{i \in N} \frac{\sum_{j \in N, j \neq i} d_{ij}^{-1}}{n-1}, \tag{3.23}$$

and is inversely related to the characteristic path length.

3.8.5 Measures of segregation

Considering a node's set of directly connected neighbours, the weighted *clustering coefficient* (Watts and Strogatz, 1998; Onnela et al., 2005) of a node i is calculated,

$$c_i = \frac{1}{n} \sum_{i \in N} \frac{\sum_{j,h \in N} (a_{ij} a_{ih} a_{jh})^{1/3}}{k_i (k_i - 1)}. \tag{3.24}$$

The clustering coefficient measures the number of actual edges connecting neighbours divided by the maximum number of possible edges between the same set of neighbours. It may be interpreted as a measure of how many local clusters exist in the network. For instance, a high clustering coefficient indicates that a node's neighbours are also well-connected to each other, i.e. they form a cluster. The *mean clustering coefficient* is a global measure of a network's clustering,

$$C = \frac{1}{n}\sum_{i\in N} C_i. \tag{3.25}$$

The *transitivity* of a network (Newman, 2003), is a normalised variant of the clustering coefficient removing bias due to nodes with low degree and is computed,

$$T = \frac{\sum_{i\in N}\sum_{j,h\in N}(a_{ij}a_{ih}a_{jh})^{1/3}}{\sum_{i\in N} k_i(k_i-1)}. \tag{3.26}$$

3.8.6 Measures of centrality

Measures of centrality may be used to identify highly connected nodes, such as network hubs. The *betweenness centrality* (Freeman, 1979) of a node i is,

$$b_i = \frac{1}{(n-1)(n-2)}\sum_{h,j\in N, h\neq j, h\neq i, j\neq i}\frac{\rho_{hj}(i)}{\rho_{hj}}, \tag{3.27}$$

where ρ_{hj} is the number of shortest paths between h and j and $\rho_{hj}(i)$ is the number of shortest paths between h and j that pass through i. Essentially, betweenness centrality measures the fraction of shortest paths in the network that pass through a given node (see Figure 3.10).

3.8.7 Measures of modularity

The *within-module degree z-score* and *module participation coefficient* (Guimerà and Amaral, 2005) measure intra- and inter-module properties by defining sub-networks (modules). For modular measures the network is fully divided into a set M of non-overlapping modules, e.g. by lobe or hemisphere, where m is the number of modules.

The *within-module degree z-score* (Guimerà and Amaral, 2005) of node i is is a statistical measure of within-module degree centrality,

$$z_i = \frac{k_i(m_i) - \bar{k}(m_i)}{\sigma^k(m_i)}. \tag{3.28}$$

The *participation coefficient* (Guimerà and Amaral, 2005) of node i,

$$y_i = 1 - \sum_{m\in M}\left(\frac{k_i(m_i)}{k_i}\right)^2, \tag{3.29}$$

is a measure of inter-module connectivity, reflecting the diversity of connections of a node to nodes in other modules (see Figure 3.10).

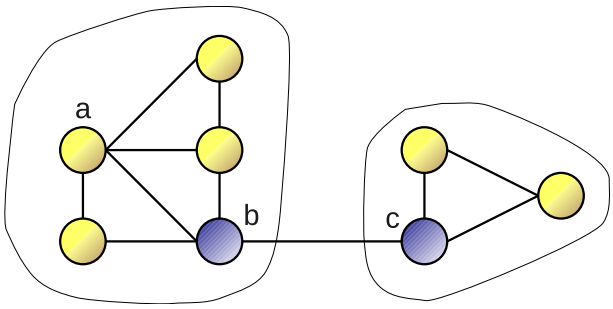

Figure 3.10: Network example with two modules. Node **a** has a betweenness centrality of seven as there are seven shortest paths passing through **a**. Likewise, node **b** has a betweenness centrality of 25, the greatest of any node in the network. Consequently, node **b** may be considered a 'hub' node. In this modular network only nodes **b** and **c** have a non-zero participation coefficient, i.e. the other nodes do not link beyond the module they belong to.

3.8.8 The small-world property

A measure of network small-worldness (Humphries and Gurney, 2008),

$$S = \frac{C/C_r}{L/L_r}, \tag{3.30}$$

where C is the clustering coefficient and L the characteristic path length of a network and C_r is the clustering coefficient and L_r the characteristic path length of a random network, usually constructed with node degree matching the network in question. Small-world networks typically have $S > 1$. For example, networks associated with ageing or disease may become either increasingly small world or random.

3.9 Network-based statistics

NBS is a non-parametric technique for identifying network differences (Zalesky et al., 2010a). NBS performs mass-univariate testing at each network connection and exploits the connectedness of network components to offer a potential gain in statistical power (through permutation test). NBS controls the family-wise error rate, in the weak sense, when the null hypothesis is tested independently at each of the $n(n-1)/2$ edges comprising the connectivity matrix. NBS can provide greater statistical power than conventional procedures for controlling the family-wise error rate, such as the false discovery rate, if the set of edges at which the null hypothesis is rejected constitutes a large component or components.

For example, to perform a two group contrast, NBS comprises fours steps as illustrated in Figure 3.11): 1) Perform a two-sample t-test at each edge independently to test the hypothesis that the value of connectivity between the two populations come from distributions with equal

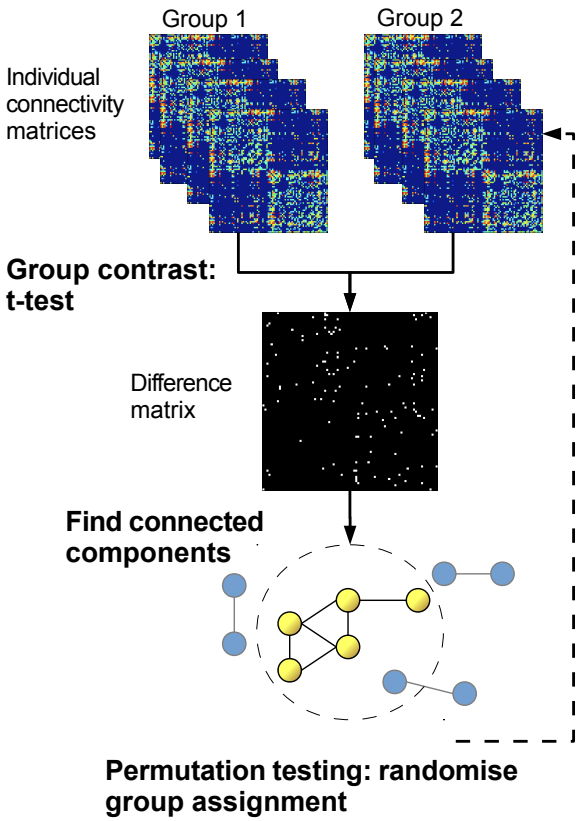

Figure 3.11: Conceptual overview of group-contrast using NBS.

means. 2) Threshold the t-statistic available at each edge to form a set of suprathreshold edges and obtain a *difference matrix*. 3) Identify any connected network components (subnetworks) in the difference matrix defined by the set of suprathreshold edges. These are referred to as observed components. Compute the magnitude of each observed component identified. 4) Repeat steps 1-3 k times, each time randomly permuting members of the two populations and storing the size of the largest component identified for each permutation. This yields an empirical estimate of the null distribution of maximal component size. A corrected p-value for each observed component cans then calculated from the null distributions.

3.10 Summation

In this chapter we have described a varied set of methods for constructing brain networks from structural and diffusion MRI data. In particular, we have focussed on methods for plausible streamline generation and have considered a number of viable network configurations and weightings. This sets the scene for a thorough assessment of the reliability of such networks.

Chapter 4

Network test-retest analysis

4.1 Overview

Whilst several previous studies have demonstrated structural brain networks obtained from healthy volunteers, only a small subset of those have assessed the reliability of the resulting networks and currently there is a lack of assessment concerning the reproducibility of these approaches. In this chapter, we report findings from experiments in which a number of factors affecting network construction were varied and assessed using repeat scans of ten healthy volunteers.

Firstly, independently of tractography and network construction, we assessed the test-retest error of two alternative registration procedures (linear and nonlinear) and the overlap agreement of two cortical parcellations. Using network test-retest measures, we then compared two alternative tractography algorithms (deterministic and probabilistic), two seeding approaches (grey and white matter), four alternative definitions of connectivity and three alternative network weightings (streamline density, streamline density with length correction and a measure of tract-averaged diffusion anisotropy). For each network configuration, we then quantified the reliability of four graph-theoretic measures using the intraclass correlation coefficient (ICC) and by comparing within- and between-subject average differences. Since these measures are an essential prerequisite for more complex analyses, such as small-world measures or the identification of network hubs, their reliability is crucial to the ultimate interpretation of such networks. We also investigated whether false connections could be reduced by an anatomically motivated filtering of streamlines by minimum length in white matter. In addition, we assessed two strategies on thresholding of connections by network weight. In the course of this work, portions of the findings in this chapter have been published previously (Buchanan et al., 2014a) and were adapted and extended below.

4.2 Participants

Ten healthy volunteers (six female) aged between 50 and 58 years underwent a dMRI protocol on two separate occasions over an interval of either two or three days (Gorgolewski et al., 2013)[1]. The subjects had a narrow age range to minimise the possible confound of increasing age on connectivity and diffusion anisotropy values. The study was approved by the local research ethics committee and informed consent was obtained from each subject.

4.3 Network construction

Following the methodology described in Chapter 3, various network configurations were constructed and subsequently assessed in a test-retest analysis. All subjects underwent the MRI protocol, as described in Section 3.2. For each T_1-weighted brain, cortical grey matter regions were identified for both the Desikan-Killiany and Destrieux atlases. Note that the brain stem was not used as a network node in the following experiments, which resulted in either 84 or 164 network nodes. We then applied the two alternative registration approaches for aligning the neuroanatomical segmentations from T_1-weighted space to dMRI space.

Following pre-processing and tractography, whole-brain networks were constructed as follows. For deterministic and probabilistic tractography algorithms we employed both WM-seeding and GM-seeding. Connections between regions were computed by identifying the streamlines connecting each pair of grey matter ROIs. The four alternative types of streamline connectivity, defined in Section 3.7.1, were computed for each set of streamlines subject to passing through at least one white matter voxel. In all cases, streamlines which did not connect between ROIs were discarded. Networks were computed for 13 different thresholds of streamline filtering by minimum contiguous length in white matter, from 0 to 6.0 mm in increments of 0.5 mm. For instance, a threshold of l mm discards any streamline which does not pass through at least l mm in white matter between grey matter ROIs. The white matter regions obtained from FreeSurfer were used as the waypoint mask.

The streamline termination constraint for curvature was empirically tested at four levels independently of network construction. From each set of resulting streamlines, *SD-weighted*, *SDL-weighted* and *FA-weighted* networks were then computed. In each case, connections were recorded in an $n \times n$ adjacency matrix, where the entry a_{ij} denotes the connection (edge) weight between node i and node j. The networks obtained by varying these factors described above were then systemically assessed in terms of test-retest performance. We note that as the sample size is small ($N = 10$), it is difficult to determine significant differences between network

[1] data available to download from GigaDB: `http://gigadb.org`

configurations due to the lack of statistical power. Nevertheless, differences and trends are discussed where appropriate.

Note that the steps for anatomical filtering of streamlines by maximum length, distance ratio, and anatomy were not used in this study. Additionally, thresholding of the network weights was not performed, except where described, as we wished to establish the test-retest reliability before using arbitrary constraints to reduce spurious streamlines. However, a post-hoc analysis of the effect of two thresholding strategies is presented in Section 4.6.9.

4.3.1 Network measures

For each connectivity matrix, four basic network measures were then computed as defined in Section 3.8. These were the *network degree* (average number of connections per node), the *network strength* (the average sum of weights per node), *network clustering coefficient* (an average measure of local connectivity), and the *characteristic path length* (mean of all shortest path lengths). Each of these four measures were computed per node for each subject. Four global network measures were then computed by taking the mean of these nodal values.

4.3.2 Thresholding of network weights

We assessed the two thresholding strategies, as described in Section 3.8.1, and assessed the effect on nodal measures. For these experiments, thresholds were computed for the three network weightings using the best performing tractography algorithm and seeding strategy. In the case of the threshold by network weight, thresholds were computed at 100 intervals between 0 and the maximum network weight. Rather than an absolute threshold on connection strength, the threshold was computed per subject at each percentile, i.e. the weights below the kth-percentile were discarded. For the threshold by proportion of subjects, thresholds were computed at 10 intervals between 0 and 1 (as there are 10 subjects in the cohort).

4.4 Test-retest statistics

For each network metric, two measures of agreement between sessions were computed, the values of ICC and a comparison of the within- versus between-subject differences.

4.4.1 The intraclass correlation coefficient

The intraclass correlation (ICC) was originally formulated for assessing multiple raters in measuring the same quantity (Shrout and Fleiss, 1979); here it was used to measure the same quantity over two sessions by estimating the proportion of between-subject variation relative to the total variation. Following the notation of Shrout and Fleiss, we computed ICC(3,1) using two-way mixed single measures using *consistency* of measurements between sessions, rather than *absolute* agreement,

$$\mathrm{ICC}(3,1) = \frac{\sigma_{BS}^2 - \sigma_\varepsilon^2}{\sigma_{BS}^2 + \sigma_\varepsilon^2}, \tag{4.1}$$

where σ_{BS}^2 is the between-subject variance and σ_ε^2 is the residual variance. The difference between consistency and absolute agreement is defined in terms of how the systematic variability due to raters (sessions) is treated.

For the unbiased formulation of the ICC used here values can range between -1 and 1. Negative values indicate greater within-subject variance than between-subject variance. Regarding test-retest analysis: ICC values < 0.5 reflect poor test-retest agreement; values between 0.5 and 0.8 reflect good agreement; and values ≥ 0.8 reflect excellent agreement.

4.4.2 Estimating within- and between-subject variation

For a paired set of N subject-specific measures, $X_1, \ldots, X_N$ and $Y_1, \ldots, Y_N$, the absolute *within-subject differences* were computed,

$$\delta_i^{WS} = |X_i - Y_i|. \tag{4.2}$$

The average *between-subject differences* (average differences of each subject against the others) were computed,

$$\delta_i^{BS} = \frac{1}{N-1} \sum_{j=1, i \neq j}^{N} |X_i - Y_j|. \tag{4.3}$$

Regional within- and between-subject differences were computed per node for each of the four network measures. Global within- and between-subject differences were computed from the global network properties. In each case, to test that $\delta^{WS} < \delta^{BS}$, a percentile bootstrap of the mean differences was used to compare the within- and between-subject components. The mean differences between the within- and between-subject components,

$$\delta = \frac{1}{N} \sum_{i}^{N} (\delta_i^{BS} - \delta_i^{WS}), \tag{4.4}$$

were computed over 5000 iterations, each time resampling with replacement from the original samples. From the distributions over these 5000 iterations, p-values were calculated as the

number of times that δ was larger (or smaller) than zero. Simultaneous probability coverage and correction for multiple comparisons was obtained by adjusting the alpha level following Wilcox (2005).

The range of δ^{BS} and δ^{WS} are dependent on the measure being assessed, but may be expressed as a percentage of the maximum value observed in X or Y. A between-subject difference $\leq$ the within-subject difference indicates poor test-retest agreement, whereas a between-subject difference $>$ the within-subject difference reflects good test-retest agreement – the larger the difference, δ, the stronger the agreement.

4.5 Test-retest variability at intermediate steps

We estimated the test-retest variability at three intermediate stages in the network mapping pipeline, namely registration, segmentation, and tractography. Regional estimates of agreement were computed from the correspondences between each pair of subject-specific ROI masks obtained from the FreeSurfer procedure. Global estimates of agreement were computed from correspondences between cerebral masks obtained from the segmentation procedure.

4.5.1 Segmentation agreement

By pairing each subject's segmentations across sessions, ROIs were compared on a voxel-wise basis in MNI-305 space[2]. In addition, the total ROI tissue volume was recorded. For each region, segmentation overlap was computed by the Dice coefficient, defined as the size of intersection of two regions divided by the mean region size, resulting in a score between 0 (no agreement) and 1 (perfect agreement).

$$D(A,B) = \frac{2|A \cap B|}{|A| + |B|}. \tag{4.5}$$

Dice coefficients were computed for the total grey matter, the total white matter, the 16 subcortical regions and all grey matter regions obtained for the two alternative cortical atlases.

4.5.2 Registration agreement

For each subject, the T_1-weighted extracted brains from both sessions were registered to the FA volume of the first session using the procedure described in Section 2.5. Spatial alignment between the aligned T_1-weighted pair was then assessed by the voxel-to-voxel correspondence

[2] `http://www.bic.mni.mcgill.ca/ServicesAtlases/MNI305` (Montreal Neurological Institute and Hospital, McGill University)

using an intensity-based normalised correlation cost function suitable for intra-modal registration (Jenkinson et al., 2002). This was calculated,

$$NC = \frac{\sum(X.Y)}{\sqrt{\sum X^2}\sqrt{\sum Y^2}}, \tag{4.6}$$

where X and Y are the reference and target images, each represented as a set of intensities. For completeness, the T_1-weighted extracted brains were also registered to the FA volume of the second session by the same procedure and results averaged to reduce any bias arising from either session. Regional estimates of between-session alignment were computed from the normalised correlation scores within a set of ROI masks constructed from the union of FreeSurfer segmentations from both sessions. Global estimates of between-session alignment were computed within the grey matter, white matter and whole-brain masks.

4.5.3 Tract density agreement

In order to provide a measure of regional tract agreement over sessions prior to network construction, we computed the ACM volumes recording the streamline density per voxel (Section 3.6.6). Using the transforms described above, the ACM volume from the second session was aligned to the ACM volume of the first session (and vice versa, with results averaged) and compared on a voxel-wise basis using normalised correlation (Equation 4.6).

4.6 Results

Before performing connectivity mapping, we assessed the test-retest properties for segmentation, registration and tractography.

4.6.1 Evaluation of neuroanatomical segmentation

FreeSurfer tissue segmentation, subcortical segmentation and cortical parcellations were compared in FreeSurfer's native MNI-305 space. Visual inspection of the segmentations for each subject indicated that the FreeSurfer morphometric procedure provided plausible brain extraction, tissue segmentation and cortical labelling. Figure 4.1 shows an example the two cortical segmentation atlases for one subject. Note that the grey and white tissue masks were identical for both cortical atlases.

Figure 4.2(a) shows the mean volumes for 16 subcortical regions and the 68 cortical regions of the Desikan-Killiany atlas with a mean volume of 5.71 cm^3. We observed that the size of grey matter regions was quite variable from region to region.

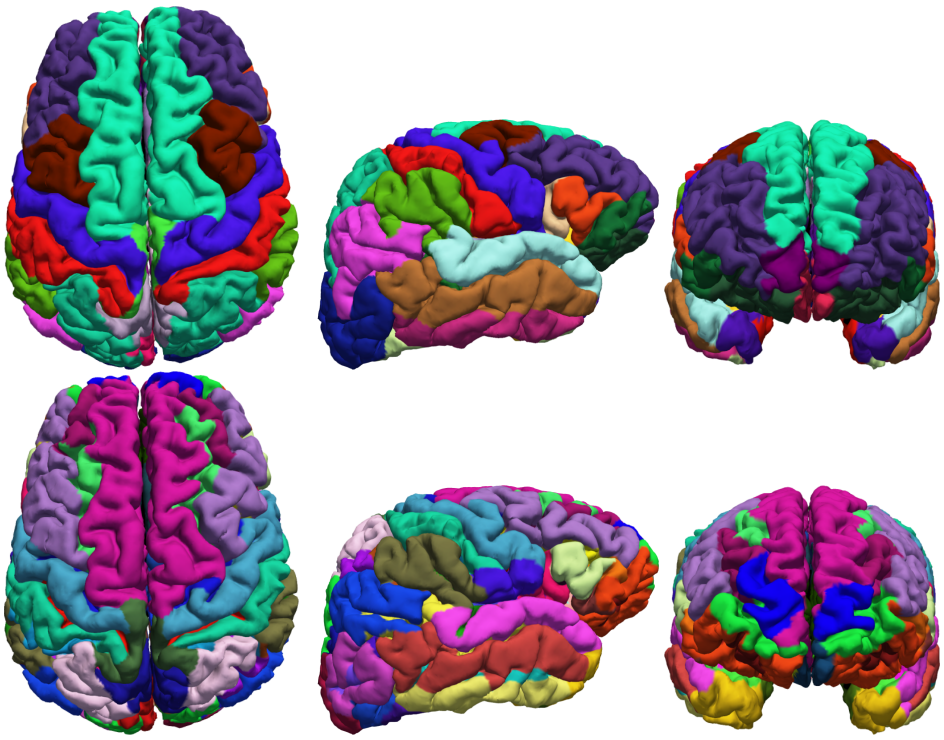

Figure 4.1: Cortical parcellations visualised on the pial surface (56 year old male), for the Desikan-Killiany atlas (top) and the Destrieux atlas (bottom). Colour codings are defined in the FreeSurfer package.

The mean Dice coefficient measuring segmentation overlap between sessions was 0.77 ± 0.02 for grey matter and 0.92 ± 0.01 for white matter. Figure 4.2(b) shows the corresponding mean Dice coefficient for each of these regions. Overall, the mean Dice coefficient was 0.83 ± 0.02 for the 16 subcortical structures assessed (excluding the brain stem). The mean Dice coefficient was 0.69 ± 0.03 for the 68 cortical regions of the Desikan-Killiany atlas. However, the mean Dice coefficient was 0.63 ± 0.04 for the 148 cortical regions of the Destrieux atlas (in the interest of space, the individual scores for these ROIs are not shown). All cortical and subcortical regions assessed showed at least some spatial overlap between sessions; no region was entirely misplaced. Inspection revealed that there was little variance in the Dice overlap between hemispheres or across the cortical lobes, although temporal, cingulate and frontal regions had slightly higher overlap agreement than occipital and parietal regions. The segmentation of sub-cortical and medial structures was found to be most reliable, in terms of test-retest agreement. In comparison, the segmentation overlap of the cortex was far more variable. The two regions showing the poorest agreement were the frontal and temporal poles. In conclusion, as the Desikan-Killiany provided better segmentation agreement than the Destrieux atlas for our cohort, it was used for all subsequent experiments.

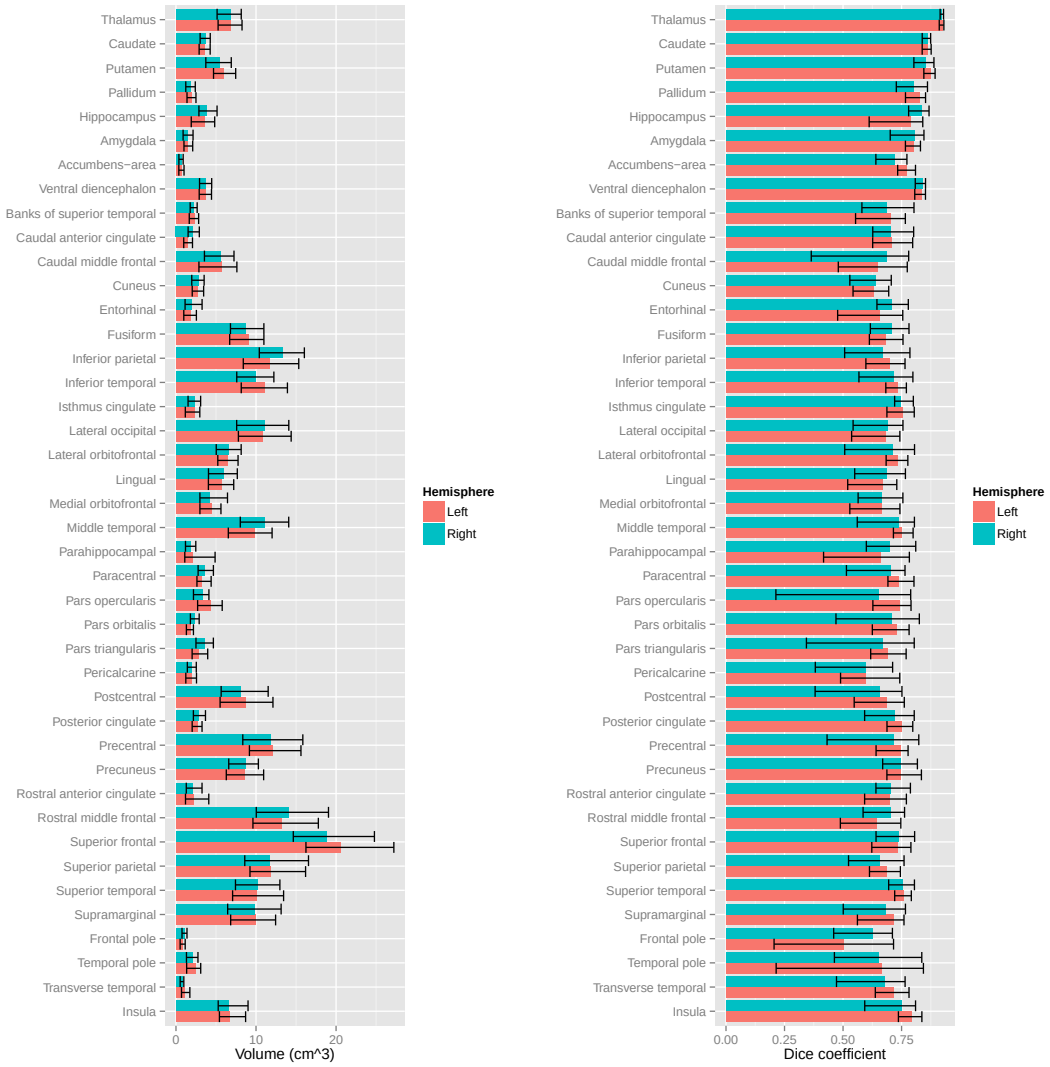

(a) (b)

Figure 4.2: a) Mean volumes of the grey matter structures obtained from FreeSurfer; b) Mean segmentation overlap between sessions, assessed by Dice coefficient. In both cases the error bars show the 95% inter-percentile range ($N = 10$).

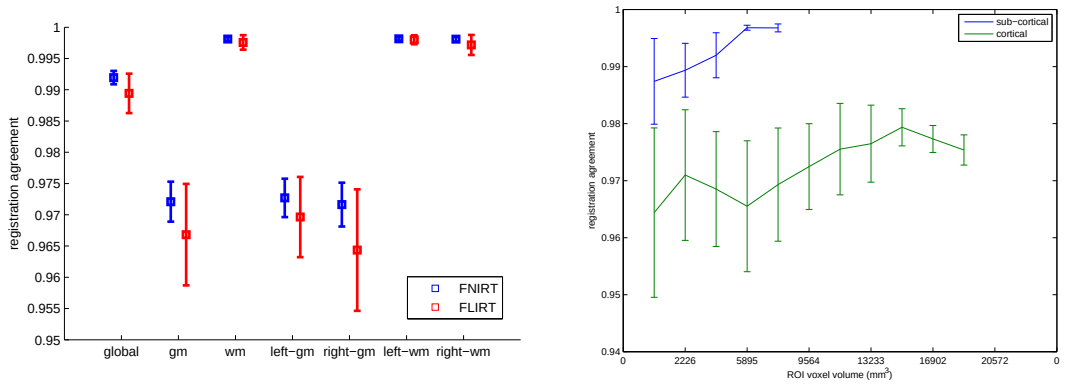

(a) (b)

Figure 4.3: a) Population registration agreement for linear (FLIRT) and nonlinear (FNIRT) registration assessed by intensity-based normalised correlation; b) Population registration agreement using FNIRT for subcortical and cortical nodes computed over ten ROI groupings ordered by volume.

4.6.2 Evaluation of registration

We compared the registration test-retest performance between sessions for both linear and nonlinear registration in grey matter, white matter and grey matter ROIs. Notably, nonlinear registration (FNIRT) was found to improve alignment over linear registration (FLIRT 12-DOF affine transform) in all brain regions as shown in Figure 4.3(a). In terms of the mean normalised correlation, a reduction in registration error of 0.0025 was obtained globally, with a 0.0053 reduction in grey matter and a 0.0005 in white matter. In particular, improvements were noted in the alignment of cortical regions. In the best cases, the mean normalised correlation error was reduced by 0.02 or greater in bilateral rostral middle frontal regions, right caudal middle frontal gyrus and right superior frontal gyrus. As a result, nonlinear registration was used for all subsequent experiments. We determined that using the *FA* parametric map as a registration target provided slightly better alignment than either a single T_2-weighted extracted brain volume or an average of the seven T_2-weighted brain volumes.

Using nonlinear registration, spatial alignment of T_1-weighted volumes was assessed between sessions by normalised correlation with an overall score of 0.992 ± 0.002 obtained within cerebral areas, 0.972 ± 0.006 in grey matter and 0.998 ± 0.001 in white matter.

Figure 4.3(b) shows population registration agreement using FNIRT for subcortical and cortical nodes computed over ten ROI groupings ordered by volume. This shows both that registration agreement for subcortical (medial) structures was typically better than cortical regions and that small cortical regions showed poorer alignment. Figure 4.4(a) shows the regional estimates

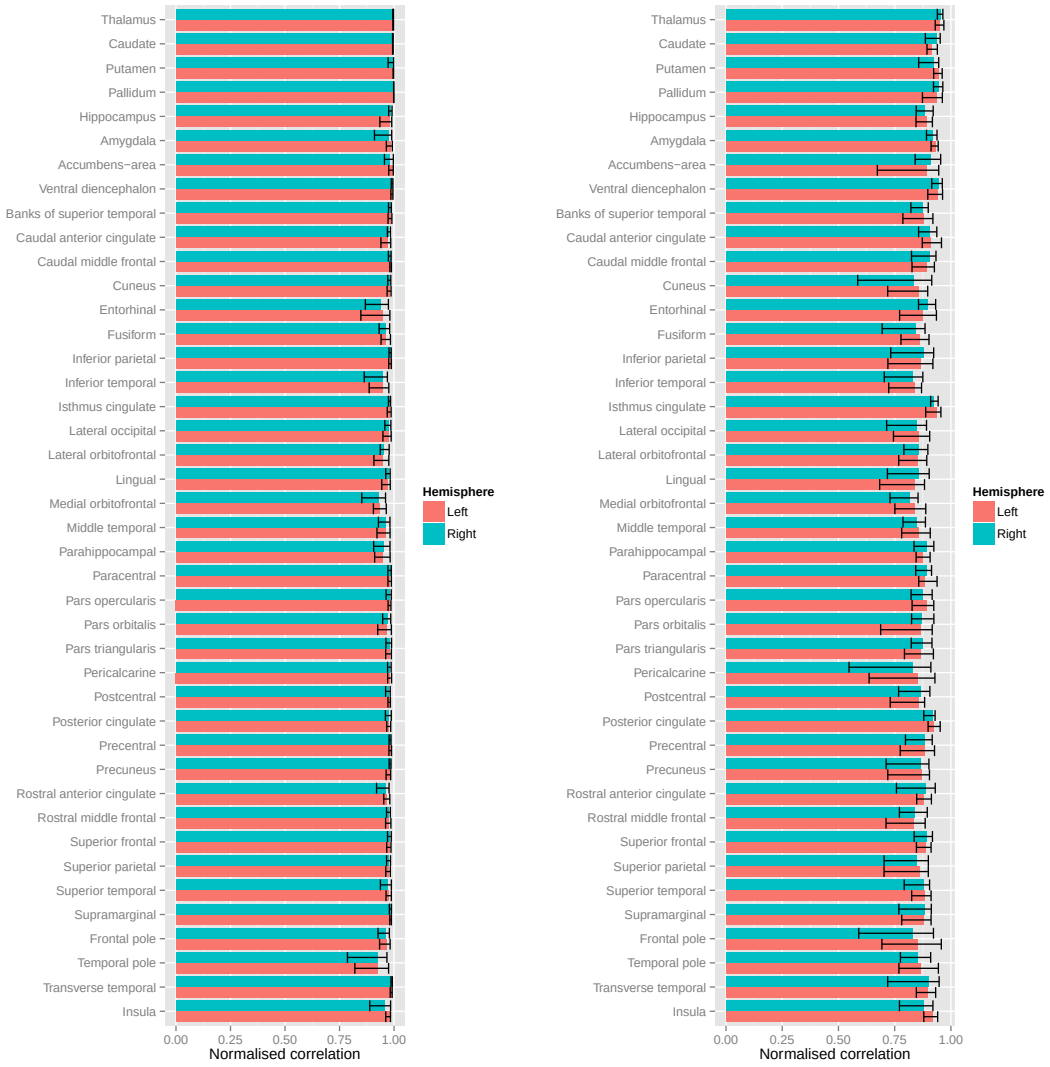

Figure 4.4: Nodal measures of test-retest agreement with the 95% inter-percentile range ($N = 10$) for: a) Registration agreement of nonlinear registration assessed by normalised correlation; b) Streamline-density agreement assessed by normalised correlation.

of alignment for each of the 84 nodes assessed by normalised correlation. In general, the registration error was very low. The poorest overall alignment was observed in the temporal poles, medial orbitofrontal and entorhinal regions, and the best alignment was observed in sub-cortical structures such as the thalamus, pallidum and putamen.

4.6.3 Evaluation of streamline-density agreement

The streamline-density agreement was assessed between sessions in diffusion space by normalised correlation with an overall score of 0.912 ± 0.02 obtained within cerebral areas, 0.860 ± 0.03 in grey matter and 0.919 ± 0.02 in white matter. Figure 4.4(b) shows the regional estimates of agreement, indicating that the regions with the poorest test-retest agreement were the frontal poles, pericalcarine cortex and the cuneus.

4.6.4 Regional variation in reliability

Figure 4.5 shows the regional test-retest agreement for segmentation, registration and streamline density. These plots show the mean values from Figures 4.2(b), 4.4(a), 4.4(b) in order to present a visualisation of possible regional variability in these measures. Broadly, all three measures of reliability showed the strongest agreement in subcortical and medial structures, and a gradual drop off in reliability was observed for more distal cortical regions. Large subcortical structures such as the thalamus and putamen were consistently reliable across measures. The poorest segmentation overlap was observed in the left frontal pole. On the whole, registration agreement was high across all nodes as shown in Figure 4.5(b), with the poorest agreement in the temporal poles which have a small cortical volume. Streamline density was observed to follow a similar pattern of regional variability to the other two measures although the poorest agreement was observed in the parietal regions as shown in Figure 4.5(c).

4.6.5 Comparison of tractography configurations

We evaluated four alternative streamline termination thresholds on curvature, independently of network construction. Thresholds were assessed at greater than 45, 60, 70 or 80 degrees. These comparisons were performed with probabilistic tracking and white matter seeding. Visual inspection of the resulting streamlines (Figure 4.6) revealed that a 45 degree constraint was too severe for computing whole-brain connections and many genuine ‘U-fibres’ were eliminated. In terms of network connectivity, the percentage of total streamlines seeded that were subsequently identified as interconnections was 9.83% for 45 degrees, 19.83% for 60 degrees, 26.18% for 70 degrees and 31.97% for 80 degrees. An empirical evaluation determined that

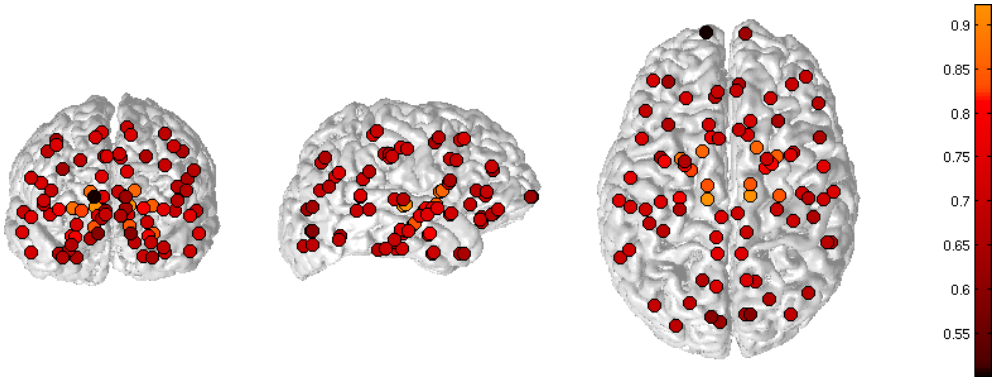

(a)

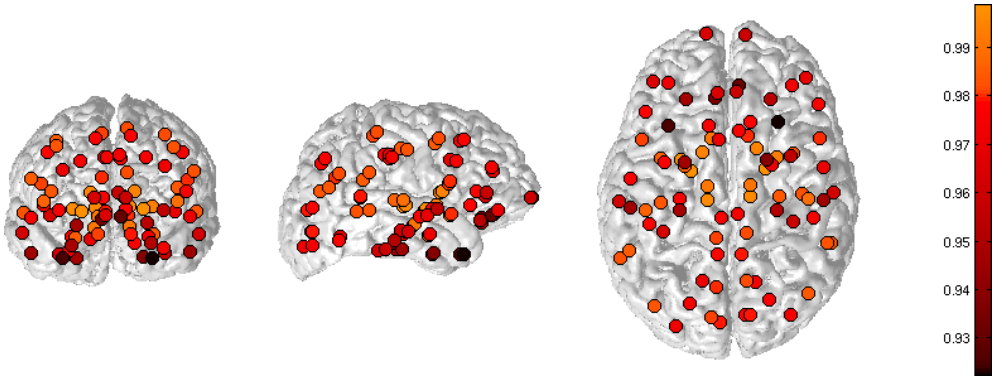

(b)

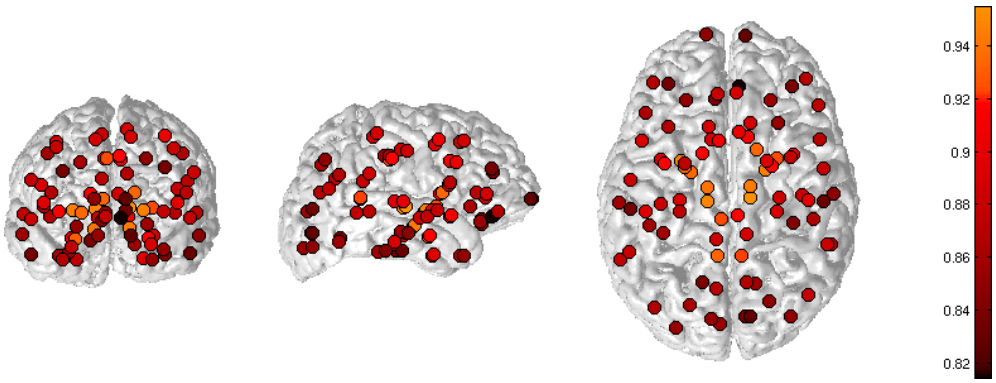

(c)

Figure 4.5: Coronal, sagittal and axial views of the regional test-retest agreement for segmentation, registration and streamline density. For each view, nodes are plotted at the region centroid where node colour indicates the mean nodal value ($N = 10$) for: a) Segmentation overlap between sessions, assessed by Dice coefficient. b) Registration agreement of nonlinear registration assessed by normalised correlation; c) Streamline-density agreement assessed by normalised correlation.

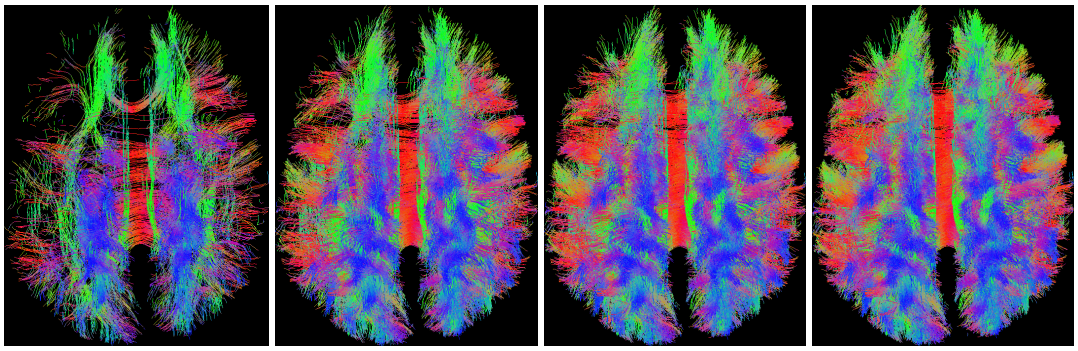

(a) 45° (b) 60° (c) 70° (d) 80°

Figure 4.6: Streamlines remaining following four alternative streamline termination thresholds on curvature for one subject (56 year old male).

using either 70 or 80 degrees permitted reconstruction of plausible U-fibres whilst removing some spurious streamlines. Likewise, an empirical investigation of the streamline termination constraint on voxel-wise FA determined that a threshold of 0.1 was necessary to ensure tracking into the target grey matter regions.

Following network construction and using a curvature constraint of 80 degrees and no waypoint constraint, approximately $39.7 \pm 3.0\%$ (mean $\pm$ standard deviation) of the total streamlines seeded were identified as interconnections between ROIs for FACT-GM. Likewise, the interconnections identified using FACT-WM, FDT-GM and FDT-WM were $32.5 \pm 4.8\%$, $31.8 \pm 3.7\%$ and $26.2 \pm 3.7\%$, respectively.

4.6.6 Networks

Following network construction, Figure 4.7 illustrates a subset (approximately 20%) of the streamlines found to interconnect between the 84 nodes, i.e. streamlines that begin and terminate in grey matter, for one subject. Figure 4.8 shows two representation of the resulting network for the same subject, in this case using probabilistic tractography and WM-seeding.

Figure 4.9 shows the mean connectivity matrices and corresponding histograms of weights generated for the three network weightings, using the networks computed by FDT with white matter seeding for illustration. In each case the networks were produced from the same set of streamlines. Both streamline density weightings (SD and SDL) follow a similar distribution, which approximates a power law; tracking results in many low weighted connections but very few strong connections. However, due to the length correction, the SDL-weighting penalised long-range links as was evident by the down-weighting of inter-hemispheric connections as shown in Figure 4.9(b). The FA-weighting produced a markedly different distribution of edge

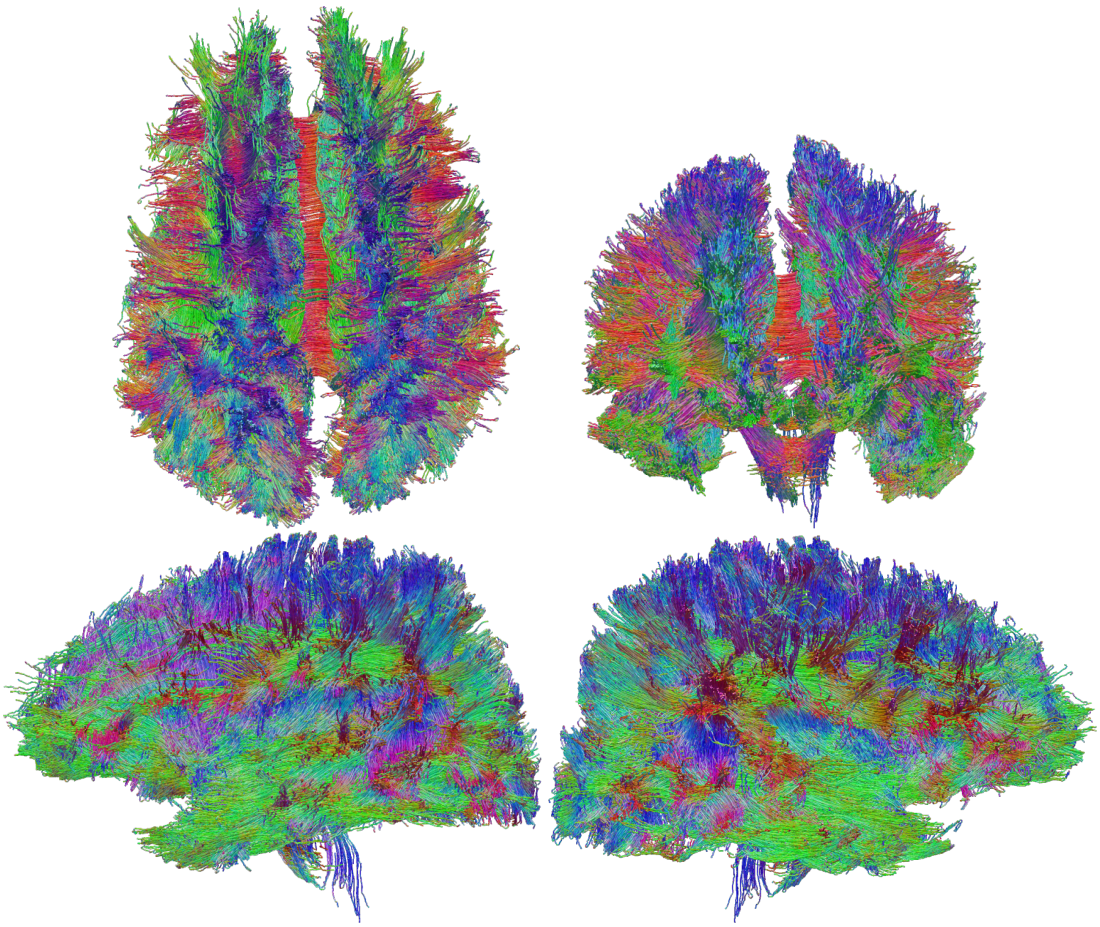

Figure 4.7: A subset of the inter-connecting streamlines for one subject (56 year old male) constructed by probabilistic tractography using white matter seeding. Streamlines have been filtered by curvature and length for visualisation, where colour indicates the x, y, z (red, green, blue) direction of each streamline segment.

weights, reflecting the mean diffusion anisotropy of interconnecting streamlines. Note that weights below 0.1 are absent due to the FA constraint applied in tracking.

Figure 4.10 shows the regional network reliability measured over different combinations of tractography algorithm, seeding approach and network weighting. In each case, the mean ICC was computed from the 84 nodes and plotted over 13 thresholds of streamline filtering by minimum length in white matter, where zero corresponds to no length constraint. Mean ICCs of node strength were used to rate the overall nodal test-retest reliability because node strength measures the sum of weights per node. An ICC score of 1 indicates perfect agreement between sessions, while a score of less than 0.5 may indicate poor agreement. In all cases, when comparing the same algorithm and weighting, the white matter seeding strategy outperformed the corresponding grey matter strategy. For instance, FDT-GM-SD obtained a mean ICC of 0.51, whereas FDT-WM-SD obtained 0.62. For most cases the probabilistic methods outperformed

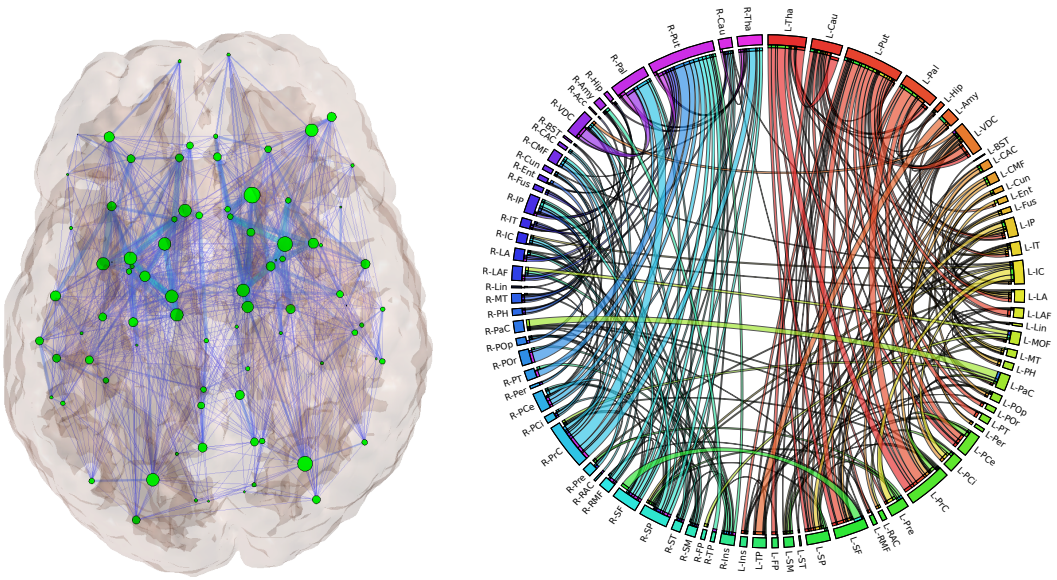

Figure 4.8: Graphical representation of the structural network for one subject (56 year old male): Left) Anatomical graph, where node size indicates node strength and edge width indicates connection weight; Right) Circle graph representation of the same network for the 500 strongest connections, where edge width indicates connection weight and node abbreviations are listed in Appendix A.

the deterministic methods, particularly for white matter seeding. However, for FA-weighting, the advantage of probabilistic over deterministic was less clear as shown in Figure 4.10(c). The best performance for both tractography algorithms was obtained when using white matter seeding and SD-weighting (mean ICC of 0.62 for FDT and 0.58 for FACT). Conversely, the poorest performance was obtained using grey matter seeding and SDL-weighting as shown in Figure 4.10(b). For many cases, the minimum streamline length threshold offered little benefit. Only in the case of SDL-weighting, as shown in Figure 4.10(a), was the mean ICC found to notably increase with length in white matter, and for this weighting the best performance was found by ensuring a white matter length $>= 3.5$ mm. This could indicate that the SDL-weighting up-weighted short unreliable streamlines while penalising longer streamlines. Additionally, in all cases using GM-seeding, except FACT-GM-SDL, a waypoint threshold > 0 mm improved performance over no waypoint threshold. Overall, the best tractography configuration was found to be the FDT algorithm with white matter seeding and SD-weighting, resulting in a mean ICC of 0.62, as shown in Figure 4.10(a). This configuration was then chosen to assess network properties in the remainder of this chapter, except where specified otherwise.

Figure 4.11 shows the mean node strength reliability as assessed by ICC for the four definitions of endpoint connectivity as described in Section 3.7.1. In each case probabilistic tractography and WM-seeding were used. Overall, the four definitions of connectivity provided similar

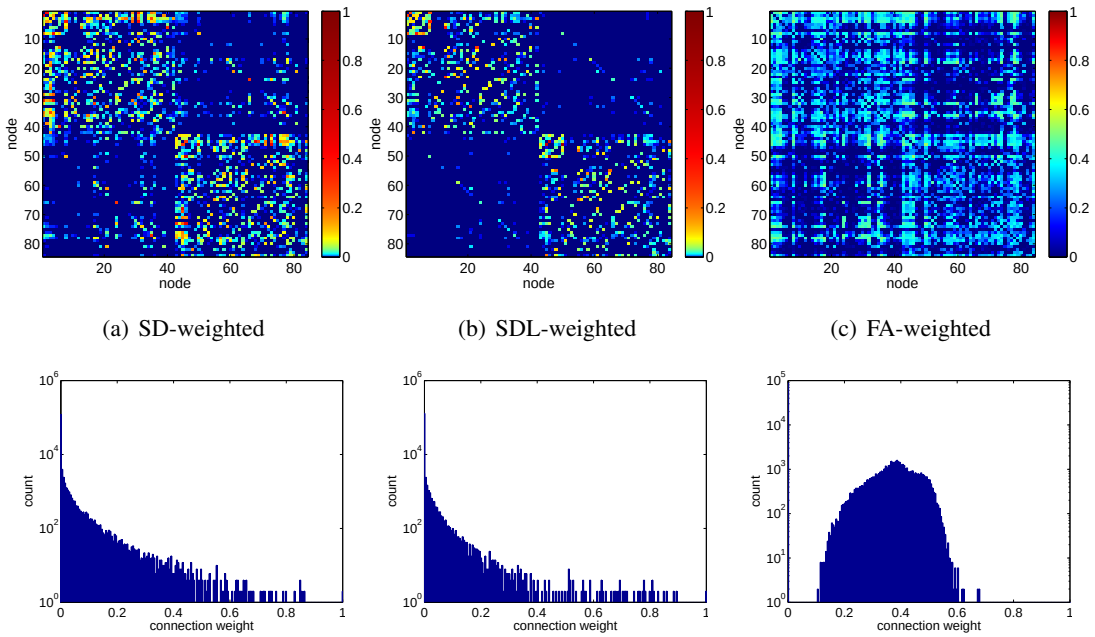

(a) SD-weighted (b) SDL-weighted (c) FA-weighted

(d) SD-weighted (e) SDL-weighted (f) FA-weighted

Figure 4.9: Top row: 84×84 mean connectivity matrices of inter-region connections averaged across all subjects (N=20) for the three network weightings and generated from the same set of streamlines: (a) streamline density; (b) streamline density with streamline length correction, (c) tract-averaged FA. Note, that both SD and SDL-weighted matrices are scaled between zero and one and (a) and (b) are log scaled. In each case, the two large rectangular patterns on the diagonal correspond to the left and right hemispheres and the node ordering is the same as Figures 4.13–4.15. Bottom row: the corresponding histograms of the connection weights for each of the three network weightings.

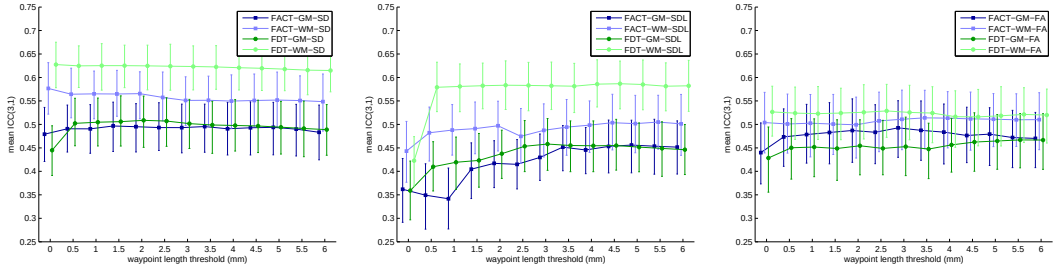

(a) SD-weighted (b) SDL-weighted (c) FA-weighted

Figure 4.10: Mean ICC values measuring test-retest reliability of node strength over thirteen thresholds of streamline filtering by white matter waypoint length for each tractography algorithm, seeding type and network weighting. 95% confidence intervals of the mean were estimated by resampling with replacement 5000 times from the 84 nodes and recomputing the mean and taking the 2.5 and 97.5 percentiles of this distribution.

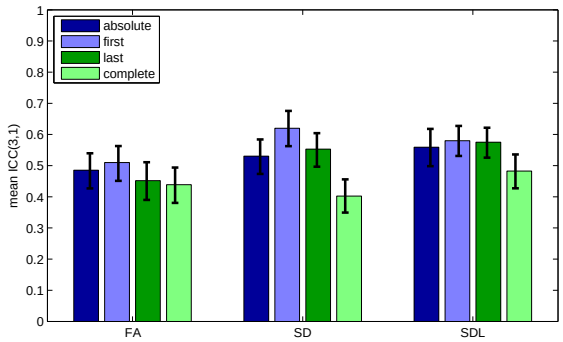

Figure 4.11: Comparison of four different connectivity configurations measuring the mean ICC for node strength using probabilistic tractography and WM-seeding. 95% confidence intervals of the mean were estimated by resampling with replacement 5000 times from the 84 nodes and recomputing the mean and taking the 2.5 and 97.5 percentiles of this distribution.

test-retest performance. However, the connectivity using 'first' ROI encountered was found to perform marginally better in each case, particularly for SD-weighting. In contrast, the connectivity measuring all potential ROIs as connected (complete) provided the poorest test-retest results in each case. As a result, 'first' connectivity was used for all subsequent experiments.

Figure 4.12 shows the relationship between network properties and tracking iterations per seed point (or number of streamlines generated) for this configuration. Figure 4.12(a) shows the population means of the four network measures. The mean network degree increased with the number of iterations, whereas the mean clustering coefficient and path length showed an inverse relationship with the number of iterations. The mean value of network strength was found to be variable when measured over few iterations but stabilised after approximately 40 iterations. Figure 4.12(b) shows the corresponding mean nodal ICCs for each network measure. The mean ICC of network strength settled to around 0.62 after approximately 10 iterations, although the other three properties showed gradual increases in reliability up to 100 iterations.

4.6.6.1 Network interpretation

Figure 4.13 shows the regionally computed mean values for the four network measures using FDT with WM-seeding, SD-weighting and a waypoint threshold set to 0.5mm. The absolute values of node strength, path length and clustering coefficient are not especially meaningful as they are dependent on the weighting scheme used, tractography configuration and choice of nodes. However, the value of one node property relative to another may offer insight into network organisation. The 95% inter-percentile range suggests that the patterns of connectivity were fairly similar across subjects. Sub-cortical structures, such as the thalamus, putamen,

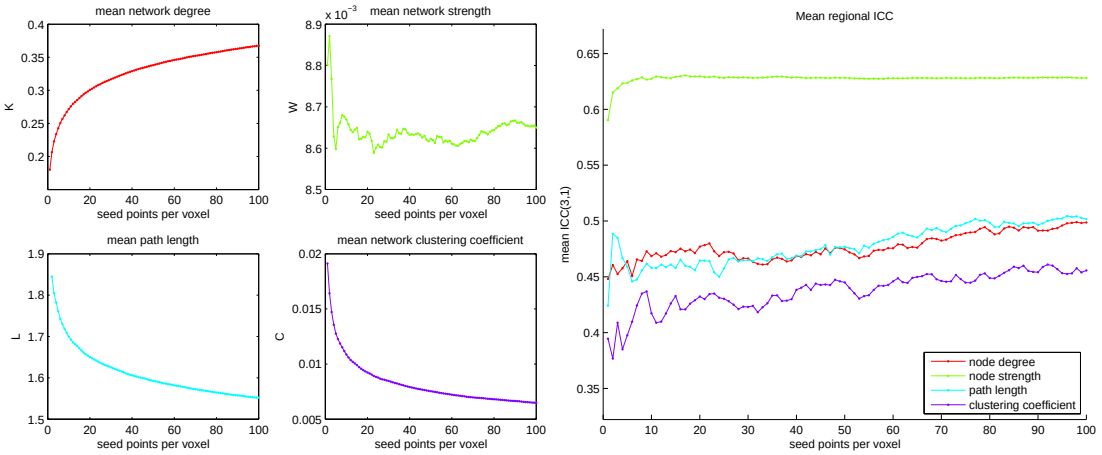

(a) (b)

Figure 4.12: The relationship between network properties and tracking iterations per seed point using FDT with white matter seeding, streamline density weighting and waypoint threshold set to 0.5 mm: (a) The population means of the four network measures; (b) The corresponding mean regional ICCs for each network measure.

and pallidum tend to be both highly interconnected and have strong total connection weightings. The entorhinal cortex, pars orbitalis and precentral gyrus show a high clustering coefficient, suggesting their direct nodal neighbours are also strongly inter-connected. However, the clustering coefficients were very variable across the cohort. The overall network degree was 31.5 ± 15.3, indicating that on average, each node is directly linked to nearly 32/84 of the other nodes. Overall network strength was 0.72 ± 0.7 reflecting the average sum of connection weights (normalised streamline density) per node. Network characteristic path length was 1.54 ± 0.2 reflecting the average weighted path length between all pairs of nodes. Finally, network clustering coefficient was 0.006 ± 0.003, indicating that on average a node's direct neighbours have a mean connection weighting of approximately 0.6% of the maximum weight.

4.6.7 Global network reliability

Table 4.1 shows a summary of the global test-retest analysis for FDT and white matter seeding (other tractography configurations are not reported) for each of the three network weightings. In each case, networks were constructed with the waypoint threshold set to 0.5 mm as this was closest to the best performance for all three weightings (Figure 4.10). All global network properties show good within-subject agreement with all ICCs > 0.59. In all cases, the within-subject differences (δ^{WS}) were smaller than the between-subject differences (δ^{BS}). For SD-weighting, ICCs were between 0.62 and 0.76, within-subject differences 3.2–11.9%,

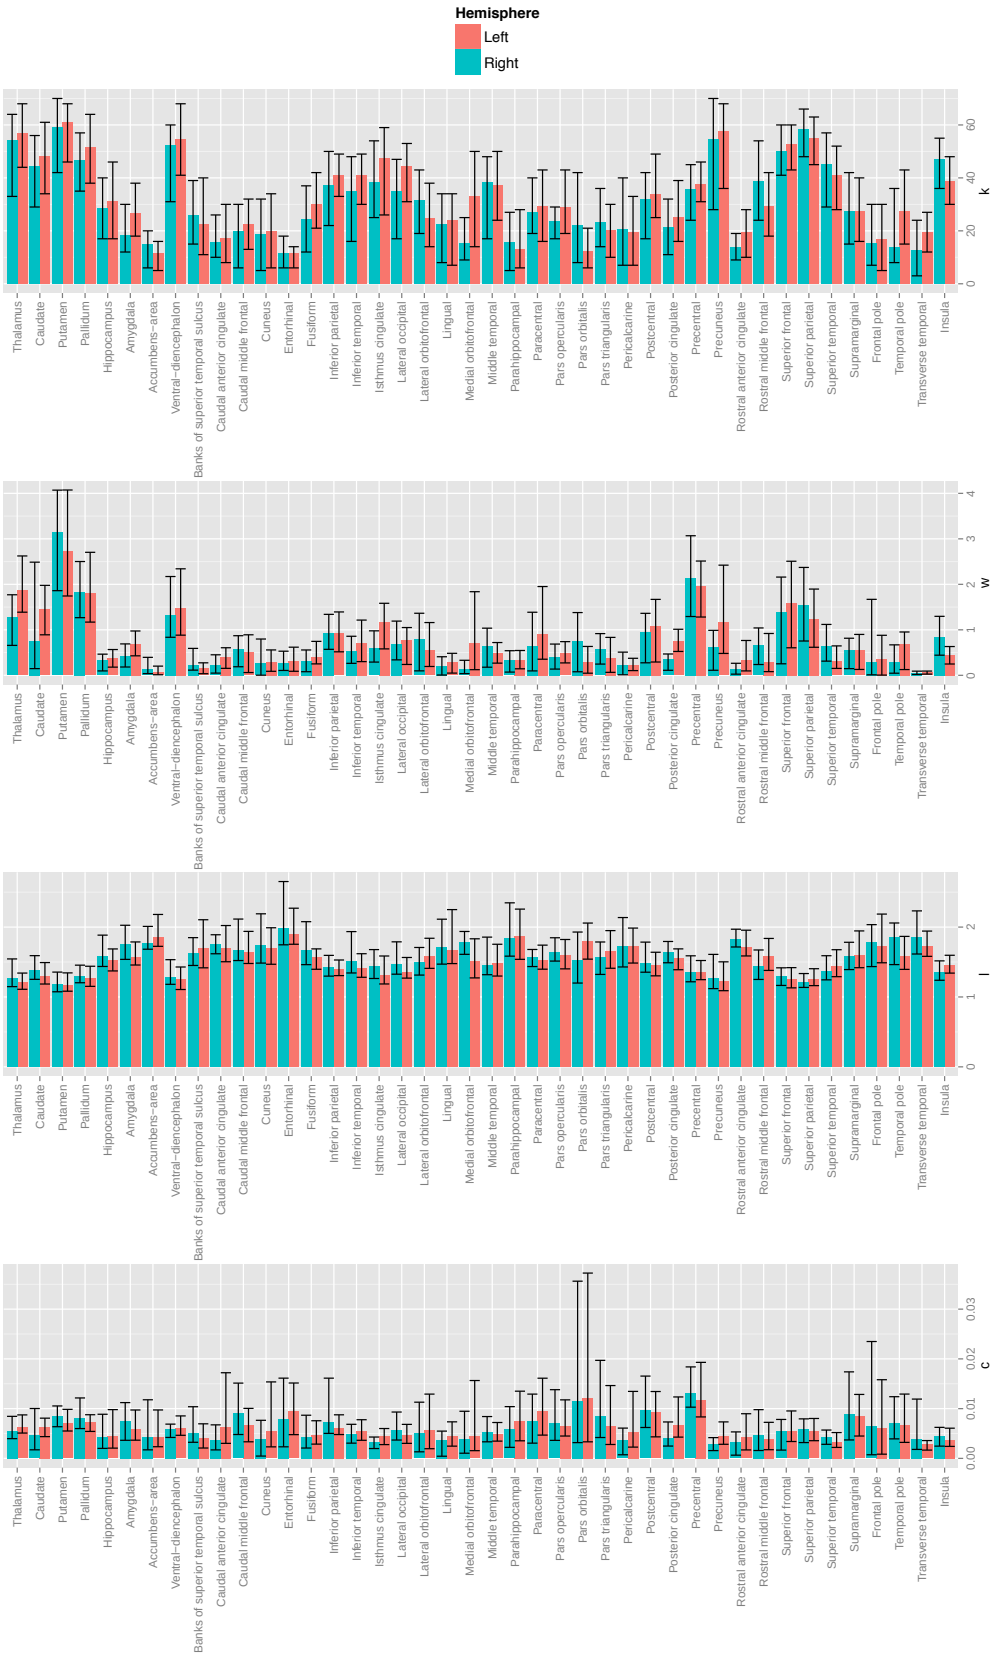

Figure 4.13: Nodal values computed from population (N=20) means, from top-to-bottom: node degree (k), node strength (w) , path length (l) and clustering coefficient (c) as defined in Section 3.8. Error bars show the 95% inter-percentile range.

FDT, WM-seeding, SD-weighted

	ICC	mean δ^{WS} %		mean δ^{BS} %		p-value
network degree	0.66	8.22	(8.3)	14.10	(11.0)	0.110
network strength	0.75	11.93	(9.9)	19.95	(14.4)	0.066
characteristic path length	0.62	3.22	(3.6)	5.83	(4.8)	0.099
network clustering coefficient	0.76	6.37	(6.2)	12.56	(9.3)	**0.018**

FDT, WM-seeding, SDL-weighted

	ICC	mean δ^{WS} %		mean δ^{BS} %		p-value
network degree	0.66	8.22	(8.3)	14.10	(11.0)	0.110
network strength	0.76	8.86	(6.1)	15.09	(10.5)	**0.025**
characteristic path length	0.59	3.06	(3.3)	5.51	(4.5)	0.097
network clustering coefficient	0.71	8.35	(7.2)	16.18	(9.6)	**0.006**

FDT, WM-seeding, FA-weighted

	ICC	mean δ^{WS} %		mean δ^{BS} %		p-value
network degree	0.66	8.22	(8.3)	14.10	(11.0)	0.110
network strength	0.67	9.82	(10.5)	17.54	(13.9)	0.070
characteristic path length	0.64	4.46	(4.8)	8.25	(7.1)	0.060
network clustering coefficient	0.69	4.77	(5.3)	8.87	(8.5)	0.073

Table 4.1: Summary of the global test-retest analysis for the three types of network weighting using FDT and white matter seeding, showing: ICC, mean within-subject differences (δ^{WS}), mean between-subject differences (δ^{BS}). The differences are expressed as a percentage, bracketing indicates the standard deviation and emboldening indicates significance ($p < 0.05$).

FDT, WM-seeding, SD-weighted

	mean ICC		mean δ^{WS} %		mean δ^{BS} %	
node degree	0.50	(0.3)	16.35	(7.0)	23.25	(8.1)
node strength	0.62	(0.2)	24.18	(8.3)	35.69	(11.0)
path length	0.53	(0.2)	5.22	(1.5)	8.08	(2.0)
clustering coefficient	0.46	(0.3)	21.82	(7.4)	28.77	(8.8)

FDT, WM-seeding, SDL-weighted

	mean ICC		mean δ^{WS} %		mean δ^{BS} %	
node degree	0.50	(0.3)	16.35	(7.0)	23.25	(8.1)
node strength	0.58	(0.2)	21.85	(7.7)	31.11	(10.5)
path length	0.51	(0.3)	4.91	(1.5)	7.57	(2.0)
clustering coefficient	0.45	(0.3)	22.92	(8.8)	31.36	(9.8)

FDT, WM-seeding, FA-weighted

	mean ICC		mean δ^{WS} %		mean δ^{BS} %	
node degree	0.50	(0.3)	16.35	(7.0)	23.25	(8.1)
node strength	0.52	(0.3)	19.02	(8.0)	27.28	(9.2)
path length	0.56	(0.2)	6.35	(1.9)	10.02	(2.4)
clustering coefficient	0.51	(0.2)	8.21	(2.2)	12.38	(2.8)

Table 4.2: Summary of the regional (nodal) test-retest analysis for all 84 nodes for the three types of network weighting using FDT and white matter seeding, showing: mean ICC, mean within-subject differences (δ^{WS}) and mean between-subject differences (δ^{BS}). The differences are expressed as a percentage and bracketing indicates the standard deviation.

between-subject differences 5.8–20% and the corresponding p-values between 0.018 and 0.11. For SDL-weighting, ICCs were between 0.59 and 0.76, within-subject differences 3.1–8.9%, between-subject differences 5.5–16.2% and the corresponding p-values between 0.006 and 0.11. For FA-weighting, ICCs were between 0.64 and 0.69, within-subject differences 4.5–9.8%, between-subject differences 8.3–17.5% and the corresponding p-values between 0.06 and 0.11. Overall, the global test-retest properties were broadly similar over the three weightings, although SDL-weighting obtained marginally better performance than the other weightings. For network strength and clustering coefficient the SDL-weighting provided the most significant differences, whereas for characteristic path length, the FA-weighting came closest to significance. In each case the results for network degree were identical because the same set of streamlines was used.

4.6.8 Regional network reliability

Table 4.2 shows a summary of the regional test-retest analysis. For SD-weighting, mean ICCs were between 0.46 and 0.62, mean within-subject differences 5.2–24.2%, mean between-subject differences 8.1–35.7%. For SDL-weighting, mean ICCs were between 0.45 and 0.58, mean within-subject differences 4.9–22.9% and mean between-subject differences 7.6–31.4%. For FA-weighting, mean ICCs were between 0.50 and 0.56, mean within-subject differences 6.4–19.0% and mean between-subject differences 10.0–27.3%. For all four measures, SD-weighting obtained marginally better regional test-retest performance than SDL-weighting in terms of mean ICCs. However, in the case of path length and clustering coefficient the FA-weighting obtained better performance than SD-weighting. Again, the results for node degree were identical across weightings as the same set of streamlines was used.

Figure 4.14 shows the regional ICCs for each of the 84 nodes, again using FDT with white matter seeding, SD-weighting and a waypoint threshold set to 0.5mm. Only 22.6% (19/84) of nodes obtained ICCs ≥ 0.5 across all four measures. 77.4% (65/84) of nodes showed poor within-subject agreement with an ICC < 0.5 across one or more measures. Furthermore, 10.7% (9/84) of nodes showed a negative ICC across one or more measures. For the unbiased formulation of the ICC used here, estimates can be negative for greater within-subject variance than between-subject variance. These nine nodes were left/right amygdala, left/right pallidum, left caudal middle frontal, left superior frontal, right inferior temporal, right middle frontal, right lateral orbitofrontal and right superior parietal.

Figure 4.15 shows the corresponding differences between the within- and between-subject components, $\delta^{BS} - \delta^{WS}$. Although, for 83.3% (70/84) of nodes the within-subject differences were smaller than the between-subject differences across all four measures, on average less than $1/8$ nodes show that the differences were significant ($p < 0.05$) for our sample. Negative differences occurred for 17.9% (15/84) of nodes in at least one measure, indicating that the within-subject difference outweighed the estimated between-subject difference. These 15 nodes were left pallidum, left hippocampus, left amygdala, left caudal anterior cingulate, right entorhinal, right fusiform, left lateral occipital, right lateral orbitofrontal, right middle frontal, left/right pars orbitalis, right pericalcarine, right paracentral, right transverse temporal and right temporal pole.

Results revealed a small variation in test-retest scores between hemispheres. For example, for nodes within the left hemisphere the mean within-subject difference (for node strength) was 22.51%, the mean between-subject difference was 34.82% and the mean ICC was 0.66. For nodes with the right hemisphere the mean within-subject difference was 25.83%, the mean between-subject was 36.56% and the mean ICC was 0.59. However, testing the mean differ-

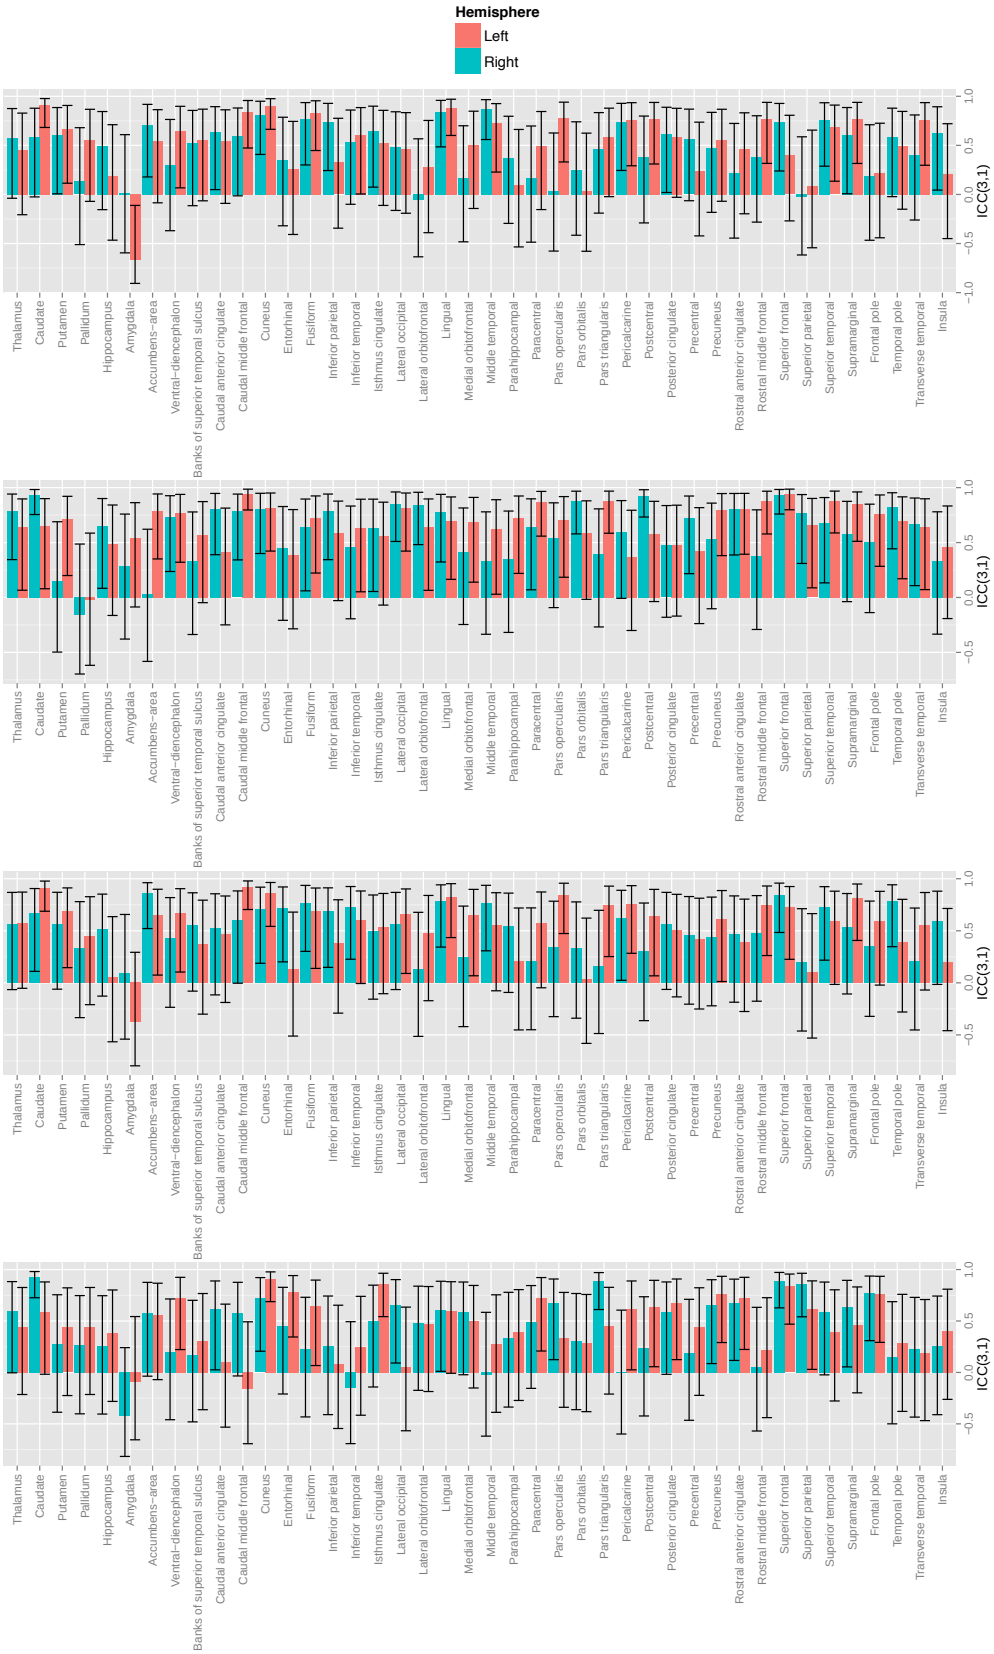

Figure 4.14: Nodal ICC values with 95% confidence intervals (N=10) from top-to-bottom: node degree, node strength, path length and clustering coefficient.

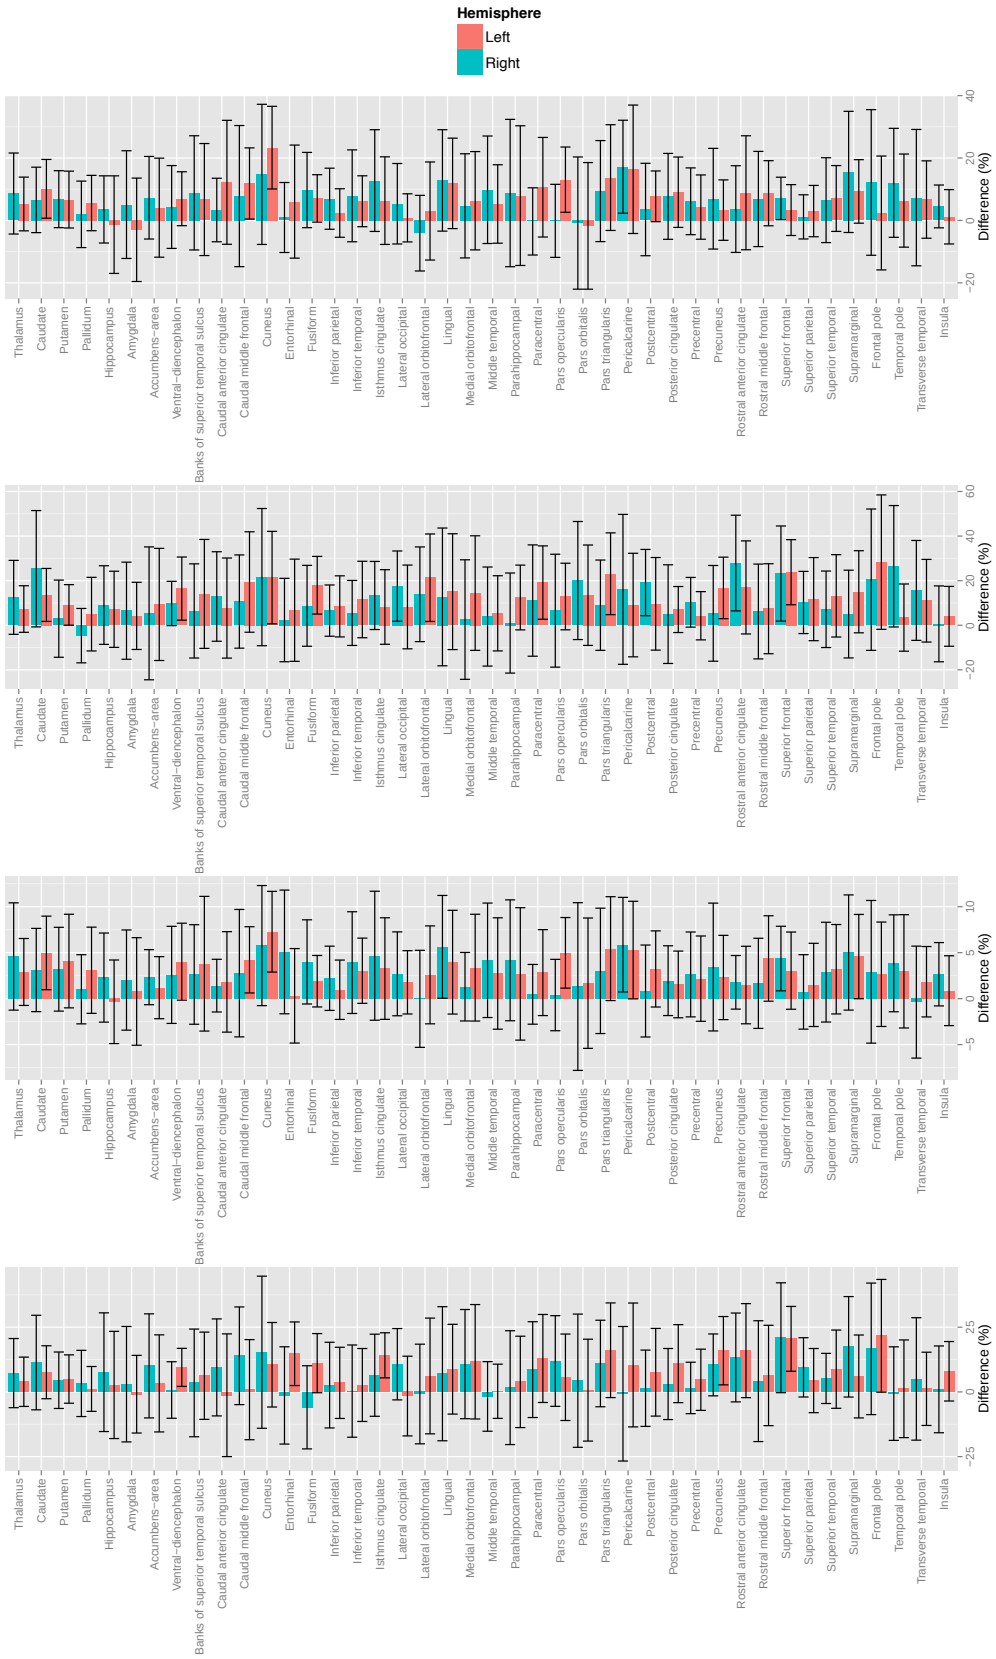

Figure 4.15: Mean nodal differences ($\delta^{BS} - \delta^{WS}$) with 95% confidence intervals of the mean estimated by bootstrap resampling (N=10), from top-to-bottom: node degree, node strength, path length and clustering coefficient.

ences δ, with a one-way ANOVA suggested that these differences were not significant for any of the four node measures.

4.6.8.1 Variation between cortical and sub-cortical nodes

Our analysis revealed a variation in regional test-retest scores between cortical and sub-cortical nodes. For cortical nodes the mean within-subject difference (of node strength) was 24.96%, the mean between-subject difference was 37.16% and the mean ICC was 0.65. However, for sub-cortical nodes the mean within-subject difference was 20.82%, the mean between-subject was 29.45% and the mean ICC was 0.50. Subsequently, we compared the 84×84 whole-brain networks with 68×68 networks constructed from cortico-cortical connections only. A percentile bootstrap contrast, which compared the mean within- versus between-subject differences (δ) between the whole-brain and cortico-cortical networks, suggested that there were no significant differences for any of the four node measures. This indicates that for our sample, the inclusion of sub-cortical connections does not reduce test-retest reliability.

4.6.9 Thresholding of network weights

Figure 4.16 shows thresholding by connection weight. Figure 4.17 shows thresholding of connections by proportion of subjects in which they were observed to occur. Figure 4.16(a) and 4.17(a) show how the mean network properties vary over a range of thresholds. For instance, node degree and node strength declines with the threshold level. Figure 4.16(b) and 4.17(b) show how corresponding mean ICCs vary over the same range of thresholds. For these experiments all three network weightings were again assessed as thresholding may effect each differently. Node degree was the same across the network weightings as the same set of streamlines was used.

Note that, for Figure 4.16(a) no value for path length was shown after more than approximately 45% of the network weights were discarded. This occurred because in one subject, this threshold level resulted in a disconnected network, i.e. not all nodes are connected and the path length becomes infinite. Arguably, both thresholding strategies showed little benefit in improving nodal reliability, particularly for node degree and node strength. However, increases in nodal ICC were observed for path length and clustering coefficient.

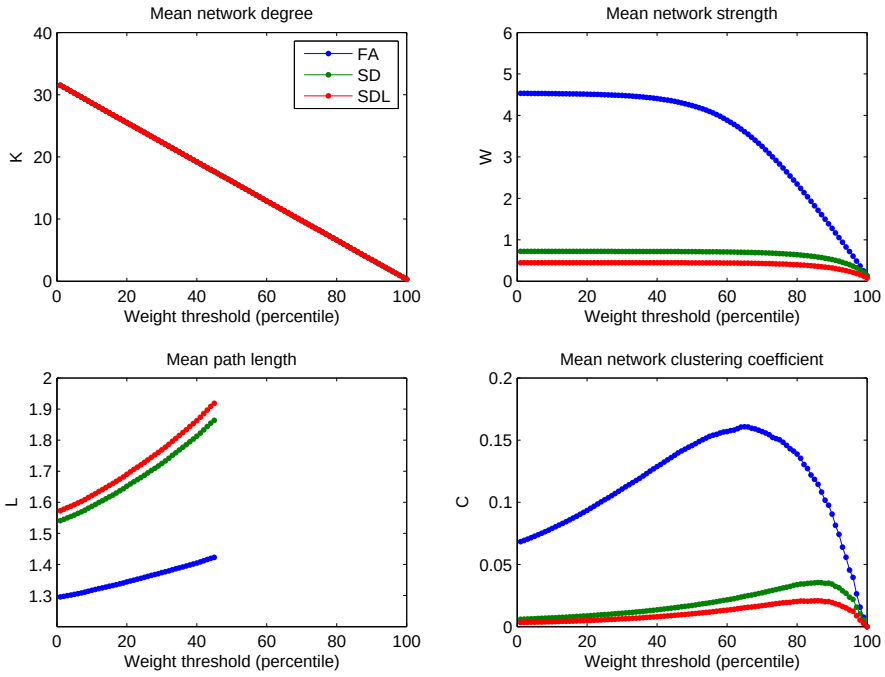

(a) Network measures

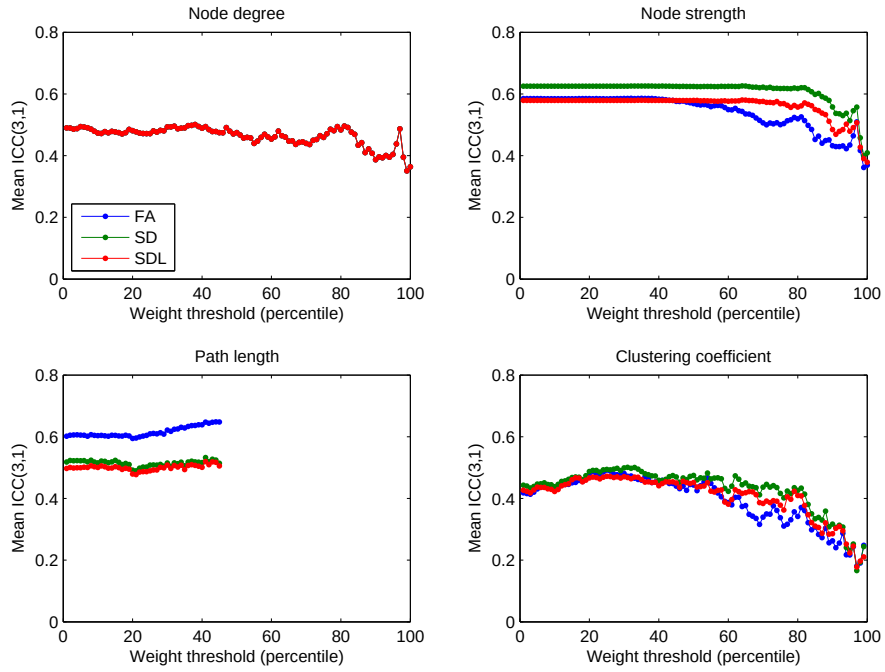

(b) Nodal ICC

Figure 4.16: Relationship between network properties (A), mean ICC (B) over a range of threshold on connection weights.

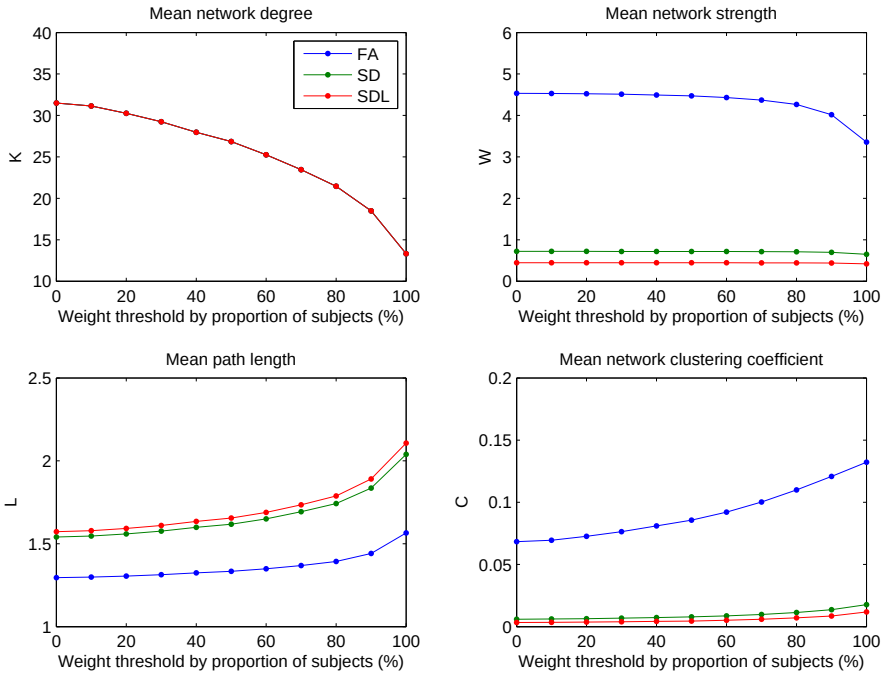

(a) Network measures

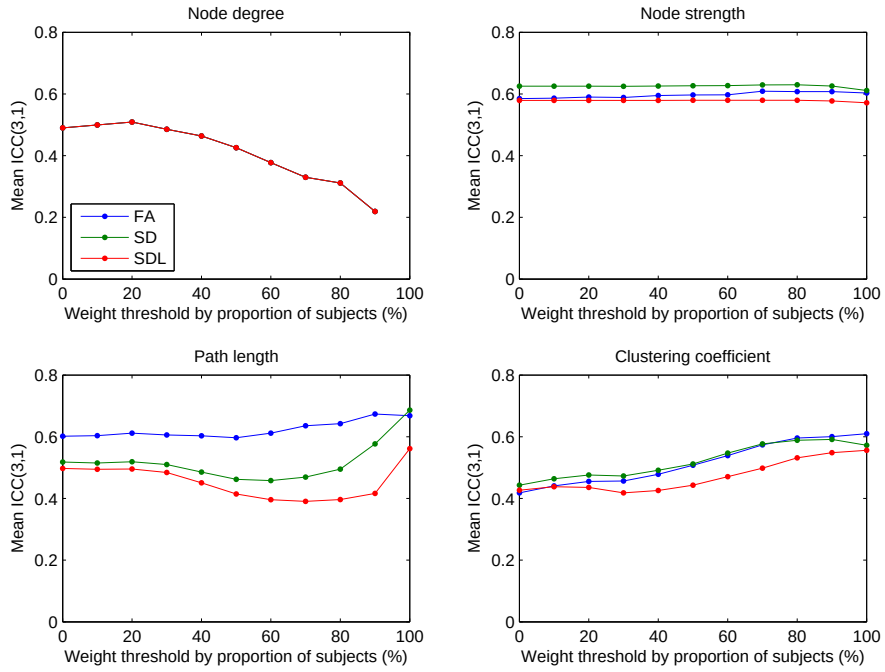

(b) Nodal ICC

Figure 4.17: Relationship between network properties (A), mean ICC (B) over a range of thresholds by the proportion of subjects for which a connection occurs.

4.7 Discussion

4.7.1 Relationship to previous work

Preceding our work, several studies have assessed aspects of network reliability using repeat scans of healthy human volunteers. Hagmann et al. (2008) assessed structural networks obtained from DSI, while Vaessen et al. (2010) assessed reproducibility over different sets of diffusion gradient directions using DTI. Bassett et al. (2010) compared reliability in both DTI and DSI, Zalesky et al. (2010b) investigated the effect of network resolution using DTI and high-angular resolution, and Cammoun et al. (2011) investigated the effect of network resolution using DSI. Finally, Cheng et al. (2012a) assessed test-retest reliability using DTI with focus on the differences between binary and weighted networks. These studies have addressed network reliability in different ways. Hagmann et al. found a within-subject network correlation of 0.78 for a single subject, and a mean between-subject correlation of 0.65 (N=5). Similarly, Cammoun et al. compared matrices directly by Pearson's correlation coefficient and found within-subject correlations ranging from 0.874 to 0.976 (N=5) and between-subject correlations of between 0.724 to 0.958 (N=20) across five different network resolutions. Vaessen et al. found within-subject coefficients of variation (CV) $< 3.8\%$ and ICCs between 0 and 0.94 for node degree, path length and clustering coefficient (N=6). Bassett et al. found within-subject CVs $< 5\%$ and ICCs > 0.72 (N=7). Cheng et al. found good test-retest agreement in both global and regional network properties (N=44). In addition, the authors concluded that binary networks and weighted networks of the same subjects have different organisational properties and that weighted networks may be more appropriate (Cheng et al., 2012a). Though it is not straightforward to compare between assays, these studies indicate that, globally, same-subject networks differ between scanning sessions and acquisition types, yet the between-subject variation is typically greater than the within-subject variation.

4.7.2 Network reliability

We obtained similar results in assessing global network properties. Using probabilistic tractography and white matter seeding, the global within-subject differences ($< 11.9\%$) were smaller than the global between-subject differences with p-values at either $p < 0.05$ or trend level and ICCs > 0.59 (Table 4.1). These findings indicate that global network properties can be estimated reliably from session to session for all three types of network weighting tested. Overall, the network weighting recording streamline density with streamline length correction (Hagmann et al., 2008) showed marginally better global test-retest performance that the other two weightings.

The comparison of tractography algorithms, seeding approaches and network weightings (Figure 4.10) offers some insights into network reliability. Firstly, in all cases white matter seeding produced networks with better test-retest reliability than grey matter seeding, in terms of the mean ICC measuring node strength. We note that these results may be partly biased due to the larger sample obtained with white matter seeding, as our setup involves seeding from all voxels in either grey or white matter, and there were approximately 15% more seed points in white matter than grey, and in areas of higher diffusion anisotropy. Nevertheless, we believe some grey matter seeding error may arise because all grey matter voxels were included even those unlikely to be involved in interfacing with white matter, potentially resulting in a greater proportion of spurious tracts. Secondly, overall, the probabilistic algorithm produced better test-retest performance than the deterministic method. This was perhaps due to the limited sample of possible streamlines produced by the deterministic method. The deterministic approach estimates a best fit of the diffusion tensor model at each voxel, whereas the probabilistic approach estimates a distribution. Concerning probabilistic tractography, the fibre model and the number of iterations are important considerations. The general recommendation is to use 1000 or more iterations per seed point. In this study, we applied 100 streamlines per seed voxel, equating to approximately 6 million streamlines per subject. Unfortunately, networks constructed with thousands of MCMC iterations are often impractical due to the computational cost involved. However, experiments using more than 100 iterations may be required to test this point.

In a recent study, Parker et al. (2014) compared network connectivity using a number of alternative segmentation, registration, fibre orientation estimation and tractography approaches from the same dMRI data ($N = 28$). Crucially, by comparing alternative pipelines, the authors found high agreement in the network similarity. And they identified replication of core connections and highly connected nodes across different network threshold and node configurations. This hints that current techniques are capable of providing a useful and representative estimation of white matter organisation.

To our knowledge, results of nodal reliability have rarely been reported in previous test-retest studies, even though reliability of individual nodes is crucial to the overall interpretation of networks. For white matter seeding, the regional findings (Figure 4.10) showed that the white matter waypoint constraints were largely ineffective in improving test-retest reliability, except in the case of SDL-weighting. However, for grey matter seeding the waypoint constraints improved the regional network reliability for all three types of weighting in comparison to unrestricted connectivity mapping. This is because grey matter seeding without any streamline filtering can produce streamlines which connect directly to a neighbouring grey matter ROI without apparently passing through any white matter. Due to partial volume effects some

of these connections may be genuine, particularly at the grey-white matter interface, but this is challenging to validate at dMRI resolution and such connections may often be spurious. However, we believe a constraint on white matter length is a more valid method of streamline filtering than thresholding purely by length. Waypoint constraints use prior knowledge of the white matter location to filter streamlines, whereas arbitrary length thresholds may remove genuine streamlines regardless of whether they pass through white matter. Our results indicated that constraining streamlines to pass through at least one white matter voxel was enough to improve network reliability.

An important consideration when quantifying interconnections is the normalisation applied to correct streamline counts for effects due to differences in grey and/or white matter volumes. However, some researchers have suggested that volume correction (e.g. SD and SDL-weightings) may overcompensate for volume-driven effects on streamline counts (Van Den Heuvel and Sporns, 2011). Streamline density weightings may down-weight the connection strengths for subjects with larger brains, potentially skewing the results of a group-wise analysis. It is not clear how such effects may be counteracted to allow representative comparison of connectivity between individuals. Alternatively, instead of streamline density it is possible to record some measure of tissue microstructure, such as diffusion anisotropy, averaged along the length of each tract (Iturria-Medina et al., 2007; Robinson et al., 2010). Such weightings somewhat circumvent the need for correction, as the weights reflect the underlying diffusion properties and are thus less affected by differences in tissue volume. Although we found that FA-weighting provided poorer test-retest performance than the SD-weighting (for most of the aspects assessed), it is possible that network weights based on properties of tissue microstructure may be more appropriate for group-wise analysis.

Some researchers have pointed out a bias arising from the tracking procedure, meaning that the number of fibres identified between a pair of regions decreases as a function of the distance (Zalesky and Fornito, 2009). This is due to the greater number of propagation steps and hence accumulated error associated with longer streamlines. An attempt to correct such biases may be made by a normalisation on streamline length. We compared networks with and without length correction (SDL and SD-weighting) finding that in almost all cases, whether seeding from grey or white matter, that the uncorrected weighting (SD) produced networks with better regional reliability (Figure 4.10), although SDL obtained marginally superior global reliability (Table 4.1). This suggests that for our data, the length correction overcompensates for accumulated errors in longer streamlines. Cheng et al. (2012a) made similar findings, suggesting that the length correction may not be necessary because the success rate for inter-region tracks was lower with longer fibres.

Concerning thresholding of network weights, it is problematic to ensure that an arbitrary thresh-

old removes spurious connections while retaining genuine patterns of connectivity. Additionally, applying the same threshold to multiple connectivity matrices is likely to results in different levels of sparsity, whereas matching sparsity across subjects is also problematic (Zalesky et al., 2010a). Thresholding has previously been reported to have a dramatic effect on the amount of false positive and false negative connections and the resulting measures of connectivity (van Wijk et al., 2010; de Reus and van den Heuvel, 2013). Trials comparing a range of suprathresholds on the connection weights indicated that some thresholds offered improvement in the overall test-retest reliability of connection (edge) statistics but marginal improvements in nodal and global properties. On the other hand, severe thresholds can be destructive. For instance, Figure 4.16(a) showed that thresholding causes some nodes to become disconnected (infinite path length). Clearly, such a brain network is anatomically implausible and no brain region could be entirely disconnected for our cohort. This highlights some serious concerns with thresholding approaches. Arguably, rather than using arbitrary thresholds it may be better to pursue anatomical constraints in order to reduce spurious connections, such as the FA, curvature and waypoint thresholds used in tractography. However, in the case of noisy data, some form of thresholding may still be necessary but should be applied with caution.

Overall, it is difficult to avoid the conclusion that our findings highlight some concerns about the regional reliability of dMRI networks, suggesting that the connections to some nodes were computed unreliably from session to session (Table 4.2, Figure 4.14, 4.15). On average, only one in eight nodes show a significant difference in between- and within-subject variation for our sample. In addition, only 22.6% (19/84) nodes have ICCs ≥ 0.5 across all four measures.

4.7.3 Sources of test-retest variation

Some test-retest variation may be due to scanner noise and inhomogeneities between sessions and some may be due to systematic variation in processing (Hagmann et al., 2010a; Van Essen and Ugurbil, 2012). Evidently, many intermediate steps are involved in generating structural networks from dMRI data. The variability at each step contributes to the variability in the following stage and in the resulting measures of connectivity. In particular, following our approach, network properties are dependent on reliable registration, node segmentation, diffusion processing and tractography.

Considering the regional test-retest agreement for segmentation, registration and streamline density (Figure 4.5), broadly, all three measures of reliability showed the strongest agreement in subcortical and medial structures. Overall, a gradual drop off in reliability was observed for more distal cortical regions. Segmentation and registration may perform more poorly in frontal and temporal regions because of EPI susceptibility artefacts and, to some extent, par-

tial voluming effects. Therefore, achieving accurate registration may be more problematic in these cortical regions in comparison to medial regions. The poorest segmentation overlap was observed in the left frontal pole, as shown in Figure 4.5(a), which was observed to be the one of the smallest and outermost nodes. The Desikan-Killiany atlas has known problems in segmenting the frontal pole because of its poorly defined cortical boundary (Desikan et al., 2006). Streamline density, which is derived from tractography, would be expected to be affected by errors in nodal segmentation and registration. Indeed, streamline density was observed to follow a similar pattern of regional variability to registration and segmentation, but with the poorest agreement in parietal regions as shown in Figure 4.5(c).

In our case, the regional network properties are reliant on the FreeSurfer morphometric procedure. It should be noted that the task of automated cortical labelling is extremely challenging. These methods have been demonstrated to show good test-retest reliability across scanner manufacturers and across field strengths (Han et al., 2006; Wonderlick et al., 2009; Reuter et al., 2012), although other studies have shown discrepancies in test-retest reliability (Morey et al., 2010; Gronenschild et al., 2012). Note that the volumetric agreement by ICC in these comparisons takes no account of spatial overlap and so ICC may not be a very appropriate assessment of segmentation. In this work, segmentation was directly assessed by voxel-wise correspondence using Dice overlap, which reports spatial agreement rather than volumetric agreement. However, overlap measures can be adversely affected by registration error between sessions. A recent study of a large cohort of 189 elderly subjects found that reliability of cortical and subcortical parameters was generally high (cortical: ICCs > 0.87, subcortical: ICCs > 0.95), but the authors suggest that fairly large cohorts are required to detect a 10% change between groups in regional cortical volumes (Liem et al., 2014).

An appropriate definition of network nodes is essential but far from trivial (Hagmann et al., 2010a; Zalesky et al., 2010b). Ideally, the node definition should be chosen such that it allows complete and unambiguous segmentation of the underlying tracts. A network formed from too few large nodes may fail to capture the true connectivity between regions. However, a network formed from too many small nodes may be influenced by systematic errors and noise, leading to spurious findings. Ideally, we wish to obtain a network resolution where the measurement of genuine structural differences is larger than the noise measurements. We chose the Desikan-Killiany atlas for our analysis as using larger and fewer nodes may minimise the effects of image noise and systematic error. We also assessed a 164 node parcellation (Destrieux et al., 2010) but segmentation was found to be poorer than with the 84 node configuration, in terms of test-retest agreement, and therefore not pursued. Given the resolution used in this study (84 nodes), our findings suggest that the measurement noise may degrade some of the genuine patterns of connectivity for several nodes. However, networks of 50–100 nodes are typical for

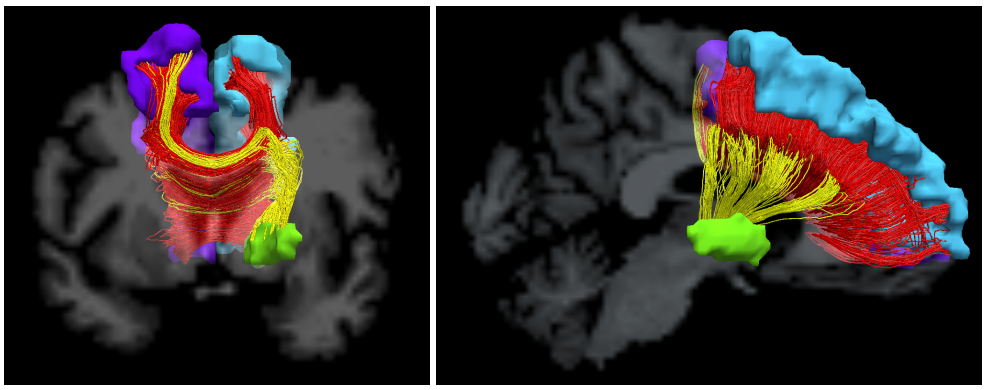

Figure 4.18: Coronal and sagittal views of anatomically plausible streamlines (red) and implausible streamlines (yellow). Implausible streamlines falsely connect from the left superior frontal gyrus to bilateral subcortical regions.

current connectome studies.

Many network studies have computed cortico-cortical connections only, even though, subcortical structures have an essential role in brain wiring. For instance, the thalamus is highly connected to many cortical regions (Behrens et al., 2003a). Therefore, we found it important to include sub-cortical structures in our analysis. Although we found a variation in the nodal properties between cortical and sub-cortical regions, the differences in reliability between whole-brain and cortico-cortical networks were not significant (Section 4.6.8.1).

Additionally, some error may reflect tractography issues in estimating the underlying axonal tracts. This may due to both ROI segmentation errors affecting seeding and methodological issues in streamline construction. Validation of tractography has been performed with good agreement with the underlying neuronal connections in the porcine and macaque brain (Parker et al., 2002; Dyrby et al., 2007), and with good agreement in a realistic artificial phantom (Fillard et al., 2011). However, tractography is know to be strongly affected by measurement noise resulting in both false positive and false negative connections (Jbabdi and Johansen-Berg, 2011; Zalesky and Fornito, 2009). Figure 4.18 illustrates a subset of false streamlines, which pass through the corpus callosum, obtained with probabilistic tractography. Bilateral cortical-subcortical connections are considered to be anatomically implausible, with the vast majority of callosal tracts believed to connect exclusively between cortical regions (Gazzaniga, 1998; Funnell et al., 2000). Such errors arise in streamline tracking due to limitations of tractography and the problems with crossing, branching and kissing fibres. Streamlines like those illustrated were not uncommon and this highlights serious limitations of our data and methods. The uncertainty in fibre directions for noisy measurements may also be a factor in the poor test-retest results. Fillard et al. (2011) suggested that for medium or low signal-to-noise datasets, an appropriate prior on the spatial smoothness of either the diffusion model or the fibres is

recommended for correct modelling. Refer to Yo et al. (2009) for a comparison of various tractography algorithms in a structural connectome context (although not using test-retest data).

4.7.4 Limitations of study

Analysis of reliability is challenging because results depend on both the variables analysed and the metric used. Here we used four dependent variables: node degree, node strength, path length and clustering coefficient. There are of course many other ways to characterise networks but these measures are the basis of graph-theoretic analyses. It is worth noting that in this study agreement was based on either global or nodal measures, which are themselves a function of elements of the adjacency matrix. It is also possible to directly assess every (non-zero) element of the adjacency matrix (see Zalesky et al. 2010a), which may offer further insight into network reliability at the level of individual connections. In terms of metrics, we relied on both the within- versus between-subject differences and the ICC. One issue with the ICC is that large samples may be required to estimate scores to acceptable precision. Shoukri et al. (2004) determined that for two repeated measures, in order to estimate a minimum acceptable ICC score of 0.8 with 95% confidence intervals of width 0.2, then 52 subjects are required (see Table 3 Shoukri et al. 2004). Similarly, to estimate a minimum acceptable ICC score of 0.6 with 95% confidence intervals of width 0.2, then 158 subjects are required. Clearly, such samples may be unrealistic in the case of MRI, and to date, no test-retest study has assessed a sample of this size, although Cheng et al. (2012a) used 44 subjects. Therefore, the studies relying on ICC (and also Pearson's correlation coefficient, which itself is a biased version of ICC because it assumes independent variances) have biased estimates. We can thus recommend that future dMRI network studies use a sample of more than 50 subjects. Furthermore, as well as ICCs, we have also quantified test-retest performance by comparing the within-subject to the mean between-subject differences. For a small sample the measurement error may be quite high, but bootstrap estimates are less affected by sample size. For most nodes the within-subject differences were smaller than between-subject differences yet most differences were not found to be significant, which can also reflect a lack of power. Nevertheless, our analysis provides insight into whether the genuine patterns of connectivity can be identified despite noisy measurements.

4.7.5 Conclusions

This chapter presented a test-retest analysis of structural brain networks obtained from MRI by comparing various factors affecting network construction. Our key findings were: 1) Probabilistic tractography was found to perform better than a deterministic method; 2) Performance

was improved when seeding from white matter, rather than grey. 3) Thresholding of network weights must be applied with caution in order to remove spurious connections while retaining genuine patterns of connectivity. We recommend that future dMRI network studies use large samples of ideally more than 50 subjects, in order to accurately estimate genuine between-subject differences above noise measurements. Overall, our results suggest that current connectome mapping techniques (at 1.5 T) are adequate for reliably measuring global network measures. However, regional network measures may not be reliable, leading to concerns about the validity of studies based on such measures particularly with small sample sizes.

Chapter 5

Brain networks in amyotrophic lateral sclerosis

5.1 Overview

Our method of network construction was used in an analysis of brain connectivity in a group of 30 amyotrophic lateral sclerosis (ALS) patients when compared with a group of age-matched healthy controls. For each subject, we identified 85 network nodes (including the brain stem) and whole-brain networks were constructed using an anatomically motivated white matter waypoint constraint and a weighting reflecting tract-averaged fractional anisotropy. An established statistical technique (NBS) was then used, without *a priori* selected regions, to identify a subnetwork (10 nodes and 12 bidirectional connections) of reduced connectivity in the ALS group compared with the controls ($p = 0.02$, corrected). These findings suggest that degeneration in ALS is strongly linked to the motor cortex. Reduced FA in three of the impaired network connections, which involved fibres of the corticospinal tract, was found to correlate with rate of disease progression. In addition, we used a novel network-tract comparison, which revealed that the connections involved in the affected network had a strong correspondence (mean overlap of 86.2%) with impaired white matter tracts identified using a standard voxel-based methods (TBSS).

This chapter begins with a brief overview of ALS, focussing on the reported changes in brain structure and connectivity associated with the disease. This is followed by a description of the participants, methods and the analyses performed. Results are then presented and, finally, the findings and limitations are discussed. In the course of this work, portions of the findings in this chapter have been published previously (Buchanan et al., 2014b) and were adapted and extended below.

5.2 Amyotrophic lateral sclerosis

ALS, the most common form of motor neurone disease, is a devastating adult-onset neurodegenerative disorder (Rowland and Shneider, 2001). The disease is characterised by chronic and usually rapid degeneration of the upper motor neurons of the motor cortex and the lower motor neurons of the brain stem and spinal cord. In sufferers, this leads to loss of motor function, progressive weakness, muscle atrophy and eventually death, often by respiratory failure. The median time from onset to death is just 39 months. However, for most patients, their cognitive function remains largely intact. It is though that the disease process causes Wallerian-like axonal degeneration (Bruijn et al., 2004), meaning that the axonal transport is gradually disrupted and the distal portion of the axon effectively becomes separated from the cell body. Eventually the axonal skeleton disintegrates, and the axonal membrane breaks apart, followed by degradation of the myelin sheath.

Approximately 5–15% sufferers are also afflicted with frontotemporal dementia (FTD) (Geser et al., 2009) and a proportion of non-demented patients present with specific cognitive impairment (Phukan et al., 2012), suggesting that the disease also affects extra-motor areas. Though the aetiology of ALS is not well understood, MRI has proved useful in probing the white matter degeneration attributed to ALS.

5.2.1 Previous MRI studies

Previous studies using dMRI have identified reduced white matter integrity in the corticospinal tract (Abe et al., 2004; Agosta et al., 2010; Blain et al., 2011; Ciccarelli et al., 2006; Douaud et al., 2011; van der Graaff et al., 2011; Sarro et al., 2011; Verstraete et al., 2010), the corpus callosum (Douaud et al., 2011; Sarro et al., 2011; Verstraete et al., 2010) and uncinate fasciculus (Sarro et al., 2011; Sato et al., 2010). Functional MRI studies have also identified abnormalities in extra-motor brain areas, including prefrontal regions (Tsermentseli et al., 2012; van der Graaff et al., 2009). And voxel-based morphometry analyses have found evidence of reduced grey matter volumes in the superior, medial and mid frontal gyri and anterior cingulate (Ellis et al., 2001; Kassubek et al., 2005) and reduced white matter volume and integrity in frontotemporal regions (Abrahams et al., 2005; Kassubek et al., 2005). Such findings have led to suggestions that ALS may be a progressive multi-system disorder (Cirillo et al., 2012; Geser et al., 2008; Rose et al., 2012).

However, many previous studies have relied on ROI approaches, which are typically constrained to a limited number of major white matter pathways or cortical areas. As the brain is a strongly interconnected system the degeneration of motor neurons may result in complex

and widespread changes. It is possible that whole-brain network analysis of the structural connectivity between brain regions (Sporns, 2011b), involving many hundreds of potential connections, may help to localise the degeneration in connectivity associated with ALS, especially for secondary or more widespread impairments in connectivity.

Whole-brain network analysis and statistical techniques, such as NBS (Zalesky et al., 2010a), can be used to identify differences in connectivity in case-control studies. NBS has previously been demonstrated in schizophrenia (Zalesky et al., 2010a, 2011) and in a similar ALS cohort (Verstraete et al., 2011, 2013). We investigated whether whole-brain network analysis, without *a priori* selected regions, can provide further insights into the white matter changes associated with the disease. This approach may provide supporting evidence of the localised impairments in motor areas surrounding the corticospinal tract and the corpus callosum, but it also has a capacity to identify secondary changes in connectivity in extra-motor areas.

In addition, we compared the results from a conventional TBSS approach (Smith et al., 2006) with the results from a network analysis of the same dMRI data. The aims of this chapter are as follows: 1) apply whole-brain network analysis with strong constraints on anatomically plausibility in order to identify white matter impairments due to ALS; 2) assess the agreement of network-based analysis with a conventional TBSS analysis; 3) assess any association between disease state and the findings from network analysis.

5.3 Participants

30 ALS patients and 30 age- and education-matched healthy controls were recruited and underwent an MRI protocol (as described in Section 3.2) and cognitive tests. The study was approved by the National Health Service Scotland Research Ethics Committee and the Department of Psychology, University of Edinburgh and informed consent was obtained from each subject.

The patient group were recruited from regional ALS services at the following sites throughout Scotland: Western General Hospital, Edinburgh; Southern General Hospital, Glasgow; and Ninewells Hospital, Dundee. All had clinical and electrophysiological evidence of combined upper and lower motor neuron involvement and fulfilled the revised El Escorial criteria for clinical definite and probable ALS (Brooks et al., 2000). 26 patients had sporadic ALS and four had a history of suspected ALS in a first degree relative. Ten patients had bulbar onset, eleven had upper limb onset and nine had lower limb onset. Exclusion criteria included the presence of another neurological disorder, the presence or history of a psychiatric disorder, and the presence of severe cardiovascular risk factors. None of the patients had evidence of dementia in

clinical notes or on initial discussion, although one patient was subsequently found to fulfil the criteria for possible behavioural variant FTD (Rascovsky et al., 2011) after a detailed interview with a caregiver. Disease severity was assessed using the ALS Functional Rating Scale-Revised (ALSFRS-R; Cedarbaum et al. 1999), and the rate of disease progression was determined using (Ellis et al., 1999; van der Graaff et al., 2011). Thirty age- and education-matched healthy controls were recruited from the University of Edinburgh, Psychology Department's Volunteer Participant Panel, staff working in the University of Edinburgh and from spouses of participating patients. None of the participants had significant neurological or psychiatric comorbidity.

5.4 Network construction

Structural networks were constructed as detailed in Chapter 3 and informed by the findings from our test-retest analysis. Whole-brain tractography was performed using probabilistic tractography as described in Section 3.6.2. Tracking was initiated from all white matter voxels and streamlines were constructed in two collinear directions until terminated by the following termination criteria designed to minimise the amount of anatomically implausible streamlines: (1) exceeding a curvature threshold of 70 degrees; (2) entering a voxel with FA below 0.1; (3) entering an extra-cerebral voxel; (4) exceeding 200 mm in length; and (5) exceeding a distance ratio metric of 10. Networks were then constructed by recording connections between all ROI pairs. The endpoint of a streamline was considered to be the first grey matter ROI encountered when tracking from the seed location. The *streamline density* weighting, recorded the interconnecting streamline density corrected for ROI size was computed (Equation 3.10). As tractography is prone to producing false connections (Van Essen and Ugurbil, 2012), we used prior knowledge of white matter anatomy to discard a proportion of spurious connections. Any implausible streamlines traversing from one cortical hemisphere to any contralateral subcortical node were discarded (Funnell et al., 2000). To further reduce spurious connections, a two-step threshold on the network weights was then applied: 1) for each subject, discard the weakest 25% of weights in the matrix by connection probability (Equation 3.10); 2) across the cohort, only retain connections which occur in at least 50% of subjects. The second step is required to discard connections which have been removed for some subjects but not others by the first step. From the remaining set of streamlines *FA-weighted* networks were then computed by recording the mean FA value along interconnecting streamlines (Equation 2.7). FA-weighted networks were constructed because FA is believed to reflect aspects of the microstructural integrity of the underlying axonal fibres. For each FA-weighted matrix, five global network measures, as described in Section 3.8, were then computed. These five measures were the *network sparsity*, the *network strength* (the average sum of weights per node), *network clustering coefficient* (an average measure of local connectivity), the *network efficiency* (the average of the inverse shortest

path length) and the *network transitivity* (a normalised variant of clustering coefficient).

5.5 Statistical analysis

The global network properties and the measures of disease state and disease progression were assessed for normality by Shapiro-Wilk test with non-normality accepted at $p < 0.05$ level of significance. Group contrasts were performed by a two-sample t-test for Normally distributed data. The two-tailed probability level was used for all global comparisons. Correlations for non-Normal variables were computed with Spearman's rank correlation coefficient.

5.5.1 Network analysis

Firstly, in order to assess any differences in global connectivity between the patient and control groups, the global network measures were tested with uncorrected $p < 0.05$ being considered statistically significant. Secondly, network connections were compared between the patient and control groups using NBS (Zalesky et al., 2010a), without *a priori* selected regions. For our 85 node networks there are 3570 possible network connections. As a result, standard statistical tests may be under-powered when corrected for multiple comparisons. NBS is an alternative approach which exploits the extent to which the connections identified by the contrast are interconnected to offer a potential gain in statistical power, for which the significance of maximally connected subnetworks are assessed rather than individual connections (Zalesky et al., 2010a).

In the NBS framework, first a two-sample one-tailed t-test was performed at each of the 3570 network connections in order to identify reduced white matter integrity in the patient group compared with the control group. Secondly, a set of suprathreshold edges and the corresponding set of maximally connected network components was computed by a network-defining threshold on the t-statistics. Permutation testing, which randomly exchanged the group to which each subject belonged, was used over 5000 iterations to estimate the distribution of component size and compute a corrected p-value for the maximally connected subnetwork(s). In NBS terminology, the 'intensity' of each maximally connected network was tested rather than the 'extent' (size of the network) as this directly assesses the magnitude of the test statistic. For completeness, we also repeated the NBS analysis to identify any increased white matter integrity in the patient group compared to the control group. Both comparisons were then repeated with correction for age. In addition, the streamlines involved in any network identified by NBS were used to compute maps of streamline density per voxel for each subject, as described in Section 3.6.6. The cerebral areas identified in these maps were then visually compared with the findings from TBSS.

5.5.2 Tract-based spatial statistics

TBSS is an established voxel-based analysis of white matter tracts using dMRI, which has been described in detail elsewhere (Smith et al., 2006). A between-group comparison of FA was performed with TBSS. Firstly, a nonlinear deformation was used to align the FA map of each subject to a white matter template in standard space. White matter masks were 'skeletonised' (morphologically thinned) in order to obtain the centre-line of the principal white matter pathways, while minimising the impact of registration error and partial volume effects. Voxel-wise FA was then compared between groups within the white matter skeleton. Finally, permutation testing assigned a corrected p-value to each voxel, for which $p < 0.05$ was considered significant. Comparisons were performed to identify both reduced and increased FA in the patient group when compared to the control group.

5.5.3 Comparison between network analysis and TBSS

We formulated a simple measure of *network-tract agreement* in order to compare the findings from the network analysis and the findings from the TBSS procedure. This assesses whether the network connections overlap with the voxels of the white matter skeleton identified by TBSS. Such a measure of agreement may be useful in validating that any impaired connections identified by a network-based contrast, such as NBS, overlap with the white matter skeleton identified by a TBSS contrast when assessing the same data. Firstly, the corrected p-value maps produced by the TBSS contrast were thresholded at $p < 0.05$. These masks were then transformed to each subject's native space, using the nonlinear transforms computed by the TBSS procedure. For each subject, a measure of overlap per connection based on streamline density was computed,

$$r_{ij} = \frac{|D_{ij}|}{|S_{ij}|}, \tag{5.1}$$

where $|S_{ij}|$ is the count of all streamlines found between nodes i and j (as in Section 3.8) and $|D_{ij}|$ is the count of streamlines which pass through at least one voxel of the p-value mask (where $D_{ij} \subset S_{ij}$). Note that the correction for ROI size (as in Eq. 3.10) is not necessary here as this cancels out when applied to both D and S. For each network connection, this results in a score between 0 and 1, reflecting the proportion of streamlines connecting node i and node j which pass through any voxel identified by TBSS. A t-test was used to compare the mean overlap for connections within the observed network against the non-zero connections outwith this network.

	N (ALS:Controls)	ALS Patients	Controls
Age (years)	30:30	58.5 (28–79)	59.1 (34–79)
Gender (male:female)	30:30	17:13	17:13
WTAR IQ	27:30	104 (81–123)	108 (87–124)

Table 5.1: Demographic data for patients with ALS and control subjects: showing mean values with ranges in parentheses. Ratios are presented for number of patients versus controls and for group gender breakdown. WTAR indicates Wechsler test of adult reading

5.6 Results

5.6.1 Participants

Demographic data and the findings from a battery of neuropsychological tests have been described previously (Bastin et al., 2013; Pettit et al., 2013). Table 5.1 shows the demographic data in terms of age, gender and premorbid intelligence quotient (IQ) as assessed by the Wechsler Test of Adult Reading (WTAR). The mean age of the 30 ALS patients, of which 17 were male, was 58.5 ± 11.2 years (mean $\pm$ standard deviation). The mean age of the 30 control subjects, of which 17 were male, was 59.1 ± 12.0 years. There were no significant differences between groups in age, gender or IQ as assessed by WTAR. In the patient group, the mean ALFRS-R score was 38.8 ± 6.76. The mean disease duration was 24.0 ± 18.35 months and the median disease duration was 19.4 as two patients had disease duration of over 60 months. The corresponding disease progression rate was 0.49 ± 0.39. Two ALS sessions were excluded from the study due to incomplete MRI data or excessive motion artefact.

5.6.2 Network analysis

Visual inspection of the segmentations for each subject indicated that the FreeSurfer procedure provided plausible brain extraction, tissue segmentation and cortical labelling. Approximately 6 million streamlines were seeded per subject. Visual assessment of the streamlines remaining following network thresholding indicated that the majority of streamlines were anatomically plausible. Out of the total streamlines seeded, $14.73 \pm 1.59\%$ were then identified as interconnections in the ALS group and $15.24 \pm 1.96\%$ in the control group. A one-tailed *t*-test found no difference in the streamline success rate between groups. Figure 5.1(a) shows the mean connectivity matrix averaged across all subjects. Figure 5.1(b) shows the histograms of network weights in both groups, indicating that there is very little variation in global connectivity between the patient and control groups. All global network measures (Table 5.2) were found to be approximately Normally distributed. Although all global measures were lower in the ALS

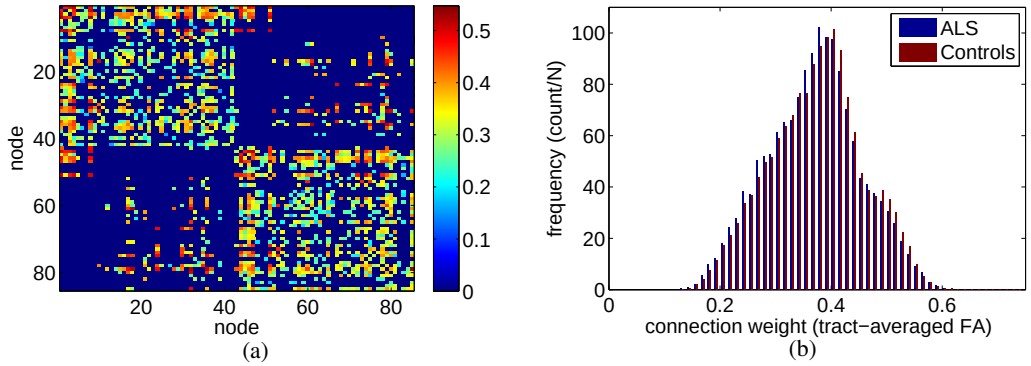

Figure 5.1: a) 85×85 connectivity matrix showing the mean weights across all subjects (N=58); b) The histograms of network weights (tract-averaged FA) across all subjects in both groups.

Network measure	ALS Patients		Controls		t-stat	p
Network sparsity	0.782	(0.01)	0.783	(0.01)	-0.451	0.654
Network strength	6.867	(0.4)	6.913	(0.5)	-0.363	0.718
Network clustering coefficient	0.227	(0.01)	0.230	(0.01)	-1.236	0.222
Network efficiency	0.229	(0.01)	0.231	(0.01)	-0.732	0.467
Network transitivity	0.200	(0.01)	0.201	(0.01)	-0.798	0.428

Table 5.2: Global network properties with mean values and standard deviations in parentheses, t-statistic and uncorrected p-value between groups.

group than in the control group, individual t-tests suggested that these differences were not significant.

However, NBS identified a subnetwork (10 nodes and 12 bidirectional connections, Figure 5.2) of impaired connectivity in the ALS group ($p = 0.020$, corrected). The t-statistic threshold was set to 2.6. This value was chosen empirically – an approach suggested by the author of NBS (Zalesky et al., 2010a).

The resulting network involves four nodes within the primary motor cortex (bilateral precentral and paracentral), left superior frontal, the left-posterior cingulate and four subcortical areas (bilateral pallidum, left thalamus, left caudate). All 12 network connections are directly linked to nodes within the primary motor cortex. Table 5.3 shows the mean FA values in both groups and the corresponding t-statistic for each of these network connections. Overall, the connections in the affected network showed a mean reduction of 0.04 ± 0.03 (approximately 10%) in terms of FA in the patient group when compared to the controls.

Figure 5.3 shows the cortical regions and the streamlines involved in the affected network for one patient. Figure 5.4 shows the areas involved in the affected network across all subjects,

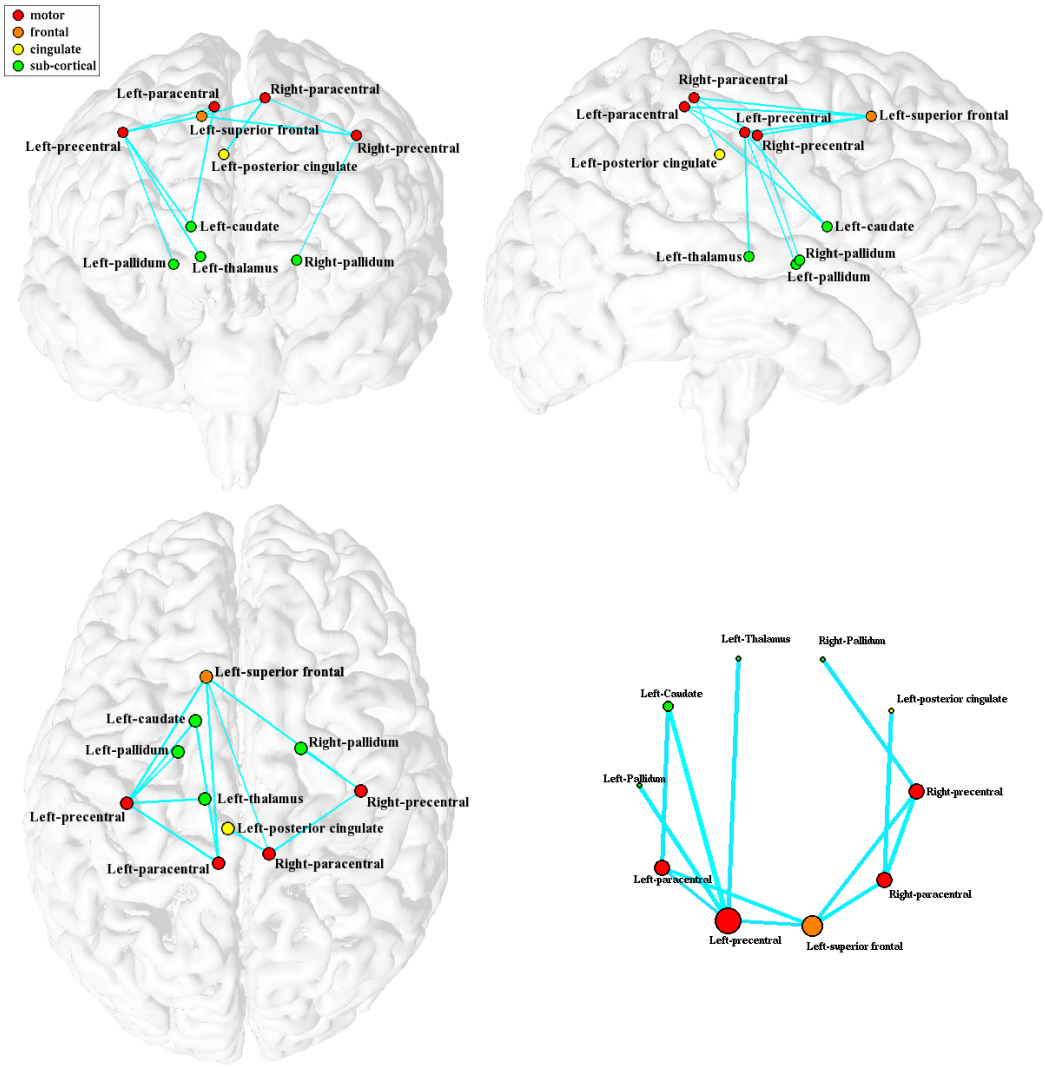

Figure 5.2: Coronal, sagittal, axial and abstract views of the impaired network nodes and interconnections identified by NBS where lobe/cortex membership is indicated by node colour.

Network connection	ALS Patients		Controls		t-stat
Left-caudate – Left-precentral	0.37	(0.08)	0.43	(0.03)	3.454
Right-paracentral – Right-precentral	0.30	(0.04)	0.33	(0.04)	3.190
Left-paracentral – Left-superior frontal	0.29	(0.04)	0.32	(0.03)	3.161
Left-pallidum – Left-precentral	0.47	(0.04)	0.50	(0.03)	3.153
Left-superior frontal – Right-paracentral	0.44	(0.1)	0.50	(0.04)	3.093
Left-posterior cingulate – Right-paracentral	0.47	(0.04)	0.50	(0.04)	2.993
Left-superior frontal – Right-precentral	0.36	(0.2)	0.47	(0.1)	2.905
Left-thalamus – Left-precentral	0.45	(0.04)	0.47	(0.02)	2.858
Right-pallidum – Right-precentral	0.46	(0.04)	0.48	(0.03)	2.843
Left-caudate – Left-paracentral	0.37	(0.04)	0.39	(0.03)	2.783
Left-precentral – Left-superior frontal	0.36	(0.04)	0.39	(0.03)	2.681
Left-paracentral – Left-precentral	0.30	(0.04)	0.33	(0.04)	2.631

Table 5.3: Mean FA values of the impaired network connections identified by NBS with standard deviations in parentheses and the t-statistic. All connections are collectively assigned a single p-value ($p = 0.020$, corrected).

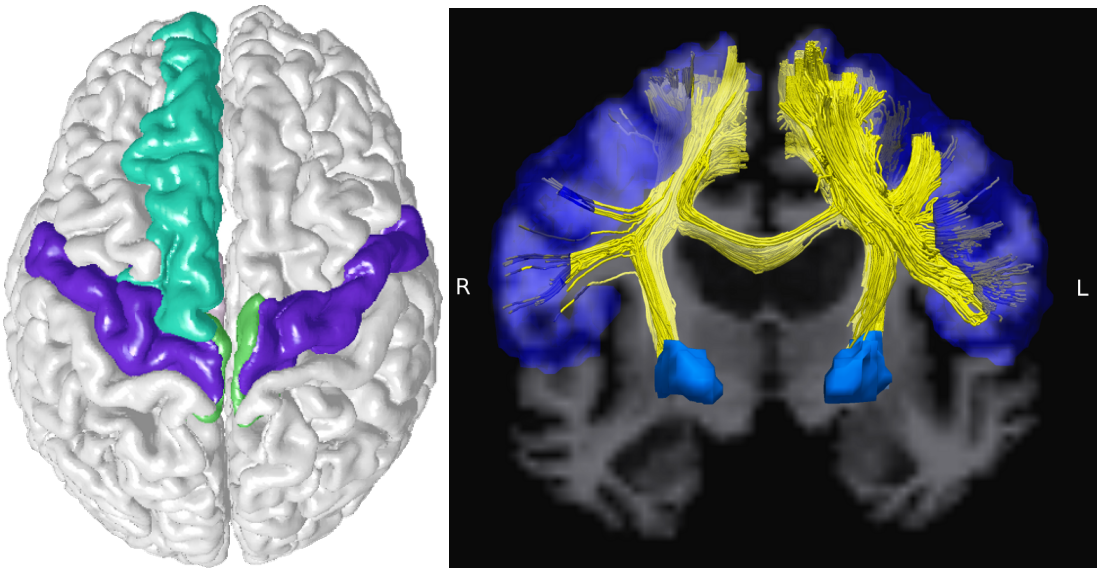

(a) (b)

Figure 5.3: a) The cortical regions (bilateral precentral, bilateral paracentral and left superior frontal gyrus) involved in the impaired network; b) Coronal view of the streamlines (yellow) involved in the impaired network for one ALS patient (63 year old male) where colouring indicates the precentral gyrus and globus pallidus.

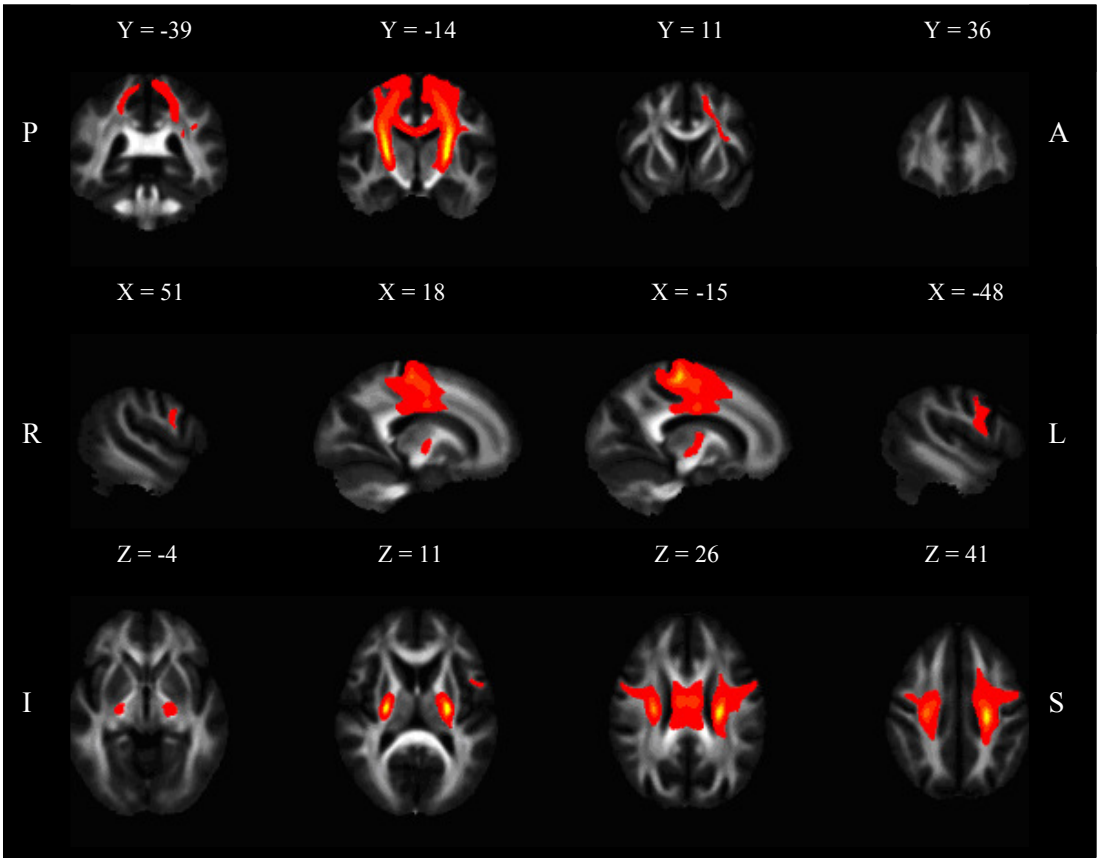

Figure 5.4: Coronal, sagittal and axial slices of the mean FA in MNI standard space, overlaid with the mean streamline density per voxel computed from the streamlines involved in the affected network and transformed to MNI space.

in terms of the mean streamline density per voxel. The converse NBS contrast testing for increased white matter integrity in the patient group when compared to the controls produced no significant results. In addition, the contrasts which included age as a covariate showed no fundamental differences in the results or topology of the observed network.

5.6.3 Comparison of network analysis to TBSS

The TBSS contrast comparing the patient and control groups found significant reductions in FA within the corticospinal tract and portions of the corpus callosum (Figure 5.5). TBSS found no areas of increased FA in the patient group when compared to the controls.

The affected areas of the white matter skeleton identified by TBSS were found to overlap with the areas involved in the affected network (Figure 5.4). Figure 5.6 shows the mean overlap proportion (Equation 5.1) with 95% inter-percentile range for the 12 network connections identified by NBS. Notably, 11 out of 12 connections had an overlap proportion > 0.75, suggesting

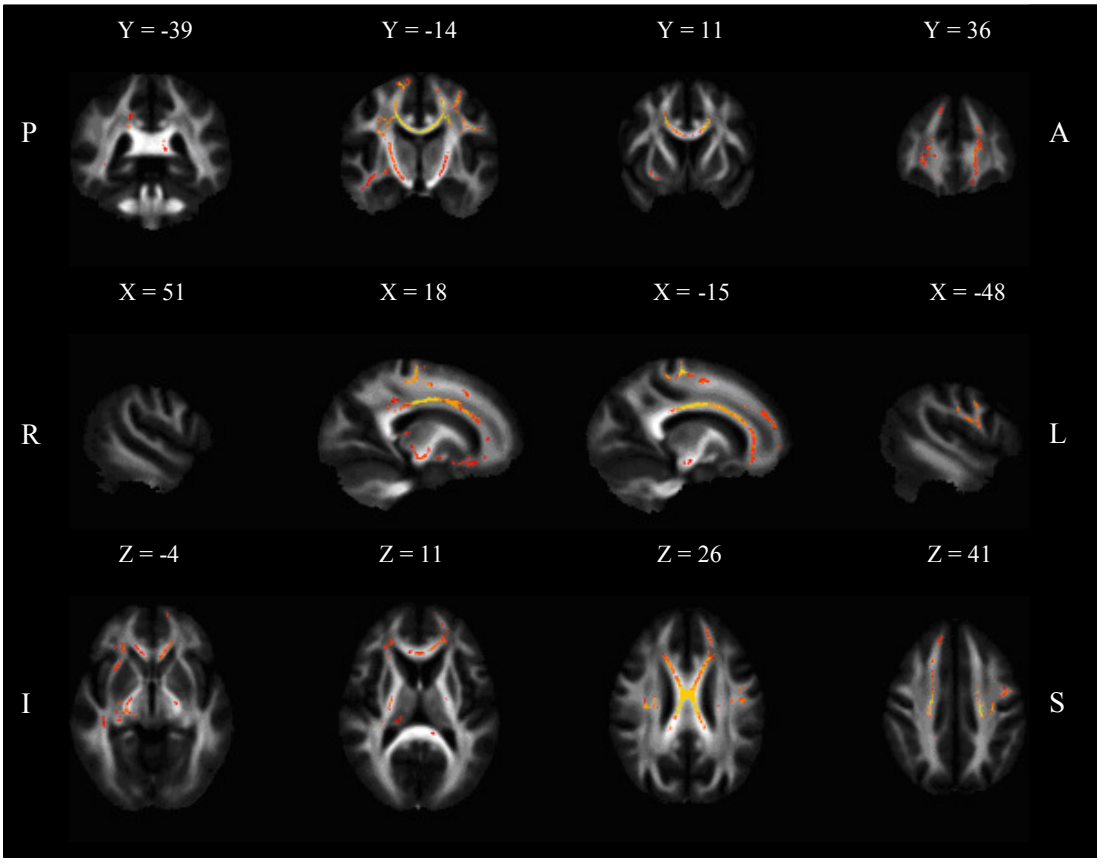

Figure 5.5: Coronal, sagittal and axial slices of the results from TBSS, showing the mean FA in MNI standard space overlaid with the voxels found to have reduced white matter integrity (FA) in the patient group compared to the controls ($p < 0.05$).

that the majority of streamlines involved in the affected network passed through the white matter identified by TBSS. Only the connection between left paracentral and left superior frontal gyrus had a lower score of 0.21, suggesting that the streamlines involved in this connection were less in agreement with the TBSS analysis. Overall, $73.6 \pm 2.2\%$ of all (whole-brain) network connections in the patient group and $73.0 \pm 2.9\%$ in the control group had at least one streamline which passed through a region identified by TBSS, suggesting that the white matter regions identified by TBSS are common to many white matter connections. However, 39.7 ± 2.3 of all network connections in the patient group and $40.1 \pm 2.0\%$ in the control group had half of the total streamlines pass through a region identified by TBSS. Crucially, the mean overlap proportion (Equation 5.1) is 0.385 ± 0.017 for connections outside the impaired network, compared to a mean proportion of 0.862 ± 0.041 for the connections within the impaired network. A t-test showed that these proportions were significantly different ($p << 0.001$).

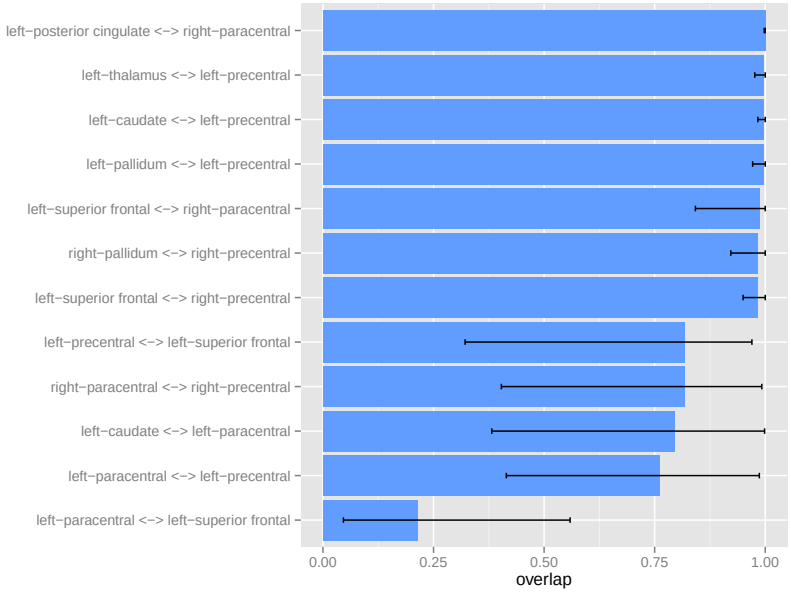

Figure 5.6: Mean network-tract overlap proportion with 95% inter-percentile range (N=58) for the 12 connections of the affected network, showing the proportion of interconnecting streamlines which pass through any white matter region identified by TBSS at $p < 0.05$.

5.6.4 Impaired connectivity correlates of disease state

Relationships between the 12 affected network connections, which showed significant differences between patient and control groups (Figure 5.2), and two clinical measures (ALSFRS-R score and disease progression) were investigated in the patient group. Correlational analyses found four uncorrected associations with the ALSFRS-R score and three associations with the disease progression rate (Table 5.4). Notably, three of the connections identified were the same for both the ALSFRS-R score and the progression rate (Figure 5.7). Relationships between the clinical measures were also assessed, finding that the ALSFRS-R score and the progression rate were inversely correlated ($\rho = -0.52, p = 0.006$).

5.7 Discussion

This study presented a whole-brain network analysis of white matter integrity in ALS, without using *a priori* selected regions. Global network properties (Table 5.2) indicate that the networks in both the patient and control groups were similar and show a level of clustering and network efficiency comparable to other network studies (Verstraete et al., 2011). Note that the absolute values of these network properties are dependent on the choice of nodes and thresholding in each study. The between-group comparison of the global network properties found no brain-

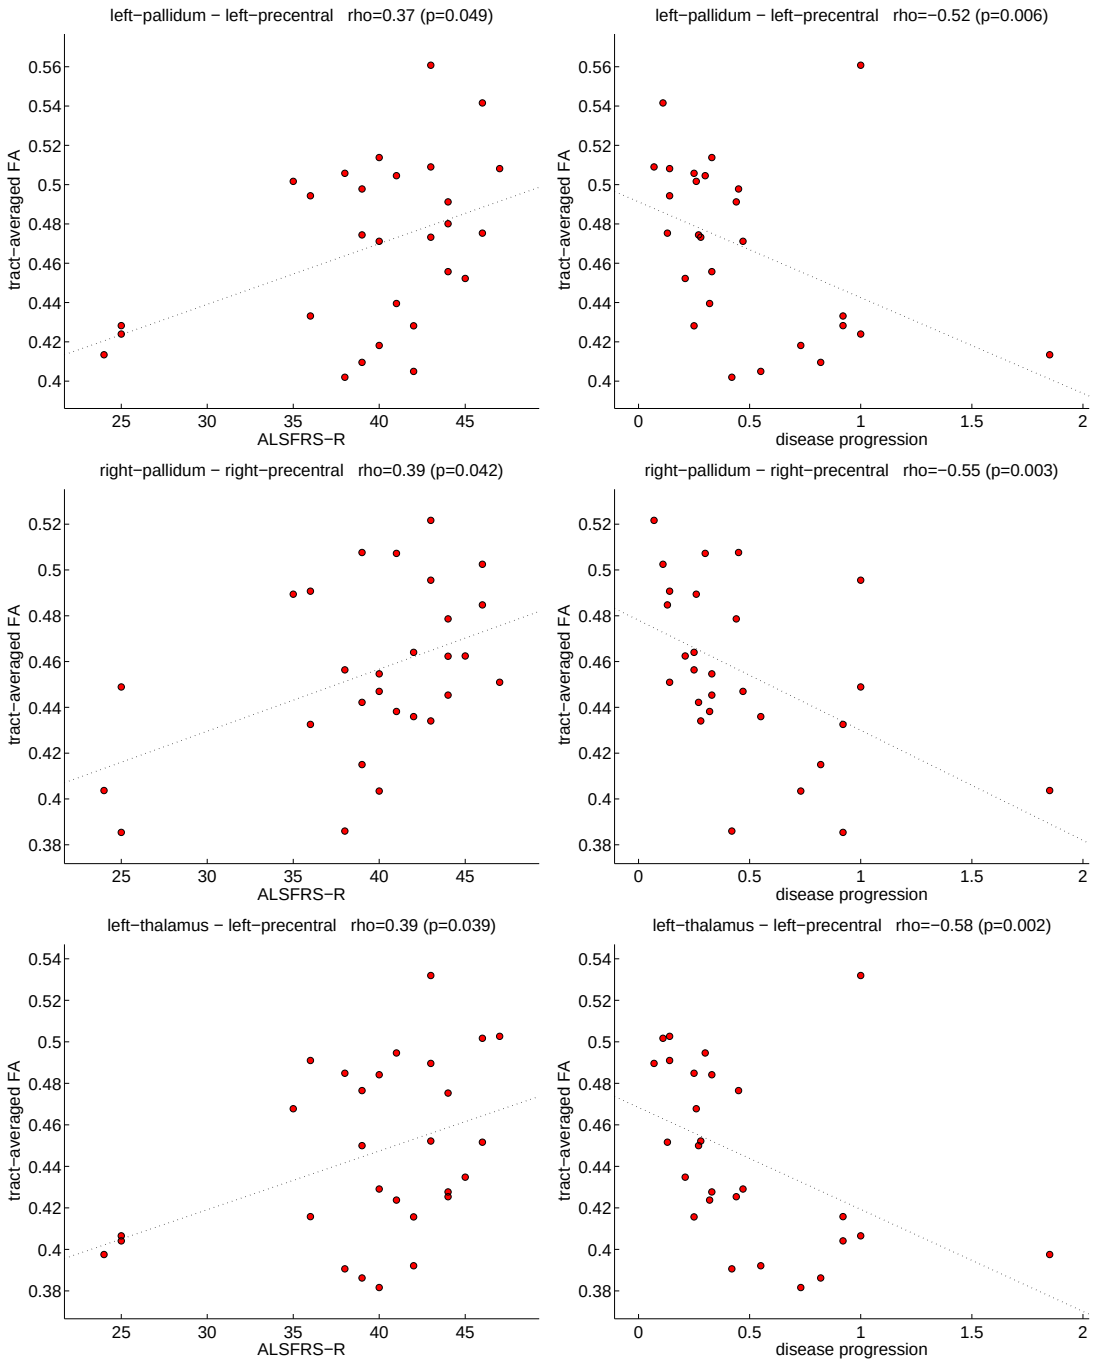

Figure 5.7: Correlations between clinical measures and the tract-averaged FA of three impaired network connections for the ALS patients ($p < 0.05$, uncorrected): Left) ALSFRS-R score; Right) the disease progression rate.

	ALSFRS-R score			*Disease progression*		
Network connection	ρ	p	q	ρ	p	q
Left-caudate – Left-paracentral	0.166	0.399	0.532	-0.035	0.862	0.862
Left-thalamus – Left-precentral	0.392	0.039	0.118	-0.578	0.002	**0.017**
Left-caudate – Left-precentral	0.383	0.044	0.118	-0.292	0.140	0.322
Left-pallidum – Left-precentral	0.375	0.049	0.118	-0.515	0.006	**0.024**
Left-paracentral – Left-precentral	-0.016	0.936	0.936	0.139	0.488	0.651
Left-paracentral – Left-superior frontal	0.104	0.597	0.717	-0.098	0.627	0.752
Left-precentral – Left-superior frontal	-0.073	0.712	0.777	0.074	0.714	0.779
Left-posterior cingulate – Right-paracentral	0.325	0.092	0.157	-0.299	0.129	0.322
Left-superior frontal – Right-paracentral	0.325	0.091	0.157	-0.261	0.188	0.322
Right-pallidum – Right-precentral	0.387	0.042	0.118	-0.552	0.003	**0.017**
Left-superior frontal – Right-precentral	0.430	0.022	0.118	-0.265	0.182	0.322
Right-paracentral – Right-precentral	0.261	0.181	0.271	-0.153	0.447	0.651

Table 5.4: Correlations between clinical measures and the tract-averaged FA of the impaired network connections, showing both uncorrected p-values and q-values adjusted by false discovery rate. Emboldening indicates significance after adjustment ($q < 0.05$).

wide impairments in connectivity for the patient group when compared to the control group as shown in Table 5.2 and Figure 5.1(b). By definition, global network properties take an average across all network connections and therefore any localised effects could be averaged away. These results suggest that, as expected, ALS is not a brain-wide disease and that global brain organisation has remained largely intact for the ALS patients. This finding is in agreement with Verstraete et al. (2011), who found no differences in global graph metrics between ALS patients and controls, but did find localised impairment in motor and frontal brain areas.

Our results suggest that in the patient group, connectivity to primary motor, prefrontal, and subcortical areas is substantially reduced, in terms of tract-averaged FA, and that these impairments are predominantly localised around the motor cortex (Figure 5.2, 5.3). Notably, although a brain-wide analysis was performed, without *a priori* selected regions, the network identified involves regions which are known to be associated with motor control and movement. The white matter pathways identified are consistent with upper motor neuron pathology.

Previous dMRI studies have shown reduced white matter integrity in the corticospinal tract and corpus callosum (Agosta et al., 2010; Cirillo et al., 2012), areas which are interlinked with several of the subcortical and motor cortex nodes within our affected network (Figures 5.2, 5.3). The network study of ALS by Verstraete et al. (2011) identified a nine node network which found a comparable pattern of motor network impairment involving connections to precentral, paracentral, pallidum, frontal areas and the cingulate cortex. In addition, voxel-based studies have identified reduced white matter changes in the corticospinal tract and the medial

portion of the corpus callosum (Abrahams et al., 2005; Filippini et al., 2010). Previous functional MRI studies have identified impairment in functional connectivity associated with ALS (Mohammadi et al., 2009). Promisingly, studies in healthy human volunteers have found some convincing associations between functional and structural connectivity (Hagmann et al., 2008; Honey et al., 2010). However, to our knowledge, this structure-function correspondence has not yet been demonstrated in an ALS network. One recent study has produced findings which suggest that the level of functional connectedness within the motor network is correlated with the rate of disease progression (Verstraete et al., 2010).

The comparison between network analysis and the findings from TBSS offers some insight. Approximately 40% of all connections had half of the total streamlines pass through a region identified by TBSS. This indicates that the white matter identified by the TBSS contrast contains a common 'hub' connecting grey matter regions, for instance, via portions of the corpus callosum. The comparison of the mean overlap proportion within and outwith the impaired network indicates a strong, but not perfect, agreement between the white matter identified by TBSS and the affected network identified by NBS. Notably, NBS identified one network connection between left paracentral and left superior frontal gyrus that was unlikely to have been concluded from the TBSS analysis (Figure 5.6). Conversely, there were several connections outside the affected network which also had a high overlap proportion (close to 1), which were not identified by NBS. Due to limitations of the data and methodological problems with tractography, some connections may contain spurious streamlines which cross the regions identified by TBSS.

The associations between the affected network connections and the disease state indicate that for the patients, the bilateral precentral to pallidum and left precentral to left thalamus connections are increasingly impaired over disease progression (Figure 5.7). Notably, all affected connections (Table 5.4) involve fibres which run through the cortical to subcortical portion of the corticospinal tract (Figure 5.3). Both the ALSFRS-R scores and disease progression rate correlate with a reduction in white matter integrity in these fibres. These findings provide supporting evidence for similar relationships found between disease progression and FA in the rostral portion of the corticospinal tract reported in a previous study (Verstraete et al., 2010).

Furthermore although measures of cognitive performance were gathered for the patient group (Pettit et al., 2013), no meaningful correlations emerged between the affected network and these measures of executive performance (data not shown). However the affected network included prefrontal regional nodes (namely the superior frontal gyrus) found to be affected in our previous analyses and related to impairments in letter fluency (Pettit et al., 2013).

Our ALS cohort includes a proportion of patients with FTD. Consequently, it might be expected that network analysis would identify impairment in extra-motor regions, as other researchers

have posited (Verstraete et al., 2011). However, with the exception of left superior frontal gyrus, our network was largely localised around the motor cortex. A post-hoc NBS analysis using a more speculative threshold on the t-statistic suggested that many extra-motor brain regions, including frontal and temporal areas could be identified by the contrast but at lower levels of significance. These findings suggest that although changes in structural connectivity may be widespread in ALS, overall the degeneration is strongly linked to the motor cortex.

5.7.1 Strengths and limitations

One strength of this study is the relatively large and well-characterised patient samples with age- and education-matched controls. However, we are aware that control subjects recruited from within the University are unlikely to be well matched for socio-economic status or IQ. To our knowledge, this is the first study to utilise probabilistic tractography and model multiple fibre populations in a network analysis of ALS. Such techniques should be better able to account for branching and crossing fibres, in comparison to deterministic tractography. In addition, previous ROI-based studies are typically limited to assessing changes in a limited number of regions. However, the network analysis used in this study offers the possibility to explore ALS as a network disease, potentially involving thousands of connections. Furthermore, in comparison to TBSS or other template-space analyses, network analysis is computed in native space rather than standard space, thereby accounting for individual differences in white matter structure and providing a more representative reconstruction of the underlying axonal wiring. Like previous network studies (Verstraete et al., 2011, 2013) we chose to use FA-weighted networks, rather than streamline density, as FA is likely to be a more representative measure of disruption in the underlying axonal fibres in case-control studies.

We note that although NBS was designed to reduce the false positive rate, tractography is known to be strongly affected by measurement noise resulting in both false positive and false negative connections (Jbabdi and Johansen-Berg, 2011; Zalesky and Fornito, 2009). Some error may reflect tractography issues in estimating the underlying axonal fibres from noisy measurements. Other errors may be due to both ROI segmentation errors affecting seeding. Additionally, thresholding of networks must be performed with caution. We believe that the two-step thresholding procedure used in this study eliminates a proportion of implausible connections without biasing the results of a group-wise analysis. However, due to the limitations of current dMRI and tractography techniques it is not possible to eliminate all erroneous connections.

In addition, when considering the correlations between clinical measures and the tract-averaged FA of three impaired network connections (Figure 5.7), we observed that there is a possibil-

ity that some data points could be outliers. An outlier which falls near the regression line will falsely increase the value of the correlation coefficient. We are aware that such outlying points could have some impact on the correlation results because the sample size is fairly small (N=28).

5.8 Conclusions

This chapter presented a whole-brain network analysis of white matter degeneration in ALS using strong constraints on the anatomical plausibility of tracts. The key findings were that while there were no brain-wide impairments in connectivity due to ALS, an impaired motor-frontal-subcortical network of reduced white matter integrity was found in the ALS patients. These findings suggest that degeneration in ALS is strongly linked to the motor cortex. The connections involved in our network analysis had a strong correspondence with tracts identified by a conventional TBSS analysis of the same data. Reduced white matter integrity in three of the impaired network connections, which involve fibres of the corticospinal tract, correlated with both the ALSFRS-R score and disease progression. Further analysis of the association between cognitive impairments and extra-motor connections merits further investigation.

Chapter 6

Brain networks in normal ageing

6.1 Overview

Our methods were used to assess structural brain connectivity in a group of 80 healthy volunteers, aged from 25 to 64 years. Such a cross-sectional analysis has the potential to find supporting evidence, in terms of brain organisation, for the widely reported age-related decline of white matter integrity. For each subject, we identified 85 network nodes and whole-brain networks were constructed using an anatomically motivated white matter waypoint constraint and a weighting reflecting tract-averaged FA. Motivated by the findings from previous research in ageing we carried out a number of statistical analyses to assess the relationships between age and brain structure in our cohort. Firstly, a structural analysis was performed to assess age-related changes in tissue volume and in global white matter integrity. Secondly, a conventional voxel-based analysis (TBSS) of the dMRI data was performed to assess the effects of age and gender within major white matter pathways. Thirdly, the relationships between age and both global and lobar network properties were assessed, including those properties which may capture a decline in connectivity, network efficiency or interhemispheric transfer. Finally, a mass-univariate network analysis (NBS) was used to identify any network connections showing changes with age and/or gender.

This chapter begins with a brief overview of the ageing brain, covering the reported age-related changes in brain structure, white matter and connectivity. This is followed by a description of the participants, methods and the analyses performed. Results are then presented and, finally, the findings and limitations are discussed.

6.2 The Ageing Brain

A huge amount of neuroscientific effort has been put into understanding the mechanisms and consequences of ageing in the brain. Research suggests that brain ageing is associated with marked chemical and structural alterations which give rise to changes in function and cognition.

6.2.1 Cognitive ageing

In general, the ability to perform working memory tasks and rapid processing becomes increasingly difficult with advancing age, even in the absence of pathology (Salthouse et al., 2003). There is strong evidence that ageing primarily affects the cognitive processes associated with frontal brain regions (Buckner, 2004; Hedden and Gabrieli, 2004; DeCarli et al., 2005), such as processing speed, working memory, executive function, problem solving and reasoning abilities (Sullivan and Pfefferbaum, 2006). The frontal ageing hypothesis (O'Sullivan et al., 2001) posits that age-related brain changes selectively impact frontal regions. However, the notion of exclusive frontal ageing remains controversial and the global disconnection hypothesis (Greenwood, 2000) predicts that changes may be more widespread. Although cognitive abilities have been found to decline in old age (60–85 years of age), there is evidence that in elderly patients, auxiliary brain regions may be recruited when performing executive functions – a pattern which is not seen in younger adults (Cabeza et al., 2002). Additionally, functional research shows a typical age-related shift from unilateral processing to bilateral processing, meaning that more regions are recruited as needed (cognitive reserve), but performance does not necessarily decline (Whalley et al., 2004). However, in very old age (more than 85 years of age) any compensatory systems may themselves become too degraded to support processing (Sullivan and Pfefferbaum, 2006).

6.2.2 Structural changes

Structural neuroimaging studies have consistently shown an age-related decline in cortical grey matter volume and an associated increase in ventricular volume and extracellular space (Good et al., 2001; Resnick et al., 2003; Sullivan et al., 2004). Grey matter atrophy has been shown to primarily affect frontal, cingulate, insular and inferior parietal regions (Good et al., 2001; Resnick et al., 2003). While there is a marked decline in cortical volume with age, post-mortem studies have shown that there is little change in the total neuronal population (Terry et al., 1987). Likewise, some studies have found a decrease in white matter volume from about 65 years of age, though others studies have found little age-related change except in localised areas of white matter showing relatively accelerated loss (Good et al., 2001). However, the

rate of both grey and white matter atrophy appears to accelerate from about 70 years of age (Resnick et al., 2003). MRI studies typically also show white matter abnormalities in the basal ganglia and hypointensities due to iron deposits in the globus pallidus and putamen (Moseley, 2002).

6.2.3 White matter changes

Several dMRI studies have identified an age-related decline in white matter integrity, which is most pronounced in frontal regions, particularly in very old age (Pfefferbaum and Sullivan, 2003; Pfefferbaum et al., 2005; Salat et al., 2005; Sullivan et al., 2006; Bastin et al., 2008, 2010). Large age-related decreases in FA have been found in white matter structures, such as the corpus callosum genu (Sullivan et al., 2006; Bastin et al., 2008, 2010), splenium (Pfefferbaum and Sullivan, 2003), anterior cingulum, and middle frontal gyrus (Pfefferbaum et al., 2005). Such age-related declines in FA may be indicative of the gradual demyelination and loss of myelinated axons, which have been observed in post-mortem studies (Sullivan et al., 2001). Histology studies have shown a brain-wide age-related decline in the number and length of myelinated fibres (Tang et al., 1997; Marner et al., 2003). Research has shown that thin, unmyelinated fibres in frontal brain regions are most susceptible to loss (Bartzokis, 2004, 2011). Some researchers have consider that these declines are manifestations of a 'disconnection syndrome' (O'Sullivan et al., 2001) which increases in severity with age. It may be that while the cortical connections of the brain remain largely intact until very old age, the overall integrity of the white matter connections deteriorates gradually from early adulthood. It is thought that this decline begins from about 20 years of age (Pfefferbaum and Sullivan, 2003), is linear and primarily impacts frontal regions. Some studies have reported that the decline is equivalent in both men and women (Sullivan et al., 2006). Longitudinal age-related decline in white matter FA has also been reported (Barrick et al., 2010). In addition, research has shown both significant negative correlation between FA and age, and a significant positive correlation between $\langle \mathrm{D} \rangle$ and age, in certain white matter structures, such as the corpus callosum genu (Bastin et al., 2008, 2010) and splenium (Pfefferbaum and Sullivan, 2003). These correlations have been found equally in both men and women. This inverse FA-diffusivity relationship is suggestive of an age-related decline of white matter integrity, possibly caused by degradation of myelin and axonal membranes and accompanied accumulation of intracellular and extracellular fluid.

Heterogeneous age-related changes in the corpus callosum (Sullivan et al., 2006; Bastin et al., 2008), the principal white matter structure for interhemispheric communication, indicate that appropriate network analysis may find age-related alterations in patterns of connectivity. Some researchers have found that brain atrophy and white matter lesions play a significant role in localised age-related changes in white matter tract shape and integrity (Bastin et al., 2010;

Maniega et al., 2014), indicating that network connectivity may be altered in such regions. Localising the age-related alterations in white matter and identifying their relationship to cognitive decline may advance our understanding of the structure-function relationship for both normal ageing and neurodegenerative disorders.

6.2.4 Networks in ageing

The changes in brain organisation due to both normal and abnormal ageing has been widely studied using network analysis and graph theory for both structural and functional MRI (Gong et al., 2009a; Robinson et al., 2010; De Boer et al., 2010; Lo et al., 2010; Wen et al., 2011; Fischer et al., 2014). Many of these studies identified marked age-related changes in connectivity, particularly in very old age. In a widely cited study, Gong et al. (2009a) investigated the effect of age and gender on the structural networks obtained from dMRI and tractography in a group of 95 subjects aged from 19 to 85 years. They identified the precuneus and posterior cingulate gyrus as highly connected hub nodes independent of age or gender. They found an overall reduction in cortical connectivity and network efficiency with increasing age. In contrast, Robinson et al. (2010) used a novel machine learning approach to identify differences in connections obtained from dMRI and tractography between two age groups (20-30 and 60-90 years of age). Particularly discriminative connections are found in the fibres between the left insula and left lingual gyrus, and between the left superior frontal gyrus and left medial orbital gyrus. In general, network studies consistently show age-related changes, both increasing and decreasing, in connections to frontal, temporal and medial brain structures. However, not all studies agree on the pathways involved.

6.3 Participants

80 healthy, right-handed, volunteers aged between 25 and 65 years were recruited by a poster campaign from staff working in the University of Edinburgh, Western General Hospital, and Royal Infirmary of Edinburgh, Scotland, UK. The study was designed to image 40 male and 40 females subjects, with uniformly distributed ages between 25 and 65 years. To provide normative data for this age range, subjects were recruited if they were native English speakers, were not on any long term medication with the exception of the contraceptive pill, had not been diagnosed with diabetes or high blood pressure, had not undergone previous cranial surgery, had alcohol consumption levels within the UK national safety guidelines (21 and 14 units per week for males and females), and were registered with a general practitioner (GP) in the UK. The study was approved by the Lothian Research Ethics Committees (REC 05/S1104/45) and all subjects gave written informed consent.

6.4 Network construction

All subjects underwent an MRI protocol, as described in Section 3.2. Following pre-processing and tractography, whole-brain networks were constructed by recording connections between all 85 ROIs (Section 3.7). The endpoint of a streamline was considered to be the first grey matter ROI encountered when tracking from the seed location. The *streamline density* weighting, recorded the interconnecting streamline density corrected for ROI size was computed (Equation 3.10). Note that while cortical grey matter volumes can be markedly different across an ageing cohort, the grey matter normalisation corrects for the bias in the resulting network measures due to between-subject variability in cortical volume. Any implausible streamlines traversing from one cortical hemisphere to any contralateral subcortical node were discarded. To reduce spurious connections, a two-step threshold on the network weights was then applied: 1) for each subject, discard the weakest 25% of weights in the matrix by connection probability (Equation 3.10); 2) across the cohort, only retain connections which occur in at least 50% of subjects. The second step is required to discard connections which have been removed for some subjects but not others by the first step. From the remaining set of streamlines *FA-weighted* networks were then computed by recording the mean FA value along interconnecting streamlines (Equation 2.7). For the context of ageing, the FA-weighting was applied because FA is believed to reflect some aspects of the microstructural integrity of the underlying axonal fibres.

For each FA-weighted matrix, seven network measures were then computed as defined in Section 3.8. This included five commonly used measures of connectivity: the *network degree* (average number of connections per node), the *network strength* (the average sum of weights per node), *network clustering coefficient* (an average measure of local connectivity), the *characteristic path length* (mean of all shortest path lengths) and the *network efficiency* (the average of the inverse shortest path length). A measure of centrality, the *betweenness centrality* was also calculated, which can be an indicator of hub nodes, which participate in many shortest paths through a network. In addition, motivated by the reported findings of recruitment and increased interhemispheric transfer in older subjects, we also included a measure of inter-module participation, the *interhemispheric participation coefficient*. First the network nodes were assigned to two modules by hemisphere and then the coefficient was calculated by Equation 3.29. Each of these seven measures were computed per node for each subject. Global network measures were computed by taking the mean of all nodal values. Lobar measures were computed by taking the mean of nodal values per lobe. The six lobar regions were frontal, temporal, parietal, occipital, cingulate, and subcortical. In addition, one further measure was computed for global networks only, the *network transitivity* (a normalised variant of the clustering coefficient).

6.5 Statistical analysis

Motivated by the reported age-related declines in grey matter volume, white matter integrity and network connectivity, a number of statistical analyses were used to assess the relationships between age and imaging properties for our subjects. For all tests, $p < 0.05$ was considered statistically significant.

6.5.1 Volumetric analysis

A volumetric analysis of the T_1-weighted tissue segmentations and the FA in cerebral white matter was performed. From the results of the FreeSurfer segmentation procedure, volumes were computed in MNI-305 space for white matter, cortical grey matter and subcortical grey matter, for each of the 80 subjects. Cerebellar regions were not included in the calculation and subcortical grey matter consisted of the regions defined in Appendix A, excluding the brain stem. Mean and median FA values in white matter were computed for each subject by: 1) aligning FreeSurfer segmentations to dMRI space, as described in Section 3.5; 2) computing the mean and median FA values from the FA voxels within the white matter mask. Each of these image properties was assessed for relationship with age using Pearson's correlation coefficient and a p-value was computed.

6.5.2 Tract-based spatial statistics

Tract-based spatial statistics was used to assess age and gender effects in the FA of major white matter pathways. Following the TBSS procedure (Smith et al., 2006), a nonlinear deformation was used to align the FA map of each subject to a white matter template in standard space. White matter masks were morphologically thinned in order to obtain the centre-line of the principal white matter pathways, while minimising the impact of registration error and partial volume effects. A one variable general linear model was used to assess the relationship between age and the FA within voxels of the white matter skeleton. This model was used to test for both an increase and decrease in FA with age. In addition a two variable model was used to test six hypotheses of age and gender effects for the same white matter voxels. Permutation testing assigned a corrected p-value to each voxel.

6.5.3 Network analysis

Following construction of FA-weighted networks we performed three levels of network analysis. Firstly, we assessed age-related change in nine network properties at a global network

level using Pearson's correlation coefficient. Secondly, we repeated the analysis at a lobar network level for seven network properties in all six lobes. Finally, we used NBS (Zalesky et al., 2010a), without *a priori* selected regions, to identify network connections showing a relationship with age or gender. NBS exploits the extent to which the connections identified by the contrast are interconnected to offer a potential gain in statistical power, for which the significance of maximally connected subnetworks are assessed rather than individual connections (Zalesky et al., 2010a). For this analysis only network connections that were present in all subjects were included. As with TBSS, a one variable general linear model was used to assess the relationship between age and the network connections, in terms of tract-averaged FA. This model was used to test for both an increase and decrease in connectivity with age. In addition a two variable model was used to test six hypotheses of age and gender effects within the whole-brain networks. Permutation testing, which randomly exchanged the group to which each subject belonged, was used over 5000 iterations to estimate the distribution of component size and compute a corrected p-value for the maximally connected subnetwork(s). In NBS terminology, the 'intensity' of each maximally connected network was tested rather than the 'extent' (size of the network) as this directly assesses the magnitude of the test statistic. As suggested by the authors of NBS, we repeated hypothesis testing over a range of t-statistic thresholds below the maximum t-statistic observed.

6.6 Results

Our methods were applied to the 80 healthy subjects (39 male). The subjects were approximately uniformly distributed in terms of age between 25 and 64 years. Overall, the mean age was 43.5 ± 10.6 years. The mean age for male subjects was 43.1 ± 10.0 years and the mean age for females subjects was 43.8 ± 11.1 years.

6.6.1 Volumetric analysis

Volumetric and structural imaging properties derived from the FreeSurfer T_1-weighted segmentations are shown in Table 6.1 for the grey matter regions, cerebral white matter and the mean and median FA within cerebral white matter. Table 6.2 shows the findings from correlational analysis of these properties with age. All grey matter regions, except from the subcortical regions in our female subjects, show a significant age-related decline in volume. Figure 6.1 shows scatter plots for all grey matter volumes ($p < 0.003$, uncorrected). However, age-related change in white matter volume was only found in our male subjects ($p = 0.041$, uncorrected). No age-related change in mean or median FA in cerebral white matter was evident.

Property	$N = 80$		Female ($N = 41$)		Male ($N = 39$)	
Grey matter volume (cm^3)	512.3	(57.5)	482.6	(40.7)	543.5	(56.4)
Cortical grey matter volume (cm^3)	454.6	(52.8)	428.3	(37.6)	482.3	(52.7)
Subcortical grey matter volume (cm^3)	57.7	(5.6)	54.3	(4.1)	61.2	(4.7)
White matter volume (cm^3)	499.2	(65.8)	459.9	(43.0)	540.5	(60.2)
Mean FA in white matter	0.323	(0.02)	0.323	(0.02)	0.324	(0.02)
Median FA in white matter	0.306	(0.02)	0.306	(0.02)	0.307	(0.02)

Table 6.1: Tissue volumes obtained from the FreeSurfer segmentation and the mean and median FA within cerebral white matter. Means are shown first with standard deviations in parenthesis.

Property	$N = 80$		Female ($N = 41$)		Male ($N = 39$)	
	r	p	r	p	r	p
Grey matter volume	**-0.423**	**<0.001**	**-0.459**	**0.003**	**-0.531**	**<0.001**
Cortical grey matter volume	**-0.422**	**<0.001**	**-0.468**	**0.003**	**-0.506**	**0.001**
Subcortical grey matter volume	**-0.369**	**0.001**	-0.267	0.101	**-0.633**	**<0.001**
White matter volume	-0.129	0.255	0.005	0.974	**-0.320**	**0.041**
Mean FA	0.012	0.915	0.046	0.779	-0.013	0.936
Median FA	-0.024	0.831	0.006	0.971	-0.046	0.773

Table 6.2: Correlations between volumetric properties and age assessed by Pearson's correlation coefficient with uncorrected p-values. Emboldening indicates significance at $p < 0.05$.

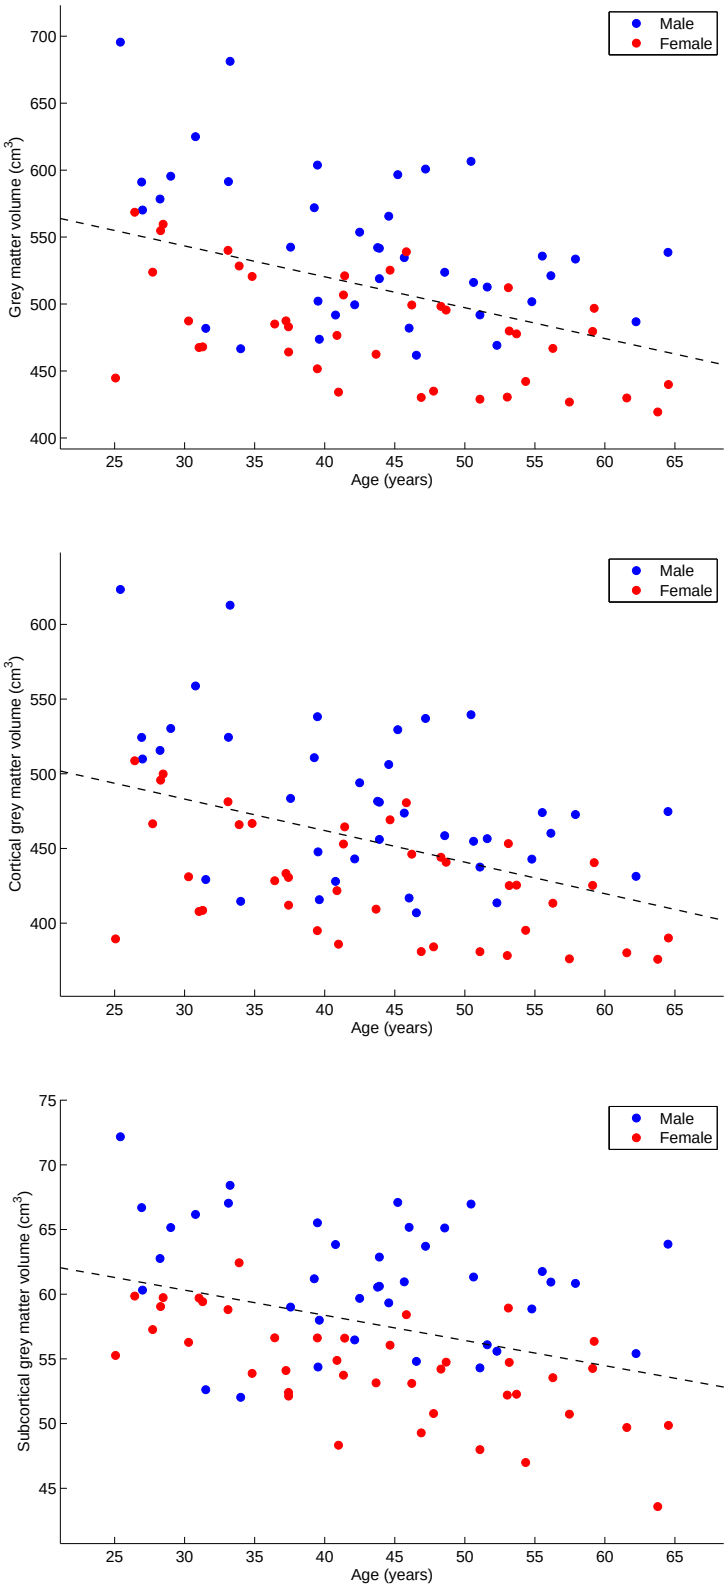

Figure 6.1: Scatter plots and linear fit of grey matter volumes against age for: total grey matter, cortical and subcortical volumes.

Hypothesis	p
H_1: FA increases with age	0.8525
H_1: FA decreases with age	0.8741

Table 6.3: Results of hypothesis testing on age using TBSS. The corrected p-values were obtained from permutation testing.

Hypothesis	p
H_1: FA increases with age, when controlled for gender	0.8849
H_1: FA decreases with age, when controlled for gender	0.8621
H_1: FA is greater in men, when controlled for age	0.7060
H_1: FA is greater in women, when controlled for age	0.7747
H_1: age has a greater effect on FA than gender	0.7688
H_1: gender has a greater effect on FA than age	0.7248

Table 6.4: Results of hypothesis testing based on gender and age using TBSS. The corrected p-values were obtained from permutation testing.

6.6.2 Tract-based spatial statistics

Table 6.3 shows the results of hypothesis testing for changes in mean FA with age for white matter voxels assessed by TBSS. In both cases, the null hypothesis must be accepted; there is no significant change in FA within these regions due to age. Table 6.4 shows the results of hypothesis testing for the effects of age and gender on FA with age for white matter voxels assessed by TBSS. In all cases, there is no significant change with age or gender.

6.6.3 Network analysis

Following connectivity mapping and constraints on the network weights, $14.52 \pm 2.27\%$ of the total streamlines seeded, were subsequently identified as interconnections across the whole cohort. The resulting mean network sparsity across all subjects was 0.794 ± 0.009. Visual assessment of the streamlines remaining following network thresholding indicated that the majority of streamlines were anatomically plausible. Figure 6.2(a) shows the mean connectivity matrix averaged across all subjects. Figure 6.2(b) shows the histograms of FA-weighted network connections with subjects divided into four distinct age groups between 25 and 65 years of age, indicating that there is very little age-related variation in global connectivity across the cohort.

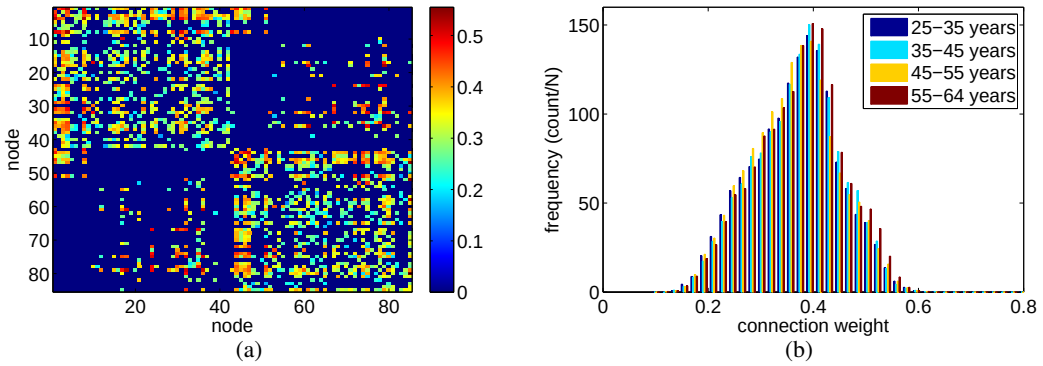

Figure 6.2: a) 85×85 connectivity matrix showing the mean weights across all subjects (N=80); b) The histograms of network weights (tract-averaged FA) with subjects placed in four distinct age groups from 25 to 65 years.

6.6.3.1 Global and lobar network analysis

Global network metrics were computed per subject, as described in Section 6.4. Table 6.5 shows the mean and standard deviations for the global measures. Table 6.6 shows the correlation coefficients between these properties and age. Just the mean betweenness centrality showed a decline in age for the male subjects ($p < 0.023$, uncorrected).

Lobar network measures were also assessed by correlation. Table 6.7 shows the mean and standard deviations for nodes within frontal and temporal lobes. Correlations were computed for all six lobes, but in the interests of space only the results for frontal and temporal regions are shown because these are arguably of the most interest in ageing. Table 6.8 shows the correlation coefficients between these lobar properties and age. After correction for multiple comparisons, none of the lobar network measures showed a relationship with age. However, some uncorrected statistics indicated a trend for age-related change in lobar network measures, for example, the betweenness centrality within the cingulate cortex (data not shown).

6.6.3.2 Network-based statistics

Table 6.9 shows the results of hypothesis testing for age-related effects in network connectivity. A number of thresholds on the t-statistic were tested for each hypothesis, but only the network (and t-statistic) corresponding to the smallest p-value was reported. In both cases, no significant age-related change in connectivity, in terms of tract-averaged FA, were found in any network components. However, the test for a decrease in FA with age comes closest to significance ($p = 0.0648$, corrected). Table 6.10 shows the results of hypothesis testing for the effects of both age and gender on connectivity. In all six cases there is no significant change

Network property	$N = 80$		Female ($N = 41$)		Male ($N = 39$)	
network degree	17.552	(0.79)	17.798	(0.68)	17.294	(0.84)
network strength	6.607	(0.51)	6.712	(0.48)	6.495	(0.52)
characteristic path length	1.178	(0.05)	1.169	(0.05)	1.188	(0.05)
network clustering coefficient	0.225	(0.01)	0.226	(0.01)	0.223	(0.01)
network efficiency	0.225	(0.01)	0.227	(0.01)	0.223	(0.01)
network transitivity	0.198	(0.01)	0.199	(0.01)	0.197	(0.01)
mean betweenness centrality	89.599	(3.83)	88.559	(3.30)	90.693	(4.08)
mean interhemispheric participation	0.216	(0.02)	0.220	(0.02)	0.212	(0.02)

Table 6.5: Global network metrics in normal ageing. Means are shown first with standard deviations in parenthesis.

Network property	$N = 80$		Female ($N = 41$)		Male ($N = 39$)	
	r	p	r	p	r	p
network degree	0.181	0.108	0.179	0.262	0.193	0.239
network strength	0.039	0.734	-0.024	0.884	0.115	0.486
characteristic path length	-0.018	0.873	0.030	0.854	-0.074	0.653
network clustering coefficient	0.016	0.885	-0.062	0.699	0.118	0.475
network efficiency	0.016	0.887	-0.030	0.851	0.074	0.656
network transitivity	0.011	0.925	-0.047	0.769	0.088	0.595
mean betweenness centrality	-0.194	0.084	-0.034	0.834	**-0.364**	**0.023**
mean interhemispheric participation	0.037	0.742	0.062	0.698	0.004	0.981

Table 6.6: Correlation between the global network properties and age using Pearson's correlation coefficient. Emboldened numbers indicate significance at $p < 0.05$ level (uncorrected).

in connectivity with age or gender. However, the test for a decrease in FA with age when controlled for gender was closest to significance ($p = 0.0548$, corrected). Although this test was not significant, a post-hoc analysis was used to identify the network connections involved. There are 2 connections in this subnetwork (left temporal pole - left fusiform and left temporal pole - left inferior temporal).

6.7 Discussion

The findings from the analysis of tissue volume (Table 6.2, Figure 6.1) are consistent with previous findings in brain ageing. The decline in grey matter volume with increasing age is one of the most robust findings in neuroimaging (Moseley, 2002; Sullivan and Pfefferbaum, 2006). As with previous studies, we found that grey matter degeneration was most pronounced in the cortex, but was also apparent in subcortical grey matter (although not significant in our female

Network property	$N = 80$		Female ($N = 41$)		Male ($N = 39$)	
frontal node degree	15.361	(0.83)	15.539	(0.83)	15.175	(0.80)
frontal node strength	5.570	(0.51)	5.639	(0.49)	5.498	(0.54)
frontal path length	1.215	(0.05)	1.208	(0.05)	1.222	(0.06)
frontal clustering coefficient	0.218	(0.01)	0.220	(0.01)	0.217	(0.02)
frontal betweenness centrality	96.090	(11.52)	94.659	(9.98)	97.594	(12.89)
frontal interhemispheric participation	0.255	(0.03)	0.258	(0.02)	0.252	(0.03)
temporal node degree	14.335	(0.78)	14.665	(0.61)	13.987	(0.79)
temporal node strength	5.019	(0.43)	5.154	(0.39)	4.876	(0.43)
temporal path length	1.303	(0.06)	1.287	(0.05)	1.319	(0.06)
temporal clustering coefficient	0.251	(0.01)	0.253	(0.01)	0.250	(0.01)
temporal betweenness centrality	36.332	(6.03)	35.696	(5.78)	37.000	(6.30)
temporal interhemispheric participation	0.050	(0.02)	0.055	(0.02)	0.045	(0.02)

Table 6.7: Frontal and temporal network metrics in normal ageing. Means are shown first with standard deviations in parenthesis.

subjects). Although, our male subjects showed a decline in white matter volume with age, this change did not survive a correction for multiple comparisons. We note that more sophisticated analyses of structure could be undertaken using techniques, such as voxel-based morphometry (Ashburner and Friston, 2000). However, as structural imaging was not the main focus of our study, this was not pursued.

The absence of a significant decline in white matter integrity in our analyses was surprising. We emphasise that no age-related change in FA was found in total cerebral white matter (Table 6.6), or by the TBSS voxel-based analysis, or by the NBS network analysis, even when using gender as a covariate. Therefore, we conclude that for this cohort there is no detectable change in white matter integrity due to age. This was unexpected when considering that age-related declines in FA, particularly in frontal brain areas, have been identified consistently by dMRI studies of normal ageing (Sullivan and Pfefferbaum, 2006). Previous dMRI studies of network connectivity in normal ageing have shown age-related changes, especially in very old age (Gong et al., 2009a; Robinson et al., 2010; De Boer et al., 2010).

However, the graph-theoretic network analysis did show an age-related decline in the *mean betweenness centrality* for the male subjects (Table 6.2, uncorrected). Betweenness centrality is a measure of node centrality, which measures the proportion of all shortest (weighted) paths in the network that traverse through a given node. Nodes with high betweenness centrality participate in a large number of shortest paths and researchers have made use of it to identify hub nodes within a network (Gong et al., 2009a; Yan et al., 2011). A decline in the mean betweenness centrality could indicate that the overall influence of hub nodes declines with age. Notably, Zhu et al. (2012) found that the betweenness centrality of certain nodes changed sig-

Network property	N=80		Female (N=41)		Male (N=39)	
	r	p	r	p	r	p
frontal node degree	0.188	0.095	0.106	0.511	0.270	0.097
frontal node strength	0.084	0.457	-0.021	0.894	0.195	0.234
frontal path length	-0.005	0.967	0.083	0.605	-0.102	0.538
frontal clustering coefficient	0.019	0.867	-0.121	0.451	0.157	0.340
frontal betweenness centrality	-0.152	0.179	-0.051	0.749	-0.255	0.117
frontal interhemispheric participation	-0.101	0.373	-0.149	0.352	-0.053	0.747
temporal node degree	0.112	0.323	0.023	0.884	0.176	0.283
temporal node strength	-0.069	0.540	-0.104	0.518	-0.024	0.886
temporal path length	0.000	0.997	0.050	0.754	-0.056	0.735
temporal clustering coefficient	-0.042	0.710	-0.088	0.585	0.032	0.846
temporal betweenness centrality	0.154	0.172	0.124	0.441	0.182	0.267
temporal interhemispheric participation	0.157	0.164	0.048	0.766	0.228	0.162

Table 6.8: Correlations between the mean lobar network properties and age using Pearson's correlation coefficient for frontal and temporal lobes.

Hypothesis	p	t-stat
H_1: network-weighting increases with age	0.2388	0.2
H_1: network-weighting decreases with age	0.0648	0.4

Table 6.9: Results of hypothesis testing on age using NBS. In each case, the p and t-stat shown correspond to the most significant network obtained by NBS for the range of t-stat thresholds tested.

Hypothesis	p	t-stat
H_1: network-weighting increases with age, when controlled for gender	0.2302	0.2
H_1: network-weighting decreases with age, when controlled for gender	0.0548	0.4
H_1: network-weighting is greater in men, when controlled for age	0.1356	0.2
H_1: network-weighting is greater in women, when controlled for age	0.2002	0.2
H_1: age has a greater effect on network-weighting than gender	0.2560	0.2
H_1: gender has a greater effect on network-weighting than age	0.1364	0.2

Table 6.10: Results of hypothesis testing based on gender and age using NBS. In each case, the p and t-stat shown correspond to the most significant network obtained by NBS for the range of t-stat thresholds tested.

nificantly with age in a structural inter-regional correlation study (Zhu et al., 2012). However, bearing in mind that for our study the statistic for mean betweenness centrality does not survive a correction for multiple comparisons, and is not seen at all in the female subjects, we should not draw any clear conclusions about node centrality from our data.

It is worth noting the differences between the TBSS and NBS approaches. Note that TBBS is performed independently of tractography and network analysis, and therefore is not affected by problems associated with network construction. TBSS can be considered a conservative technique as only certain primary white matter areas are assessed. Arguably, the difference in p-values between the TBSS and NBS analysis may indicate that NBS at the network level is more sensitive (or perhaps speculative) to alterations in connectivity between cortical areas than TBSS.

There are a number of reasons our method failed to reveal any expected changes due to age. Our cohort does not include subjects older than 64 years of age. However, only after about 70 years of age do studies show pronounced declines in white matter volume and integrity (Pfefferbaum et al., 2005; Salat et al., 2005; Sullivan et al., 2006). In contrast to neurodegenerative disease or brain injury, the effects of normal ageing are subtle and occur gradually. Crucially, there is no evidence of severed axonal fibres occurring due to normal ageing (Sullivan and Pfefferbaum, 2006). Age-related white matter changes are thought to be small, gradual and typically without lesions for normal ageing. As a result, the locus of cognitive decline may be challenging to detect with conventional MRI analysis. Also, it is possible that our current techniques in dMRI, tractography or network construction result in a signal-to-noise ratio below that required to observe the expected alterations in normal ageing.

6.8 Conclusions

This chapter presented an analysis of brain structure, white matter integrity and network connectivity in a group of volunteers undergoing normal ageing. Whole-brain networks were constructed using a validated approach with strong constraints on the anatomical plausibility of tracts. While a significant age-related decline in grey matter volume was found from structural MRI, consistent with the literature of brain ageing, there were no meaningful brain-wide impairments due to ageing found by a voxel-based analysis of white matter integrity or at any level of network analysis. However, network analysis when corrected for gender did find a trend level decline in white matter connectivity between temporal nodes in the left hemisphere, but this was not found to be significant. The lack of evidence for white matter changes due to ageing is unexpected. However, it is possible that our current techniques in network construction are too limited to observe the relatively subtle changes due to normal ageing.

One conclusion that we can draw from these results is that for healthy subjects younger than 65 years of age, cortical volume is a more sensitive indicator of age-related changes in the brain than any of the measures of white matter integrity that was assessed. However, for subjects older than 70 years of age we would expect to observe measurable changes in white matter connectivity, especially in frontal and temporal regions (Pfefferbaum et al., 2005; Salat et al., 2005) and at the network level as reported in the literature (Gong et al., 2009a; Robinson et al., 2010; De Boer et al., 2010). Ultimately, improvements in the signal-to-noise of network approaches may increase the power of analyses assessing network properties and cognitive impairment. These analyses could be key to understanding the structure-function relationship in both normal and abnormal ageing.

Chapter 7

Discussion and conclusions

7.1 Overview

This chapter presents a discussion of the findings from the preceding results chapters. We focus on the relative advantages and limitations of a network approach based on the described techniques and cohorts, and we attempt to place in context our findings from the evaluation of these methods. This is followed by a discussion of the current challenges facing connectome approaches and suggestions for future work. Finally, we conclude with some final remarks.

7.2 Summary of findings

We evaluated several key steps and representations in the construction of structural networks using repeat scans of healthy volunteers. To our knowledge, this work was the first to directly compare probabilistic and deterministic tractography and alternative seeding strategies in a test-retest network analysis. Although a small sample was used in the test-retest study, our findings showed that global network measures were estimated reliably, but there were concerns about the reliability of nodal measures. We obtained broadly similar results to previous test-retest analyses (Hagmann et al., 2008; Vaessen et al., 2010; Bassett et al., 2010; Cammoun et al., 2011; Cheng et al., 2012a). Although reliability has been assessed in various ways, these studies indicate that, globally, the between-subject variation is typically greater than the within-subject variation. To our knowledge, the test-retest study (Buchanan et al., 2014a) was amongst the first to quantify the reliability of nodal network properties. Our findings indicated that connections to some nodes were computed unreliably from session to session. For some network connections, the measurement noise could be greater than the genuine signal. Notably, a more recent study made similar findings on the reliability of nodal measures (Andreotti et al.,

2014).

The comparison of tractography algorithms, seeding approaches and network weightings offered some insights into appropriate methods of network construction. In all the cases assessed, white matter seeding produced networks with better test-retest reliability than grey matter seeding, in terms of the mean ICC measuring node strength. We believe some grey matter seeding error may arise because all grey matter voxels were included, even those unlikely to interface with white matter, potentially resulting in a greater proportion of spurious streamlines. Overall, the probabilistic tractography algorithm produced better test-retest performance than the deterministic method. This was perhaps due to the limited sample of possible streamlines produced by the deterministic method. Deterministic methods are known to be somewhat simplistic, leading to inherent tracking artefacts (Wiegell et al., 2000; Tuch et al., 2002). In addition, we determined that in the not uncommon case a streamline passes through multiple ROIs before terminating, a definition of endpoint connectivity which considered the first ROI encountered was the most reliable, when compared to absolute endpoints or complete connectivity between ROIs along the streamline. Streamline termination is a well-known weakness of current tractography techniques (Jbabdi and Johansen-Berg, 2011) and there is no means to prevent some streamlines from either ending prematurely or extending past the true endpoint. Note that the above findings were made with a small sample and further analysis may be required to test these points.

Following evaluation of network methods, we applied network analysis to an ALS cohort in a case-control study and to a cross-sectional study of normal ageing. To our knowledge, this work was the first to utilise probabilistic tractography and model multiple fibre populations in a network analysis of ALS. The key findings from the network analysis of ALS were that while there were no brain-wide impairments in connectivity due to ALS, there was an impaired motor-frontal-subcortical network of reduced white matter integrity in the ALS patients. The connections involved in the network analysis had a strong correspondence with tracts identified by a conventional TBSS analysis of the same data. Reduced white matter integrity in three of the impaired network connections, which involve fibres of the cortico-spinal tract, correlated with the rate of disease progression. However, no convincing white matter impairments were found in several frontal areas believed to be associated with FTD in ALS. It could be that the limitations of the data and methods mean that the resulting network analysis is not sensitive enough to detect the structural changes thought to be associated with FTD and cognitive impairment.

In addition, it is worth noting that a preliminary case-control analysis was also performed in a cohort of Multiple Sclerosis (MS) patients using the same set of network methods (findings not published). These analyses showed widespread global changes in white matter connec-

tivity, broadly compatible with demyelination and white matter impairment reported in MS (Kutzelnigg et al., 2005). However, localising these alterations in network connectivity using NBS proved to be problematic. In the case of such widespread white matter changes, network analysis may not be especially appropriate.

Concerning analyses of normal ageing, no age-related change in FA was found in total cerebral white matter, or by a TBSS voxel-based analysis, or by a NBS network analysis. Undoubtedly, the effects of ageing on white matter are more subtle than changes associated with neurodegenerative disorders such as ALS or MS. We note that using NBS, the t-statistics observed in the analysis of ageing (Table 6.9) are small and far below those observed in the ALS analysis (Table 5.3). Indeed, age-related white matter changes are thought to be small until the formation of white matter lesions in advanced age (Sullivan and Pfefferbaum, 2006). It is possible that current techniques in dMRI, tractography or network construction result in a signal-to-noise ratio below that required to observe the expected age-related alterations in brain organisation. However, it was evident that the FA-weighted connectivity matrices and histograms of network weights were near identical for both the ALS cohort (Figure 5.1) and the ageing cohort (Figure 6.2). Note that the cohorts span a similar age range. This suggests that, on average, global patterns of structural connectivity are remarkably similar between subjects.

An important consideration is the type of network weighting. We chose to use FA-weighted networks for both the ALS and ageing cohorts. Although FA-weighting provided marginally poorer test-retest reliability than the SD-weighting at the global and nodal network level, there are concerns about the ROI volume correction on SD-weighting, which may unduly bias a group contrast. Some researchers have suggested that ROI volume correction may overcompensate for volume-driven effects on streamline counts (Van Den Heuvel and Sporns, 2011). Such weightings may be affected by between-subject differences in tissue volume rather than streamline density. It is not clear how such effects may be counteracted to allow representative comparison of connectivity between individuals. We believe that FA is a more appropriate measure of between-subject connectivity as it somewhat circumvents the need for correction by ROI size and it is thought to reflect aspects of the underlying integrity of white matter.

In the test-retest analysis, we sought to understand network reliability without applying arbitrary thresholds on network weights. However, given the proportion of false streamlines observed, we found it necessary to apply stronger constraints for the ALS and ageing cohorts. These constraints focussed on thresholding of network weights. Thresholding resulted in discarding the lowest 25% of network weights per subject and only retaining connections which occurred in the majority of subjects.

7.3 Strengths and limitations

Network analysis has the potential to localise pathways of connectivity at a macroscopic resolution. In comparison to conventional voxel-based or ROI-based analyses, many thousands of connections can be identified with high-throughput network methods. However, the application of network resolution was not especially revealing for our analyses, perhaps with the exception of the connection between left paracentral and left superior frontal gyrus identified by the ALS NBS analysis (Figure 5.2), which was not identified by TBSS. It could be argued that there is little justification for using a complex set of network methods, with many sources of potential error, over simpler and well-defined techniques such as TBSS, if no further information on localisation can be obtained. On the other hand, we believe that as dMRI acquisitions and image processing techniques improve, network analysis will permit higher resolution probing of connectivity than by spatially limited approaches, such as TBSS. Additionally, in comparison to TBSS or other template-space analyses, network analysis is computed in native space rather than standard space, thereby accounting for individual differences in white matter structure and providing a more representative reconstruction of the underlying axonal wiring.

The most concerning aspect of the described network methods is the apparent lack of reliability and repeatability in computing certain regional network properties. Some test-retest variation may be due to scanner noise and inhomogeneities between sessions and some may be due to systematic variation in processing. Evidently, many intermediate but fundamental steps are involved in generating structural networks from dMRI data. The variability at each step contributes to the errors in the following stage and in the resulting measures of connectivity. As discussed in Chapter 4, network measures are dependent on reliable registration, node segmentation, diffusion processing and tractography.

As with many neuroimaging analyses, network approaches rely on accurate registration of image volumes. Accurate alignment of neuroanatomical nodes with diffusion data is an open research problem (Klein et al., 2009). We observed that registration error of even one voxel can disrupt the correct mapping of fibres. Nonlinear (deformable) transforms improve local alignment (over linear transforms) particularly in cortical regions, though several regions remain susceptible to errors in alignment.

EPI suffers from inherent geometric distortions from magnetic field inhomogeneities due to the rapid sampling of the gradient echo train. There are some principled methods to reduce these distortions, at least in part (Jones et al., 2013). For certain applications these distortions may be tolerable. However, we believe these distortions are a problem for whole-brain quantitative studies which must rely in mapping diffusion data to specific anatomical ROIs. In this work, it was apparent that such distortions can adversely affect the proper registration between dMRI

and T_1-weighted data.

Concerning thresholding of network weights, it is problematic to ensure that an arbitrary threshold removes spurious connections while retaining genuine patterns of connectivity. Implausible streamlines like those illustrated in Figure 4.18, highlight serious concerns about the tractography configuration used. In this particular case false streamlines were eliminated by discarding any bilateral cortico-subcortical connections. However, there were many other network connections for which similarly questionable streamlines were observed. To our knowledge, there is no means to entirely reject implausible streamlines because of the lack of complete ground-truth data for white matter connectivity. Thresholding has previously been reported to have a dramatic effect on the amount of false positive and false negative connections and the resulting measures of connectivity (van Wijk et al., 2010; de Reus and van den Heuvel, 2013). As a result, any application of current network methods must be used with caution and should account for a proportion of spurious streamlines remaining, even after network thresholding. We believe that the two-step thresholding procedure used in this study eliminates a proportion of implausible connections without biasing the results of a group-wise analysis. However, due to the limitations of current dMRI and tractography techniques it is not possible to eliminate all false connections. This highlights the need for even more stringent constraints and priors on dMRI data and tractography.

NBS has proved useful in mass univariate testing of network connections by exploiting the connectedness of network components to improve statistical power. Unfortunately, as noted by the authors of NBS, network analysis requires selecting an arbitrary threshold on the t-statistics which then affects the extent of any observed network (Zalesky et al., 2010a). However, Smith and Nichols (2009) suggest an alternative threshold-free approach for a similar mass univariate test, which merits further investigation in the network analysis setting.

Another crucial aspect of analysis is the sample size. One strength of this study is the relatively large and well-characterised patient samples for both the ALS and normal ageing cohorts. In contrast, we have noted the limitation of the small sample size used in the test-retest analysis. In some cases, this has made it difficult to draw clear conclusions about the effect of the different steps used in network construction. Our work (Buchanan et al., 2014a) and other recent research (Liem et al., 2014) point towards the benefit of using large cohorts.

7.4 Challenges and Future Work

As our findings indicate, there are many avenues for future improvements in connectome technologies. For instance, improvements in neuroanatomical segmentation, registration, fibre ori-

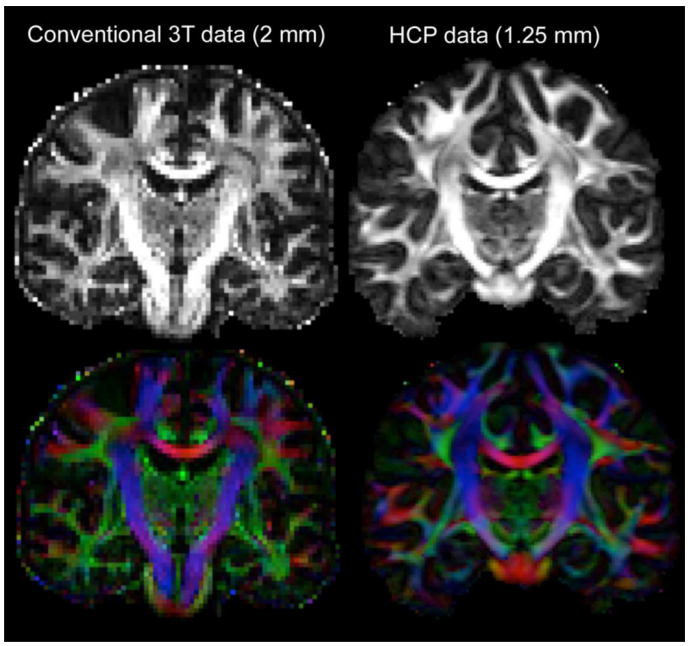

Figure 7.1: Coronal slices of FA and colour-encoded principal diffusion direction images from the Human Connectome Project dMRI data, compared with a conventional 2 mm data set of a different subject. (Adapted from Van Essen et al. 2013.)

entation estimation, tractography and post-processing methods will undoubtedly lead to reduced artefacts and improve the signal-to-noise of the resulting network measures.

MRI acquisition, field strength, spatial resolution and angular resolution are crucial factors. Many studies have demonstrated useful network connectivity gathered at 3 T and several studies have shown useful results with appropriate acquisition at 1.5 T. However, our findings suggest that methods at 1.5 T have limitations in detecting population differences in case-control or cross-sectional studies. We believe that high-resolution dMRI acquisitions could be hugely beneficial to future connectome analyses. For example, the state-of-the-art acquisitions and dMRI processing techniques developed within the Human Connectome Project (Van Essen et al., 2013), as illustrated in Figure 7.1, may lead to more convincing macroscale connectomes.

Until recently most network studies have been acquired with DTI, despite its known limitations. There has been a gradual shift towards more complex acquisitions, such as HARDI and DSI, for which the fibre crossing problem becomes more tractable. Vaessen et al. (2010) assessed reproducibility over different sets of diffusion gradient directions using DTI. Bassett et al. (2010) compared reliability in both DTI and DSI. Zalesky et al. (2010b) investigated the effect of network resolution using DTI and high-angular resolution. Cammoun et al. (2011) investigated the effect of network resolution using DSI. These studies indicate that networks can be estimated from various dMRI data, but that richer and higher-resolution acquisitions seem more in agreement with the underlying white matter connectivity. Over time, higher

resolution imaging will permit characterisation of cerebral structures at sub-millimetre scales. However, all acquisitions have inherent limits and reconstruction errors are inescapable. Researchers have shown that, in general, it is preferable to use isotropic voxels to minimise partial voluming effects or any bias arising from using a through-plan dimension much smaller than the in-plane voxel size (Jones et al., 2002).

An appropriate definition of network nodes is essential but far from trivial (Hagmann et al., 2010a; Zalesky et al., 2010b). Zalesky et al. (2010b) showed that the estimates of various organisational parameters such as clustering, path length, and efficiency were consistent across different parcellation scales at the same resolution, i.e. the same number of nodes. However, these parameters vary considerably as a function of spatial scale. Their findings indicate that any comparison of network parameters across studies must be made with reference to the spatial scale of the parcellation. For current network methods affected by noise, a compromise must be made between network resolution and an acceptable level of measurement noise. If issues with the reliability of nodal measures and individual network connections cannot easily be overcome with richer clinical data then low-resolution lobar networks, in the manner of Cabeen et al. (2013), may be less susceptible to the errors of individual nodes. In addition, it is not clear if anatomical divisions of the cortex, such as the sulcal and gyral boundaries used by FreeSurfer, or functional subdivisions are more appropriate for the the connectome mapping task (Hagmann et al., 2010a).

Methods in tractography are critical to understanding structural connectivity and the techniques have advanced steadily since the conception of tractography. We demonstrated that probabilistic tractography and a two fibre model performed better than the deterministic method used, in terms of the mean ICC for node strength. We believe this is because probabilistic methods should be able to account for branching and crossing fibres, in comparison to deterministic tractography. However, due to limitations of dMRI data, it is not possible to eliminate multi-fibre problems altogether, as illustrated in Figure 4.18. It remains unclear how to select the optimal fibre reconstruction method. Fillard et al. (2011) suggested that for medium or low signal-to-noise datasets, an appropriate prior on the spatial smoothness of either the diffusion model or the fibres is recommended for correct modelling. This merits further investigation.

There is a need to simplify and standardise the gathering, storage and analysis of connectome data. Tools such as the *brain connectivity toolbox* (Rubinov and Sporns, 2010), NBS (Zalesky et al., 2010a) and the Connectome Mapping Toolkit (Gerhard et al., 2011) are beginning to provide researchers with the necessary tools for analysing and visualising networks. Ultimately, as network methods become accepted and integrated this may facilitate answering fundamental neuroscientific questions relating to healthy and pathological brain organisation. Developments from other network sciences, such as biology and social science, may also provide inspiration.

As mentioned in Chapter 2, meaningful visualisation of networks remains an open problem. There is a need to visualise thousands of network connections coherently without losing essential information.

Another important, but largely missing, aspect of connectome mapping is a robust set of tools to validate the resulting networks and prevent the generation of unreliable data, i.e. in the case that no data is better than erroneous data. As with all neuroimaging analyses, it is crucial to inspect data at intermediate stages and if necessary reject the outcome. Typically, checks of data are performed manually, but there is an incentive to develop a set of automated and objective measures in order to achieve this. There is a pressing need to validate tractography results against the underlying axonal ground-truths, as some researchers have previously addressed in part (Dyrby et al., 2007; Hagmann et al., 2008). A full and proper characterisation of this anatomical knowledge may lead to more appropriate tractography methods and constraints. Ultimately, with sufficient prior knowledge, any network found to deviate from population norms could be treated with caution and possibly rejected from subsequent analysis.

7.5 Concluding remarks

This thesis presented methods for constructing white matter structural networks from MRI data. We sought to determine whether genuine patterns of connectivity can be identified despite noisy measurements. Our key findings were: 1) Probabilistic tractography was found to perform better than a deterministic method; 2) Performance was improved when seeding from white matter, rather than grey. 3) Thresholding of network weights must be applied with caution in order to remove spurious connections while retaining genuine patterns of connectivity. Our findings suggest that current connectome mapping techniques (at 1.5 T) are adequate for reliably measuring global network measures. However, regional network measures may not be as reliable, leading to concerns about the validity of studies based on such measures particularly with small sample sizes. We recommend that future dMRI network studies use large samples of ideally more than 50 subjects. Whilst current network methods are capable of characterising the genuine between-subject differences in connectivity, it is challenging to measure subtle white matter changes, for example, due to normal ageing.

We conclude that there are many future challenges facing connectome mapping technologies and future work should be undertaken to address these concerns. Future improvements to acquisitions and methods are likely to achieve network resolutions of many thousands of nodes, marking a shift from the macroscopic to the microscopic scale. Ultimately, these developments are likely to be a benefit to the understanding of the structure-function relationship and the basis of neural connectivity.

Appendix A

List of grey matter structures

A.1 Sub-cortical regions

List of sub-cortical regions (all entries except brain stem have left/right hemisphere counterparts):

Region	Abbreviation
Brain stem	BS
Thalamus	Tha
Caudate	Cau
Putamen	Put
Pallidum	Pal
Hippocampus	Hip
Amygdala	Amy
Accumbens-area	Acc
Ventral-diencephalon	VDC

A.2 Desikan-Killiany cortical regions

List of 34 cortical regions (all entries have left/right hemisphere counterparts):

Region	Abbreviation	Lobe
Banks of superior temporal sulcus	BST	temporal
Caudal anterior cingulate	CAC	cingulate
Caudal middle frontal	CMF	frontal
Cuneus	Cun	occipital
Entorhinal	Ent	temporal
Fusiform	Fus	temporal
Inferior parietal	IP	parietal
Inferior temporal	IT	temporal
Isthmus cingulate	IC	cingulate
Lateral occipital	LA	occipital
Lateral orbitofrontal	LAF	frontal
Lingual	Lin	occipital
Medial orbitofrontal	MOF	frontal
Middle temporal	MT	temporal
Parahippocampal	PH	temporal
Paracentral	PaC	frontal
Pars opercularis	POp	frontal
Pars orbitalis	POr	frontal
Pars triangularis	PT	frontal
Pericalcarine	Per	occipital
Postcentral	PCe	parietal
Posterior cingulate	PCi	cingulate
Precentral	PrC	frontal
Precuneus	Pre	parietal
Rostral anterior cingulate	RAC	cingulate
Rostral middle frontal	RMF	frontal
Superior frontal	SF	frontal
Superior parietal	SP	parietal
Superior temporal	ST	temporal
Supramarginal	SM	parietal
Frontal pole	FP	frontal
Temporal pole	TP	temporal
Transverse temporal	TT	temporal
Insula	Ins	-

Bibliography

Abe, O., Yamada, H., Masutani, Y., Aoki, S., Kunimatsu, A., Yamasue, H., Fukuda, R., Kasai, K., Hayashi, N., Masumoto, T., Mori, H., Soma, T., and Ohtomo, K. (2004). Amyotrophic lateral sclerosis: diffusion tensor tractography and voxel-based analysis. *NMR Biomed.*, 17:411–416.

Abrahams, S., Goldstein, L. H., Suckling, J., Ng, V., Simmons, A., Chitnis, X., Atkins, L., Williams, S. C. R., and Leigh, P. N. (2005). Frontotemporal white matter changes in amyotrophic lateral sclerosis. *Journal of neurology*, 252:321–331.

Aganj, I., Lenglet, C., Jahanshad, N., Yacoub, E., Harel, N., Thompson, P. M., and Sapiro, G. (2011). A Hough transform global probabilistic approach to multiple-subject diffusion MRI tractography. *Medical Image Analysis*, 15(4):414–425.

Agosta, F., Pagani, E., Petrolini, M., Caputo, D., Perini, M., Prelle, A., Salvi, F., and Filippi, M. (2010). Assessment of white matter tract damage in patients with amyotrophic lateral sclerosis: a diffusion tensor MR imaging tractography study. *American Journal of Neuroradiology*, 31(8):1457–1461.

Alemán-Gómez, Y., Melie-García, L., and Valdés-Hernandez, P. (2005). IBASPM: Toolbox for automatic parcellation of brain structures. In *Presented at the 12th Annual Meeting of the Organization for Human Brain Mapping*, pages 11–15.

Alexander, D. C. (2008). A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features. *Magnetic Resonance in Medicine*, 60:439–448.

Andersson, J. L. R., Jenkinson, M., and Smith, S. (2007). Non-linear registration aka Spatial normalisation. Technical Report TR07JA2, Oxford Centre for Functional Magnetic Resonance Imaging of the Brain, University of Oxford.

Andreotti, J., Jann, K., Melie-Garcia, L., Giezendanner, S., Dierks, T., and Federspiel, A. (2014). Repeatability Analysis of Global and Local Metrics of Brain Structural Networks. *Brain connectivity*, 4:203–20.

Araque, A., Parpura, V., Sanzgiri, R. P., and Haydon, P. G. (1999). Tripartite synapses: Glia, the unacknowledged partner. *Trends in Neurosciences*, 22(5):208–215.

Arfanakis, K., Cordes, D., Haughton, V. M., Carew, J. D., Elizabeth Meyerand, M., and Meyerand, M. E. (2002). Independent component analysis applied to diffusion tensor MRI. *Magnetic Resonance in Medicine*, 47(2):354–363.

Ashburner, J. and Friston, K. J. (2000). Voxel-based morphometry–the methods. *NeuroImage*, 11:805–821.

Bai, F., Shu, N., Yuan, Y., Shi, Y., Yu, H., Wu, D., Wang, J., Xia, M., He, Y., and Zhang, Z. (2012). Topologically convergent and divergent structural connectivity patterns between patients with remitted geriatric depression and amnestic mild cognitive impairment. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 32:4307–18.

Barabasi, A.-L. and Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439):11.

Barrick, T. R., Charlton, R. A., Clark, C. A., and Markus, H. S. (2010). White matter structural decline in normal ageing: A prospective longitudinal study using tract-based spatial statistics. *NeuroImage*, 51:565–577.

Bartzokis, G. (2004). Age-related myelin breakdown: A developmental model of cognitive decline and Alzheimer's disease. *Neurobiology of Aging*, 25(1):5–18.

Bartzokis, G. (2011). Alzheimer's disease as homeostatic responses to age-related myelin breakdown. *Neurobiology of Aging*, 32(8):1341–1371.

Basser, P. J., Mattiello, J., and LeBihan, D. (1994). Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of magnetic resonance. Series B*, 103:247–254.

Basser, P. J. and Pajevic, S. (2000). Statistical artifacts in diffusion tensor MRI (DT-MRI) caused by background noise. *Magnetic Resonance in Medicine*, 44(1):41–50.

Basser, P. J., Pajevic, S., Pierpaoli, C., Duda, J., and Aldroubi, A. (2000). In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44(4):625–32.

Basser, P. J. and Pierpaoli, C. (1996). Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of magnetic resonance Series B*, 111(3):209–219.

Bassett, D. S., Brown, J. A., Deshpande, V., Carlson, J. M., and Grafton, S. T. (2010). Conserved and variable architecture of human white matter connectivity. *NeuroImage*, 54(2):1262–1279.

Bassett, D. S. and Gazzaniga, M. S. (2011). Understanding complexity in the human brain. *Trends in Cognitive Sciences*, 15(5):200–209.

Bastin, M., Piatkowksi, J., Storkey, A., Brown, L. J., Maclullich, A. M., and Clayden, J. D. (2008). Tract shape modelling provides evidence of topological change in corpus callosum genu during normal ageing. *NeuroImage*, 43(1):20–28.

Bastin, M. E., Munoz Maniega, S., Ferguson, K. J., Brown, L. J., Wardlaw, J. M., MacLullich, A. M., and Clayden, J. D. (2010). Quantifying the effects of normal ageing on white matter structure using unsupervised tract shape modelling. *NeuroImage*, 51(1):1–10.

Bastin, M. E., Pettit, L. D., Bak, T. H., Gillingwater, T. H., Smith, C., and Abrahams, S. (2013). Quantitative tractography and tract shape modeling in amyotrophic lateral sclerosis. *Journal of magnetic resonance imaging : JMRI*, 38:1140–1145.

Bear, M. F., Connors, B. W., and Paradiso, M. A. (2007). *Neuroscience*, volume 2. Lippincott Williams & Wilkins.

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

Behrens, T. E. J., Johansen-Berg, H., Woolrich, M. W., Smith, S. M., Wheeler-Kingshott, C. A. M., Boulby, P. A., Barker, G. J., Sillery, E. L., Sheehan, K., Ciccarelli, O., Thompson, A. J., Brady, J. M., and Matthews, P. M. (2003a). Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. *Nature Neuroscience*, 6(7):750–757.

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

Blain, C. R. V., Brunton, S., Williams, V. C., Leemans, A., Turner, M. R., Andersen, P. M., Catani, M., Stanton, B. R., Ganesalingham, J., Jones, D. K., Williams, S. C. R., Leigh, P. N., and Simmons, A. (2011). Differential corticospinal tract degeneration in homozygous 'D90A' SOD-1 ALS and sporadic ALS. *Journal of neurology, neurosurgery, and psychiatry*, 82:843–849.

Bonilha, L., Nesland, T., Rorden, C., and Fridriksson, J. (2014). Asymmetry of the structural brain connectome in healthy older adults. *Frontiers in Psychiatry*, 4(JAN).

Brooks, B. R., Miller, R. G., Swash, M., and Munsat, T. L. (2000). El Escorial revisited: revised criteria for the diagnosis of amyotrophic lateral sclerosis. In *Amyotrophic lateral sclerosis and other motor neuron disorders : official publication of the World Federation of Neurology, Research Group on Motor Neuron Diseases*, volume 1, pages 293–299.

Bruijn, L. I., Miller, T. M., and Cleveland, D. W. (2004). Unraveling the mechanisms involved in motor neuron degeneration in ALS. *Annual review of neuroscience*, 27:723–749.

Buchanan, C. R., Pernet, C. R., Gorgolewski, K. J., Storkey, A. J., and Bastin, M. E. (2014a). Test-retest reliability of structural brain networks from diffusion MRI. *NeuroImage*, 86:231–243.

Buchanan, C. R., Pettit, L. D., Storkey, A. J., Abrahams, S., and Bastin, M. E. (2014b). Reduced structural connectivity within a prefrontal-motor-subcortical network in amyotrophic lateral sclerosis. *Journal of Magnetic Resonance Imaging*.

Buckner, R. L. (2004). Memory and executive function in aging and ad: Multiple factors that cause decline and reserve factors that compensate. *Neuron*, 44:195–208.

Bullitt, E., Gerig, G., Pizer, S. M., Lin, W., and Aylward, S. R. (2003). Measuring tortuosity of the intracerebral vasculature from MRA images. *IEEE transactions on medical imaging*, 22:1163–1171.

Busbice, T., Gleeson, P., Khayrulin, S., Cantarelli, M., Dibert, A., Idili, G., Palyanov, A., and Larson, S. (2013). The NeuroML C. elegans Connectome. *Frontiers in Neuroinformatics*.

Bystron, I., Blakemore, C., and Rakic, P. (2008). Development of the human cerebral cortex: Boulder Committee revisited. *Nature reviews. Neuroscience*, 9:110–122.

Cabeen, R. P., Andreyeva, K., Bastin, M. E., and Laidlaw, D. H. (2013). A Diffusion MRI resource of 80 age-varied subjects with neuropsychological and demographic measures. In *Proceedings 21st Scientific Meeting International Society for Magnetic Resonance in Medicine*, page 2138.

Cabeza, R., Anderson, N. D., Locantore, J. K., and McIntosh, A. R. (2002). Aging gracefully: compensatory brain activity in high-performing older adults. *NeuroImage*, 17:1394–1402.

Calamante, F., Tournier, J. D., Jackson, G. D., and Connelly, A. (2010). Track-density imaging (TDI): Super-resolution white matter imaging using whole-brain track-density mapping. *NeuroImage*, 53:1233–1243.

Cammoun, L., Gigandet, X., Meskaldji, D., Thiran, J. P., Sporns, O., Do, K. Q., Maeder, P., Meuli, R., and Hagmann, P. (2011). Mapping the human connectome at multiple scales with diffusion spectrum MRI. *Journal of Neuroscience Methods*, 6(2007):1–12.

Catani, M., Thiebaut de Schotten, M., Slater, D., and Dell'Acqua, F. (2013). Connectomic approaches before the connectome. *NeuroImage*, 80:2–13.

Cedarbaum, J. M., Stambler, N., Malta, E., Fuller, C., Hilt, D., Thurmond, B., and Nakanishi, A. (1999). The ALSFRS-R: A revised ALS functional rating scale that incorporates assessments of respiratory function. *Journal of the Neurological Sciences*, 169:13–21.

Cheng, H., Wang, Y., Sheng, J., Kronenberger, W. G., Mathews, V. P., Hummer, T. A., and Saykin, A. J. (2012a). Characteristics and variability of structural networks derived from diffusion tensor imaging. *NeuroImage*, 61:1153–64.

Cheng, H., Wang, Y., Sheng, J., Sporns, O., Kronenberger, W. G., Mathews, V. P., Hummer, T. A., and Saykin, A. J. (2012b). Optimization of seed density in DTI tractography for structural networks. *Journal of Neuroscience Methods*, 203:264–272.

Chung, M. K., Adluru, N., Dalton, K. M., Alexander, A. L., and Davidson, R. J. (2011). Scalable Brain Network Construction on White Matter Fibers. *Society*, pages 1–6.

Ciccarelli, O., Behrens, T. E., Altmann, D. R., Orrell, R. W., Howard, R. S., Johansen-Berg, H., Miller, D. H., Matthews, P. M., and Thompson, A. J. (2006). Probabilistic diffusion tractography: a potential tool to assess the rate of disease progression in amyotrophic lateral sclerosis. *Brain : a journal of neurology*, 129:1859–1871.

Cirillo, M., Esposito, F., Tedeschi, G., Caiazzo, G., Sagnelli, A., Piccirillo, G., Conforti, R., Tortora, F., Monsurrò, M. R., Cirillo, S., and Trojsi, F. (2012). Widespread Microstructural White Matter Involvement in Amyotrophic Lateral Sclerosis: A Whole-Brain DTI Study. *American Journal Of Neuroradiology*, pages 1–7.

Clayden, J. D., Storkey, A. J., and Bastin, M. E. (2007). A probabilistic model-based approach to consistent white matter tract segmentation. *IEEE Transactions on Medical Imaging*, 26(11):1555–1561.

Conturo, T. E., Lori, N. F., Cull, T. S., Akbudak, E., Snyder, A. Z., Shimony, J. S., McKinstry, R. C., Burton, H., and Raichle, M. E. (1999). Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences of the United States of America*, 96(18):10422–10427.

Cook, P. A., Bai, Y., Nedjati-Gilani, S., Seunarine, K. K., Hall, M. G., Parker, G. J., and Alexander, D. C. (2006). Camino: Open-source diffusion-MRI reconstruction and processing. In *14th Scientific Meeting of the International Society for Magnetic Resonance in Medicine*, page 2759.

Côté, M. A., Girard, G., Boré, A., Garyfallidis, E., Houde, J. C., and Descoteaux, M. (2013). Tractometer: Towards validation of tractography pipelines. *Medical Image Analysis*, 17:844–857.

Crick, F. and Jones, E. (1993). Backwardness of human neuroanatomy. *Nature*, 361(6408):109–110.

Crofts, J. J., Higham, D. J., Bosnell, R., Jbabdi, S., Matthews, P. M., Behrens, T. E. J., and Johansen-Berg, H. (2010). Network analysis detects changes in the contralesional hemisphere following stroke. *NeuroImage*, 54(1):161–169.

De Boer, R., Schaap, M., Van Der Lijn, F., Vrooman, H. A., De Groot, M., Vernooij, M. W., Ikram, M. A., Van Velsen, E. F. S., Van Der Lugt, A., Breteler, M. M. B., and Niessen, W. J. (2010). Statistical analysis of structural brain connectivity. *Medical Image Computing and Computer-Assisted Intervention*, 13(Pt 2):101–108.

de Reus, M. A. and van den Heuvel, M. P. (2013). Estimating false positives and negatives in brain networks. *NeuroImage*, 70:402–409.

DeCarli, C., Massaro, J., Harvey, D., Hald, J., Tullberg, M., Au, R., Beiser, A., D'Agostino, R., and Wolf, P. A. (2005). Measures of brain morphology and infarction in the framingham heart study: Establishing what is normal. *Neurobiology of Aging*, 26:491–510.

Desikan, R. S., Ségonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., Buckner, R. L., Dale, A. M., Maguire, R. P., Hyman, B. T., Albert, M. S., and Killiany, R. J. (2006). An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage*, 31(3):968–980.

Destrieux, C., Fischl, B., Dale, A., and Halgren, E. (2010). Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. *Neuroimage*, 53(1):1–15.

Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271.

Douaud, G., Filippini, N., Knight, S., Talbot, K., and Turner, M. R. (2011). Integration of structural and functional magnetic resonance imaging in amyotrophic lateral sclerosis. *Brain : a journal of neurology*, 134:3470–3479.

Dyrby, T. B., Sø gaard, L. V., Parker, G. J., Alexander, D. C., Lind, N. M., Baaré, W. F. C., Hay-Schmidt, A., Eriksen, N., Pakkenberg, B., Paulson, O. B., and Jelsing, J. (2007). Validation of in vitro probabilistic tractography. *NeuroImage*, 37(4):1267–1277.

Einstein, A. (1905). On the movement of small particles suspended in a stationary liquid demanded by the molecular-kinetic theory of heat. [English translation]. *Annalen der Physik*, 17:549–560.

Ellis, C. M., Simmons, A., Jones, D. K., Bland, J., Dawson, J. M., Horsfield, M. A., Williams, S. C., and Leigh, P. N. (1999). Diffusion tensor MRI assesses corticospinal tract damage in ALS. *Neurology*, 53:1051–1058.

Ellis, C. M., Suckling, J., Amaro, E., Bullmore, E. T., Simmons, A., Williams, S. C., and Leigh, P. N. (2001). Volumetric analysis reveals corticospinal tract degeneration and extramotor involvement in ALS. *Neurology*, 57:1571–1578.

Embleton, K. V., Morris, D. M., Haroon, H. A., Ralph, M. A. L., and Parker, G. J. (2007). Anatomical Connectivity Mapping. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, volume 15.

Filippini, N., Douaud, G., Mackay, C. E., Knight, S., Talbot, K., and Turner, M. R. (2010). Corpus callosum involvement is a consistent feature of amyotrophic lateral sclerosis. *Neurology*, 75:1645–1652.

Fillard, P., Descoteaux, M., Goh, A., Gouttard, S., Jeurissen, B., Malcolm, J., Ramirez-Manzanares, A., Reisert, M., Sakaie, K., Tensaouti, F., Yo, T., Mangin, J.-F., and Poupon, C. (2011). Quantitative evaluation of 10 tractography algorithms on a realistic diffusion MR phantom. *NeuroImage*, 56(1):220–234.

Fischer, F. U., Wolf, D., Scheurich, A., and Fellgiebel, A. (2014). Association of structural global brain network properties with intelligence in normal aging. *PLoS ONE*, 9.

Fischi-Gómez, E., Vasung, L., Meskaldji, D.-E., Lazeyras, F., Borradori-Tolsa, C., Hagmann, P., Barisnikov, K., Thiran, J.-P., and Hüppi, P. S. (2014). Structural Brain Connectivity in School-Age Preterm Infants Provides Evidence for Impaired Networks Relevant for Higher Order Cognitive Skills and Social Cognition. *Cerebral cortex (New York, N.Y. : 1991)*.

Fischl, B., Salat, D. H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., Kouwe, A. V. D., Killiany, R., Kennedy, D., Klaveness, S., Montillo, A., Makris, N., Rosen, B., and Dale, A. M. (2002). Whole Brain Segmentation: Neurotechnique Automated Labeling of Neuroanatomical Structures in the Human Brain. *Neuron*, 33(3):341–355.

Fischl, B., Salat, D. H., Van Der Kouwe, A. J. W., Makris, N., Ségonne, F., Quinn, B. T., and Dale, A. M. (2004a). Sequence-independent segmentation of magnetic resonance images. *NeuroImage*, 23 Suppl 1:S69–84.

Fischl, B., Van Der Kouwe, A., Destrieux, C., Halgren, E., Ségonne, F., Salat, D. H., Busa, E., Seidman, L. J., Goldstein, J., Kennedy, D., Caviness, V., Makris, N., Rosen, B., and Dale, A. M. (2004b). Automatically parcellating the human cerebral cortex. *Cerebral Cortex*, 14(1):11–22.

Freeman, L. C. (1979). Centrality in social networks conceptual clarification. *Social networks*, 1(3):215–239.

Friman, O., Farnebäck, G., and Westin, C.-F. (2006). A Bayesian approach for stochastic white matter tractography. *IEEE transactions on medical imaging*, 25:965–978.

Funnell, M. G., Corballis, P. M., and Gazzaniga, M. S. (2000). Cortical and subcortical interhemispheric interactions following partial and complete callosotomy. *Archives of neurology*, 57(2):185–189.

Gazzaniga, M. S. (1998). The split brain revisited. *Scientific American*, 279:50–55.

Gerhard, S., Daducci, A., Lemkaddem, A., Meuli, R., Thiran, J.-P., and Hagmann, P. (2011). The Connectome Viewer Toolkit: An Open Source Framework to Manage, Analyze, and Visualize Connectomes. *Frontiers in neuroinformatics*, 5(June):15.

Geser, F., Brandmeir, N. J., Kwong, L. K., Martinez-Lage, M., Elman, L., McCluskey, L., Xie, S. X., Lee, V. M.-Y., and Trojanowski, J. Q. (2008). Evidence of multisystem disorder in whole-brain map of pathological TDP-43 in amyotrophic lateral sclerosis. *Archives of neurology*, 65:636–641.

Geser, F., Martinez-Lage, M., Kwong, L. K., Lee, V. M.-Y., and Trojanowski, J. Q. (2009). Amyotrophic lateral sclerosis, frontotemporal dementia and beyond: the TDP-43 diseases. *Journal of neurology*, 256:1205–1214.

Gigandet, X., Hagmann, P., Kurant, M., Cammoun, L., Meuli, R., and Thiran, J.-P. (2008). Estimating the Confidence Level of White Matter Connections Obtained with MRI Tractography. *PLoS ONE*, 3(12):9.

Gong, G., He, Y., Concha, L., Lebel, C., Gross, D. W., Evans, A. C., and Beaulieu, C. (2009a). Mapping Anatomical Connectivity Patterns of Human Cerebral Cortex Using In Vivo Diffusion Tensor Imaging Tractography. *Cerebral Cortex*, 19(3):524–536.

Gong, G., Rosa-Neto, P., Carbonell, F., Chen, Z. J., He, Y., and Evans, A. C. (2009b). Age- and gender-related differences in the cortical anatomical network. *Journal of Neuroscience*, 29(50):15684–15693.

Good, C. D., Johnsrude, I. S., Ashburner, J., Henson, R. N., Friston, K. J., and Frackowiak, R. S. (2001). A voxel-based morphometric study of ageing in 465 normal adult human brains. *NeuroImage*, 14:21–36.

Gorgolewski, K. J., Storkey, A., Bastin, M. E., Whittle, I. R., Wardlaw, J. M., and Pernet, C. R. (2013). A test-retest fMRI dataset for motor, language and spatial attention functions. *GigaScience*, 2:6.

Greenwood, P. M. (2000). The frontal aging hypothesis evaluated. *Journal of the International Neuropsychological Society*, 6(06):705–726.

Greicius, M. D., Supekar, K., Menon, V., and Dougherty, R. F. (2009). Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cerebral Cortex*, 19(1):72–8.

Greve, D. N. and Fischl, B. (2009). Accurate and robust brain image alignment using boundary-based registration. *NeuroImage*, 48(1):63–72.

Gronenschild, E. H. B. M., Habets, P., Jacobs, H. I. L., Mengelers, R., Rozendaal, N., Van Os, J., and Marcelis, M. (2012). The Effects of FreeSurfer Version, Workstation Type, and Macintosh Operating System Version on Anatomical Volume and Cortical Thickness Measurements. *PLoS ONE*, 7(6):e38234.

Guimerà, R. and Amaral, L. A. N. (2005). Cartography of complex networks: modules and universal roles. *Journal of statistical mechanics Online*, 2005(P02001):nihpa35573.

Hagmann, P. (2005). *From diffusion MRI to brain connectomics*. PhD thesis, ITS Institut de traitement des signaux.

Hagmann, P., Cammoun, L., Gigandet, X., Gerhard, S., Ellen Grant, P., Wedeen, V., Meuli, R., Thiran, J. P., Honey, C. J., and Sporns, O. (2010a). MR connectomics: Principles and challenges. *J Neurosci Methods*, 194(1):34–45.

Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C. J., Wedeen, V. J., and Sporns, O. (2008). Mapping the Structural Core of Human Cerebral Cortex. *PLoS Biology*, 6(7):15.

Hagmann, P., Kurant, M., Gigandet, X., Thiran, P., Wedeen, V. J., Meuli, R., and Thiran, J.-P. (2007). Mapping Human Whole-Brain Structural Networks with Diffusion MRI. *PLoS ONE*, 2(7):9.

Hagmann, P., Sporns, O., Madan, N., Cammoun, L., Pienaar, R., Wedeen, V. J., Meuli, R., Thiran, J. P., and Grant, P. E. (2010b). White matter maturation reshapes structural connectivity in the late developing human brain. *Proceedings of the National Academy of Sciences*, 107(44):1–6.

Hagmann, P., Thiran, J.-P., Jonasson, L., Vandergheynst, P., Clarke, S., Maeder, P., and Meuli, R. (2003). DTI mapping of human brain connectivity: statistical fibre tracking and virtual

dissection. Technical Report 3, Signal Processing Institute, Swiss Federal Institute of Technology, 1015 Lausanne, Switzerland. patric.hagmann@epfl.ch.

Han, X., Jovicich, J., Salat, D., Van Der Kouwe, A., Quinn, B., Czanner, S., Busa, E., Pacheco, J., Albert, M., Killiany, R., Maguire, P., Rosas, D., Makris, N., Dale, A., Dickerson, B., and Fischl, B. (2006). Reliability of MRI-derived measurements of human cerebral cortical thickness: the effects of field strength, scanner upgrade and manufacturer. *NeuroImage*, 32(1):180–194.

Heckemann, R. A., Hajnal, J. V., Aljabar, P., Rueckert, D., and Hammers, A. (2006). Automatic anatomical brain MRI segmentation combining label propagation and decision fusion. *NeuroImage*, 33(1):115–126.

Hedden, T. and Gabrieli, J. D. E. (2004). Insights into the ageing mind: a view from cognitive neuroscience. *Nature reviews. Neuroscience*, 5:87–96.

Herculano-Houzel, S. (2009). The human brain in numbers: a linearly scaled-up primate brain. *Frontiers in human neuroscience*, 3:31.

Honey, C., Sporns, O., Hagmann, P., Cammoun, L., Gigandet, X., and Meuli, R. (2008). The Structural Core of Human Cerebral Cortex and its relation to the brains default network. In *Proceedings 16th Scientific Meeting International Society for Magnetic Resonance in Medicine*, page 839.

Honey, C. J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J. P., Meuli, R., and Hagmann, P. (2009). Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences of the United States of America*, 106(6):2035–40.

Honey, C. J., Thivierge, J. P., and Sporns, O. (2010). Can structure predict function in the human brain? *NeuroImage*, 52:766–776.

Humphries, M. D. and Gurney, K. (2008). Network Small-World-Ness: A Quantitative Method for Determining Canonical Network Equivalence. *PLoS ONE*, 3(4):10.

Ingalhalikar, M., Smith, A., Parker, D., Satterthwaite, T. D., Elliott, M. a., Ruparel, K., Hakonarson, H., Gur, R. E., Gur, R. C., and Verma, R. (2014). Sex differences in the structural connectome of the human brain. *Proceedings of the National Academy of Sciences of the United States of America*, 111(2):823–8.

Irimia, A., Chambers, M. C., Torgerson, C. M., and Van Horn, J. D. (2012). Circular representation of human cortical networks for subject and population-level connectomic visualization. *NeuroImage*, 60:1340–1351.

Iturria-Medina, Y., Canales-Rodríguez, E. J., Melie-García, L., Valdés-Hernández, P. A., Martínez-Montes, E., Alemán-Gómez, Y., and Sánchez-Bornot, J. M. (2007). Characterizing brain anatomical connections using diffusion weighted MRI and graph theory. *NeuroImage*, 36(3):645–660.

Iturria-Medina, Y., Sotero, R. C., Canales-Rodríguez, E. J., Alemán-Gómez, Y., and Melie-García, L. (2008). Studying the human brain anatomical network via diffusion-weighted MRI and Graph Theory. *NeuroImage*, 40(3):1064–1076.

Jabr, F. (2012). The Connectome Debate: Is Mapping the Mind of a Worm Worth It? *Scientific American*.

Jahanshad, N., Aganj, I., Lenglet, C., Joshi, A., Jin, Y., Barysheva, M., McMahon, K. L., de Zubicaray, G. I., Martin, N. G., Wright, M. J., and Others (2011). Sex differences in the human connectome: 4-Tesla high angular resolution diffusion imaging (HARDI) tractography in 234 young adult twins. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pages 939–943. IEEE.

Jahanshad, N., Nir, T. M., Toga, A. W., Jack, C. R., Bernstein, M. a., Weiner, M. W., and Thompson, P. M. (2015). Seemingly unrelated regression empowers detection of network failure in dementia. *Neurobiology of aging*, 36 Suppl 1:S103–12.

Jbabdi, S. and Johansen-Berg, H. (2011). Tractography: where do we go from here? *Brain Connectivity*, 1(3):169–183.

Jenkinson, M., Bannister, P., Brady, M., and Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2):825–841.

Jenkinson, M. and Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, 5(2):143–156.

Jones, D. K. (2004). The Effect of Gradient Sampling Schemes on Measures Derived from Diffusion Tensor MRI: A Monte Carlo Study. *Magnetic Resonance in Medicine*, 51(4):807–815.

Jones, D. K. (2010). *Diffusion MRI: Theory, Methods, and Applications*. Oxford University Press, USA.

Jones, D. K. and Cercignani, M. (2010). Twenty-five pitfalls in the analysis of diffusion MRI data. *NMR in Biomedicine*, 23(7):803–820.

Jones, D. K., Knösche, T. R., and Turner, R. (2013). White matter integrity, fiber count, and other fallacies: The do's and don'ts of diffusion MRI.

Jones, D. K., Williams, S. C. R., Gasston, D., Horsfield, M. A., Simmons, A., and Howard, R. (2002). Isotropic resolution diffusion tensor imaging with whole brain acquisition in a clinically acceptable time. *Human brain mapping*, 15(4):216–230.

Jones, E. G. (1999). Colgi, Cajal and the Neuron Doctrine. *Journal of the history of the neurosciences*, 8:170–178.

Kassubek, J., Unrath, A., Huppertz, H.-J., Lulé, D., Ethofer, T., Sperfeld, A.-D., and Ludolph, A. C. (2005). Global brain atrophy and corticospinal tract alterations in ALS, as investigated by voxel-based morphometry of 3-D MRI. *Amyotrophic lateral sclerosis and other motor neuron disorders : official publication of the World Federation of Neurology, Research Group on Motor Neuron Diseases*, 6:213–220.

Klein, A., Andersson, J., Ardekani, B. A., Ashburner, J., Avants, B., Chiang, M. C. M., Christensen, G. E. G., Collins, D. L. D., Gee, J., Hellier, P., Song, J. H. J., Jenkinson, M., Lepage, C., Rueckert, D., Thompson, P., Vercauteren, T., Woods, R. R. P., Mann, J. J. J., and Parsey, R. R. V. (2009). Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. *NeuroImage*, 46(3):786–802.

Koh, D. M. and Collins, D. J. (2007). Diffusion-weighted MRI in the body: Applications and challenges in oncology. *American Journal of Roentgenology*, 188:1622–1635.

Korgaonkar, M. S., Fornito, A., Williams, L. M., and Grieve, S. M. (2014). Abnormal Structural Networks Characterize Major Depressive Disorder: A Connectome Analysis.

Kuceyeski, A., Maruta, J., Niogi, S. N., Ghajar, J., and Raj, A. (2011). The generation and validation of white matter connectivity importance maps. *NeuroImage*, 58(1):109–121.

Kutzelnigg, A., Lucchinetti, C. F., Stadelmann, C., Brück, W., Rauschka, H., Bergmann, M., Schmidbauer, M., Parisi, J. E., and Lassmann, H. (2005). Cortical demyelination and diffuse white matter injury in multiple sclerosis. *Brain*, 128:2705–2712.

Latora, V. and Marchiori, M. (2001). Efficient behavior of small-world networks. *Physical Review Letters*, 87(19):198701.

Lazar, M., Weinstein, D. M., Tsuruda, J. S., Hasan, K. M., Arfanakis, K., Meyerand, M. E., Badie, B., Rowley, H. A., Haughton, V., Field, A., and Alexander, A. L. (2003). White matter tractography using diffusion tensor deflection. *Human Brain Mapping*, 18(4):306–321.

Le Bihan, D., Breton, E., Lallemand, D., Grenier, P., Cabanis, E., and Laval-Jeantet, M. (1986). MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology*, 161(2):401–407.

Le Bihan, D. and Johansen-Berg, H. (2012). Diffusion MRI at 25: Exploring brain tissue structure and function. *NeuroImage*, 61:324–341.

Le Bihan, D., Mangin, J. F., Poupon, C., Clark, C. A., Pappata, S., Molko, N., and Chabriat, H. (2001). Diffusion tensor imaging: concepts and applications. *Journal of magnetic resonance imaging : JMRI*, 13:534–546.

Leemans, A., Jeurissen, B., Sijbers, J., and Jones, D. (2009). ExploreDTI: a graphical toolbox for processing, analyzing, and visualizing diffusion MR data. In *Proceedings 17th Scientific Meeting, International Society for Magnetic Resonance in Medicine*, volume 17, page 3537.

Li, L., Rilling, J. K., Preuss, T. M., Glasser, M. F., and Hu, X. (2011). The effects of connection reconstruction method on the interregional connectivity of brain networks via diffusion tractography. *Human Brain Mapping*, 00(January).

Li, Y., Liu, Y., Li, J., Qin, W., Li, K., Yu, C., and Jiang, T. (2009). Brain Anatomical Network and Intelligence. *PLoS Computational Biology*, 5(5):17.

Liem, F., Mérillat, S., Bezzola, L., Hirsiger, S., Philipp, M., Madhyastha, T., and Jäncke, L. (2014). Reliability and statistical power analysis of cortical and subcortical FreeSurfer metrics in a large sample of healthy elderly. *NeuroImage*, 108:95–109.

Lo, C.-Y., Wang, P.-N., Chou, K.-H., Wang, J., He, Y., and Lin, C.-P. (2010). Diffusion tensor tractography reveals abnormal topological organization in structural cortical networks in Alzheimer's disease. *Journal of Neuroscience*, 30(50):16876–16885.

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412:150–157.

Maldjian, J. A., Laurienti, P. J., Kraft, R. A., and Burdette, J. H. (2003). An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fMRI data sets. *NeuroImage*, 19(3):1233–1239.

Maniega, S. M. n., Hernández, M. C. V., Clayden, J. D., Royle, N. A., Murray, C., Morris, Z., Aribisala, B. S., Gow, A. J., Starr, J. M., Bastin, M. E., and Others (2014). White matter hy-

perintensities and normal-appearing white matter integrity in the aging brain. *Neurobiology of aging*, 36(2):909–918.

Marcus, D. S., Harwell, J., Olsen, T., Hodge, M., Glasser, M. F., Prior, F., Jenkinson, M., Laumann, T., Curtiss, S. W., and Van Essen, D. C. (2011). Informatics and Data Mining Tools and Strategies for the Human Connectome Project. *Frontiers in neuroinformatics*, 5(June):12.

Marner, L., Nyengaard, J. R., Tang, Y., and Pakkenberg, B. (2003). Marked loss of myelinated nerve fibers in the human brain with age. *The Journal of Comparative Neurology*, 462(2):144–152.

McRobbie, D. W., Moore, E. A., Graves, M. J., and Prince, M. R. (2006). *MRI from Picture to Proton*. Cambridge University Press.

Mohammadi, B., Kollewe, K., Samii, a., Krampfl, K., Dengler, R., and Munte, T. F. (2009). Changes of resting state brain networks in amyotrophic lateral sclerosis. *Experimental neurology*, 217:147–153.

Morey, R. A., Selgrade, E. S., Wagner, H. R., Huettel, S. A., Wang, L., and McCarthy, G. (2010). Scan-rescan reliability of subcortical brain volumes derived from automated segmentation. *Human brain mapping*, 31(11):1751–1762.

Mori, S., Crain, B. J., Chacko, V. P., and Van Zijl, P. C. (1999). Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology*, 45(2):265–269.

Mori, S. and Zhang, J. (2006). Principles of diffusion tensor imaging and its applications to basic neuroscience research. *Neuron*, 51:527–39.

Moseley, M. (2002). Diffusion tensor imaging and aging - a review. *NMR in Biomedicine*, 15(7-8):553–560.

Moseley, M. E., Cohen, Y., Kucharczyk, J., Mintorovitch, J., Asgari, H. S., Wendland, M. F., Tsuruda, J., and Norman, D. (1990). Diffusion-weighted MR imaging of anisotropic water diffusion in cat central nervous system. *Radiology*, 176:439–445.

Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM review*, pages 167–256.

Onnela, J. P., Saramäki, J., Kertész, J., and Kaski, K. (2005). Intensity and coherence of motifs in weighted complex networks. *Physical Review E*, 71(6):65103.

O'Sullivan, M., Jones, D. K., Summers, P. E., Morris, R. G., Williams, S. C. R., and Markus, H. S. (2001). Evidence for cortical "disconnection" as a mechanism of age-related cognitive decline. *Neurology*, 57(4):632.

Park, C. H., Kim, S. Y., Kim, Y. H., and Kim, K. (2008). Comparison of the small-world topology between anatomical and functional connectivity in the human brain. *Physica A*, 387(23):5958–5962.

Parker, C. S., Deligianni, F., Cardoso, M. J., Daga, P., Modat, M., Dayan, M., Clark, C. A., Ourselin, S., and Clayden, J. D. (2014). Consensus between Pipelines in Structural Brain Networks. *PloS one*, 9(10):e111262.

Parker, G. J. M., Haroon, H. A., and Wheeler-Kingshott, C. A. M. (2003). A framework for a

streamline-based probabilistic index of connectivity (PICo) using a structural interpretation of MRI diffusion measurements. *Journal of Magnetic Resonance Imaging*, 18(2):242–254.

Parker, G. J. M., Stephan, K. E., Barker, G. J., Rowe, J. B., MacManus, D. G., Wheeler-Kingshott, C. A. M., Ciccarelli, O., Passingham, R. E., Spinks, R. L., Lemon, R. N., and Turner, R. (2002). Initial demonstration of in vivo tracing of axonal projections in the macaque brain and comparison with the human brain using diffusion tensor Imaging and fast marching tractography. *NeuroImage*, 15(4):797–809.

Pettit, L. D., Bastin, M. E., Smith, C., Bak, T. H., Gillingwater, T. H., and Abrahams, S. (2013). Executive deficits, not processing speed relates to abnormalities in distinct prefrontal tracts in amyotrophic lateral sclerosis. *Brain : a journal of neurology*, 136:3290–3304.

Pfefferbaum, A., Adalsteinsson, E., and Sullivan, E. V. (2005). Frontal circuitry degradation marks healthy adult aging: Evidence from diffusion tensor imaging. *NeuroImage*, 26(3):891–899.

Pfefferbaum, A. and Sullivan, E. V. (2003). Increased brain white matter diffusivity in normal adult aging: relationship to anisotropy and partial voluming. *Magnetic Resonance in Medicine*, 49(5):953–961.

Pfefferbaum, A., Sullivan, E. V., Hedehus, M., Lim, K. O., Adalsteinsson, E., and Moseley, M. (2000). Age-related decline in brain white matter anisotropy measured with spatially corrected echo-planar diffusion tensor imaging. *Magnetic Resonance in Medicine*, 44:259–268.

Phukan, J., Elamin, M., Bede, P., Jordan, N., Gallagher, L., Byrne, S., Lynch, C., Pender, N., and Hardiman, O. (2012). The syndrome of cognitive impairment in amyotrophic lateral sclerosis: a population-based study. *Journal of neurology, neurosurgery, and psychiatry*, 83:102–108.

Poupon, C., Clark, C. A., Frouin, V., Régis, J., Bloch, I., Le Bihan, D., and Mangin, J. (2000). Regularization of diffusion-based direction maps for the tracking of brain white matter fascicles. *NeuroImage*, 12:184–195.

Purves, D., Augustine, G., Fitzpatrick, D., Hall, W., LaMantia, A., and White, L. (2012). *Neuroscience*. Sinauer Associates, Incorporated, fifth edition.

Rascovsky, K., Hodges, J. R., Knopman, D., Mendez, M. F., Kramer, J. H., Neuhaus, J., Van Swieten, J. C., Seelaar, H., Dopper, E. G. P., Onyike, C. U., Hillis, A. E., Josephs, K. A., Boeve, B. F., Kertesz, A., Seeley, W. W., Rankin, K. P., Johnson, J. K., Gorno-Tempini, M. L., Rosen, H., Prioleau-Latham, C. E., Lee, A., Kipps, C. M., Lillo, P., Piguet, O., Rohrer, J. D., Rossor, M. N., Warren, J. D., Fox, N. C., Galasko, D., Salmon, D. P., Black, S. E., Mesulam, M., Weintraub, S., Dickerson, B. C., Diehl-Schmid, J., Pasquier, F., Deramecourt, V., Lebert, F., Pijnenburg, Y., Chow, T. W., Manes, F., Grafman, J., Cappa, S. F., Freedman, M., Grossman, M., and Miller, B. L. (2011). Sensitivity of revised diagnostic criteria for the behavioural variant of frontotemporal dementia. *Brain*, 134(9):2456–2477.

Resnick, S. M., Pham, D. L., Kraut, M. A., Zonderman, A. B., and Davatzikos, C. (2003). Longitudinal magnetic resonance imaging studies of older adults: a shrinking brain. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 23:3295–3301.

Reuter, M., Schmansky, N. J., Rosas, H. D., and Fischl, B. (2012). Within-subject template estimation for unbiased longitudinal image analysis. *NeuroImage*, 61(4):1402–1418.

Robinson, E. C., Hammers, A., Ericsson, A., Edwards, A. D., and Rueckert, D. (2010). Identifying population differences in whole-brain structural networks: a machine learning approach. *NeuroImage*, 50(3):910–919.

Robinson, E. C., Valstar, M., Hammers, A., Ericsson, A., Edwards, A. D., and Rueckert, D. (2008). Multivariate statistical analysis of whole brain structural networks obtained using probabilistic tractography. *Medical Image Computing and Computer-Assisted Intervention*, 11(Pt 1):486–493.

Romero-Garcia, R., Atienza, M., Clemmensen, L. H., and Cantero, J. L. (2012). Effects of network resolution on topological properties of human neocortex. *NeuroImage*, 59(4):3522–3532.

Rose, S., Pannek, K., Bell, C., Baumann, F., Hutchinson, N., Coulthard, A., McCombe, P., and Henderson, R. (2012). Direct evidence of intra- and interhemispheric corticomotor network degeneration in amyotrophic lateral sclerosis: An automated MRI structural connectivity study. *NeuroImage*, 59:2661–2669.

Rowland, L. P. and Shneider, N. A. (2001). Amyotrophic lateral sclerosis. *New England Journal of Medicine*, 344(22):1688–1700.

Rubinov, M. and Sporns, O. (2010). Complex network measures of brain connectivity: uses and interpretations. *NeuroImage*, 52(3):1059–1069.

Salat, D. H., Tuch, D. S., Greve, D. N., van der Kouwe, A. J. W., Hevelone, N. D., Zaleta, A. K., Rosen, B. R., Fischl, B., Corkin, S., Rosas, H. D., and Others (2005). Age-related alterations in white matter microstructure measured by diffusion tensor imaging. *Neurobiology of Aging*, 26(8):1215–1227.

Salthouse, T. A., Atkinson, T. M., and Berish, D. E. (2003). Executive functioning as a potential mediator of age-related cognitive decline in normal adults. *Journal of experimental psychology. General*, 132:566–594.

Sarro, L., Agosta, F., Canu, E., Riva, N., Prelle, A., Copetti, M., Riccitelli, G., Comi, G., and Filippi, M. (2011). Cognitive Functions and White Matter Tract Damage in Amyotrophic Lateral Sclerosis: A Diffusion Tensor Tractography Study. *American Journal of Neuroradiology*, 32:1866–1872.

Sato, K., Aoki, S., Iwata, N. K., Masutani, Y., Watadani, T., Nakata, Y., Yoshida, M., Terao, Y., Abe, O., Ohtomo, K., and Others (2010). Diffusion tensor tract-specific analysis of the uncinate fasciculus in patients with amyotrophic lateral sclerosis. *Neuroradiology*, 52(8):729–733.

Scholz, J., Klein, M. C., Behrens, T. E. J., and Johansen-Berg, H. (2009). Training induces changes in white-matter architecture. *Nature Neuroscience*, 12(11):1370–1371.

Ségonne, F., Dale, A. M., Busa, E., Glessner, M., Salat, D., Hahn, H. K., and Fischl, B. (2004). A hybrid approach to the skull stripping problem in MRI. *NeuroImage*, 22(3):1060–1075.

Shattuck, D. W., Mirza, M., Adisetiyo, V., Hojatkashani, C., Salamon, G., Narr, K. L., Poldrack, R. A., Bilder, R. M., and Toga, A. W. (2008). Construction of a 3D probabilistic atlas of human cortical structures. *NeuroImage*, 39(3):1064–1080.

Shoukri, M. M., Asyali, M. H., and Donner, A. (2004). Sample size requirements for the design of reliability study: review and new results. *Statistical Methods in Medical Research*, 13:251–271.

Shrout, P. E. and Fleiss, J. L. (1979). Intraclass correlations: uses in assessing rater reliability. *Psychological Bulletin*, 86(2):420–428.

Shu, N., Liu, Y., Li, J., Li, Y., Yu, C., and Jiang, T. (2009). Altered Anatomical Network in Early Blindness Revealed by Diffusion Tensor Tractography. *PLoS ONE*, 4(9):13.

Shu, N., Liu, Y., Li, K., Duan, Y., Wang, J., Yu, C., Dong, H., Ye, J., and He, Y. (2011). Diffusion Tensor Tractography Reveals Disrupted Topological Efficiency in White Matter Structural Networks in Multiple Sclerosis. *Cerebral Cortex*, 21(November):2565–2577.

Skudlarski, P., Jagannathan, K., Anderson, K., Stevens, M. C., Calhoun, V. D., Skudlarska, B. A., and Pearlson, G. (2010). Brain connectivity is not only lower but different in schizophrenia: a combined anatomical and functional approach. *Biological Psychiatry*, 68(1):61–69.

Skudlarski, P., Jagannathan, K., Calhoun, V. D., Hampson, M., Skudlarska, B. A., and Pearlson, G. (2008). Measuring brain connectivity: diffusion tensor imaging validates resting state temporal correlations. *NeuroImage*, 43(3):554–561.

Sled, J. G., Zijdenbos, A. P., and Evans, A. C. (1998). A nonparametric method for automatic correction of intensity nonuniformity in MRI data. *IEEE Transactions on Medical Imaging*, 17(1):87–97.

Smith, S. M. (2002). Fast robust automated brain extraction. *Human Brain Mapping*, 17(3):143–155.

Smith, S. M., Jenkinson, M., and Johansen-Berg, H. (2006). Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. *Neuroimage*, 31:1487–1505.

Smith, S. M. and Nichols, T. E. (2009). Threshold-free cluster enhancement: Addressing problems of smoothing, threshold dependence and localisation in cluster inference. *NeuroImage*, 44:83–98.

Sporns, O. (2011a). *Networks of the Brain*. MIT Press.

Sporns, O. (2011b). The human connectome: a complex network. *Annals Of The New York Academy Of Sciences*, 1224(1):109–125.

Sporns, O., Tononi, G., and Kötter, R. (2005). The Human Connectome: A Structural Description of the Human Brain. *PLoS Computational Biology*, 1(4):e42.

Stehling, M. K., Turner, R., and Mansfield, P. (1991). Echo-planar imaging: magnetic resonance imaging in a fraction of a second. *Science*, 254:43–50.

Sullivan, E. V., Adalsteinsson, E., Hedehus, M., Ju, C., Moseley, M., Lim, K. O., and Pfefferbaum, A. (2001). Equivalent disruption of regional white matter microstructure in ageing healthy men and women. *Neuroreport*, 12:99–104.

Sullivan, E. V., Adalsteinsson, E., and Pfefferbaum, A. (2006). Selective age-related degradation of anterior callosal fiber bundles quantified in vivo with fiber tracking. *Cerebral Cortex*, 16(7):1030–9.

Sullivan, E. V. and Pfefferbaum, A. (2006). Diffusion tensor imaging and aging. *Brain*, 30(6):749–761.

Sullivan, E. V., Rosenbloom, M., Serventi, K. L., and Pfefferbaum, A. (2004). Effects of age and sex on volumes of the thalamus, pons, and cortex. *Neurobiology of Aging*, 25:185–192.

Tang, Y., Nyengaard, J. R., Pakkenberg, B., and Gundersen, H. J. (1997). Age-induced white matter changes in the human brain: a stereological investigation. *Neurobiology of aging*, 18:609–15.

Terry, R. D., DeTeresa, R., and Hansen, L. A. (1987). Neocortical cell counts in normal human adult aging. *Annals of neurology*, 21:530–539.

Tournier, J. D., Calamante, F., and Connelly, A. (2007). Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution. *NeuroImage*, 35:1459–1472.

Tournier, J. D., Yeh, C. H., Calamante, F., Cho, K. H., Connelly, A., and Lin, C. P. (2008). Resolving crossing fibres using constrained spherical deconvolution: Validation using diffusion-weighted imaging phantom data. *NeuroImage*, 42:617–625.

Tsermentseli, S., Leigh, P. N., and Goldstein, L. H. (2012). The anatomy of cognitive impairment in amyotrophic lateral sclerosis: More than frontal lobe dysfunction. *Cortex*, 48:166–182.

Tuch, D. S., Reese, T. G., Wiegell, M. R., Makris, N., Belliveau, J. W., and Wedeen, V. J. (2002). High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magnetic Resonance in Medicine*, 48(4):577–582.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., and Joliot, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage*, 15(1):273–289.

Vaessen, M. J., Hofman, P. A. M., Tijssen, H. N., Aldenkamp, A. P., Jansen, J. F. A., and Backes, W. H. (2010). The effect and reproducibility of different clinical DTI gradient sets on small world brain connectivity measures. *NeuroImage*, 51(3):1106–1116.

van den Heuvel, M. P. and Hulshoff Pol, H. E. (2010). Exploring the brain network: A review on resting-state fMRI functional connectivity.

Van Den Heuvel, M. P., Mandl, R. C. W., Kahn, R. S., and Hulshoff Pol, H. E. (2009). Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain. *Human Brain Mapping*, 30(10):3127–3141.

Van Den Heuvel, M. P. and Sporns, O. (2011). Rich-Club Organization of the Human Connectome. *Journal of Neuroscience*, 31(44):15775–15786.

Van Den Heuvel, M. P., Stam, C. J., Boersma, M., and Hulshoff Pol, H. E. (2008). Small-world and scale-free organization of voxel-based resting-state functional connectivity in the human brain. *NeuroImage*, 43(3):528–539.

van der Graaff, M. M., de Jong, J. M. B. V., Baas, F., and de Visser, M. (2009). Upper motor neuron and extra-motor neuron involvement in amyotrophic lateral sclerosis: A clinical and brain imaging review. *Neuromuscular Disorders*, 19:53–58.

van der Graaff, M. M., Sage, C. A., Caan, M. W. A., Akkerman, E. M., Lavini, C., Majoie, C. B., Nederveen, A. J., Zwinderman, A. H., Vos, F., Brugman, F., van den Berg, L. H., de Rijk, M. C., van Doorn, P. A., Van Hecke, W., Peeters, R. R., Robberecht, W., Sunaert, S., and de Visser, M. (2011). Upper and extra-motoneuron involvement in early motoneuron disease: a diffusion tensor imaging study. *Brain : a journal of neurology*, 134:1211–1228.

Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., and Ugurbil, K. (2013). The WU-Minn Human Connectome Project: An overview. *NeuroImage*, 80:62–79.

Van Essen, D. C. and Ugurbil, K. (2012). The future of the human connectome. *NeuroImage*, 62(2):1–12.

van Wijk, B. C. M., Stam, C. J., and Daffertshofer, A. (2010). Comparing brain networks of different size and connectivity density using graph theory. *PLoS ONE*, 5.

Várkuti, B., Cavusoglu, M., Kullik, A., Schiffler, B., Veit, R., Yilmaz, O., Rosenstiel, W., Braun, C., Uludag, K., Birbaumer, N., and Sitaram, R. (2011). Quantifying the Link between Anatomical Connectivity, Gray Matter Volume and Regional Cerebral Blood Flow: An Integrative MRI Study. *PLoS ONE*, 6(4):15.

Verstraete, E., van den Heuvel, M. P., Veldink, J. H., Blanken, N., Mandl, R. C., Pol, H. E. H., and van den Berg, L. H. (2010). Motor network degeneration in amyotrophic lateral sclerosis: A structural and functional connectivity study. *PLoS ONE*, 5.

Verstraete, E., Veldink, J. H., Mandl, R. C. W., Van Den Berg, L. H., and Van Den Heuvel, M. P. (2011). Impaired Structural Motor Connectome in Amyotrophic Lateral Sclerosis. *PLoS ONE*, 6(9):10.

Verstraete, E., Veldink, J. H., van den Berg, L. H., and van den Heuvel, M. P. (2013). Structural brain network imaging shows expanding disconnection of the motor system in amyotrophic lateral sclerosis. *Human brain mapping*, 00:1–11.

Vogelstein, J., Bogovic, J., Carass, A., and Gray, W. (2010). Graph-Theoretical Methods for Statistical Inference on MR Connectome Data. *Organization Human Brain Mapping*, pages 56307–56307.

Vogelstein, J. T., Gray, W. R., Vogelstein, R. J., and Priebe, C. E. (2011). Graph Classification using Signal Subgraphs: Applications in Statistical Connectomics. *Submitted to IEEE PAMI*, page 12.

Wang, R., Benner, T., Sorensen, a. G., and Wedeen, V. J. (2007). Diffusion Toolkit : A Software Package for Diffusion Imaging Data Processing and Tractography. *Proc. Intl. Soc. Mag. Reson. Med.*, 15:3720.

Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442.

Wedeen, V. J., Hagmann, P., Tseng, W.-Y. I., Reese, T. G., and Weisskoff, R. M. (2005). Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magnetic Resonance in Medicine*, 54(6):1377–1386.

Wedeen, V. J., Wang, R. P., Schmahmann, J. D., Benner, T., Tseng, W. Y. I., Dai, G., Pandya, D. N., Hagmann, P., D'Arceuil, H., and de Crespigny, A. J. (2008). Diffusion spectrum magnetic resonance imaging (DSI) tractography of crossing fibers. *NeuroImage*, 41:1267–1277.

Wee, C.-Y., Yap, P.-T., Li, W., Denny, K., Browndyke, J. N., Potter, G. G., Welsh-Bohmer, K. A., Wang, L., and Shen, D. (2011). Enriched white matter connectivity networks for accurate identification of MCI patients. *NeuroImage*, 54(3):1812–1822.

Wen, W., Zhu, W., He, Y., Kochan, N. A., Reppermund, S., Slavin, M. J., Brodaty, H., Craw-

ford, J., Xia, A., and Sachdev, P. (2011). Discrete neuroanatomical networks are associated with specific cognitive abilities in old age. *Journal of Neuroscience*, 31(4):1204–1212.

Whalley, L. J., Deary, I. J., Appleton, C. L., and Starr, J. M. (2004). Cognitive reserve and the neurobiology of cognitive aging. *Ageing Research Reviews*, 3:369–382.

White, J. G., Southgate, E., Thomson, J. N., and Brenner, S. (1986). The Structure of the Nervous System of the Nematode Caenorhabditis elegans. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 314:1–340.

Wiegell, M. R., Larsson, H. B., and Wedeen, V. J. (2000). Fiber crossing in human brain depicted with diffusion tensor MR imaging. *Radiology*, 217:897–903.

Wilcox, R. R. (2005). *Introduction to robust estimation and hypothesis testing*. Academic Press, second edition.

Wonderlick, J. S., Ziegler, D. A., Hosseini-Varnamkhasti, P., Locascio, J. J., Bakkour, A., Van Der Kouwe, A., Triantafyllou, C., Corkin, S., and Dickerson, B. C. (2009). Reliability of MRI-derived cortical and subcortical morphometric measures: effects of pulse sequence, voxel geometry, and parallel imaging. *NeuroImage*, 44(4):1324–1333.

Xue, R., van Zijl, P. C., Crain, B. J., Solaiyappan, M., and Mori, S. (1999). In vivo three-dimensional reconstruction of rat brain axonal projections by diffusion tensor imaging. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 42:1123–1127.

Yan, C., Gong, G., Wang, J., Wang, D., Liu, D., Zhu, C., Chen, Z. J., Evans, A., Zang, Y., and He, Y. (2011). Sex- and brain size-related small-world structural cortical networks in young adults: a DTI tractography study. *Cerebral Cortex*, 21(2):449–458.

Yo, T.-S., Anwander, A., Descoteaux, M., Fillard, P., Poupon, C., and Knösche, T. R. (2009). Quantifying brain connectivity: a comparative tractography study. *Medical Image Computing and Computer-Assisted Intervention*, 12(Pt 1):886–893.

Ystad, M., Hodneland, E., Adolfsdottir, S., Haász, J., Lundervold, A. J., Eichele, T., and Lundervold, A. (2011). Cortico-striatal connectivity and cognition in normal aging: a combined DTI and resting state fMRI study. *NeuroImage*, 55(1):24–31.

Zalesky, A. and Fornito, A. (2009). A DTI-derived measure of cortico-cortical connectivity. *IEEE Transactions on Medical Imaging*, 28(7):1023–1036.

Zalesky, A., Fornito, A., and Bullmore, E. T. (2010a). Network-based statistic: identifying differences in brain networks. *NeuroImage*, 53(4):1197–1207.

Zalesky, A., Fornito, A., Harding, I. H., Cocchi, L., Yücel, M., Pantelis, C., and Bullmore, E. T. (2010b). Whole-brain anatomical networks: does the choice of nodes matter? *NeuroImage*, 50(3):970–983.

Zalesky, A., Fornito, A., Seal, M. L., Cocchi, L., Westin, C.-F., Bullmore, E. T., Egan, G. F., and Pantelis, C. (2011). Disrupted axonal fiber connectivity in schizophrenia. *Biological Psychiatry*, 69(1):80–89.

Zhu, W., Wen, W., He, Y., Xia, A., Anstey, K. J., and Sachdev, P. (2012). Changing topological patterns in normal aging using large-scale structural networks. *Neurobiology of Aging*, 33:899–913.