Homogeneity based segmentation and
enhancement of Diffusion Tensor Images
A white matter processing framework

About the cover
The cover has been designed by Paulo Rodrigues.
It is inspired by the artwork of Tim Burton, the author's favourite film director, writer and artist. It is an amalgamate of several elements. On the front, a human brain is shown as illustrated by Gray. Behind it, a spot of water representing diffusion. Hovering the brain, the doors of perception. On the back, a spiral represents a PhD: slowly unravelling certain aspects of life and science, but sometimes one can go downward the spiral, into nothingness.

Colophon

This thesis was typeset by the author using LaTeX2_ε. The main body of the text was set using a 11-points Times Roman font.

Advanced School for Computing and Imaging

This work was carried out in the ASCI graduate school. ASCI dissertation series number 226.

The cover has been designed by Paulo Reis Rodrigues.

Homogeneity based segmentation and enhancement of Diffusion Tensor Images
A white matter processing framework

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de

Technische Universiteit Eindhoven, op gezag van de

rector magnificus, prof.dr.ir. C.J. van Duijn, voor een

commissie aangewezen door het College voor

Promoties in het openbaar te verdedigen

op dinsdag 19 April 2011 om 16.00

door

Paulo Reis Rodrigues

geboren te Vila Franca, Portugal

Dit proefschrift is goedgekeurd door de promotor:

prof.dr.ir. B.M. ter Haar Romeny

Copromotor:
dr. A. Vilanova

Contents

1 Introduction

In expanding the field of knowledge we but increase the horizon of ignorance.

Henry Miller

The human brain is one of the most intriguing organs in the human body. This complex collection of neurons defines the intelligent creatures that we are, who we are, how we are, our feelings, our dreams, our thoughts. Even more intriguing, its self awareness, its self introspection, its sentiency. This fascinating capability has been driving thinkers, philosophers, scientists, for millennia.

Fortunately, we live in exciting times when new imaging techniques provide an unprecedented look inside the structure and function of the brain [32]. Besides collecting data about the brain's many subsystems, these techniques have large societal benefits, as to aid in the diagnosis of many brain related diseases.

Magnetic resonance imaging (MRI) is a medical imaging technique (see figure 1.1(a)) which, through the use of strong magnetic fields and radio waves, is able to produce three-dimensional images of brain structures with a high degree of detail (see figure 1.1(b)). Its popularity is due to not only the greater tissue contrast of the produced images than computed tomography (CT) does, but also to the fact that it does so in a safe way, i.e. without the use of radiation (X-rays) or radioactive tracers (PET, SPECT).

(a)   (b)
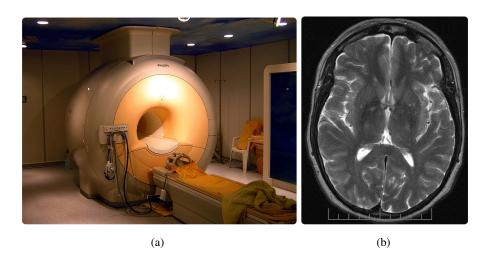
Figure 1.1: (a) A modern high field (3T) clinical MRI scanner, and (b) an MRI T2 brain axial image from an anonymized patient provided by Aaron G. Filler, MD, PhD for educational, research, and teaching purposes.

Diffusion weighted MRI (DW-MRI) is a recent imaging technique where the acquisition is made sensitive to the microscopic Brownian movement of water molecules, so called diffusion. In fibrous tissues this movement is hindered by the local structure enforcing the water molecules to diffuse mainly along these

fibers. By probing this diffusion at different directions, DW-MRI allows a non-invasive evaluation of the structural integrity of fibrous tissues, such as brain's white matter. Therefore, DW-MRI constitutes a valuable tool for understanding the brain, its mechanisms, functions and pathologies.

In the development of diffusion tensor imaging (DTI), first introduced by Basser [10], the movement of water molecules is measured in different directions and the average diffusion is then described by a 3×3 positive definite matrix, the so called diffusion tensor (DT). Here, different measures can be extracted to characterize tissue at microscopic detail. Fractional anisotropy, for instance, describes how directionally dependent (i.e. how anisotropic) the diffusion ensemble is, enabling researchers to evaluate changes in areas of neural degeneration and demyelination in diseases like Multiple Sclerosis [89].

From such a DT image, tractography methods reconstruct the tissue's fiber bundles, enabling the evaluation of the complex anatomical connectivity patterns within the brain (Conturo et al. [27], Basser et al. [11], Mori et al. [95]), but also the fibrous structure of muscle tissue, e.g., of the heart (Zhukov and Barr [150]) or the skeletal muscle (Heemskerk et al. [58]). These tractography algorithms can be used to estimate several white matter tracts - for instance, the fibers of the *corticospinal tract*, along which motor information travels between the cortex and the spinal cord. Figure 1.2(a) shows one of Gray's [54] remarkable illustrations of the anatomy of the human brain, depicting a major fiber tract, the *corpus callosum* in the center, connecting both hemispheres of the brain. In comparison, figure 1.2(b) shows the white matter of a brain, and its structure, as reconstructed through diffusion tractography.

In the past few years, there has been a worldwide strongly increasing interest in brain connectivity. The Human Connectome Project (HCP) [1] is the first large scale attempt at collecting and sharing data, aiming at constructing a map of the complete structural and functional connections of the brain, *in vivo*. The Blue Brain Project (BBP) [2] approaches the challenge of understanding the brain, its function and dysfunction, in a reverse-engineering way, by collecting data and detailed simulations such as the first model of a neocortical column with 10000 digitizations of real neurons. The Visualization of Brain Connectivity Data [3] addresses the visualisation of brain connectivity information, with interest in the joint visualisation of DTI and fMRI/EEG data.

DTI has also been used in many studies to explore the anatomical basis of human cognition and its disorders. Most studies are at an exploratory stage, aiming at providing an understanding of the underlying diseases mechanisms and eval-

---

[1] http://www.humanconnectomeproject.org
[2] http://bluebrain.epfl.ch/
[3] http://www.narcis.info/research/RecordID/OND1316395/Language/en

(a)   (b)
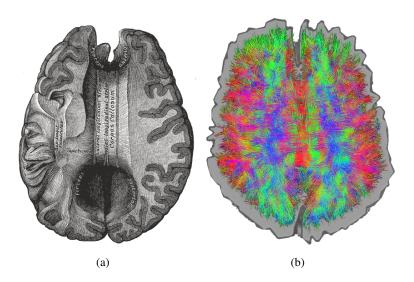
Figure 1.2: Human brain's white matter, depicting the *corpus callosum* in the center, from above: a) as illustrated by Gray [54]; b) estimated through DTI based tractography [117], where color encodes orientation of the local diffusion pattern (see section 2.5), superimposed over illustration of a brain axial slice.

uating differences across various subjects. Concha et al. [26] conducted a study demonstrating that DTI is capable of detecting, stating and following the microstructural degradation of white matter following corpus callosotomy in epilepsy patients. Brunenberg et al. [17] studied the potential of different DW-MRI modelling techniques to analyse the connectivity of *subthalamic nucleus* (STN) subterritories in order to avoid negative cognitive and emotional side-effects after deep brain stimulation (DBS) for Parkinson's disease. Several studies have also been conducted to investigate disorders such as: infarct (Moseley et al. [97]), schizophrenia (Kanaan et al. [71]), Alzheimer (Chua et al. [23]), multiple sclerosis (Melhem et al. [89]), language disorders (Klingberg et al. [77]), ageing (Salat et al. [119]), among others (Catani [19]).

DTI is also a valuable tool in tissue characterization, surgical planning, and treatment follow-up in patients with cerebral neoplasms (Field et al. [44]). Questions such as *"How to access the lesion, most safely?"* or *"How close is the tumour to the pathways interconnecting vital functional areas like vision, language or motor system?"* can be appraised thanks to DTI.

These multitude of applications and procedures usually involve the virtual dissection of white matter tracts of interest. The isolation of these bundles is usually manually defined and therefore require a great deal of neuroanatomical know-

ledge. This also makes DW-MRI based tractography heavily operator dependent and can take up to several hours of work per patient. The connectivity of the full brain can be extracted, however the amount of data to analyse and visualize is so vast that cluttering problems must be dealt with (see figure 2.13). Furthermore, as MRI technology progresses, higher resolution data will be available, hence this problem will become more urgent.

Several techniques have been proposed with the intent to automatically group individual fibers in coherent structures (Moberts et al. [94]). However, these methods operate over derived structures, thus not using the full tensor information, becoming sensitive to the choice of tractography method.

An alternative is the direct segmentation of the tensor image in volumetric regions containing fiber bundles that correspond to larger anatomical structures [7, 118, 138, 151]. Segmentation is needed to determine regions of interest where subsequent quantitative analysis and visualisation can be applied. These techniques provide tools to extract shape, size and other structural characteristics potentially useful for the analysis of pathologies, or the study of cognitive development in different populations such as premature neonates [133, 146], without the additional tractography step.

Several algorithms have been proposed for scalar image segmentation, however, how to extend these into tensor images still holds challenging questions. Frequently, new metrics in the space of tensors are introduced, and typically these algorithms are not automatic, i.e. they require the tuning of several parameters in order to achieve the desired results. Additionally, the limited interaction possibilities with the user prevents the added value of specialists' knowledge.

Although its great potential, a major limitation of DTI based tractography is that the calculated trajectory of the fiber may not follow the true tract, mainly due to partial volumes effects. The spatial resolution of a DW-MRI image, with typical voxel resolutions on the order of $1 - 2\,\mathrm{mm}^3$, is much lower than individual fiber bundles 10^{-6} m, leading to multiple fiber orientations within a voxel. A diffusion tensor fails to capture this complex structure leading to ambiguous fiber trajectories. Therefore, high angular resolution diffusion imaging (HARDI) techniques, pioneered by Tuch [130], were introduced. These are able to better capture the intra-voxel diffusion pattern compared to DTI. However, HARDI acquisitions in general produces very noisy diffusion patterns due to the low signal-to-noise ratio (SNR) from the scanners at high sensitivity to diffusion. Furthermore, it still exhibits limitations in areas where the diffusion pattern is asymmetrical (bifurcations, splaying fibers, etc.). To overcome these limitations the post-processing of the data is a crucial step.

1.1   Organization and Contributions of this Manuscript

This thesis focuses on the development of techniques able to automatically perform the virtual dissection of white matter structures. To segment such structures in a tensor field, the similarity of diffusion tensors must be assessed for partitioning data into regions, which are homogeneous in terms of tensor characteristics. In this thesis, the concept of **tensor homogeneity** is explored in order to achieve new methods for segmenting, filtering and enhancing diffusion images.

We start by first describing the brain, its main building elements, from the axons to the major fiber tracts connecting the different parts of the grey matter. When imaging the brain it is important to understand what we are imaging, the underlying anatomy and connectivity. In chapter 2 we address the brain, its structure, how it can be probed through imaging techniques, in particular through DW-MRI. Finally how can we exploit this information, i.e. process it, in order to identify white matter structures.

Since we do not know the "ground truth" of the human brain anatomical connectivity, one of the most challenging, yet extremely important aspects of DW-MRI is validation. In order to validate the accuracy of novel techniques in the modelling and processing of DTI/HARDI, software and hardware phantoms are often created. However, these models are an over-simplified version of the real underlying fiber configurations in the brain white matter. Different models are used in literature without justification which one is more appropriate than other, nor how the properties of the derived models/features change using one model or another. In **chapter 3** we analyse the most common synthetic data models, the multi-tensor model [2] and Söderman and Jöhnson's model [123], and hardware phantom and *in-vivo* data as well. This study is aimed to help scientists choosing the most appropriate synthetic data model when conducting DTI and HARDI experiments.

We start by studying the space of diffusion tensors, with special focus on the different measures that can be used to define (dis)similarities between tensors, i.e. to define homogeneity in a tensor field. Many different measures have been proposed to compute similarities and distances between tensors. Essential for algorithms such as segmentation, registration and quantitative analysis of DTI datasets, these measures are classified and summarized in **chapter 4**. This evaluation led to the development of a novel approach to semi-automatically define the similarity measures that better suit the segmentation task at hand, in **chapter 5**.

Chapter 6 explores the intrinsic hierarchical nature of the brain tissue: axons, fiber bundles, fiber tracts. Here, a multi-resolution watershed framework is presented, where the tensor field's homogeneity is used to automatically achieve a hierarchical representation of white matter structures in the brain, allowing the simultaneous segmentation of different structures, with different sizes.

HARDI, the successor of DTI, is able to more accurately model the diffusion pattern in areas of complex fiber heterogeneity, however, at the cost of poor image quality (low SNR). Moreover, HARDI based methods can not reconstruct asymmetric diffusion patterns such as bifurcations and splaying fibers. The processing of HARDI data is paramount, and the contextual (neighbourhood) information plays an important role. The processing of HARDI data is based on modelling the stochastic process of water diffusion within tissues, inferring the homogeneity characteristics of the diffusion field. Convolutions with these kernels are then performed in the coupled spatial and angular domain. However, these approaches have high computational complexity of an already complex HARDI data processing. In **chapter 7**, an accelerated framework for HARDI data denoising, regularization and enhancement is presented.

Although HARDI has proven to better characterize complex intra-voxel structures than its predecessor DTI, its higher acquisition times and significantly lower signal-to-noise ratios established DTI as more attractive for use in clinical research. In **chapter 8** we use contextual information derived from DTI data, to obtain similar crossing information as from HARDI data. We show that with extrapolation of the contextual information the obtained crossings are the same as the ones from the HARDI data, and the robustness to noise is considerably better.

Software contribution
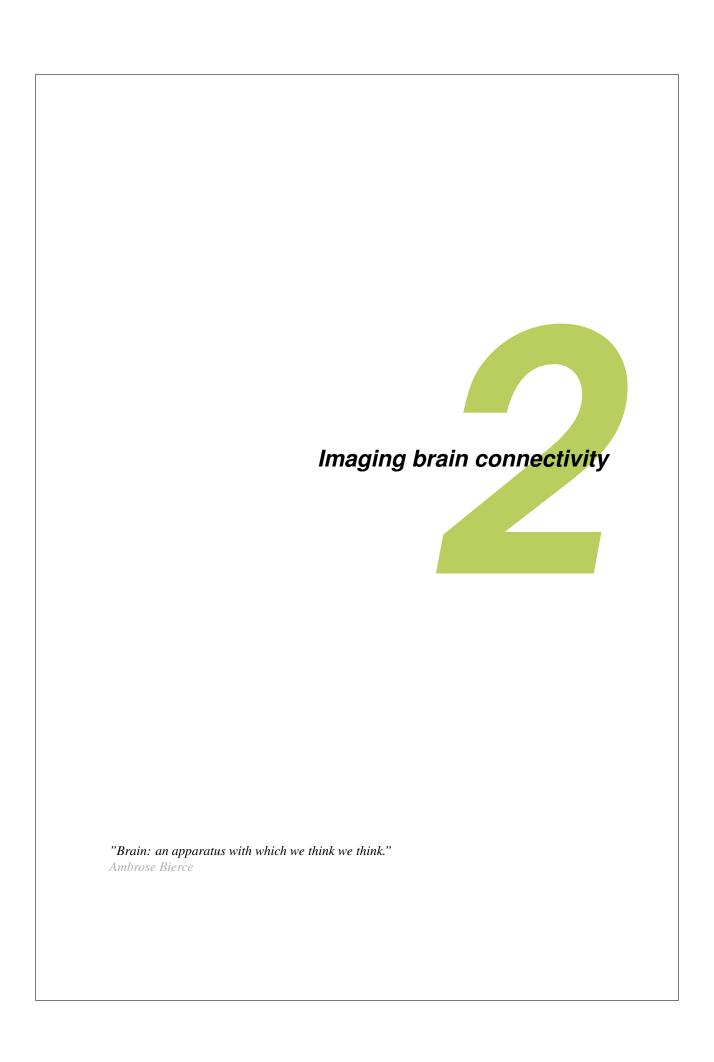
All the methods presented in this thesis were developed in a framework common to our research group - named DTItool [117]. Lead by Dr. A. Vilanova, several researchers contributed to this framework unifying a collection of state-of-the-art algorithms dedicated to medical image processing and visualization. This tool has been developed in collaboration with Vesna Prckovska[4], Tim Peeters[5] and other members of the BMIA[6] group, in particular the collaborations with Markus van Almsick and Remco Duits. The tool and the implemented algorithms are available on request.

---

[4]http://www.vesnaprckovska.net
[5]http://www.timpeeters.com
[6]http://www.bmia.bmt.tue.nl

2

Imaging brain connectivity

"Brain: an apparatus with which we think we think."
Ambrose Bierce

Contents

2.1   A brief history of Neuroimaging

In a process called pneumoencephalography, the cerebrospinal fluid involving the brain is drained and replaced with air. This changes the relative density of the brain and its surroundings, causing it to show better on an x-ray. This incredibly unsafe and very painful procedure can be considered the beginning of neuroimaging, in the start of the 20^{th} century. This method is now fully abandoned, and right so.

In 1927 Egas Moniz pioneered cerebral angiography with a suitable contrast medium as to provide images of blood vessels within the brain, allowing the detection of abnormal blood vessels with great precision.

In the early 1970s, computerized axial tomography (CAT or CT scanning), was introduced by Cormack and Hounsfield, providing more detailed anatomical images of the brain, for both diagnostic and research purposes, granting them the 1979 Nobel Prize for Physiology or Medicine. Soon after, in the early 1980s, single photon emission computed tomography (SPECT) and positron emission tomography (PET) of the brain was developed thanks to the use of radionuclides. These techniques can show the amount of brain activity in the various regions,

thanks to the measurements of blood flow, oxygen and glucose metabolism of the working brain, although with a low spatial resolution.

Magnetic resonance imaging (MRI) was developed, roughly at the same time, by Mansfield and Lauterbur, awarded the Nobel Prize for Physiology or Medicine in 2003. Using strong magnetic fields and radio frequency fields, this technique provides detailed images, with much greater contrast between the different soft tissues of the body than CT, with low invasiveness, lack of radiation exposure, thus of special interest for neurological imaging. During the 1980s many improvements and applications of MR appeared. Exploiting the magnetic properties of haemoglobin and capillary responses to increased metabolic need of active areas, changes in blood flow associated with neural activity can be measured with functional magnetic resonance imaging (fMRI). Images can now be created reflecting which brain structures are activated while performing different tasks.

Diffusion weighted MRI (DW-MRI) is a recent imaging technique, where the acquisition is made sensitive to the microscopic movement of water molecules (diffusion) restricted by the local structure. Its basic principles were introduced in the mid-1980s [80, 90, 127]. These measurements thus allow the evaluation of the structural integrity of fibrous tissues, such as the white matter, in the brain. In the 1990s, Peter Basser and his co-workers [10], established diffusion tensor imaging (DTI) as a viable imaging method, granting Basser the 2008 International Society for Magnetic Resonance in Medicine Gold Medal.

2.2 Human Brain: from nervous tissue to architecture

The human cerebrum is an electrochemical machine that is continuously processing information about its surroundings, gathered through the senses. This information is processed according to brain's previous experience (memory) and results in an appropriate response or action. How the brain works is very much related to the collective behaviour of the billions of cells in it.

2.2.1 Nervous tissue

The neuron is a fundamental component in comprehending the brain. The human brain is composed by billions of neurons and supporting cells. They are responsible for cognitive and memory functions - they define what we are. The cortex of the human brain is extremely complex, containing in the order of one hundred billion (10^{11}) neurons [41], processing and producing electrical and chemical currents. As clearly illustrated by Santiago Ramón y Cajal (1852-1934) [145], in figure 2.1, the neurons are connected to each other in an extremely intricate network.
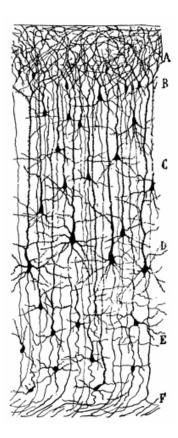
Figure 2.1: This drawing first appeared in volume two, part two of Cajal's *Textura del Sistema Nervioso del Hombre y de los Vertebrados*, published in Madrid in 1904. Using a method called Golgi staying [52], Cajal produced images such as this, where we can see the six layers of the mouse neocortex. Cajal's drawings remain fundamental for neuroscience by showing that the nervous system is a complex network composed of individual neurons. Cajal received the Nobel Prize in Physiology or Medicine in 1906.

Each neuron (see figure 2.2) consists of a body (called the soma) and tentacles (the dendrites) which connect to thousands of other neurons, seeking and receiving information. The surface of the cell body integrates the information arriving at its dendrites. If the excitation is sufficient, it triggers impulses that are conducted away along an axon. Each neuron has a single axon leaving its cell body, although an axon can branch to stimulate more cells.

Typically an axon radius varies from $0.2\,\mu$m to $20\,\mu$m, and can reach up to a meter length, as from sensory neurons in the feet to neurons in the spinal cord. There are three types of neurons:
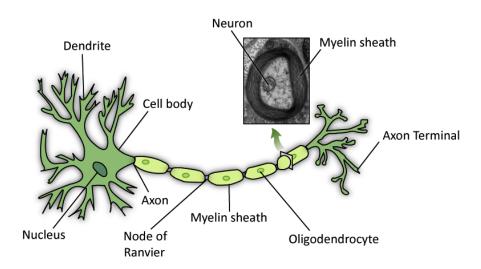
Figure 2.2: Structure of a typical neuron. The dendrites extending from the cell body receive the information. A single axon transmits the impulse away. Many axons are wrapped by a myelin sheath. *Figure adapted from the originals generated and deposited into the public domain by the US National Cancer Institute's Surveillance, Epidemiology and End Results (SEER) Program and by the Electron Microscopy Facility at Trinity College.*

- **Sensory neurons**, or afferent neurons, translate physical input from the environment - such as light, sound, temperature, pressure, taste, smell - into electrical signals. The impulses are carried into the central nervous system (CNS), i.e. the brain and spinal cord.

- **Motor neurons**, or efferent neurons, carry impulses from the CNS to effectors, i.e. muscles or glands.

- **Interneurons**, or association neurons, are located in the brain and spinal cord, and help provide complex reflexes and higher cognitive functions such as learning and memory.

The cerebral tissue contains many other cells that serve a variety of functions such as supplying the neurons with nutrients, removing waste from neurons, and providing immune functions, which are called neuroglia. Two of the most important types of neuroglia are the Schwann cells and oligodendrocytes, which produce myelin sheaths surrounding axons of many neurons (see figure 2.2). Schwann cells produce myelin in the peripheral nervous system (composed by sensory and motor neurons), while oligodendrocytes produce myelin in the CNS.

During development, the myelin sheaths are formed by successive wrappings of these cells around the axons. These multiple compact membrane layers facilitate far more rapid conduction of impulses. The myelin sheath is interrupted at regular intervals of 1 to 2 millimiters by small gaps of 1 to 2 μm known as nodes of Ranvier. In the CNS, myelinated axons form the **white matter**, and the unmyelinated dendrites and cell bodies form the **gray matter**.

2.2.2   Neuroanatomy

The complex connectivity network of the neurons might seem largely random, however, the brain is extremely well organized, at all levels, and every day we learn more: from regular neuronal patterns in cortical columns, precise sensory maps, energy minimizing distances in connectivity, etc.

The first basic structures are the two cerebral hemispheres. Each hemisphere primarily receives sensory input from the opposite side of the body and conveys motor control primarily over that side. There are two of every brain organ located on each hemisphere (except the pineal gland).

Incidentally, René Descartes [33] thought the unique pineal gland, located almost in the middle of the brain, might be involved with the mind or the soul; he called it the "seat of the soul".

Much of the neural activity of the brain occurs within a layer of gray matter, just a few millimetres thick on its outer surface. This layer, called the **cerebral cortex** or **neocortex**, is densely packed with neuron bodies and dendrites. The cortex is a highly convoluted structure whose ridges and valleys are dubbed, respectively, gyri and sulci. The neocortex is the center of our most impressive capabilities, such as learning, memory, language, and consciousness. The other parts of our brains are similar to those in other mammals, however the neocortex is comparatively much larger than in other mammals. The human cerebral cortex is "new" in an evolutionary sense, hence the name neocortex.

The exterior of the cerebral cortex is divided into four lobes: frontal, parietal, temporal and occipital (see figure 2.3). The neocortex is organized into regions that process specific types of input or are specialized in specific cognitive functions. These functions fall into three general categories: motor, sensory and associative. The primary **motor** cortex, illustrated in figure 2.3, is located along the *gyrus* on the posterior border of the frontal lobe, just in front of the central *sulcus*. Each point on its surface is associated with the movement of a different part of the body, as illustrated by the *cortical homunculus*, discovered and described by Wilder Penfield. Similarly, just behind the central *sulcus*, on the anterior edge of the parietal lobe, lies the primary **somatosensory** cortex. Each point in this
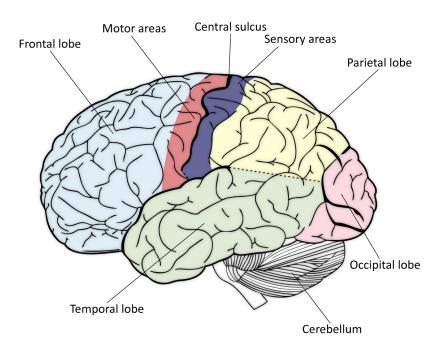
Figure 2.3: **The cerebrum.** This diagram shows the lobes of the cerebrum and indicates some of the known specialized regions. *Based on Figure 728 from Gray's Anatomy [54].*

area receives the input from sensory neurons. Large parts of the somatosensory cortex is dedicated to fingers, lips and tongue given our need for manual dexterity and speech. The area of the cortex that is not dedicated to motor or sensory functions is referred to as **association** cortex. Higher mental activities take place here, reaching its greatest extent in humans, where it takes 95% of the cortex's area.

Millions of axons interconnect the myriad of neurons in the neocortex forming the **white matter**, depicted in figure 2.4. Following the ratios provided by Miller et al. [93], assuming a brain volume of 1250 mL, white matter accounts for about $44 - 48\%$ of the brain volume between age 20 and 50. The white matter is a hierarchically ordered tissue: from aligned microscopic axons to large **bundles** running together between various gray matter regions. These coherent groups are called white matter **fiber tracts** (*fasciculi*), and are broadly classified as **association**, **commissural** or **projection** fibers:

Association fibers are confined to one hemisphere and interconnect cortical areas within that hemisphere. They are further classified as either *short* or *long* as-
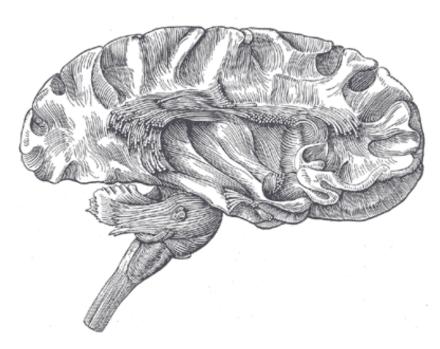
Figure 2.4: Dissection of cortex and brain-stem showing association fibers. *Figure 752 from Gray's Anatomy [54].*

sociation fibers. Short association fibers lie immediately beneath the gray matter, connecting nearby specialized regions within a gyrus or between gyri by looping around the sulci. Long association fibers interconnect functionally specialized areas of the cortex. Some of the long association fibers are the *cingulum*, *superior longitudinal and arcuate fasciculi*, and *interior occipitofrontal and uncinate fasciculi* (see figure 2.5(a) and later on figure 5.6).

Commissural fibers interconnect across the midline, mostly corresponding areas within the two hemispheres. The largest commissure is the *corpus callosum* located in the center of the brain (see figure 2.5(b)(b) and later on figure 5.5), connecting both halves. It is composed of axons of variable diameter and conduction velocities. There are several additional smaller commissures such as the *anterior commissure*, the *posterior commissure*, and the *hippocampal commissure*.

Projection fibers unite the cortex with various lower parts of the brain and the spinal cord. These fibers are classified into two groups on the basis of the direction in which the fibers conduct. Afferent fibers are those on the way to

the cortex, while efferent fibers are those originating in regions of the cortex and which are proceeding to the basal ganglia, brain stem, and spinal cord. The *corticospinal tract*, for instance, as illustrated in figure 2.5(b) mostly contains motor axons travelling between the cortex and the spinal cord.

(a)   (b)
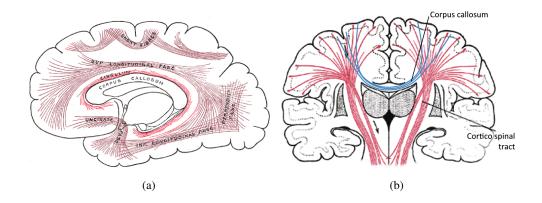
Figure 2.5: Diagram showing principal systems of fiber tracts in the cerebrum. a) Association fibers, with the *cingulum* highlighted; b) Projection fibers, such as the *cortico spinal tract*, combined with the *corpus callosum*, the largest commissure in the brain. *Based on Figures 751 and 764 from Gray's Anatomy [54].*

Additionally, there are two fluid systems in the brain: the cerebrospinal fluid (CSF) in the ventricles and around the brain; and the vascular system providing blood.

Let's take a moment to summarize the scales of brain tissue in order to help understand the discussion in the following sections:

- Most of axon fibers diameters, in human white matter, are less than 10^{-6} m;

- The packing density of axon fibers is of order 10^{11} m^{-2};

- Fiber tracts, i.e. coherent fiber bundles, vary in diameter from several centimetres down to a few microns;

- The voxel dimensions in a typical biomedical 1.5-3T MRI is of order 10^{-3} m.

2.3   Principles of diffusion

Diffusion is a process arising in nature, which results in particle or molecular mixing, due to collisions between atoms or molecules in a fluid. The physical law that explains this process is called Fick's first law [43] relating particle concentration C to the diffusive flux $\mathbf{J}$ through the relationship

$$\mathbf{J} = -D\nabla C \qquad \boxed{2.1}$$

where the constant of proportionality D is the so called "diffusion coefficient". Fick's law embodies the idea that particles go from areas of higher to lower concentration resulting in an even concentration in the whole fluid. The diffusion coefficient D depends on the temperature and the intrinsic properties of the medium. Its sensitivity on the local microstructure enables its use as a probe of the physical properties of the biological tissue.

Robert Brown (a Scottish botanist), while studying pollen particles floating in water under the microscope, observed this seemingly random movement of particles, the so called Brownian motion. In the beginning of the 20^{th} century, Einstein [42] established a relationship between the mean-squared displacement of particles, characterizing the Brownian motion, and the classical diffusion coefficient D in Fick's Law, given by

$$< x^2 >= 2Dt \qquad \boxed{2.2}$$

where $< x^2 >$ is the mean-square displacement of the particles during a diffusion time t, and D is the same diffusion coefficient appearing in Fick's law 2.1. Thus, the larger the diffusion coefficient D, the greater the distance of a particle is expected to travel on average during the same diffusion time.

The Brownian motion can be described by the diffusion displacement probability density function (PDF), also called diffusion propagator, $P(\mathbf{r}, t)$, where $\mathbf{r} = \mathbf{R} - \mathbf{R}_0$ is the net displacement of a particle, initially located at position $\mathbf{R}_0$, with $\mathbf{R}$ being the displacement after time t. Using Fick's law of diffusion, the diffusion process can be approximated as:

$$\frac{\partial}{\partial t}P(\mathbf{r}, t) = D\nabla^2 P(\mathbf{r}, t), \qquad \boxed{2.3}$$

The solution to this equation is the propagator given by Basser et al [10], i.e. a Gaussian PDF:

$$P(\mathbf{r}, t) = \frac{1}{\sqrt{(4\pi Dt)^3}} \exp\left(\frac{-\mathbf{r}^2}{4Dt}\right). \qquad \boxed{2.4}$$

In tissues with a fibrous structure - such as the axons in the white matter - diffusion of water molecules occurs faster along fibers' longer axis and slower in the orthogonal direction. When the movement of molecules is not restricted, and therefore equal in all directions, diffusion is described as being **isotropic** (figure 2.6(b)), whereas when there is a preferred direction of movement (as illustrated in figure 2.6(a) diffusion is said to be **anisotropic**. Diffusion is an isotropic case can be described by D in any direction $\mathbf{r}$ whereas in an anisotropic case D depends on the direction $\mathbf{r}$.
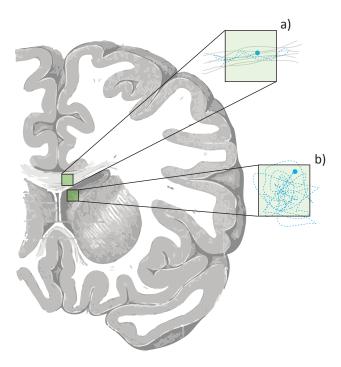
Figure 2.6: Within neural tissue, the movement of water molecules is hindered by the local fibrous structure (a), whereas in the ventricles the molecules diffuse in all directions equally (b).

2.4 Diffusion weighted imaging

Because diffusion is influenced by the geometrical structure of the environment, diffusion weighted magnetic resonance imaging (DW-MRI) provides a unique opportunity to non-invasively probe the structure of tissues. By capturing the average diffusion of the water molecules within biological tissue, DW-MRI man-

ages to describe the structure of tissues such as the white matter in the brain, or muscle fibers.

MRI exploits the fact that the human body is mainly constituted by water molecules, and each molecule has two hydrogen protons. When the scanner applies a powerful magnetic field, the magnetic moments of some of these protons change, aligning with the direction of the magnetic field. A radio frequency is then briefly applied, producing an electromagnetic field, causing the flip of the spin of the aligned protons in the body. After the field is turned off, the protons decay to the original state and the difference in energy between the two states is released as a radio frequency photon. These photons produce the electromagnetic field detected by the scanner.

Additional magnetic fields are applied in order to make the field strength depend on the position within the scanned subject, thus making the frequency of the released photons dependent on the position. An image can be constructed since the protons in different tissues return to their equilibrium state at different rates.

Stejskal and Tanner [125] in 1965, proposed an imaging sequence used to measure the diffusion of water molecules in a given direction $\mathbf{g}$. This sequence, called pulse gradient spin echo (PGSE), illustrated in figure 2.7, applies two gradient pulses in direction $\mathbf{g}$.
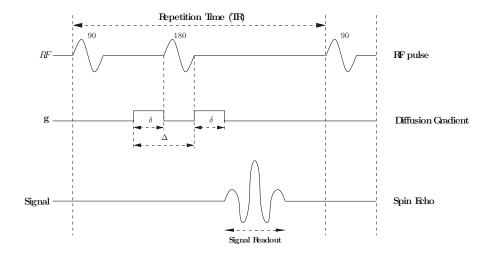
Figure 2.7: Scheme of the pulse gradient spin echo sequence, proposed by Stejskal and Tanner [125], adapted from [82].

The first $90°$ pulse causes a phase shift of the spins, thus encoding their position in function of the frequency. After time Δ the $180°$ pulse combined with the second gradient pulse cause the shift to be cancelled for static spins. In the

meanwhile, some protons underwent Brownian motion, therefore the refocus will not be perfect and the measured MRI signal will be attenuated resulting in signal attenuation measured in T2-weighted images, i.e. diffusion weighted (DW) images. The faster the water molecules diffuse, the more dephased they will be and the weaker the recorded signal. Assuming a Gaussian PDF, the relation between this attenuation and the amount of diffusion can be expressed through [125]:

$$\frac{S_{\mathbf{g}}}{S_0} = \exp^{\gamma^2 G^2 \delta^2 (\Delta - \delta/3) D} = \exp^{-bD} \qquad \boxed{2.5}$$

where S_0 is the signal intensity without the diffusion weighting $b = 0$, $S_{\mathbf{g}}$ is the signal with the gradient $\mathbf{g}$, γ is the gyromagnetic ratio, G is the strength of the gradient pulse, δ is the duration of the pulse, Δ is the time between the two pulses, and D is the diffusion-coefficient. The so called b value, proposed by Le Bihan et al. [81], which is proportional to the square of the gradient strength, is used to characterize the level of sensitivity to diffusion. A typical value is b $= 1000\,\text{s/mm}^2$.

2.4.1 Diffusion tensor imaging

Basser et al. [10] proposed the use of a second order symmetric positive-definite tensor to model the diffusion properties of biological tissues. The diffusion propagator in a homogeneous anisotropic environment can be well described by a Gaussian PDF, establishing the DT model as follows:

$$P(\mathbf{r}, t) = \frac{1}{\sqrt{(4\pi t)^3 |\mathbf{D}|}} \exp\left(\frac{-1}{4t} \mathbf{r}^{\mathrm{T}} \mathbf{D}^{-1} \mathbf{r}\right), \qquad \boxed{2.6}$$

where $|\mathbf{D}|$ is the determinant of tensor $\mathbf{D}$.

The diffusion coefficient D is related to each direction $\mathbf{r} \in \mathbb{R}^3$:

$$D = \mathbf{r}^T \mathbf{D} \mathbf{r}. \qquad \boxed{2.7}$$

In this model, the scalar diffusion coefficient D is replaced by a positive semi-definite matrix $\mathbf{D}$ representing diffusion, the diffusion tensor. This diffusion tensor is a 3×3 symmetric positive definite matrix that characterizes diffusion in 3D, and it is usually represented by an ellipsoid (see figure 2.8). The scalar components of a tensor $\mathbf{D}$ are denoted by:

$$\mathbf{D} = \begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{xy} & D_{yy} & D_{yz} \\ D_{xz} & D_{yz} & D_{zz} \end{pmatrix}. \qquad \boxed{2.8}$$

Since $\mathbf{D}$ is symmetric it only has six different values, and therefore, it has only six unknown coefficients that we need to estimate. DTI needs at least six DW images ($S_{\mathbf{g}}$) and one unweighted diffusion image ($S_0, b = 0\,\mathrm{s/mm^2}$) [10], typically called B_0 image, to solve the system of equations to obtain the tensor. The typical acquisition setting [69] consists of 20 to 60 DW images acquired with non-collinear diffusion gradient directions and $b = 1000\,\mathrm{s/mm^2}$ and a single B_0 image.

Using a tensor representation of the diffusion and the Stejskal-Tanner equation 2.5 we obtain:

$$S_i = S_0 \exp\left(-b\mathbf{g}_i^{\mathrm{T}}\mathbf{D}\mathbf{g}_i\right) \qquad \boxed{2.9}$$

where $D_i = \mathbf{g}_i^{\mathrm{T}}\mathbf{D}\mathbf{g}_i$ is called apparent diffusion coefficient (ADC) in the direction of $\mathbf{g}_i$ and $\mathbf{D}$ is a diffusion tensor.

The shape of diffusion can be easily visualized with ellipsoidal glyphs (squished or stretched spheres). Figure 2.8 illustrates diffusion as anisotropic (cigar shaped), as a planar shaped, but it may also be spherical, as in isotropic diffusion.
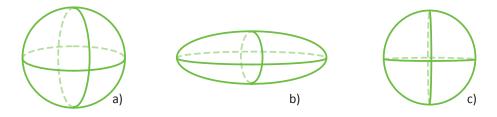
Figure 2.8: The three stereotypes of Gaussian diffusion in 3D, visualized with ellipsoidal isoprobability surface glyphs with a) isotropic, b) linear or c) planar shape. In DTI, all diffusion shapes are spanned by an interpolation between these three types.

The DT can be decomposed, by eigenanalysis, into eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$ and corresponding eigenvectors $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$. The first vector gives the principal direction of diffusion, the other two span an orthogonal plane to it and the eigenvalues quantify the diffusivity in these directions. When $\lambda_1 \gg \lambda_2$, $\mathbf{e}_1$ is aligned with the preferred diffusion direction of the water molecules in that voxel, and λ_1 is its diffusivity (figure 2.9).

The analysis and processing of the diffusion tensor field is explained in more detail in section 2.5.

2.4.2 Beyond DTI: high angular resolution diffusion imaging

The diffusion tensor model provides good results where, within a voxel, there is only one fiber population, i.e. fibers are aligned along a single direction. However, when several fiber populations intersect, DT fails to identify the different
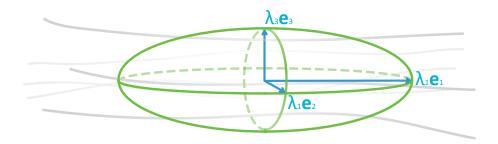
Figure 2.9: Diffusion within a coherent arrangement of fiber bundles is represented by a tensor whose main eigenvector $\mathbf{e}_1$ coincides with the orientation of the bundles. The eigenvectors and eigenvalues define the tensor shape.

fiber directions simultaneously. This problem is due to limitations of the resolution of current MRI machinery, and the protocol with limited number of directions. While the radius of an axon varies from $0.2\ \mu m$ to $20\ \mu m$, the resolution of DW images ranges from $1\ mm^3$ to $8\ mm^3$ in clinical settings. As a consequence, one voxel may contain distinct fiber bundles with crossing, kissing and splaying geometrical configurations, as illustrated in figure 2.10, which DTI inadequately characterizes by averaging the several fiber orientations. Several studies have been conducted aiming at quantifying the amount of multi-fiber voxels within a brain. Alexander et al. [2] classified 5% of voxels within the brain as complex structure (non-Gaussian). Tuch [130] showed that $2/3$ of the white matter has more complex intra-voxel structures. According to Behrens et al. [14], $1/3$ of the white matter voxels contains crossing fibers. In a more recent work, Jeurissen et al. [63] reports 90% of the voxels, within the white matter, to have a multi-fiber configuration - a much higher proportion than previously reported. It is clear that in these areas we need modelling techniques able to provide higher angular resolution.

Approaches based on high angular resolution diffusion imaging (HARDI) were pioneered by Tuch [130]. In HARDI more sophisticated models are employed to reconstruct more complex fiber structures and to better capture the intra-voxel diffusion pattern. Figure 2.11 illustrates this relation between underlying fibrous structure, and respective DTI and HARDI reconstructions. Some of the proposed models include high-order tensors [102], mixture of Gaussians [64, 130], spherical harmonic (SH) transformations [50], diffusion orientation transform (DOT) [104], orientation distribution function (ODF) [35] using the Q-ball imaging [131], and the spherical deconvolution approach [128] by estimating the fiber orientation distribution (FOD).
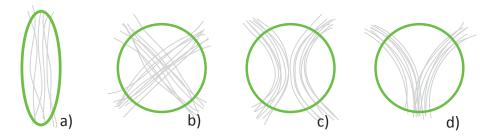
Figure 2.10: Examples of complex intra-voxel fiber configurations that cannot be resolved using the Gaussian diffusion model of DTI. From left to right: single fiber, crossing, kissing, and splaying fibers.

It is important to note that all of the diffusion weighted MRI modelling techniques model functions that reside on a sphere. For simplicity we will refer to them as spherical distribution function (SDF). Whereas the physical meaning of these SDFs can be different (a probability density function (PDF), iso-surface of a PDF, ODF, FOD, etc.), in all cases they characterize the intra-voxel diffusion process, i.e. the underlying fiber distribution within a voxel.
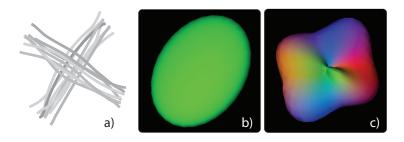
Figure 2.11: A more complex intra-voxel structure, such as crossing fibers (a), is not well captured by the DT model (b), whereas a HARDI model such as Qball (c) is able to identify the two fiber populations.

The HARDI acquisition schemes typically use a higher number N of gradient directions $\mathbf{g}_i$ (usually $60 \leq N \leq 200$) than DTI and scans are made using b-values of over 3000 s/mm^2 [113]. These are needed to be able to capture the more complex profiles, however as a consequence, HARDI produces, in general, noisy diffusion patterns due to the low SNR. Another important limitation is that HARDI acquisition schemes take too long time. In a clinical setting for instance, imaging time is often critical, and HARDI is just not feasible. Due to the simplicity and post-processing speed, DTI is the most widespread DW-MRI technique. Notwithstanding, profuse research is being done in new models that are capable

of characterizing the diffusion propagator, but keeping DT's simplicity.

2.5 Processing DTI

The processing and visualisation of diffusion tensor images present several challenges, given its multivalued nature and the complex interrelationships between the different tensors. Several approaches have been proposed through the last decade, where most of them first reduce the dimensionality of the data by extracting relevant information from the DT. The following sections present an overview of the different processing methods, where we discuss two important characteristics: the dimensionality to which the tensor is reduced; and the ability to show local or global information, i.e. the complex inter-voxel relationships.

2.5.1 Anisotropy measures / Biomarkers

Medical researchers and practitioners are well trained in reading scalar images, i.e. gray-level images, as an X-ray for example. The diffusion tensor is a rich formalism able to provide a considerable amount of information. It is not surprising that DTs are too complex structures to interpret and analyse. After formalizing DTI, Basser [9] defined a set of scalar, rotationally invariant measures to quantify different characteristics of the DT, and therefore the underlying diffusion characteristics. In the following, we briefly describe some of the most frequently used scalar measures:

- Mean diffusivity (MD) is the average of the DT eigenvalues or trace. MD is low within the white matter, whereas, for example in the ventricles, it is high due to the unrestricted diffusion of the water molecules. This measure of overall diffusion rate can be used to delineate the area affected by a stroke, as demonstrated by Van Gelderen [132].

$$MD(\mathbf{D}) = <\mathbf{D}> = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3} = \frac{Trace(\mathbf{D})}{3}$$

- Fractional anisotropy (FA) is one of the most used indices in clinical applications. This rotationally invariant, dimensionless measure, expresses the anisotropy of the tensor ranging from 0, when the tensor is completely isotropic, to 1, when the diffusion is bound to a single axis.

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

- Geometrical diffusion measures, linear C_l, planar C_p and spherical C_s anisotropy, proposed by Westin et al. [142], characterize the tensor as cigar-shaped (C_l), disk-shaped (C_p) or spherical-shaped (C_s), as shown in figure 2.8.

$$
\begin{aligned}
C_l &= \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + \lambda_3} \\
C_p &= \frac{2(\lambda_2 - \lambda_3)}{\lambda_1 + \lambda_2 + \lambda_3} \\
C_s &= \frac{3\lambda_3}{\lambda_1 + \lambda_2 + \lambda_3}
\end{aligned}
$$

Chapter 4 presents an extensive analysis of the many measures defined in literature, as well as distances and similarity measures used to compare DTs.

This data can be visualised slice per slice (as in figure 2.12(b)), using volume rendering techniques, or with common techniques used for 3D scalar fields.

Colour can be used to indicate the orientation of the underlying tensor. Applying the standard RGB colouring of the principal eigenvector $\mathbf{e}_1$, by mapping the vector components to RGB, allows the delineation of basic neuroanatomic features, as shows in figure 2.12(a). This orientational information can be combined with other measures to better clarify differences in the tissue. The RGB map can be weighted by anisotropy (FA) (figure 2.12(b)), to highlight white matter structures. This visualisation has nice results since main fiber tracts are aligned in the X, Y or Z directions, thus clearly visible in red, green or blue.

2.5.2 Glyphs

As introduced in section 2.4.1, the diffusion tensor can be represented by a graphical object, the tensor glyph. In the simplest form, a sphere can be deformed according to DT's eigenvalues and oriented along its eigenvectors. The first use of these ellipsoids for DT glyphs, was done by Pierpaoli et al [109], where an array of glyphs was put together, showing a 2D slice of DTI data (see figure 2.12(c)).

Many different glyph based techniques have been presented in literature. Laidlaw et al. [79], for instance, developed a method based on oil painting and brush strokes to enhance the diffusion patterns, and used it to visualize sections of mice spinal cords. Kindlmann [75] created a class of tensor glyphs, based on super-quadrics, where a sphere indicates isotropic diffusion, cylinders are used for linear and planar anisotropy, and the intermediate forms of anisotropy are represented by shapes close to a box. With these glyphs, tensors' orientations are better depicted than with ellipsoids, specially for the more planar shapes.
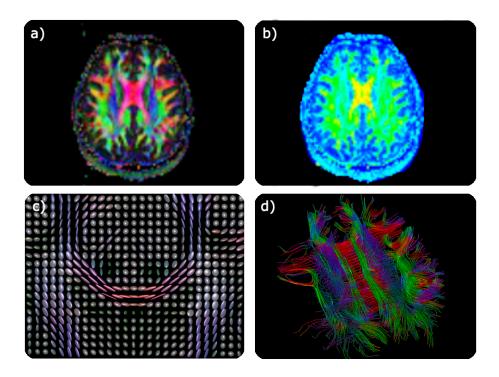
Figure 2.12: Several examples of common processing methods of a DTI dataset of a healthy brain: (a) axial slice with RGB colouring of the principal eigenvector; (b) axial slice with fractional anisotropy ranging from low (blue) anisotropy to high (red) anisotropy; (c) DT ellipsoids; (d) DTI based tractography, with RGB colouring of the principal eigenvector. *Images realized with the developed framework DTItool.*

It is important to note, that these glyph based methods, although fully representing the diffusion tensor, do not express the relationships between voxels, i.e., they do not expose the contextual features across the tensor field (figure 2.12(c)), they just depict the local properties.

2.5.3 Tractography

Up to now we saw how with diffusion MRI techniques we can obtain a representation of water diffusivity per voxel of the human brain. However, as we have seen in section 2.2, the neuron pathways that constitute the white matter are of major interest in analysing brain's connectivity.

Diffusion MRI constitutes a powerful non-invasive tool to analyse the structure of the white matter within a voxel, but also to investigate the anatomy of the brain and its connectivity. In DW-MRI based tractography, the axonal paths are estim-

ated either by following the main direction of diffusion or by using a probabilistic model of the diffusion. For reviews of tractography techniques refer to the works of Mori et al. [96] and Lori et al. [84].

There are two main philosophies in tractography.

Simple **deterministic**, or streamline, tractography traces fiber bundles by following the principal direction of diffusion from point to point in the image volume [11, 27, 68, 86, 95, 142, 144]. Figure 2.12(d) shows an example of deterministic fiber tracking using the developed framework DTItool.

Probabilistic tractography algorithms use a probability density function to model the uncertainty in the fiber orientation in each voxel. The algorithm runs repeated streamline processes with orientations drawn from the model. The fraction of fibers that pass through a voxel provides a connectivity index, between two regions in the brain, reflecting fiber organization [15, 56, 60, 83, 99, 105, 147]. These methods are computationally expensive, and thus with less appeal for clinical applications. Another drawback of this approach is the fact that any two points in space are connected and therefore it is necessary to establish a criterion for when points are considered not to be connected.

It often happens that in some parts of the trajectory the local diffusion profile does not support the presence of a fiber. This can be due to noise, or the presence of a high number of fiber populations with heterogenous orientations. A possible solution for this problem of local perturbations may be provided by global tractography methods, optimising a global criterion for connectivity. **Geodesic** fiber tracking (GT), first proposed by Parker et al. [106], interprets brain fibers as minimal distance paths (geodesics) for a metric derived from the diffusion profile. Succinctly, a distance field from a seed region is constructed. This is done by solving a partial differential equation (PDE), the so called Eikonal equation. Solving this equation in a heterogeneous and highly anisotropic medium, as is the human brain, is a technically challenging problem. There have been a few attempts at solving this problem [62, 83]. A drawback of this approach is the fact that any two points in space are connected, thus, it is necessary to define not only start but also end points.

In these approaches, tracking is initiated from a region of interest (ROI) from which a series of points are taken as starting point for tracing fibers. The definition of these regions, therefore, influences the obtained bundles. This is usually done by the user, thus compromising reproducibility. The anatomical connectivity of the full brain can be extracted, however the amount of data to analyse and visualise is such that cluttering problems must be dealt with (see figure 2.13). Furthermore, as MRI technology progresses, higher resolution data will be available, hence this problem will get greater importance.

Many of these tractography methods are based on DT images (fields), thus they reflect the same limitation in handling complex structures like crossing, kiss-
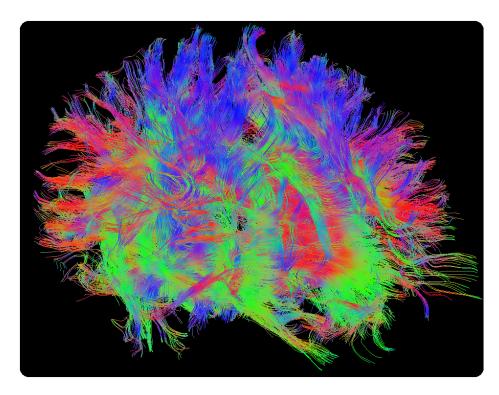
Figure 2.13: Result of a full brain deterministic tractography. Using the developed framework DTITool, the fiber bundles are reconstructed if $FA > 0.2$. Fibers are coloured with the typical RGB mapping of the principal eigenvector. $128 \times 128 \times 60$ dataset provided by Poupon et al. [111].

ing or splaying fibers. Several authors provide tractography algorithms based on multiple-fiber reconstruction using HARDI models [37, 105, 115, 130].

We would like to point out that the underlying physical basis of these functions is a gross simplification of the actual complex distribution of different barriers to diffusion, and therefore, the relationship between any estimation of fiber bundles with the true distribution of fibers still requires a great deal of validation and verification.

2.6 *In vivo* virtual dissection

Delineation and analysis of brain's architecture has been an active area of research for more than a century. As we referred to in the beginning of this chapter, several techniques have been developed to aid in the extraction and visualisation of this complex network of connections in the brain. These techniques were in-

vasive and therefore inappropriate for the study of living human subjects in a clinical environment. With the advent of MR imaging, specially diffusion weighted imaging, *in vivo* **dissection** of the white matter became possible.
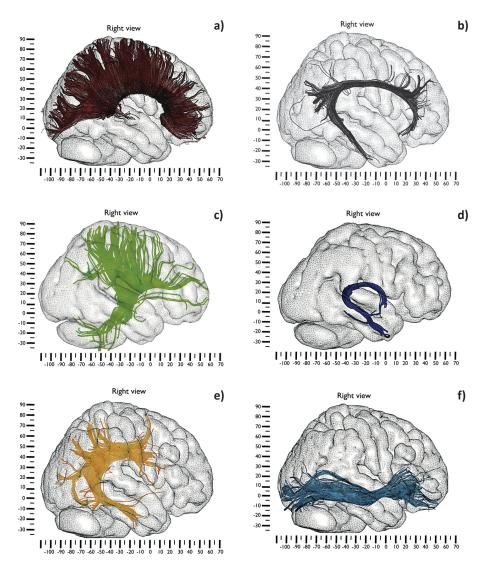
Figure 2.14: In vivo dissection using diffusion tensor streamline tractography. Major white matter tracts are shown dissected in accordance with anatomical knowledge: a) *corpus callosum*, b) *cingulum*, c) *corona radiata*, d) *fornix*, e) *arcuate fasciculus*, f) *Inferior fronto-occipital fasciculus*. Adapted from Catani et al. [21]

Catani et al. [20] conducted a remarkable study where they performed the vir-

tual dissection of several white matter structures such as *corpus callosum*, *superior longitudinal fasciculus*, *cingulum*, and the *fornix*, among others. Although no quantitative validation was performed, they showed the reconstructions of several known fiber bundles. The study was conducted over a DTI image acquired with a 1.5 T scanner, with typical settings: 64 gradient directions, b-value $1300\,\mathrm{s/mm^2}$ and with isotropic $(2.5 \times 2.5 \times 2.5\,\mathrm{mm^3})$ resolution. The total acquisition time was 14 min. In this DTI image, a deterministic fiber tracking algorithm was employed, in order to reconstruct the fiber bundles starting at defined regions - the "seedpoints". These regions were carefully set based on classical neuroanatomical works. Afterwards, the obtained bundles were pruned in order to get rid of unwanted fiber bundles.

Figure 2.14 shows some of these structures reconstructed using diffusion tensor imaging tractography, as in Catani's atlas [21].

Following this work, Wakana et al. [136] constructed an atlas of white matter based on tractography, showing the reconstructions of several structures. Later they performed a quantitative analysis showing a higher reproducibility [137].

This virtual *in vivo* dissection through tractography produces very interesting results that can be used to understand brain's function and to aid in the diagnosis and treatment of several disorders of neurological nature. However, it obviously requires considerable expertise knowledge in defining the proper seeding regions to obtain the desired fiber bundles, and then pruning the results to discard unwanted bundles. This cumbersome, time consuming, and prone to errors process, calls for an automatic *in vivo* virtual dissection methodology.

2.6.1   Clustering

Several authors proposed fiber clustering methods to automatically group individual fiber bundles into coherent tracts. There are two main issues to deal with when clustering fibers. First is which clustering method to use, and second how to assess similarity between fibers.

Moberts et al. [94] presented a framework to evaluate the various clustering algorithms and conducted an interesting comparative survey aiming at physician's quality criteria.

Several fiber clustering algorithms have been proposed. Corouge et al. [28] use a method that propagates cluster labels from fiber to neighbour fiber, assigning each unlabeled fiber to the cluster of its closest neighbour, if it is below a certain threshold. Shimony et al. [122] use a fuzzy c-means algorithm where a fiber is assigned to a cluster based on a confidence function. Zhang and Laidlaw [149] use a hierarchical approach. It starts from initially individual clusters, and at each stage the algorithm groups the two most similar clusters. From a hierarchical

clustering algorithm a dendogram is constructed, and the number of resulting clusters is defined by the parameter: at which level of the dendogram is cut. Several authors use spectral manifold learning techniques in order to produce a mapping from the fiber tracts to a high dimensional Euclidean space, where regular clustering algorithms are applied [16, 100, 129]. Several approaches also incorporate *a priori* knowledge of anatomical structures such as [85, 139].

A key factor in these algorithms is the choice of similarity measure between clusters. Most fiber similarity measures are based on the Euclidean distance between some parts of the fibers. Corouge et al. [28] defines distances based on pairs of each point in a fiber to the closest point on the other fiber. According to Brun et al. [16] two fibers are similar if their start and end points are near. Zhang and Laidlaw [149] define a distance based on the average distance from any point of the shorter fiber to the closest point on the longer fiber. More recently, Wasserman et al. [139] devised a framework using the inner product based on Gaussian processes, between fibers. This metric facilitates the combination of fiber tracts, and operations like tract membership to a bundle or bundle similarity.

Although these algorithms present interesting results in delineating white matter bundles, they involve some choices of measures and thresholds, that will affect their outcome. Atlases can be used as priori knowledge, and thus avoiding user's bias, however this knowledge does not always apply. In an unhealthy brain we do not know what to expect, and prior assumptions can lead to undesired results. Furthermore, one key shortcoming of clustering algorithms is the fact that they operate over derived structures - the reconstructed fiber bundles. Clustering results are thus intimately dependent on the choice of fiber tracking algorithm and parameters.

2.6.2 Segmentation

An alternative to clustering fibers is the direct segmentation of the image into volumetric regions. The assumption here is that tensors will belong to the same bundle if they are similar to each other.

Several algorithms have been proposed over the past years for the segmentation of tensor fields. Zhukov et al. [151] proposed a level-set method over a scalar field derived from anisotropy measures. However this method fails to distinguish between regions with same anisotropy but different direction.

Level-set methods using the full tensor information have been proposed by Zhizhou and Vemuri [138] and Rousson et al. [118], however, these iterative gradient descent based solutions seek a local solution and therefore are highly sensitive to initialization and parameter settings.

Watershed based methods, such as proposed by Rittner and Lotufo [91], are well known by their over-segmentation results. More recent and more efficient methods like the globally optimal graph-cuts have been applied to DTI by Weldeselassie and Hanarneh [141], however they provide a binary partition of the data, into one object and the background.

More recent work, such as Niethammer et al. [98], focusses on the specific problem of segmenting a tubular structure such as the *cingulum*.

As previously stated, several segmentation techniques require the notion of **homogeneity** within a tensor field, i.e., a measure which indicates when a tensor is considered to be similar enough to belong to the same group. Clearly, the segmentation results are highly dependent on the choice of measure, independently of the used segmentation method. So here again, the problem of how to define a distance between DT imposes itself.

2.7 Summary
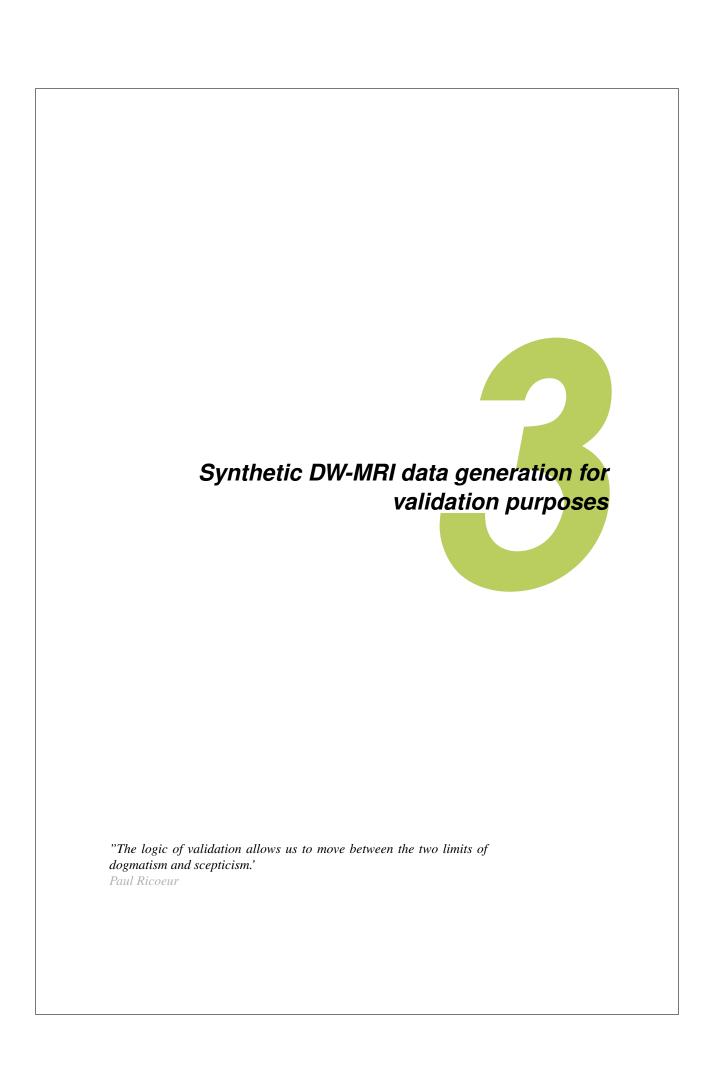
DW-MRI and tensor-based tractography have been proven capable to provide valuable biomarkers for a wide range of applications from characterizing brain disorders and contributing to their diagnosis, to analyse the differences on white matter and consequences in brain function. These procedures usually involve the virtual dissection of white matter tracts of interest. The manual isolation of these bundles requires a great deal of neuroanatomical knowledge and can take up to several hours of work.

The connectivity of the full brain can be extracted, however the amount of data to visualise is such that cluttering problems must be dealt with. Furthermore, as MRI technology progresses, higher resolution data will be available, hence this problem will acquire greater importance.

Several clustering techniques and segmentation techniques have been introduced with success in different application domains. However the automatic identification of white matter structures remains a difficult problem. The main problem lies in the fact that it is a task of the user to choose thresholds, similarity measures, parameters, all depending on the particular task at hand - *how is the dataset acquired and in which bundles is he/she interested?*

This thesis focuses on the development of techniques able to automatically perform the identification of white matter structures. To segment such structures in a tensor field, the similarity of diffusion tensors must be assessed for partitioning data into regions, which are homogeneous in terms of tensor characteristics. This concept of tensor **homogeneity** is explored in order to achieve new methods for segmenting, filtering and enhancing diffusion images.

Synthetic DW-MRI data generation for validation purposes

"The logic of validation allows us to move between the two limits of dogmatism and scepticism.'

Paul Ricoeur

Contents

3.1 Overview

In this chapter, we cover some basic techniques for creating synthetic datasets in order to validate modelling or processing techniques for DW-MRI data. Furthermore, a comparison and evaluation of the similarities between the generated synthetic, hardware phantom and real data is performed. The findings of this work can be used as a guideline for researchers when selecting the most appropriate synthetic data model for evaluating their research.

3.2 Introduction

There is a wide range of utilizations of DTI and newly developed HARDI techniques, from depicting the local structure of the probability density function of the water diffusion, segmentation of white matter structures, to the regularization schemes for the noisy local and global structures. Thorough validation of these methods is needed to fully evaluate their value. There are two main validation strategies: synthetic (either by software simulations or by hardware phantoms) or independent anatomical data. In simulations and phantoms the "ground truth" is known since it is defined upon creation of the artificial image. However, these

datasets hold an over-simplified version of the true neuronal environment. In contrast, in a real image of brain's white matter, we need to derive a "gold standard", which, given the nature of the data, is difficult to obtain.

In literature, two methods are most often used for generating artificial DW-MRI data: the multi-tensor model and the Söderman and Jönsson's model of a restricted diffusion inside a cylinder. In this chapter we analyse these strategies and compared them to acquired hardware phantom data and *in-vivo* data.

3.3 Data

3.3.1 Synthetic data generation

The purpose of the software simulations is to construct mathematical models of the neuroanatomical environment and its DW-MRI acquisition, in order to generate a synthetic image where the underlying fiber configuration is known.

The multi-tensor model was used by, e.g., Pierpaoli et al. [110]; Alexander et al. [2], Hess et al. [59], Tuch [130], Alexander and Barker [1] and Descoteaux et al. [34]. Söderman and Jöhnson's [123] model of restricted diffusion inside a cylinder was used by, e.g., Özarslan et al. [104], Barmpoutis et al. [8], von dem Hagen et al. [135] and others. Next, we describe both methods in detail.

3.3.1.1 Multi-tensor model

The multi-tensor model assumes that single fiber responses can be described by a second order tensor, and that the tissue in a voxel with more than one fiber is simply composed by a superposition of multiple second order tensors that describe the underlying probability density function under assumption of no exchange of molecules between the different fiber compartments.

For a given set of gradient directions $\mathbf{g}_i$, in our case generated by the static repulsion algorithm [70], the corresponding signal value will be:

$$S(b, \mathbf{g}_i) = \sum_{k=1}^{n} p_k \exp^{b\mathbf{g}_i^T \mathbf{D}_k \mathbf{g}_i} \qquad \boxed{3.1}$$

where b is the b-value and p_k are the different weights for the diffusion tensors $\mathbf{D}_k$ that compose the signal by the Stejskal-Tanner equation, introduced in chapter 2.

In this chapter, we investigate the properties of two fiber bundles crossing at certain angle. Therefore, we set the parameters according to the different areas of the domain we want to simulate:

- **For the crossing voxels:** we set the weights to be equal $p_1 = p_2 = 0.5$. For the second order tensor $\mathbf{D}_k$, we fix the eigenvalues to $[\lambda_1, \lambda_2, \lambda_3] =$
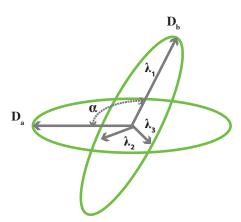
Figure 3.1: Illustration of two fibers crossing at angle α simulated by the multi-tensor model.

$[300, 300, 1700] \times 10^{-6} \text{ mm}^2/\text{s}$. We rotate the tensors such that they form the required angle of crossing;

- **For the single fiber voxels:** We set $n = 1$, $p_1 = 1$ and we use the same eigenvalues for the tensor $\mathbf{D}_1$ as in the previous scenario;

- **For the isotropic voxels:** we simulate the signal by fixing the eigenvalues to $[\lambda_1, \lambda_2, \lambda_3] = [700, 700, 700] \times 10^{-6} \text{ mm}^2/\text{s}$.

The specific choice of values for the eigenvalues is according to eigenvalues of tensors found in anisotropic regions of human brain as reported in the work of Descoteaux et al. [34].

3.3.1.2 Söderman and Jöhnson's model

This model is based on the MRI signal attenuation from molecules diffusing restricted inside a cylinder of radius ρ and length L, with the free diffusion coefficient given by D_0. For crossing fibers configurations, presence of more than one cylinder is assumed and the resulting signal attenuation is averaged from the independent signals within each cylinder. Multiple fiber orientations can be modelled assuming that the diffusing molecules are constrained within these cylinders with no possibility for exchange between the cylinders. The simulations from this model employ the exact form of the MR signal attenuation from particles diffusing inside cylindrical boundaries. Therefore it has the advantage over the multi-tensor model that does not enforce mono-exponential decay.

The signal attenuation is given by:

$$S(\rho, q, \vartheta, \Delta) = \sum_{n=0}^{\infty} \sum_{k=1}^{\infty} \sum_{m=0}^{\infty} \frac{2 K_{nm} \rho^2 (2\pi q \rho)^4 \sin^2(2\vartheta) \gamma_{km}^2}{[(n\pi\rho/L)^2 - (2\pi q \rho \cos\vartheta)^2]^2}$$

$$\times \frac{[1 - (-1)^n \cos(2\pi q L \cos\vartheta)] [J_m'(2\pi q \rho \sin\vartheta)]^2}{L^2 [\gamma_{km}^2 - (2\pi q \rho \sin\vartheta)^2]^2 (\gamma_{km}^2 - m^2)}$$

$$\times \exp\left(-\left[\left(\frac{\gamma_{km}}{\rho}\right)^2 + \left(\frac{n\pi}{L}\right)^2\right] D_0 \Delta\right) \qquad \boxed{3.2}$$

Where J_m is the m-th order Bessel function, γ_{km} is the k-th solution to $J_m'(\gamma) = 0$ with the convention $\gamma_{10} = 0$, and $K_{nm} = \delta_{n0}\delta_{m0} + 2[(1 - \delta_{n0}) + (1 - \delta_{m0})]$, ϑ is dependent on the gradient vector $\mathbf{g}_i$, since it is the angle between the cylinder and the applied diffusion gradient, and $q = |\mathbf{q}|$ is the magnitude of the q-space vector $\mathbf{q} = \gamma \delta \mathbf{g}_i / 2\pi$, where γ is the gyromagnetic ratio, and can be re-expressed as: $q = \sqrt{\frac{b}{4\pi^2(\Delta - \delta/3)}}$.

Söderman and Jöhnson's model is significantly more complex and computationally expensive than the multi-tensor model. However, it models the MRI signal decay more accurately, especially at high b-values. On the other hand, the multi-tensor model is fast, but simply models the signal based on a second order tensor and the assumption of mono-exponential decay.

We fix the parameters required for Söderman and Jöhnson's's equation, as in the work of Özarslan et al. [104] to $L = 5$ mm, $\rho = 5$ μm, $D_0 = 2.02 \times 10^{-3}$ mm^2/s, and truncate the series at $n = 1000$ and $k, m = 10$.

3.3.1.3   Noise simulation

For more realistic simulations, noise is added to the simulated DW-MRI signal from equation 3.1 or equation 3.2.

This signal is considered to be the magnitude of a complex signal and the noise is the result of Gaussian noise with standard deviation of $\sigma = S_0/\zeta$ added to the real and imaginary parts of the signal. That is, to each $S(\mathbf{g}_i)$, we add a random complex number with real and imaginary parts coming from a zero mean Gaussian distribution $G(\sigma)$ with $\sigma = 1/\zeta(S_0 = 1)$ and take the modulus to obtain the noisy synthetic data. A typical value is $\zeta = 20$ so that the unweighed signal S_0 has a signal-to-noise ratio (SNR) of 20.

3.3.1.4   Synthetic data fields

For the validation of processing algorithms (i.e. filtering, enhancement of diffusion profiles, noise removal and others), using both of the above described mod-

els, we create two fiber bundles crossing under different angles. The different fiber bundles in the volume are parametrically defined by specifying the center line of the bundle, seen as a cylinder with a given radius r (see figure 3.2). In each voxel of the volume, the respective fiber(s) orientation(s) is taken in order to simulate the respective signal attenuation.
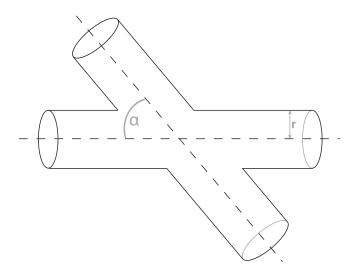
Figure 3.2: Illustration of tubes crossing at angle α and radius r.

We create the crossing tubes under angles of $30°, 50°$ and $65°$ with 72 and 120 number of gradients and b-values, SNRs and other parameters (mainly required for Söderman and Jöhnson's model) as in table 3.1. We used the same fiber configurations and parameters as in the acquired hardware phantom described in the following section.

parameters				
b	(s/mm^2)	1000	2000	4000
Δ	(ms)	32.44	39.18	47.18
δ	(ms)	25.34	32.08	40.08
t	(ms)	24.0	28.5	33.8
snr		15.3	13.3	11.9

Table 3.1: Parameters of our acquisition protocol.

3.3.2 Hardware phantom

Hardware phantoms aim at constructing a physical model of the neuroanatomical environment of different fiber populations, specially fiber crossings, in order to generate a real DWI dataset where the underlying ground truth is known. The benefits of employing hardware phantoms over the synthetic data are mainly due to the real MRI acquisitions which include real noise without the need for additional noise modelling.

The phantoms used in this work are constructed following a method described by Pullens et al. [114]. A bundle of yarn strands, each consisting of 22, 10 μm fibers, is wrapped into a shrink wrap tube. Crossing phantoms were constructed by interdigitating 25 bundles of 400 yarns, such that they form an 'X' shape, as in figure 3.3. The phantom was transferred to a 0.03 g/l MnCl*4H2O solution and shrunk by heating the solution to 95 °C. For resistive coil loading, 2.4g/l NaCl was added. The setup is shown in figure 3.4, where fiber bundles of three sets of crossings under angles of $30°, 50°, 65°$ are presented.

The hardware phantom data was acquired with exactly the same acquisition parameters as the *in-vivo* data, and therefore presents similar noise levels. The verisimilitude of this phantom is further increased by embedding it in a solution that produces a $T2$ image similar to the real acquisition. However, due to the high thickness of the fibers and their loose packing, the anisotropy is lower than in biological tissue. Moreover, it is important to keep in mind that the phantoms are still a highly over-simplified version of the brain white matter fibrous tissue.
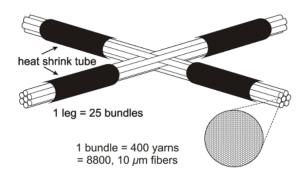
Figure 3.3: Schematic overview of the phantom construction [114].

3.3.3 In-vivo human brain data

Several datasets were generously acquired by Alard Roebroeck, Maastricht University, for us to evaluate the different methods for data generation. This data has also been used to validate the new processing algorithms addressed in this thesis.
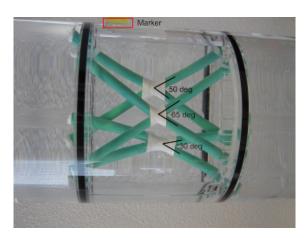
Figure 3.4: The setup: container with the phantoms [114].

DW-MRI acquisition was performed on a volunteer VP (25 years, female), using a twice refocused spin-echo echo-planar imaging sequence on a Siemens Allegra 3T scanner (Siemens, Erlangen, Germany), with FOV 208×208 mm and voxel size $[2.0 \times 2.0 \times 2.0]$ mm^3. Informed consent was obtained prior to the measurement. Ten horizontal slices were positioned through the body of the *corpus callosum* and *centrum semiovale*. Datasets were acquired with the gradient sampling schemes generated with static repulsion algorithm with the diffusion-weighted volumes interleaved with zero weighted volumes every 12^{th} scanned diffusion gradient directions. Gradient schemes of 72 and 120 directions were used in combination with b-values of 1000, 2000, 4000 s/mm^2, and gradient pulse duration δ, and gradient spacing Δ, given in table 3.1. In the same session, an anatomical data set (192 slices, voxel size $1 \times 1 \times 1$ mm^3) was acquired using the ADNI-MPRAGE protocol (i.e. a T1-weighted 3D anatomical scan).

3.3.4   Fiber Cup hardware phantom

The Fiber Cup diffusion phantom [112] was specially designed for the benchmarking of tractography techniques. It was constructed as to resemble a coronal slice of a human brain, and therefore holds different fiber tracts crossing at different angles. Figure 3.5 depicts this configuration.

In order to create large bundles, hydrophobic acrylic fibers with diameter of the same order as myelinated axons were used. A polyurethane positive and negative mould of the target bundles were manufactured in order to strongly tighten the fibers together. Several layers of fibers were interleaved and stacked to build the different fiber crossing configurations. The density of fibres was close to 1900 fibers/mm^2.

Since the standard DW-MRI acquisition scheme is extremely sensitive to phase inhomogeneities, great care was taken to insure the absence of any air bubbles. A special filling process under vacuum was devised and an ultrasound beam is used to destroy any remaining air bubbles. The container was filled using pure water.

The DW-MRI images were acquired with a 3T Tim Trio MRI systems of the NeuroSpin centre, equipped with a whole body gradient coil (40 mT/m, 200 T/m/s), and using a 12-channel receive only head coil, in combination with the whole body transmit coil of the MRI system.

Two datasets were acquired at two different spatial resolutions: 3 mm isotropic and 6 mm isotropic. The settings were as follows:

- for the 3 mm isotropic acquisition: field of view FOV $=$ 19.2 cm, matrix 64×64, slice thickness TH $=$ 3 mm, repetition time TR $=$ 5 s, 2 repetitions, read bandwidth RBW $=$ 1775 Hz/pixel, partial Fourier factor 6/8, parallel reduction factor GRAPPA $=$ 2. Three diffusion sensitizations at b-values $b = 650, 1500, 2000$ s/mm^2 corresponding to the echo times TE $= 77, 94, 102$ ms respectively.

- for the 6 mm isotropic acquisition: field of view FOV $=$ 38.4 cm, matrix 64×64, slice thickness TH $=$ 6 mm, repetition time TR $=$ 5 s, 1 repetition, read bandwidth RBW $=$ 1775 Hz/pixel, partial Fourier factor 6/8, parallel reduction factor GRAPPA $=$ 2, . Three diffusion sensitizations at b-values $b = 650, 1500, 2650$ s/mm^2 corresponding to the echo times TE $= 77, 94, 102$ ms respectively.

The diffusion sensitization was applied along a set of 64 orientations, uniformly distributed over the sphere.

This phantom provides a rich architecture of fiber tracts crossing, splaying and also curve shaped fibers. However it presents some limitations, specially in its artificiality. In order to increase the SNR level, the voxel size had to be too coarse (3 and 6 mm isotropic). Furthermore, given the used fiber thickness (20pm), the anisotropy is lower than in reality.

3.4 Analysis

3.4.1 Maxima detection

It is generally assumed that the maxima of the reconstructed profiles by the HARDI techniques are a good estimate of the underlying fiber direction [88]. Therefore, calculating the maxima is an important step for evaluation. Since these profiles are functions described by the SH, the maxima extraction is challenging due to the difficulties in finding analytical solutions.

Figure 3.5: The Fiber Cup [112] ground truth, with 16 fiber candidates highlighted with different colours.

In our implementation, maxima detection is performed using a finite difference algorithm on a discrete grid on the reconstruction profiles, which are basically functions on a sphere. This grid is created by tessellating an icosahedron by a certain order that dictates the total number of points on the spherical mesh. If a mesh point is above the threshold and it is above all its neighbours, the mesh point direction is kept as a maxima. This threshold avoids selecting small peaks that may appear due to noise. In each of these regions, the local maximum is determined, as illustrated in figure 3.6.

This method is straightforward and easy to implement, but can be inaccurate if the tessellation order is chosen too low. A high tessellation order improves the accuracy, but also increases the computational cost. Another drawback is the *ad-hoc* selection of the threshold, that is, if the threshold is too high some of the maxima will not be detected as illustrated in figure 3.6.

3.4.2   Data analysis

In this study we compare:

- The multi-tensor model (MT) section 3.3.1.1;

- The Söderman and Jöhnson's model (S) section 3.3.1.2;
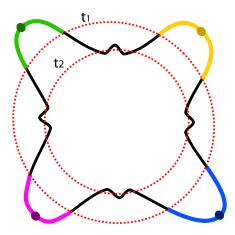
Figure 3.6: The maxima algorithm uses a threshold to limit the region growing on the spherical function. This leads to isolated regions from which the local maxima are extracted. With threshold t_1 the 4 maxima are correctly extracted, whereas with threshold t_2 the spurious noisy peaks are incorrectly extracted.

- A hardware phantom (HW) 3.3.2;

- An *in-vivo* human brain dataset (H) 3.3.3.

Simulation Data: The synthetic tubes of three crossings under $30°, 50°, 65°$, were analysed by defining two regions of interest (ROI) in the crossing and single fiber areas.

Hardware phantom: The hardware phantom contains the 3 different crossing configurations, $30°, 50°, 65°$ as illustrated in figure 3.4. Two ROIs were manually defined (similar to the ROIs in the synthetic tubes), to capture the crossing voxels, and the voxels in the bundle's legs, i.e. in the single fiber regions.

In-vivo human brain dataset: ROIs were manually defined to capture crossing voxels between the *corpus callosum* and the *corona radiata*, and also voxels in the base of the *corpus callosum* as linear shaped diffusion voxels.

For the analysis we define ROIs of about 60 voxels each, in the crossing and linear areas of the computer generated synthetic, hardware phantom and real data, all simulated or scanned under identical conditions.

3.4.3 Reconstruction techniques and measures

From the simulated and acquired signals, we estimate diffusion tensors (DT) and Qballs with Laplace Beltrami smoothing with $\lambda = 0.006$ as presented in the work of Descoteaux et al. [35]. From these reconstruction profiles, we calculate several scalar measures on the four types of data. From the diffusion tensors we calculate fractional anisotropy *FA* and mean diffusivity *MD* [134], C_l and C_p [143], as well as *K1*, *K2* and *K3* [74]. From Qball we calculate generalized anisotropy *GA* [103] and fractional multi-fiber index *FMI* [**?**]. We also quantify the angular difference of the simulated linear direction and estimated main eigenvector in the DTI analysis, as well as the angular error and the standard deviation of the simulated and recovered crossing angles in Qball.

Additionally we give a qualitative analysis on the similarity of the profiles by the four different types of data.

3.5 Results

3.5.1 Quantitative results

In figure 3.7 the average normalized signal values for the analysed voxels in the different datasets in the linear and crossing part are summarized. The signal decays with increasing b-value in most of the cases. Only in the hardware phantoms the signal values start to increase at $b = 4000\,\mathrm{s/mm^2}$, suggesting that scanning under high b-values is not recommended for this type of data. To avoid clutter in presenting the results, in figure 3.8 we only summarize the results of the average values for DTI and HARDI scalar measures. For the DTI measures we use typical b-values of $1000\,\mathrm{s/mm^2}$ and for the HARDI measures significantly higher b-values of $4000\,\mathrm{s/mm^2}$ as recent literature in HARDI suggests. Given that the results for 72 and 120 gradient directions were essentially the same, we present only the latter and only the datasets with $65°$ crossing fibers since qualitatively they are the most similar to the selected region of the *centrum semiovale* from the *in-vivo* brain data. For the HARDI measures we use 4^{th} order of the SH coefficients calculated from the Qball model.

The results in bold, in figure 3.8 are the ones most similar to the real data. We observe good similarity between the real and the multi tensor (MT) model in many of the results. The results from the Söderman and Jöhnson (S) model become similar to the *in-vivo* ones at high b-values, which can be useful for the analysis of HARDI data.
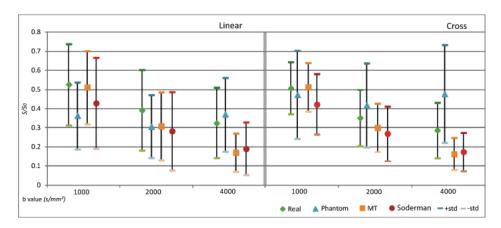
Figure 3.7: Average signal values and standard deviation at different b-values and in different synthetic, hardware phantom and real data configurations. We observe signal decay with the increase of the b-values in all the datasets and in both linear and crossing areas.

b value (s/mm²)		Linear				Cross			
		real	phantom	MT	S	real	phantom	MT	S
1000	FA	**0.80**	0.40	**0.80**	0.88	**0.37**	**0.43**	0.55	0.67
	MD	**0.0008**	0.0012	**0.0008**	0.0011	**0.0008**	0.0010	**0.0007**	0.0009
	CI	**0.59**	0.21	**0.60**	0.71	**0.10**	**0.20**	0.27	0.37
	Cp	**0.14**	**0.15**	0.05	0.05	0.31	0.23	0.27	0.28
	K1	**0.0023**	0.0037	**0.0023**	0.0032	0.0023	0.0030	0.0021	0.0028
	K2	**0.0011**	0.0007	**0.0011**	0.0019	0.0004	0.0006	0.0006	0.0011
	K3	**0.94**	0.53	**0.99**	0.99	-0.37	0.29	0.52	0.68
	Ang. Diff	-	17.58	**1.61**	1.77	-	-	-	-
4000	GA	**0.44**	0.07	**0.49**	0.63	**0.18**	**0.08**	0.32	0.41
	FMI	0.16	1.25	0.22	0.26	**0.88**	**0.96**	0.26	0.26
	Ang. err;std	-	-	-	-	-	13.58;7.28	**10.15;5.34**	10.6;7.76

Figure 3.8: Results from DTI and HARDI measures at different b-values and in different synthetic, hardware phantom and real data configurations.

3.5.2   Qualitative results

From the qualitative observation of the data (see figure 3.9), we conclude that the noise in the hardware phantoms is more prominent compared to the real data. The synthetic data however, has an underestimation of the noise.

3.6   Conclusion

In this chapter we presented an analysis of different artificial DWI phantom datasets compared to *in-vivo* brain data scanned under the same acquisition paramet-
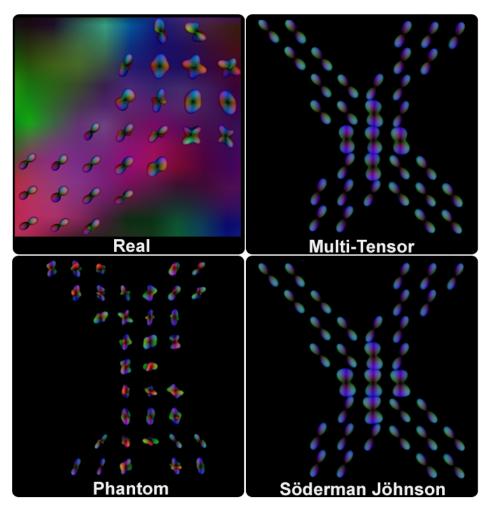
Figure 3.9: Qball glyphs with 4^{th} order of SH representation for the different datasets with acquisition parameters $b = 1000\,\text{s/mm}^2$, $NG = 120$.
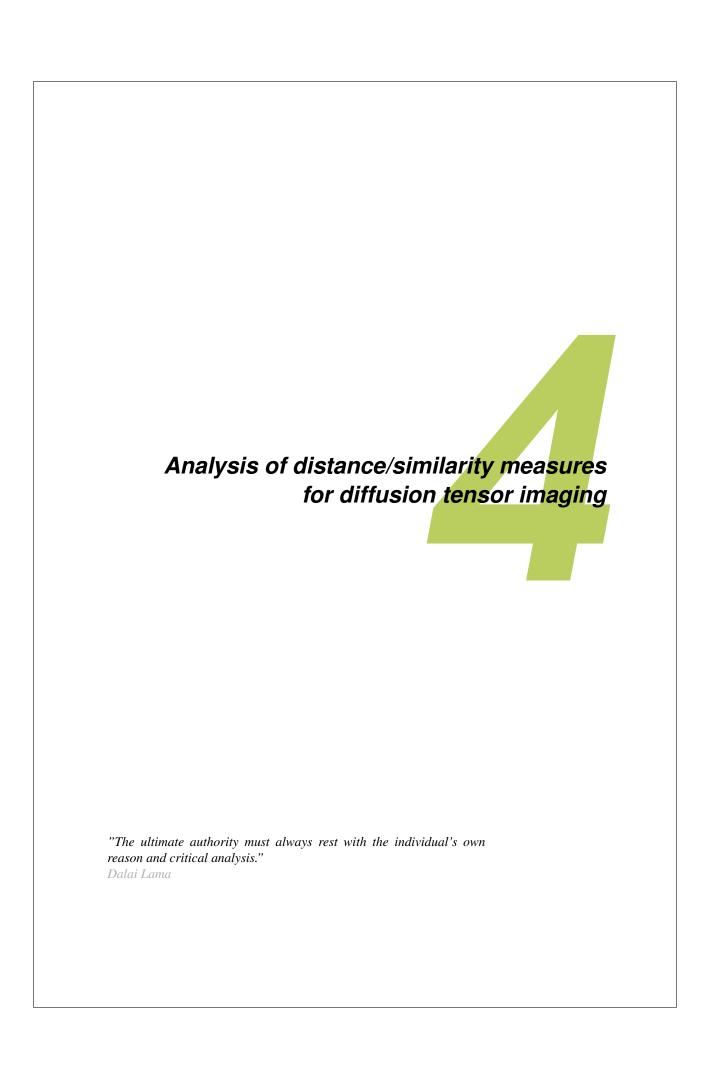
ers. Two different gradient sampling schemes and three different b- values were used. Furthermore, the additional errors imposed by the models for synthetic data generation were examined. It appears that the Söderman and Jöhnson's model imposes less errors compared to the multi-tensor model. However, in the analysis of the measures derived from the DTI data, in most of the cases the multi-tensor model exhibited similar behaviour to the *in-vivo* data at the same b-value. Additionally, this model is simpler and has low computational costs. Together, it is a good choice for software synthetic data generation for validation of DTI-based algorithms. The qualitative results suggest that the noise level for the selected b-

values is much lower in the Söderman and Jöhnson's model and the multi-tensor model than for the scanned data. In this analysis, only one specific region of crossings in the *in-vivo* brain was examined and these results might vary in other selections. Future work should address a different range of acquisition parameters and different selections of regions from *in-vivo* data. The Söderman and Jöhnson's model is frequently used for rat studies in high-field pre-clinical scanners, which allows for a significantly shorter gradient pulse duration δ compared to a clinical scanner. Future work should investigate this type of real data and compare to the synthetic simulations generated by the Söderman and Jöhnson's model. Future work should evaluate recent methods for generating fiber crossing phantoms out of real acquisition data, such as Caan et al. [18]. At the end of this chapter, we would like to point out that as a common practice, throughout this thesis, the validation procedure is as follows:

1. Initial evaluation against a dataset, synthetically generated using the MT or the Söderman and Jöhnson's model

2. Proof-of-concept evaluation against a hardware phantom dataset (optional)

3. Qualitative evaluation using *in-vivo* dataset.

Analysis of distance/similarity measures for diffusion tensor imaging

"The ultimate authority must always rest with the individual's own reason and critical analysis."
Dalai Lama

Contents

4.1 Overview

Many different measures have been proposed to compute similarities and distances between diffusion tensors. These measures are commonly used for algorithms such as segmentation, registration and quantitative analysis of DTI datasets. The results obtained from these algorithms are extremely dependent on the chosen measure. The measures presented in literature can be of complete different nature, and it is often difficult to predict the behaviour of a given measure for a specific application. In this chapter, we classify and summarize the different measures that have been presented in literature. We also present a framework to analyse and compare the behaviour of the measures according to several selected properties. We expect that this framework will help in the selection of a measure for a given application and to identify when the generation of a new measure is needed. This framework will also allow the comparison of new measures with existing ones.

The work presented in this chapter was done in co-authorship with T.H.J.M. Peeters [?].

4.2 Introduction

In chapter 2, we showed various ways of processing and visualizing diffusion tensor imaging data. Several algorithms can be used to extract the white matter fiber tracts or the fibrous structure of the heart tissue. Numerous studies try to identify the subtleties of neurological disorders or to investigate the brain's development by evaluating differences across subjects. These multitude of applications and procedures usually involve the virtual dissection of white matter tracts of interest [118, 138, 151] and also to register different DTI data sets [3, 92, 148]. It is often also necessary to derive statistical properties of diffusion tensors to identify differences, e.g., between healthy and pathology areas [89]. In all these methods, there is the common need to define the (dis)similarity between tensors, i.e., to define homogeneity in a tensor field.

In segmentation and registration, similarity measures are applied to match DTs in voxels in a certain region, and between regions of different data sets. In quantitative analysis or DT statistics, distance or similarity measures of DTs in neighbouring voxels can be used to classify the amount of variability in a selected voxel [109] or volume of interest. The results of these applications are highly dependent on the choice of measure.

Alexander et al. [3] listed several measures and analysed their results for segmentation. However, since then, various people have introduced new measures for comparing DTs. These measures are of different nature and it is difficult to

predict which measure will give better, or similar results. Numerous measures exist, and there is a need for an overview that compares and classifies them in a structured way. This comparison can help to support researchers in choosing an appropriate measure, and being able to predict the behaviour of the measures for their concrete application.

In this chapter, we provide this analysis and improve the intuition in the behaviour of the measures. The intrinsic characteristics of a measure are analysed without having a specific application in mind. This allows an evaluation of the nature of the measure in itself. It is beyond the scope of this chapter to make an application-oriented analysis, (e.g., finding the best measure for DTI adult brain registration). However, this chapter aims to help in making a first selection of the possible measures that could be used for such applications by looking at the characteristics of the problem and the characteristics of the measures. We expect that it will also help to identify when a new measure is necessary, and compare its behaviour with existing ones.

First, we present the notations used in this chapter. In section 4.4, we describe the properties that will be used for the analysis of the measures. In Section 4.5, we give an overview of existing measures from literature. In section 4.6, we explain how we evaluate the properties of the measures and show some simple results to illustrate our methods. Section 4.7 presents the results of the experiments. Finally, in section 4.8, we summarize and discuss the results of this chapter.

4.3   Notation

In this chapter we often work with equations that include more than one diffusion tensor. Therefore, we extend the previously defined notation. Diffusion tensors are denoted by capital bold letters, for example, $\mathbf{D}, \mathbf{A}, \mathbf{B} \in Sym_3^+$. Eigenvalues of tensor $\mathbf{D}$ are $\lambda_1^{\mathbf{D}} \geq \lambda_2^{\mathbf{D}} \geq \lambda_3^{\mathbf{D}} \geq 0$ and the corresponding eigenvectors are $\mathbf{e}_1^{\mathbf{D}}, \mathbf{e}_2^{\mathbf{D}}$ and $\mathbf{e}_3^{\mathbf{D}}$. We will denote the trace $(\sum_{i=1}^3 \mathbf{D}_{ii}) = (\sum_{i=1}^3 \lambda_i^{\mathbf{D}})$ of $\mathbf{D}$ with $\mathrm{tr}(\mathbf{D})$. The determinant of $\mathbf{D}$ will be denoted by $|\mathbf{D}|$.

With *measure* we refer to a function m that has two tensors $\mathbf{A}, \mathbf{B}$ as input, and returns a non-negative scalar value:

$$m : Sym_3^+ \times Sym_3^+ \mapsto \mathbb{R}_0^+. \qquad \boxed{4.1}$$

If a measure returns a larger value for more similar $\mathbf{A}$ and $\mathbf{B}$, then we call the measure a *similarity measure*. If it returns a larger value for less similar $\mathbf{A}$ and $\mathbf{B}$, we call it a *distance measure*. We will denote similarity measures with s and distance measures with d.
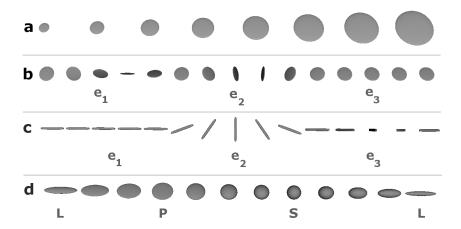
Figure 4.1: Ellipsoidal glyphs showing smoothly varying diffusion tensors (DTs). (**a**) a DT with planar shape where the size is increased smoothly, i.e., MD increases; (**b**) a DT with planar shape is rotated smoothly, until π, first around $\mathbf{e}_1$, then around $\mathbf{e}_2$ and finally around $\mathbf{e}_3$; (**c**) a DT with elongated shape is rotated smoothly as in (b); (**d**) DTs where the shape changes from elongated (L), to planar (P), to spherical (S) and back to elongated (L).

4.4 Properties

In this section, we present a list of properties that can be evaluated for the different measures. Diffusion tensors can be classified by their size, orientation and shape. We evaluate the measures according to their sensitivity to changes in these properties. These changes are illustrated in figure 4.1. We also include as properties how robust the measures are to noise, and whether a measure is a metric or not.

4.4.1 Size

We understand as the size of a DT the mean diffusivity $MD = \mathrm{tr}(\mathbf{D})/3$. This is illustrated in figure 4.1(a). We consider a measure to be size-invariant if it is invariant to isotropic scaling, i.e., if it fulfils:

$$m(t\mathbf{A}, u\mathbf{B}) = m(\mathbf{A}, \mathbf{B}), \qquad \boxed{4.2}$$

where t and u are positive scalar values.

4.4.2 Orientation

A measure m is rotation invariant, if the value of m does not change when the input tensors are rotated:

$$m(R^T \mathbf{A} R, P^T \mathbf{B} P) = m(\mathbf{A}, \mathbf{B}), \qquad \boxed{4.3}$$

where R and P are rotation matrices. The orientation invariance can be divided in two parts. One is whether the measure is sensitive, in general, to the difference in orientation between tensors. Orientation changes are illustrated in figures 4.1(b) and 4.1(c).

The other invariance, included in the previous, is invariance to image rotation. If we define a DTI image as $f : \mathbb{R}^3 \mapsto Sym_3^+$ in most of the cases we want our measure to be invariant with respect to rigid body transformations of f (i.e., rotation and translation). In the case of DTI images, the image transformation also has to be applied to the tensor. From these transformations the rotation is the only one that affects the tensor. Being invariant to image rotation means that we want to fulfil equation 4.3 when $R = P$. If the image f is transformed with other non-orthogonal transformations (e.g., non-uniform scaling, skewing), it is not clear how this should affect the tensor and, therefore, we do not consider this further in this chapter.

4.4.3 Shape

The shape of a diffusion tensor can be defined as *elongated*, *planar*, *spherical* (see figure 2.8) or as an interpolation between these types. The shape is given by the ratio between the different eigenvalues. A graphical representation of interpolation between different tensor shapes is shown in figure 4.1(d). A measure m is shape-invariant if the value of $m(\mathbf{A}, \mathbf{B})$ does not change when changing the shape (i.e., the ratio between eigenvalues) of $\mathbf{A}$, $\mathbf{B}$, or both.

4.4.4 Robustness

Measures are never completely insensitive for noise, however, if small changes in the input produce small changes in output, then we consider the measures to be robust under noise. The measure m is robust under noise when:

$$|m(\mathbf{A} + \mathbf{E}1, \mathbf{B} + \mathbf{E}2) - m(\mathbf{A}, \mathbf{B})| \leq \varepsilon, \qquad \boxed{4.4}$$

where ε is a very small scalar value and the components $\mathbf{E}\{1, 2\}_{ij}$ of noise tensors $\mathbf{E}1, \mathbf{E}2 \in Sym_3^+$ have also very small values.

4.4.5 Metric

A distance measure d is a semi-metric if, for two tensors $\mathbf{A}$ and $\mathbf{B}$, it satisfies the following conditions:

$$\mathbf{A} = \mathbf{B} \quad \Leftrightarrow \quad d(\mathbf{A}, \mathbf{B}) = 0, \tag{4.5}$$

$$d(\mathbf{A}, \mathbf{B}) \quad = \quad d(\mathbf{B}, \mathbf{A}). \tag{4.6}$$

Condition 4.5 is important because it allows us to distinguish between equal and non-equal tensors. Condition 4.6 is necessary if we do not want the results to depend on the order in which we deal with the DTs in a volume. If the measure has to be a metric, it also has to fulfil:

$$d(\mathbf{A}, \mathbf{B}) \leq d(\mathbf{A}, \mathbf{C}) + d(\mathbf{C}, \mathbf{B}). \tag{4.7}$$

Condition 4.7 is important in applications where, e.g., one needs to take the mean or do interpolation between tensors [5, 108].

4.5 Measures

In this section, we present a classification of similarity and distance measures for diffusion tensors that have been used in literature. This classification is based on the nature of the derivation of the measure: measures based on scalar indices; measures that make use of the angles between eigenvectors; measures based on linear algebra; measures based on imposing the preservation of positive definiteness of the tensor, i.e., Riemannian geometry; measures considering the DTs as a representation of a probability density function; and, finally, measures that combine different measures from the previous classes.

4.5.1 Scalar indices

Given a scalar index $g : Sym_3^+ \mapsto \mathbb{R}_0^+$, the simplest way to obtain a difference between two DTs $\mathbf{A}$ and $\mathbf{B}$ is by using the absolute difference $|g(\mathbf{A}) - g(\mathbf{B})|$ of the scalar index of the two tensors. There exist numerous scalar indices that can be chosen for g. Two well-known examples are fractional anisotropy (FA) and linear anisotropy (c_l). For a selection of scalar indices, see table 4.1 and refer to Westin et al. [143] and Vilanova et al. [134]. These indices reduce the 6D information in a DT to a scalar value. In the computation of the scalar value, only the eigenvalues, which are rotation invariant, are used, thus they do not depict the directional variation of the diffusion anisotropy. The measures created from

Name	Abbrev.		Equation
Mean diffusivity	MD	$=$	$\mathrm{tr}(\mathbf{D})/3 = (\lambda_1 + \lambda_2 + \lambda_3)/3$
Fractional anisotropy	FA	$=$	$\dfrac{\sqrt{(\lambda_1-\lambda_2)^2+(\lambda_2-\lambda_3)^2+(\lambda_1-\lambda_3)^2}}{\sqrt{2(\lambda_1^2+\lambda_2^2+\lambda_3^2)}}$
Relative anisotropy	RA	$=$	$\dfrac{\sqrt{(\lambda_1-\lambda_2)^2+(\lambda_2-\lambda_3)^2+(\lambda_1-\lambda_3)^2}}{\sqrt{2}(\lambda_1+\lambda_2+\lambda_3)}$
Linear anisotropy	c_l	$=$	$(\lambda_1 - \lambda_2)/(\lambda_1 + \lambda_2 + \lambda_3)$
Planar anisotropy	c_p	$=$	$2(\lambda_2 - \lambda_3)/(\lambda_1 + \lambda_2 + \lambda_3)$
Isotropy	c_s	$=$	$3\lambda_3/(\lambda_1 + \lambda_2 + \lambda_3)$
Volume ratio	VR	$=$	$\lambda_1\lambda_2\lambda_3/MD^3$

Table 4.1: Scalar indices for diffusion tensors [134, 143].

scalar indices will be denoted by ds, with the short name of the index as subscript, e.g. ds_{FA}, ds_{C_l}, ds_{MD}. Thus

$$ds_{FA}(\mathbf{A}, \mathbf{B}) = |FA(\mathbf{A}) - FA(\mathbf{B})|. \qquad \boxed{4.8}$$

When using ds, most information is lost. Each DT is represented by one scalar value, while six scalar values are needed to represent the full DT. Thus, the measures based on scalar indices can be rather limited.

More scalar indices can be derived from tensors. For example, several authors in DTI literature recognized the benefit of *tensor invariants* as measures of the diffusion tensor shape that do not require diagonalization. Kindlmann [73] used these invariants, like the mean, variance and skewness, which are invariant to rotation, to measure the shape gradients in tensor fields. However, using them for constructing a distance measure will give similar results to ds and will not solve the problem that just one aspect is being shown. Thus, we do not treat them separately here.

4.5.2   Angular difference

Angular difference d_{ang_i}, $i \in \{1, 2, 3\}$ of the eigenvectors $\mathbf{e}_i^{\mathbf{D}}$ is often used as a distance between tensors. It measures changes in orientation [152]:

$$d_{ang_i}(\mathbf{A}, \mathbf{B}) = \arccos(\mathbf{e}_i^{\mathbf{A}} \cdot \mathbf{e}_i^{\mathbf{B}}). \qquad \boxed{4.9}$$

Using d_{ang_1} only makes sense for tensors where the diffusion is mainly linear. If the tensors have a planar shape then d_{ang_3} can be used. For tensors with a spherical shape, any d_{ang_i} can be considered random and should not be used.

4.5.3 Linear algebra

A class of measures deals with the diffusion tensor components as vectors elements. A typical distance measure is the L^n-norm of the component-wise difference of two vectors:

$$d_{Ln}(\mathbf{A}, \mathbf{B}) = \sqrt[n]{\sum_{i=1}^{3} \sum_{j=1}^{3} (\mathbf{A}_{ij} - \mathbf{B}_{ij})^n}. \qquad \boxed{4.10}$$

In DTI literature, the L^2-norm, d_{L2}, is most commonly used for computing a distance measure (see Batchelor et al. [13]), therefore, we only treat d_{L2} in this chapter. d_{L2} is the same as the *Frobenius distance* [152] which is computed by $d_F(\mathbf{A}, \mathbf{B}) = \sqrt{\mathrm{tr}((\mathbf{A} - \mathbf{B})^2)}$. One can also compute the *scalar product* of two tensors by summing the products of components of the tensors [3]. The result can be used as a similarity measure s_{sp}:

$$s_{sp}(\mathbf{A}, \mathbf{B}) = \sum_{i=1}^{3} \sum_{j=1}^{3} \mathbf{A}_{ij} \mathbf{B}_{ij}. \qquad \boxed{4.11}$$

Measures s_{sp} and d_{Ln} treat the DTs as simple vectors and ignore the matrix or tensor nature of them. Another class of measures use the fact that we have matrices. Pierpaoli and Basser [109] propose to use the sum of the squared vector dot products of the eigenvectors weighted by the product of the eigenvalues as a tensor scalar product [66]: s_{tsp} includes the collinearity of the orientation of the tensors weighted by their eigenvalues. The value is maximized if the tensors are aligned.

$$s_{tsp}(\mathbf{A}, \mathbf{B}) = \sum_{i=1}^{3} \sum_{j=1}^{3} \lambda_i^{\mathbf{A}} \lambda_j^{\mathbf{B}} (\mathbf{e}_i^{\mathbf{A}} \cdot \mathbf{e}_j^{\mathbf{B}})^2. \qquad \boxed{4.12}$$

where $\cdot$ indicates the dot product between vectors. This measure is also called tensor dot product [12]. It is used to construct the *lattice index*, which we show in section 4.5.6. Jonasson et al. [66] use the normalized tensor scalar product s_{ntsp} in order to make it invariant to scaling of the tensors:

$$s_{ntsp}(\mathbf{A}, \mathbf{B}) = \frac{s_{tsp}(\mathbf{A}, \mathbf{B})}{\mathrm{tr}(\mathbf{A}) \mathrm{tr}(\mathbf{B})}. \qquad \boxed{4.13}$$

Instead of applying the above mentioned measures to the tensors directly, they can also be applied to the *deviatoric* of the DTs (see e.g. Alexander et al. [3]). The deviatoric $\tilde{\mathbf{D}}$ of tensor $\mathbf{D}$ represents the non-isotropic part of $\mathbf{D}$. It expresses just

the shape and orientation of the DT, independent of the size. It can be computed as follows:

$$\tilde{\mathbf{D}} = \mathbf{D} - \frac{1}{3}\mathrm{tr}(\mathbf{D})\mathbf{I}, \qquad \boxed{4.14}$$

where $\mathbf{I}$ is the identity matrix. Note that $\tilde{\mathbf{D}}$ is not always a positive definite tensor. This means that it can have negative eigenvalues, and some of the measures will also give negative values.

4.5.4   Riemannian geometry

If we constrain the matrices to positive definite matrices we get another class of measures based on Riemannian geometry. Batchelor et al. [13] introduced a geodesic-based distance d_g that measures the distance between two tensors in the space of positive definite tensors:

$$d_g(\mathbf{A}, \mathbf{B}) = N(\mathbf{A}^{-\frac{1}{2}}\mathbf{B}\mathbf{A}^{-\frac{1}{2}}), \qquad \boxed{4.15}$$

where

$$N(\mathbf{D}) = \sqrt{\sum_{i=1}^{3}(\log(\lambda_i^{\mathbf{D}}))^2}. \qquad \boxed{4.16}$$

This measures the distances along geodesics in the manifold of symmetric positive defined matrices. Pennec et al. [108] introduced a similar framework with the same distance measure, and extended it with methods for filtering and regularization of tensor fields. The disadvantage of this approach is that it is computationally expensive.

Arsigny et al. [5] introduced a new *Log-Euclidian* framework. It has similar theoretical properties as the framework by Pennec et al., but with simpler and faster calculations. They derive the following Log-Euclididan distance measure d_{LE}:

$$d_{LE}(\mathbf{A}, \mathbf{B}) = \sqrt{\mathrm{tr}((\log(\mathbf{A}) - \log(\mathbf{B}))^2)}. \qquad \boxed{4.17}$$

This measure is equivalent to the d_{L2} of the logarithm of the matrices. The details of its computation and derivation can be found in Arsigny et al. [5].

4.5.5   Statistics

A diffusion tensor can be interpreted as the covariance matrix of a Gaussian distribution describing the local diffusion. Thus, a natural family of dissimilarity measures between DTs would be the statistical divergence that measures the overlap of probability density functions. Given a diffusion tensor $\mathbf{D}$, the displacement

r of water molecules at time t is a random variable with the following probability density function (PDF):

$$P(\mathbf{r}|t, \mathbf{D}) = \frac{1}{\sqrt{(2\pi)^n |2t\mathbf{D}|}} e^{-(\mathbf{r}^T \mathbf{D}^{-1} \mathbf{r})/(4t)},$$

where $||$ is the determinant and n is the dimensionality of the square matrix $\mathbf{D}$.

Wang and Vemuri [138] proposed to use the square-root of the J-divergence (symmetrized Kullback-Leibler) as a new definition of DT distance d_{KL}:

$$d_{KL}(\mathbf{A}, \mathbf{B}) = \frac{1}{2} \sqrt{\mathrm{tr}(\mathbf{A}^{-1}\mathbf{B} + \mathbf{B}^{-1}\mathbf{A}) - 2n}, \qquad \boxed{4.18}$$

where the dimensionality n is 3 for DTs.

In probability theory, class separability can be measured by the overlap between the corresponding PDFs. Therefore, the overlap of PDFs can also be used as a similarity measure between tensors. The calculation of the overlap cannot be done analytically and often approximations are being used. The Chernoff bound [38] gives us the upper bound of the probability error, $P(error)$, of a Bayesian classifier for two classes, w_1 and w_2, given their PDFs $P(w_1)$ and $P(w_2)$. For normal distributions we have:

$$P(error) \leq P^\beta(w_1) P^{1-\beta}(w_2) e^{-k\beta},$$

where β is a parameter that needs to be optimized to find the Chernoff bound. A special case is the Bhattacharyya bound where $\beta = 1/2$. This bound is never looser than the optimal Chernoff bound and can be directly calculated. For DTs, it becomes the following similarity measure:

$$s_{Bhat}(\mathbf{A}, \mathbf{B}) = e^{-\frac{1}{2} \ln\left(\frac{1}{2}|\mathbf{A}+\mathbf{B}|/\sqrt{|\mathbf{A}||\mathbf{B}|}\right)} \qquad \boxed{4.19}$$

4.5.6 Composed

As mentioned in section 4.5.1, a scalar measure in itself can give very limited information of the difference between DTs (e.g., FA just gives information about the anisotropy). Usually, a measure that reflects the changes of a combination of these properties is necessary. Therefore, several authors have tried to combine simple measures to obtain a more complete measure. Often, the measures that are combined have quite different natures and therefore *ad-hoc* normalizations and weighting factors are needed.

Pollari et al. [92] introduced shape-dependent similarity measures, s_l, s_p, s_s, which are used depending on the DT shape, linear, planar and spherical:

$$\begin{aligned}
s_l(\mathbf{A}, \mathbf{B}) &= |\mathbf{e}_1^{\mathbf{A}} \cdot \mathbf{e}_1^{\mathbf{B}}| = \cos(d_{ang_1}(\mathbf{A}, \mathbf{B})), \\
s_p(\mathbf{A}, \mathbf{B}) &= |\mathbf{e}_3^{\mathbf{A}} \cdot \mathbf{e}_3^{\mathbf{B}}| = \cos(d_{ang_3}(\mathbf{A}, \mathbf{B})), \\
s_s(\mathbf{A}, \mathbf{B}) &= 1 - \frac{|\mathrm{tr}(\mathbf{A}) - \mathrm{tr}(\mathbf{B})|}{\max(\mathrm{tr}(\mathbf{A}), \mathrm{tr}(\mathbf{B}), 1)}, \\
s_{T_2}(a, b) &= 1 - \frac{|g^a - g^b|}{\max(g^a, g^b, 1)},
\end{aligned}$$

where a and b are the voxels with tensors $\mathbf{A}$ and $\mathbf{B}$. g^a, g^b are the grey-levels in a, b in the T_2 MRI data. Using s_l, s_p, s_s and s_{T_2}, Pollari et al. introduce a DT distance measure for registration of DTI brain data sets that looks at the overlap between diffusion shapes and weights this with the most reliable information for that shape:

$$I(a, b) = \begin{aligned}&\hat{c}_l^{\mathbf{A}} \hat{c}_l^{\mathbf{B}} s_l(\mathbf{A}, \mathbf{B}) + \hat{c}_p^{\mathbf{A}} \hat{c}_p^{\mathbf{B}} s_p(\mathbf{A}, \mathbf{B}) + \\ &\gamma * \hat{c}_s^{\mathbf{A}} \hat{c}_s^{\mathbf{B}} \left(s_s(\mathbf{A}, \mathbf{B}) + s_{T_2}(a, b) \right) / 2\end{aligned} \qquad \boxed{4.20}$$

where γ is $\frac{1}{2}$ in all of their experiments because they want to give less weight to isotropic voxels. The anisotropy measures are defined as: $\hat{c}_l = \frac{\lambda_1 - \lambda_2}{\lambda_1}$, $\hat{c}_p = \frac{\lambda_2 - \lambda_3}{\lambda_1}$, $\hat{c}_s = \frac{\lambda_3}{\lambda_1}$, which is a variation of the measures proposed by Westin [143] listed in table 4.1. Because we are analyzing measures for DTs only, in section 4.7 we use a modified similarity measure s_{pnl} that disregards the s_{T_2} term of equation 4.20:

$$s_{pnl}(\mathbf{A}, \mathbf{B}) = \begin{aligned}&\hat{c}_l^{\mathbf{A}} \hat{c}_l^{\mathbf{B}} s_l(\mathbf{A}, \mathbf{B}) + \hat{c}_p^{\mathbf{A}} \hat{c}_p^{\mathbf{B}} s_p(\mathbf{A}, \mathbf{B}) + \\ &\gamma * \hat{c}_s^{\mathbf{A}} \hat{c}_s^{\mathbf{B}} s_s(\mathbf{A}, \mathbf{B}).\end{aligned} \qquad \boxed{4.21}$$

We also use $\gamma = \frac{1}{2}$, although a more precise analysis of the robustness of this measure to the changes of γ would be needed.

Pierpaoli and Basser [109] introduced the *lattice index* as an intervoxel anisotropy measure that takes the DTs in neighbouring voxels into account. For the computation of the lattice index they defined a measure s_{LI} that gives a similarity between two tensors:

$$s_{LI}(\mathbf{A}, \mathbf{B}) = \frac{\sqrt{3}}{\sqrt{8}} \frac{\sqrt{s_{tsp}(\tilde{\mathbf{A}}, \tilde{\mathbf{B}})}}{\sqrt{s_{tsp}(\mathbf{A}, \mathbf{B})}} + \frac{3}{4} \frac{s_{tsp}(\tilde{\mathbf{A}}, \tilde{\mathbf{B}})}{\sqrt{s_{tsp}(\mathbf{A}, \mathbf{A})} \sqrt{s_{tsp}(\mathbf{B}, \mathbf{B})}}, \qquad \boxed{4.22}$$

with s_{tsp} as defined in equation 4.12 and $\tilde{\mathbf{A}}, \tilde{\mathbf{B}}$ as in equation 4.14. Because $s_{tsp}(\tilde{\mathbf{A}}, \tilde{\mathbf{B}})$ can be negative, s_{LI} can give negative or imaginary values which do not fulfil the basic description of a measure as we defined it. Therefore, we do not use s_{LI} in the further analysis.

(a)   (b)
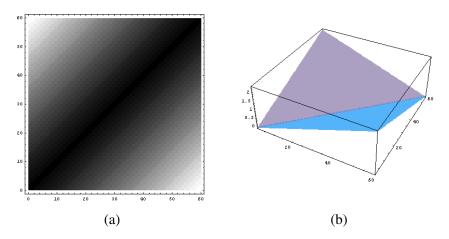
Figure 4.2: Size comparison plot of ds_{MD} for tensors with planar shape. On the axes from left to right, and from bottom to top, the size of the tensors increase while the shape and orientation are invariant. See figure 4.1(a). (a) shows a grey-value plot; (b) shows the height field.

4.6 Methods

For analysing the properties of the measures, we want to show the behaviour of each measure for the different properties in a global way. So, we show the results of each measure for sets of pairs of DTs where one property is changed. We change each property gradually and analyse the behaviour of the measures.

In order to do this analysis, we use plots as shown in figure 4.2. The axes of the plots have smoothly varying DTs and in the plot we show the similarity or difference of corresponding DTs. In figure 4.2(a) the results of the measure are shown as a grey-scale image. Figure 4.2(b) shows the same results as a height field, which gives a more clear impression about the evolution of the measure.

Furthermore, we compared the different measures by means of the root mean square difference (RMSD) of their normalized results. This allows us to grasp the similarities between the measures.

4.6.1 Size

Size is simple to evaluate because it can be captured with only one scalar value (mean diffusivity MD, see table 4.1). Figure 4.2 shows a size comparison plot for ds_{MD}. From left to right and bottom to top, we increase the size of the tensor by multiplying the eigenvalues of a base tensor with linearly increasing values. The size experiments use the an elongated tensor (with $[\lambda_1, \lambda_2, \lambda_3] = [1, 0.1, 0.1]$)

(a)   (b)
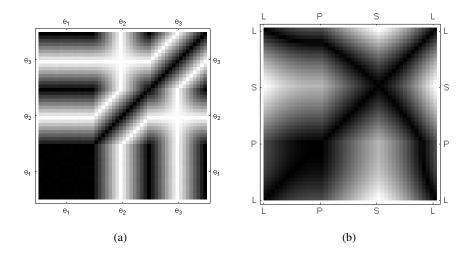
Figure 4.3: **(a)** Comparison plot for tensor with elongated shape rotated around $\mathbf{e}_1$, $\mathbf{e}_2$ and $\mathbf{e}_3$ showing d_{ang_1}. See figure 4.1(c); **(b)** ds_{FA} comparison plot for tensors with shape changing from elongated (L) to planar (P) to spherical (S) to elongated (L). See figure 4.1(d).

which is enlarged by multiplying all components of the tensor with values from 0 to 60. This is illustrated in figure 4.1(a). It can be seen from figure 4.2 that tensors with the same size (on the diagonal of the plot) have zero distance, and tensors of which the sizes differ have larger distances. Some measures (e.g., ds_{FA}) are invariant to scaling. So this plot will not be used for those measures.

4.6.2   Orientation

For orientation, we consider the sensitivity of the measure to rotation of the tensors, i.e. rotation around any axis. For tensors with an elongated shape, the measure should be invariant to rotations around $\mathbf{e}_1$. For tensors with a planar shape, the measure should be invariant to rotations around $\mathbf{e}_3$. For tensors with a spherical shape, the measure should be invariant to any rotation. We created plots for multiple types of tensors (representing linear, planar or spherical diffusion), where on both axes we gradually rotate the tensor around $\mathbf{e}_1$ until π. Then, on the middle part of the horizontal and vertical axes of the plots, we rotate around $\mathbf{e}_2$ until π. Finally, in the top and right of the two axes we rotate around $\mathbf{e}_3$ until π. The tensors on the axes of the plots are illustrated in figures 4.1(b) and 4.1(c). Figure 4.3(a) shows results for d_{ang_1}.

The orientation experiments were done with an elongated tensor with eigen-

values $[\lambda_1, \lambda_2, \lambda_3] = [1, 0.1, 0.1]$. Thus, for rotation around $\mathbf{e}_1$, the distance does not change. This can be seen in the image because in the lower-left part, the distances stay zero. When rotating the tensor around $\mathbf{e}_2$ and $\mathbf{e}_3$, it can be seen that the distance between measures gradually increases for a rotation up to $\pi/2$ and then decreases again until it is zero at π.

Furthermore, we tested the rotation invariance of the measures to the situation when we rotate the volume. We did this by applying the same rotation to every tensor in a set, and then computing the root mean square difference (RMSD) of these results to the corresponding ones without rotation.

4.6.3 Shape

We consider that DTs can have elongated, planar or spherical shape or a shape that is an interpolation of these shapes. In order to study the behaviour of the measure under changes in shape, we start with DTs with elongated shape ($\lambda_1 \gg \lambda_2 \simeq \lambda_3$ with $[\lambda_1, \lambda_2, \lambda_3] = [5/2, 1/4, 1/4]$ mm^2/s), and then gradually change the shape to planar ($\lambda_1 \simeq \lambda_2 \gg \lambda_3$ with $[\lambda_1, \lambda_2, \lambda_3] = [10/7, 10/7, 1/7]$ mm^2/s), spherical ($\lambda_1 \simeq \lambda_2 \simeq \lambda_3$ with $[\lambda_1, \lambda_2, \lambda_3] = [1, 1, 1]$ mm^2/s), and back to elongated. This is illustrated in figure 4.1(d). In order to make sure that we are only evaluating shape, we do not change the size (MD $= 1$) and orientation of the tensors in the same plot. Results for ds_{FA} are shown in figure 4.3(b). As can be seen from the black areas in the plot which are not in the diagonal, tensors with different shapes can have the same value for FA. This is a known property of FA.

4.6.4 Robustness

From the results of the previous methods, we can deduce whether a measure is sensitive to small changes for one of the properties. In addition, we introduce a small variation to the set of tensors in our experiments. To each component of the input tensors (on both axes) for making the size, shape and orientation plots, we add a uniformly distributed random value. Then we analyse this robustness by computing the root mean square difference between the plots with and without the added noise. We consider the measures robust to noise if its plots do not show sharp changes or discontinuities, and the computed root mean square differences are relatively small.

4.6.5 Metric

The conditions that need to be fulfilled for a measure to be a metric can be derived from its definition. Thus, no experiments are needed in order to evaluate this

property. However, we will summarize whether the properties in equations 4.5–4.7 are fulfilled for each of the measures.

4.7 Experiments

In this section, we analyse and categorize behaviour of the different measures using the methods described in the previous section. The behaviour of the measures is summarized in table 4.2.

4.7.1 Size

We can observe four different behaviours for the measures with respect to the relation between the sizes of the tensors and the respective change in the measure output.

In table 4.2, we list these behaviours as *invariant* or variant in three different ways: *proportional*, *multiplicative* and *fractional* .

All scalar measures listed in table 4.1, except MD, are *invariant* to scaling of one or both input tensors with a positive scalar μ. d_{ang_i}, s_{ntsp}, and s_{LI} are also invariant to scaling.

Measures ds_{MD} and d_{L2} show a behaviour as illustrated in figure 4.2. Scaling $\mathbf{A}$ and $\mathbf{B}$ will proportionally change the outcome of ds_{MD} and d_{L2}. In case the scaling factor is the same:

$$m(\mu\mathbf{A}, \mu\mathbf{B}) = (\mu r) \times m(\mathbf{A}, \mathbf{B}) \qquad \boxed{4.23}$$

where r is a scalar proportion ratio related to the measure m. This behaviour is listed as *proportional* in table 4.2.

Measures s_{sp} and s_{tsp} have behaviour as shown in figure 4.4(a). They return bigger values as the size of the tensors is bigger:

$$m(\mu_1\mathbf{A}, \mu_2\mathbf{B}) = \mu_1\mu_2 m(\mathbf{A}, \mathbf{B}) \qquad \boxed{4.24}$$

As a consequence of this multiplicative behaviour, $m(\mu\mathbf{A}, \mu\mathbf{A})$ is not constant but depends on μ, therefore there is no upper limit for the value of these measures. We list this behaviour as *mult* in table 4.2.

The remaining measures are listed as *frac*. This means that they behave as shown in figure 4.4(b). The relation between the output of the measure and the ratio of the size of the two tensors is proportional, and it is invariant to equal scaling of both tensors with scalar μ:

$$m(\mu\mathbf{A}, \mu\mathbf{B}) = m(\mathbf{A}, \mathbf{B}). \qquad \boxed{4.25}$$

(a) Linear tensors with measure s_{tsp}.

(b) Planar tensors with measure d_g.
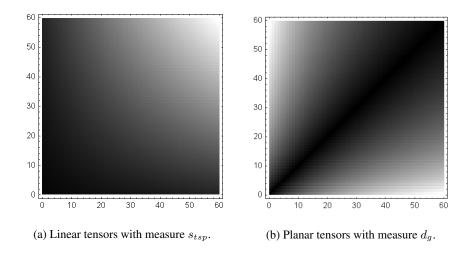
Figure 4.4: Comparison plots with tensors changing size.

4.7.2   Orientation

Measure d_{ang_1} only works well for tensors with elongated shape. All scalar-index based measures (ds) are invariant to rotation. All measures except ds and d_{ang_i} have similar behaviour under rotation, which we list as *smooth*. For tensors with elongated shape, they have the same behaviour as d_{ang_1}, which is shown in figure 4.3(a). Results for d_{KL} for tensors with planar shape is shown in figure 4.5. It is similar to that in figure 4.3(a). Except for d_{ang_i}, all measures are invariant to rotations if at least one of the two tensors that are being compared to has spherical shape, i.e.:

$$m(\mathbf{S}, \mathbf{A}) = m(\mathbf{S}, R^T \mathbf{A} R) \qquad \boxed{4.26}$$

for spherical tensor $\mathbf{S}$, and $\mathbf{A} \in Sym^+(3)$, and rotation R. For tensors whose shape is not purely elongated, planar or spherical, the resulting plots are a weighted average of the plots of the respective tensor types. This is shown in figure 4.6 for d_{LE}. All measures, except d_{ang_i} are invariant to rotations of both tensors, thus, for any rotation R:

$$m(\mathbf{A}, \mathbf{B}) = m(R^T \mathbf{A} R, R^T \mathbf{B} R). \qquad \boxed{4.27}$$

In order to refine the classification of these measures, we compared their results by computing the root mean square difference ($RMSD$) between them. Measures d_{L2}, d_g, d_{LE}, d_{KL} and d_{Bhat} (since s_{Bhat} it is not a distance, we inverted the result, $d_{Bhat} = 1 - s_{Bhat}$, before the comparison) are similar to each other
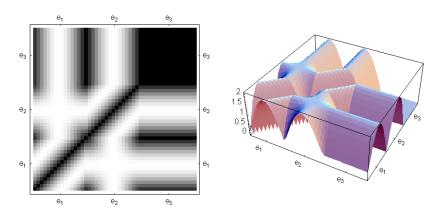
Figure 4.5: Comparison plot of planar tensors rotated around $\mathbf{e}_1$, $\mathbf{e}_2$ and $\mathbf{e}_3$ for d_{KL}. See figure 4.1(b).

($RMSD \approx 0$). We can define another subgroup with the measures s_{sp}, s_{tsp} and s_{ntsp}. These measures give the same result, $RMSD = 0$.

4.7.3 Shape

Of the measures that we analysed, only ds_{MD} is invariant to shape changes. The behaviour for ds_{FA} is shown in figure 4.3(b). We list it as *non descriptive* in table 4.2 because tensors with different shape can have a distance of zero, although it is meant to characterize tensor's shape. The other ds measures show similar behaviour where tensors that differ can have a distance of zero depending on which anisotropy measure is used. d_{ang_i} can give random values depending on the shape of the diffusion. This measure should only be applied to tensors that represent anisotropic diffusion. Thus, we list *elongated* in table 4.2. The behaviour of d_{L2} is shown in figure 4.7. The diagonal is black, and the greatest distance occurs between elongated and spherical tensors. We consider this behaviour as good and therefore we list *ok* in table 4.2.

Measures s_{sp}, s_{tsp} and s_{ntsp} all behave similar to what is shown in figure 4.8. These measures give a high similarity when comparing a tensor with elongated shape to itself. However, tensors with planar or spherical shapes are less self-similar, i.e. have a lower value for $m(\mathbf{A}, \mathbf{A})$. Thus, the similarity between a tensor and itself depends on its shape. Because of this behaviour, we cannot convert these similarity measures to distance measures which fulfil metric condition cf. equation 4.5. We list this in table 4.2 as *not self-similar (1)*.

The behaviour under shape changes for s_{pnl} is shown in figure 4.9. The values on the diagonal are higher than the values next to it because tensors are self-
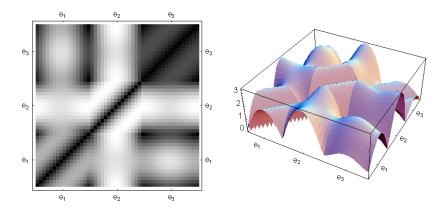
Figure 4.6: Comparison plot for tensors rotated around $\mathbf{e}_1$, $\mathbf{e}_2$ and $\mathbf{e}_3$ for d_{LE}. The tensors do not have pure elongated, planar or spherical shape, but eigenvalues $\lambda_1 = 1.0, \lambda_2 = 0.5, \lambda_3 = 0.1$.
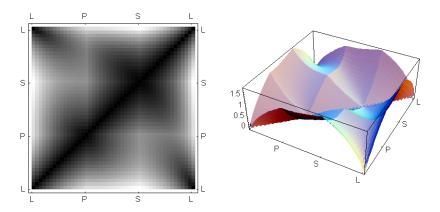
Figure 4.7: Comparison plot of d_{L2} for shapes changing from elongated to planar to spherical to elongated.

	Equation	Eqn	Cite	Size	Orientation	Shape	Robustness	Metric
ds_{FA}	$\|FA(\mathbf{A}) - FA(\mathbf{B})\|$	4.8	[143]	invariant	invariant	non descriptive	shape	no
ds_{MD}	$\|MD(\mathbf{A}) - MD(\mathbf{B})\|$	4.8	[143]	proportional	invariant	invariant	size	no
d_{ang_1}	$\arccos(\mathbf{e}_1^A \cdot \mathbf{e}_1^B)$	4.9	[152]	invariant	linear	elongated	elongated only	no
d_{L2}	$\sqrt[n]{\sum_{i=1}^3 \sum_{j=1}^3 (\mathbf{A}_{ij} - \mathbf{B}_{ij})^n}$	4.10	[13]	proportional	smooth	ok	all	yes
s_{sp}	$\sum_{i=1}^3 \sum_{j=1}^3 \mathbf{A}_{ij}\mathbf{B}_{ij}$	4.11	[3]	mult	smooth	not self-similar (1)	all	no
s_{tsp}	$\sum_{i=1}^3 \sum_{j=1}^3 \lambda_i^A \lambda_j^B (\mathbf{e}_i^A \mathbf{e}_j^B)^2$	4.12	[66]	mult	smooth	not self-similar (1)	all	no
s_{ntsp}	$\frac{s_{tsp}(\mathbf{A},\mathbf{B})}{\mathrm{tr}(\mathbf{A})\mathrm{tr}(\mathbf{B})}$	4.13	[66]	invariant	smooth	not self-similar (1)	shape/orientation	no
s_{pnl}	$\hat{c}_l^A \hat{c}_l^B s_l(\mathbf{A}, \mathbf{B}) + \hat{c}_p^A \hat{c}_p^B s_{sp}(\mathbf{A},\mathbf{B})$ $+\gamma * \hat{c}_s^A \hat{c}_s^B s_{ss}(\mathbf{A}, \mathbf{B})$	4.21	[92]	frac	smooth	not self-similar (2)	shape/orientation	no
d_g	$N(\mathbf{A}^{-\frac12}\mathbf{B}\mathbf{A}^{-\frac12})$	4.15	[13]	frac	smooth	sensitive	sensitive	yes
d_{LE}	$\sqrt{\mathrm{tr}((\log(\mathbf{A}) - \log(\mathbf{B}))^2)}$	6.3	[5]	frac	smooth	sensitive	sensitive	yes
d_{KL}	$\frac12\sqrt{\mathrm{tr}(\mathbf{A}^{-1}\mathbf{B} + \mathbf{B}^{-1}\mathbf{A}) - 2n}$	4.18	[138]	frac	smooth	sensitive	sensitive	yes
s_{Bhat}	$e^{-\frac12 ln\left(\frac{\|\frac{\mathbf{A}+\mathbf{B}}{2}\|}{\sqrt{\|\mathbf{A}\|\|\mathbf{B}\|}}\right)}$	4.19	[38]	frac	smooth	sensitive	sensitive	no

Table 4.2: Overview of properties for similarity and distance measures for diffusion tensors. The ds measures that are not listed have the same properties as ds_{FA}.
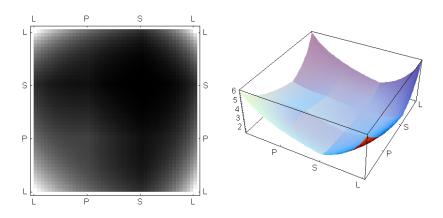
Figure 4.8: Comparison plot for changing shapes for s_{tsp}.
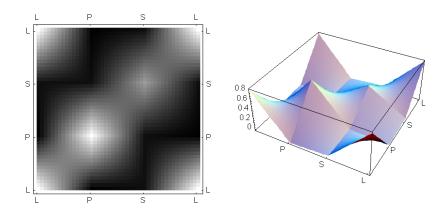
Figure 4.9: Comparison plot for changing shapes for S_{pnl}.
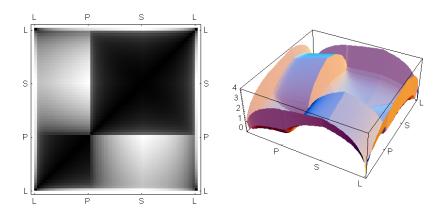
Figure 4.10: Comparison plot for shapes changing from elongated (L) to planar (P) to spherical (S) to elongated (L). The plot shows d_{KL}. $\lambda_1/\lambda_3 = 100$ for tensors with "pure" elongated and planar shape.

similar. However, the actual values on the diagonal are not all the same. Thus, for s_{pnl} the similarity between a tensor **D** and itself also depends on the shape of **D**. We list this in table 4.2 as *not self-similar (2)*.

The plots for measures based on Riemannian geometry and statistics (see sections 4.5.4 and 4.5.5) show steep edges in areas where at least one of the eigenvalues is very small. This is shown for d_{KL} in figure 4.10 where $\lambda_1/\lambda_3 = 100$ for tensors with "pure" elongated and planar shape. This behaviour is listed as *sensitive* in table 4.2. In medical data, the chance to be exactly on the very steep part is small because the fractions between eigenvalues are not that large. However, it is always possible that two similar tensors are on opposite sides of this edge, which will result in a large difference. Also, noise in medical data can change the fractions of the eigenvalues in such a way that the tensors come closer to the steep edges, i.e., small variations in the shape results in large variation in the measures.

4.7.4 Robustness

We repeated the experiments of the previous sections after adding noise as described in section 4.6.4 to the input tensors. The noise consists of uniformly distributed random values $\varepsilon \in [-0.01, 0.01]$, which are added independently to the components of the tensors, and independently for each tensor. We then compare the root mean square difference (RMSD) between the output of the normalized plots with and without noise. The results are shown in table 4.3. The more robust the measures are to noise, the lower the values in the table.

In table 4.2, the robustness of the measures is summarized. Some measures

	Eqn	shape	orientation	size
ds_{FA}	4.8	0.007	0.372	0.279
ds_{MD}	4.8	0.316	0.332	0.002
d_{ang_1}	4.9	0.409	0.006	0.338
d_{L2}	4.10	0.005	0.013	0.002
s_{sp}	4.11	0.009	0.024	0.002
s_{tsp}	4.12	0.009	0.024	0.002
s_{ntsp}	4.13	0.003	0.024	0.246
s_{pnl}	4.21	0.008	0.018	0.237
d_g	4.15	0.012	0.044	0.007
d_{LE}	6.3	0.012	0.042	0.007
d_{KL}	4.18	0.013	0.056	0.007
s_{Bhat}	4.19	0.014	0.048	0.006

Table 4.3: Root mean square difference between the sets of tensors with and without small variations.

prove to be robust within only one or two of the properties (shape, orientation and size), i.e. they have one or two relatively low values in table 4.3. We classify them as such. For example measure ds_{FA} is robust to changes in *shape* only. If the plots do not have steep parts, i.e. high discontinuities, thus the values in table 4.3 are small, we consider the noise robustness of the measures to be good for *all*.

Measure d_{L2} proves to be the most robust measure. Measure d_{ang_1} only takes the main diffusion direction into account. If the shape of the tensor is not elongated, this direction can change randomly when small changes are made to the tensors. Thus, d_{ang_1} does not behave well under noise. From the other measures, only the shape plots for d_g, d_{LE}, d_{KL}, and s_{Bhat} show steep edges. These edges appear where the shapes of the tensors are very elongated or planar. Thus, in these areas the measures are very sensitive to noise. Figure 4.11 shows this behaviour for d_{KL}.

4.7.5 Metric

All measures are symmetric, this can be also seen in the plots, since they are symmetric by the diagonal.

We have converted the similarity measures into distance measures and evaluate whether those distance measures can be metrics. Similarity measures with *multiplicative* as size change behaviour or *not self-similar* as shape change behaviour have a similarity $s(\mathbf{A}, \mathbf{A})$ that depends on the size or shape of tensor
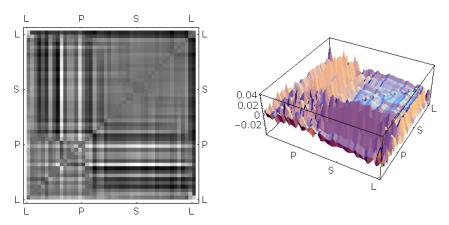
Figure 4.11: Comparison plot showing the difference between the response of d_{KL} to a set with and without random noise. $RMSD = 0.013$

A. Thus they cannot directly be translated into a distance measure that always fulfils equation 4.5. Measures ds, d_{ang_i} and s_{ntsp} are invariant to one or more of the properties of section 4.4. Thus, there are many tensors $\mathbf{A} \neq \mathbf{B}$ for which $d(\mathbf{A}, \mathbf{B}) = 0$ which invalidates metric condition 4.5.

It is clear that distance measures d_{L2}, d_g, d_{LE}, d_{KL} fulfil equations 4.5 and 4.6. They also fulfil the triangle inequality of equation 4.7 if the tensors $\mathbf{A}$ and $\mathbf{B}$ are infinitesimally close [5, 108, 138]. Therefore they are *metrics*. Because this is sufficient for those applications that need the distances to be a metric, we list them as *yes* for the metric property in table 4.2. s_{Bhat} is a similarity measure, so it cannot be a metric. However, it can be used to construct a metric, as is shown by Comaniciu and Meer [25].

4.8 Discussion
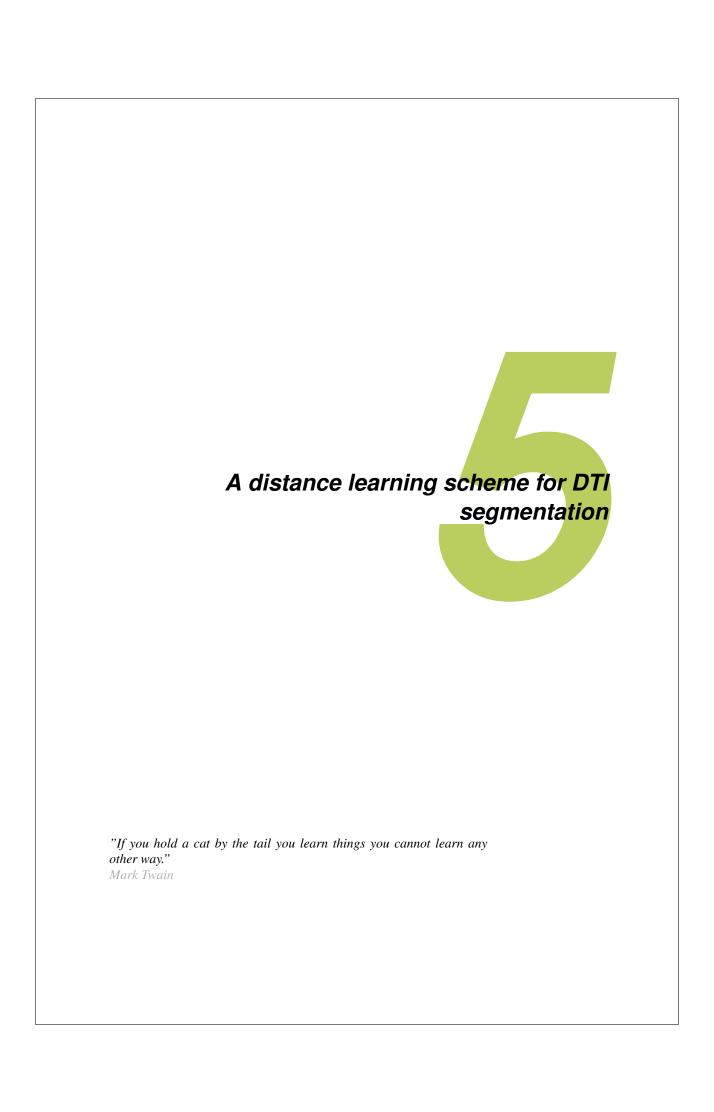
Depending on the application, different distance or similarity measures can be used. Using the previous analysis of properties we can identify from a practical point of view the differences and similarities between the different measures. It turns out that the behaviour of s_{sp} and s_{tsp} is similar, even though s_{sp} deals with the tensor as if it is a vector. The L^2 norm distance d_{L2} is relatively simple, but shows good behaviour. Also, all measures listed in sections 4.5.4 and 4.5.5 give practically the same results. Except for s_{Bhat}, the similarity measures S cannot easily be converted into metrics. Thus if that is a requirement for the application those measures are ruled out. This also rules out the ds measures and d_{ang_i}. Measure d_{L2} can be a good measure in that case. When using measures d_g, d_{LE}, d_{KL} and s_{Bhat} one has to be aware that the measures are sensitive to small shape changes close to the degenerate cases.

Throughout a full brain, all diffusion properties vary. In order to take all properties into account when registering brains, no measure should be chosen that is invariant to any of them. Also, if the weighting for all DTs used in the registration must be the same, the similarity measures that list *multiplicative* for size or *not self-similar* for shape should not be used because, even for equal DTs, the computed similarities vary depending on size and shape. For interpolation of DTs the triangle inequality condition must be satisfied, therefore only measures that are metrics can be used. Work has been done in the comparison of the different interpolation methods as in Arsigny et al. [5], Pennec et al. [108] and Kindlmann et al. [76].

We created an overview of existing distance and similarity measures for matching diffusion tensors and classified the measures. Such an overview, including introduced measures, was not previously available. We evaluated the properties of these measures and summed them up in table 4.2. When researchers want to use a similarity or distance measure for their concrete application, they can define which properties their measure should have to and then study the measures that fulfil their requirements. When new measures are introduced, it will be beneficial to classify them and see for which properties they differ from already existing measures, and how they differ. So in which sense they improve existing measures.

This chapter aims to help in making the first selection of these measures. The next step is to test what measure performs better in a concrete application, e.g., white matter segmentation. If the goal is to segment the brain using DTs, the choice of measure depends on which properties are of importance for a given area. For example, segmenting the thalamic nuclei requires dependency of orientation for the measure used [152], while white and grey matter can be distinguished using the tensor shape.

5

A distance learning scheme for DTI segmentation

"If you hold a cat by the tail you learn things you cannot learn any other way."
Mark Twain

Contents

5.1   Overview

In this chapter we discuss segmentation techniques for diffusion tensor imaging data. In particular, we partition the data into regions which are homogeneous in terms of tensor characteristics. Various distance measures have been proposed in literature for analysing the similarity between diffusion tensors, but selecting a measure suitable for the task at hand is difficult and often done by trial-and-error. In this chapter, we propose a novel approach to semi-automatically select the similarity measure or combination of measures that better suit the data. We use a linear combination of known distance measures, jointly capturing multiple aspects of tensor characteristics, for comparing DTs with the purpose of image segmentation. The parameters of our adaptive distance measure are tuned for each individual segmentation task on the basis of user-selected ROIs using the concept of kernel target alignment (KTA). Experimental results support the validity of the proposed method.

5.2   Introduction

As we saw in section 2.5, many applications, such as neurological disorder studies or surgical planning, involve the virtual dissection of white matter fiber tracts (e.g., the *corticospinal tract*). Fiber clustering techniques can be used for this

purpose, however, not without a caveat for its sensitivity to the underlying fiber tracking parameters. The direct segmentation of the tensor field into volumetric regions stands as an interesting alternative to fiber clustering to avoid this problem.

Inherent to segmentation methods is the notion of (in)homogeneity among the tensor field, i.e. tensors that belong to the same tract will be similar to each other. These techniques require then the notion of (dis)similarity of two DTs, i.e. a measure that indicates when tensors are considered to be similar enough to belong to the same region.

In chapter 4 we extensively analysed the different (dis)similarity measures and we observed that different measures capture different tensor characteristics, and thus the segmentation results are, clearly, dependent on the choice of measure. This choice depends on the application at hand. Furthermore, these measures are of different nature and sometimes lack physiological significance, afflicting even more the problem of measure selection. Usually an *ad-hoc* definition of parameter values and choice of similarity measures is used.

The contribution presented in this chapter lies in the semi automatic assessment of tensor similarity based on the data and segmentation task. The resulting measure is a suitable parametrised measure. The results of the presented method can then be used in any segmentation algorithm where homogeneity properties are used, e.g., region growing.

This problem of *metric learning* and parameter estimation has been addressed before in the machine learning and pattern recognition literature [61,120]. We extended these methods for the particular problem of diffusion tensor segmentation. With the proposed preprocessing distance learning algorithm, the parameters for a region-based segmentation algorithm are inferred from the data. A seeding region is selected (by the user) and the algorithm will segment the spatially connected 3D section with the diffusion tensors that are similar to the initial chosen region and dissimilar to the rest. The initially flexible learning scheme adapts itself to the task at hand. This technique can be used for different (region-based) segmentation algorithms. For illustration purposes, we present the results using a region growing based segmentation.

5.3   Methods

The main goal of this work is to assess what distance better expresses the homogeneity characteristics of a structure defined in a tensor field, e.g., DTI of the
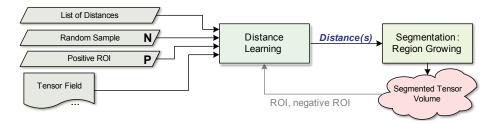
Figure 5.1: Global gist of the distance/parameter learning and segmentation

white matter. We consider a linear combination of distances also as a distance.

$$d(\mathbf{A}, \mathbf{B}) = \sum_{i=1}^{l} w_i m_i(\mathbf{A}, \mathbf{B}), \text{ and } \sum_{i=1}^{l} w_i = 1, \ w = (w_1, ..., w_l), \ \forall_m w_m \geq 0$$

(5.1)

where d is the distance resulting of l combined distances m_i, whose contributions are defined by parameter w_i. We describe the general framework in section 5.3.1.

We define a distance/parameter learning scheme, whose results are then used to drive a region growing segmentation algorithm (see figure 5.1). This distance learning algorithm infers the distance(s) that best discriminates a selected region of interest (ROI) from the entire image volume represented by a random sample of DTs. The optimal combination of distances will then be used in the segmentation algorithm and a spatially connected volume of tensors will be obtained. Then the user will be able to further improve the process by adding additional negative ROIs, i.e. examples of tensors that are different from the target region and provide complementary information.

Figure 5.2 shows the details of the distance learning algorithm. From the tensor field volume data we define a labelled set $S = (\mathbf{D}_i, l_i)$ of n DTs $\mathbf{D}$ with a label l. The set S is defined as the union of two subsets of DTs: P, a set of representative DTs from a user defined ROI (positive ROI), where $l = +1$; and N, a set of representative DTs for the whole volume (negative ROI), where $l = -1$.

Distance matrices are constructed by calculating the distance between all pairs of tensors in the set S. Each row, and column, is considered as a feature vector with the distance from a tensor to all others in the training set. From these feature vectors, symmetric matrices, referred to as kernel matrices (i.e. Gram matrices), are calculated. For a uniform behaviour of the algorithm, without minding the scale, a normalisation of the individual kernel matrices is performed. Then, with a linear combination of the different kernel matrices, one per considered distance, we define a new kernel matrix $\mathbf{K}$ with a set of unknown parameters (the weights).

Using a gradient descent based method, the weights are estimated in order to maximize the kernel target alignment measure described in section 5.3.3. This
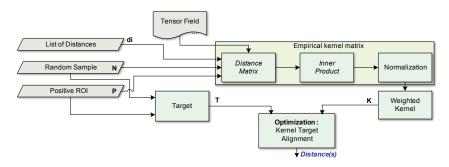
Figure 5.2: Detailed scheme of the distance learning algorithm

maximum gives the best alignment between the kernel matrix $\mathbf{K}$ and an equally sized label matrix, i.e. which combination of distances provides the best discrimination for the considered data.

In the following, we describe the optimization of the kernel target alignment for the distance learning. In section 5.4 results and evaluation of the distance learning algorithm are presented.

5.3.1 Distances

Distance measures convey different aspects of a diffusion tensor. While some capture changes in individual degrees of freedom (e.g., difference in anisotropy), others use the full tensor information. Thus their use is sometimes redundant, i.e. different measures describe common tensor attributes. There are measures that use the full tensor, like *Riemannian* based measures, that have a mathematical nature which does not have a direct intuition of the physiological meaning. Thus, the results are not predictable. Other measures like the ones presented by Kindlmann et al [74] decompose tensor variations into changes in shape and orientation, covered by three invariant gradients and three rotation tangents. In this work, a tunable difference measure between two DTs is introduced. This measure uses a weighted sum of the individual measures, however, the definition of these weights depends on the task at hand. The method presented in this chapter would be able to help in this definition. Notwithstanding, this measure is only good for very small differences, since the invariant gradient and rotation tangent coordinate frame are not accurately defined for a large difference between tensors. Therefore, we do not use these measures.

In order to show the flexibility of our framework, we evaluate the following set of different distances (chapter 4): difference of fractional anisotropy FA (ds_{FA}), difference of mean diffusivity MD ds_{MD}, angular difference d_{ang1}, Frobenius distance d_{L2}, geometric distance d_g, Log-Euclidean distance d_{LE} and the sym-

metrized Kullback-Leibler distance d_{KL}.

These measures are chosen because they fulfil $d(\mathbf{D}_1, \mathbf{D}_1) = 0$, and are symmetric and positive. Other similarity measures presented in chapter 4 could be used, however they must be converted into a distance measure, see Haasdonk et al [55]. The distance learning algorithm does not require the distance to be a true metric - the triangle inequality is not a prerequisite.

5.3.2 Empirical kernel matrices

The main idea of kernel methods is to map the input data (here the input data is the distance between tensors) to a feature space provided with a dot product. The mapped data is then dichotomized.

Let's define a feature vector representing a single object, in our case a DTI voxel. For a set $S = P \cup N$ of L objects o_j, the feature vector representing o_i, $\mathbf{f_i} = [d(o_i, o_1), ..., d(o_i, o_L)]$, is computed by evaluating a distance measure m between o_i and all other objects in S. For measure m, a kernel matrix $\mathbf{K}$ for a set of L feature vectors can be regarded as a matrix of pairwise similarity, measured by their pairwise inner-product.

As presented in Pekalska et al. [107], a kernel $\mathbf{K}$ can be defined as a mapping of the feature vectors $\mathbf{f_i}$. The kernel matrix is then the inner-product between the feature vectors

$$K_{ij} \equiv < \mathbf{f}_i, \mathbf{f}_j > = \sum_k d(\mathbf{D}_i, \mathbf{D}_k) d(\mathbf{D}_j, \mathbf{D}_k) \qquad \boxed{5.2}$$

where K_{ij} is the element in row i and column j of the kernel $\mathbf{K}$.

Each element in the kernel matrix effectively depends on all tensors in the labelled set S. The kernel has high values for similar classes, but close to 0 for inter-class tensors. For geometric interpretation, consider that the inner product is proportional to the angle between two vectors. Now we have a kernel matrix, i.e. the set of all possible inner products, and it is symmetric and positive definite.

For a uniform behaviour of the algorithm, i.e. without minding the scale of the used measures, a normalisation must be performed. We can normalise kernel matrices in such a way that the features lie on the surface of a unit hypersphere. This normalisation [53] can be done directly in the kernel as follows:

$$\widetilde{\mathbf{K}}_{ij} = \frac{\mathbf{K}_{ij}}{\sqrt{\mathbf{K}_{ii}\mathbf{K}_{jj}}} \qquad \boxed{5.3}$$

Henceforth, a normalised kernel $\widetilde{\mathbf{K}}$ from a distance measure m will be referred as $\mathbf{K}^m$.

5.3.3 Alignment

Christiani et al. [29] proposed a method to assess the quality of a binary clustering. This measure, referred to as kernel target alignment (KTA), depicts how good a kernel is with respect to a given set of labelled objects (the target) with the notion of good clustering, i.e. high similarity within clusters and low similarity between clusters. This notion is captured using the Frobenius inner product between these matrices.

The Frobenius product between two matrices $\mathbf{V}, \mathbf{P}$ is defined as

$$< \mathbf{V}, \mathbf{P} >_F = \sum_{ij} v_{ij} p_{ij}$$

The *alignment* between two arbitrary kernels $\mathbf{K}_1$ and $\mathbf{K}_2$ is

$$A(\mathbf{K}_1, \mathbf{K}_2) = \frac{< \mathbf{K}_1, \mathbf{K}_2 >_F}{\sqrt{< \mathbf{K}_1, \mathbf{K}_1 >_F < \mathbf{K}_2, \mathbf{K}_2 >_F}} \qquad \boxed{5.4}$$

A target matrix is constructed from the set of n tensors S. We define a vector of labels $\mathbf{y} \in \{-1, +1\}^n$ where 1 is the label for the positive set P, and -1 for the negative set N. The target is then calculated using the matrix product $\mathbf{T} = \mathbf{y}^T \mathbf{y}$ and the alignment can now be expressed as

$$A(\mathbf{K}, \mathbf{T}) = \frac{< \mathbf{K}, \mathbf{T} >_F}{n\sqrt{< \mathbf{K}, \mathbf{K} >_F}}, \text{ since } < \mathbf{T}, \mathbf{T} >_F = n^2 \qquad \boxed{5.5}$$

5.3.3.1 Linear combination of kernels

In machine learning, the problem of learning an adequate distance metric for the input space of data from a set of similar/dissimilar objects has been addressed in many studies in recent years like Igel et al. [61].

So far, we have a set of normalized kernels $\mathbf{K}^m$, one for each m measure. However, some kernels, i.e. some measures, may be more discriminative than others. Therefore, we introduce new weight parameters $w_m, m = 1, .., l$, with l as the number of distances to evaluate, associated to each measure m, and a new kernel will be constructed from the linear combination of the individual kernels:

$$\mathbf{K}(\mathbf{w}) = \sum_{m=1}^{l} w_m \mathbf{K}^m, \text{ and } \sum_{m=1}^{l} w_m = 1, \ \mathbf{w} = (w_1, ..., w_l), \ \forall m \ w_m \geq 0$$

$$\boxed{5.6}$$

We use the A measure defined in equation 5.5 with kernel $\mathbf{K}(\mathbf{w})$. We want to find the weights $\mathbf{w}$ that maximize $A(\mathbf{K}(\mathbf{w}), \mathbf{T})$, i.e. the $\mathbf{w}$ that gives the best discrimination between P and N. If the measures are not orthogonal to each other

and do not represent specific characteristics of the tensor, a clear interpretation of the resulting weights cannot be given. Furthermore there is not a unique solution. However, we still expect that the method will give a good balance of the measures and they will give good results although we cannot associate a clear interpretation to the measures.

5.3.4 Parameter tuning using a gradient-descent based method

The selection of weights is achieved by maximizing the alignment between the linear combination of kernels $\mathbf{K}(\mathbf{w})$ and the target matrix $\mathbf{T}$

$$
argmax_{\mathbf{w}}(A(\mathbf{K}(\mathbf{w}), \mathbf{T})) = argmax_{\mathbf{w}} \left(\frac{< \mathbf{K}(\mathbf{w}), \mathbf{T} >_F}{n\sqrt{< \mathbf{K}(\mathbf{w}), \mathbf{K}(\mathbf{w}) >_F}} \right) \quad \boxed{5.7}
$$

Expressing $\mathbf{K}(\mathbf{w})$ as a linear combination of the individual kernels $\mathbf{K}^m$

$$
argmax_{\mathbf{w}} A(\mathbf{K}(\mathbf{w}), T) = \frac{< \displaystyle\sum_{m=0}^{l} w_m \mathbf{K}^m, T >_F}{n\sqrt{< \displaystyle\sum_{m=0}^{l} w_m \mathbf{K}^m, \displaystyle\sum_{m=0}^{l} w_m \mathbf{K}^m >_F}} \quad \boxed{5.8}
$$

For the gradient-based adaptation of the kernel weights, the partial derivatives of the weighted kernel are computed with respect to the weights $\mathbf{w}$.

Rewriting $A(\mathbf{w})$ (equation 5.5) as

$$
A(\mathbf{w}) = \frac{f(\mathbf{w})}{n\sqrt{h(\mathbf{w})}} \quad \boxed{5.9}
$$

with

$$
f(\mathbf{w}) = f = < \mathbf{K}(\mathbf{w}), \mathbf{T} >_F = \sum_{ij}^{n} \sum_{m=0}^{l} w_m K_{ij}^m T_{ij}
$$

$$
h(\mathbf{w}) = h = < \mathbf{K}, \mathbf{K} >_F = \sum_{ij}^{n} \sum_{m=0}^{l} w_m K_{ij}^m \sum_{m=0}^{l} w_m K_{ij}^m
$$

The partial derivative of the kernel alignment A with respect to the weights w_m is given by:

$$\frac{\partial A(\mathbf{w})}{\partial w_m} \quad = \quad \frac{2hf' - fh'}{2nh\sqrt{h}} \qquad \boxed{5.10}$$

where

$$f' = \frac{\partial f(\mathbf{w})}{\partial w_m}, h' = \frac{\partial h(\mathbf{w})}{\partial w_m}$$

The derivative of the Frobenius product of the weighted kernel is

$$
\begin{aligned}
\frac{\partial f(\mathbf{w})}{\partial w_m} \quad &= \quad \frac{\partial < \mathbf{K}, \mathbf{T} >_F}{\partial w_m} \\
&= \quad \frac{\partial < \sum_k w_k \mathbf{K}^k, \mathbf{T} >_F}{\partial w_m} \\
&= \quad \frac{\partial \sum_k w_k < \mathbf{K}^k, \mathbf{T} >_F}{\partial w_m} \\
&= \quad < \mathbf{K}^m, \mathbf{T} >_F
\end{aligned}
$$

$$
\begin{aligned}
\frac{\partial h(\mathbf{w})}{\partial w_m} \quad &= \quad \frac{\partial < \mathbf{K}, \mathbf{K} >_F}{\partial w_m} \\
&= \quad \frac{\partial < \sum_k w_k \mathbf{K}^k, \sum_l w_l \mathbf{K}^l >_F}{\partial w_m} \\
&= \quad \frac{\partial \sum_k w_k \sum_l w_l < \mathbf{K}^k, \mathbf{K}^l >_F}{\partial w_m} \\
&= \quad \sum_k \sum_l w_l \frac{\partial w_k}{\partial w_m} < \mathbf{K}^k, \mathbf{K}^l >_F + w_k \frac{\partial w_l}{\partial w_m} < \mathbf{K}^k, \mathbf{K}^l >_F \\
&= \quad \sum_k \sum_l w_l \delta_m^k < \mathbf{K}^k, \mathbf{K}^l >_F + w_k \delta_m^l < \mathbf{K}^k, \mathbf{K}^l >_F \\
&= \quad \sum_k w_k < \mathbf{K}^k, \mathbf{K}^m >_F + \sum_l w_l < \mathbf{K}^l, \mathbf{K}^m >_F \\
&= \quad 2 \sum_k w_k < \mathbf{K}^k, \mathbf{K}^m >_F
\end{aligned}
$$

where δ_j^i is the Kronecker delta function

$$\delta_j^i = \begin{cases} 1, & \text{if } i = j \\ 0, & \text{if } i \neq j \end{cases}$$

The following parametrization takes care of the boundaries of the weights w_m, as expressed in equation 5.6, ensuring positiveness and normalization with the introduction of new parameters v_m:

$$w_m = \frac{\exp(v_m)}{\sum_l \exp(v_l)} \qquad \boxed{5.11}$$

Taking into account this parametrization, the partial derivatives are now computed with respect to the parameters v_m. Using the chain rule:

$$\frac{\partial A(\mathbf{v})}{\partial v_m} = \frac{\partial A(\mathbf{v})}{\partial w_m} \frac{\partial w_m}{\partial v_m}$$

$$= \frac{\partial A(\mathbf{v})}{\partial w_m} \frac{\exp(v_m) \sum_{l \neq m} \exp(v_l)}{(\sum_l \exp(v_l))^2} \qquad \boxed{5.12}$$

With the gradient descent based method we can estimate the maximum of our alignment function with respect to the parameters $\mathbf{v}$. The weights associated to the measures $\mathbf{w}$ can be easily calculated following equation 5.11.

5.3.5   Region Growing

We apply our method to region growing segmentation as a proof of concept of the presented distance learning method. The weights $\mathbf{w}$ that result from the previous distance learning method are used to define the distance that will drive the region growing segmentation algorithm. The algorithm starts growing from the initially selected ROI. During the growing process, the assignment of voxels is controlled by a homogeneity predicate based on the KTA obtained distance.

One class classifier

We define a predicate based on the **One Class Classifier** (OCC) methodology [126], normally used to distinguish a single class of 'normal' data points, from 'abnormal' ones: the outlier class. In our case, the normal class consists of all tensors belonging to the structure of interest to be segmented, and the outlier class consists of all the other DTs.

We start by defining two distance sample sets: a positive set D_P and a negative set D_N. D_P is the pairwise distances between the DTs from the positive set P (equation 5.13), while the negative distance set D_N is defined by the distances between the positive DT set P and the negative set N (equation 5.14).

$$D_{\mathbf{P}} = \{d(\mathbf{A}, \mathbf{B}) \mid \mathbf{A}, \mathbf{B} \in P \wedge \mathbf{A} \neq \mathbf{B}\} \qquad \boxed{5.13}$$

$$D_N = \{d(\mathbf{A}, \mathbf{B}) \mid \mathbf{A} \in P \wedge \mathbf{B} \in N\} \qquad \boxed{5.14}$$

Once the distance sample sets are defined, two histograms $H_P(d)$ and $H_N(d)$ are created for each distance set respectively. The optimal distance threshold t is then estimated by minimizing the miss-classification error of both positive and negative distances, as illustrated in figure 5.3.
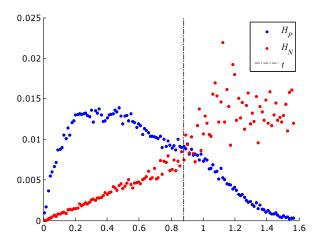
Figure 5.3: Example of the positive and negative distance histograms, for measure angular distance d_{ang}, with 256 sample distances, and the estimated optimal distance threshold.

The miss-classification error of the positive distances R_P is defined as the total area under $H_P(d)$ above threshold t (equation 5.15), which is equal to 1 minus the total area under $H_P(d)$ below t (equation 5.16). The miss classification error of the negative samples R_N is defined as the total area under $H_N(d)$ below t (equation 5.17):

$$R_P(t) \quad = \quad \sum_{i>t}^{\inf} H_P[i] \qquad\qquad \boxed{5.15}$$

$$= \quad \sum_{i=0}^{t} 1 - H_P[i] \qquad\qquad \boxed{5.16}$$

$$R_N(t) = \sum_{i=0}^{t} H_N[i] \qquad\qquad \boxed{5.17}$$

Assuming an equal cost for miss-classification of both classes, adding equation 5.16 and equation 5.17 yields total miss-classification error R_T (equation 5.18). The optimal upper distance threshold t is obtained through the minimization of this R_T.

$$R_T(t) = \sum_{i=0}^{t} (1 + H_N[i] - H_P[i]) \qquad\qquad \boxed{5.18}$$

The predicate is based on a voting system. Given a candidate DT $\mathbf{O}_i$, $M(\mathbf{O}_i)$ is defined by the elements of P whose distance to $\mathbf{O}_i$ is smaller or equal to t. The candidate is accepted if the relative amount of these distances is greater or equal to the defined quantile α:

$$PR(\mathbf{O}_i) = \begin{cases} true & \text{if } \frac{1}{|P|}|M(\mathbf{O}_i)| \geq \alpha \\ false & \text{otherwise} \end{cases} \qquad \boxed{5.19}$$

where $M(\mathbf{O}_i) = \{\mathbf{X} \mid \mathbf{X} \in P \wedge d(\mathbf{X}, \mathbf{O}_i \leq t)\}$, and $|P|$ denotes the number of elements in the set P.

In our experiments we use a majority vote, i.e. $\alpha = 0.5$.

The region growing algorithm, starting from the initially selected ROI, accepts the spatially connected voxels, whose respective tensors satisfy the predicate PR.

5.4   Results

The synthetic image shown in figure 5.4 was designed so that the regions, despite having distinct DTs, share some properties with other regions also. E.g., R1 has the same anisotropy as R2 and R3. With this synthetic data, we intent to illustrate the behaviour of the presented algorithm. To segment R1 an adequate distance must be chosen, for example d_{FA} would segment R1, R2 and R3 together.
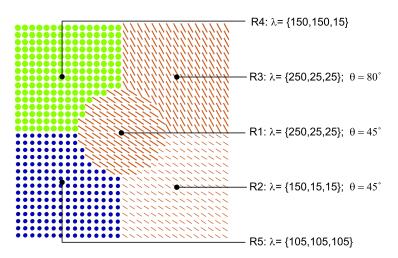
Figure 5.4: Superquadric glyphs [75] showing the five distinct regions in a 30×30 tensor synthetic image. DTs have λ as eigenvalues and the main eigenvector is rotated with angle θ (in the plane of the image depicted).

Choosing a ROI in R1, and randomly sampling 45 DTs, our algorithm estimates $w_{d_{L2}} = 1.0$ as the best discriminating distance. With these parameters, the region growing algorithm successfully segments only R1. Choosing a ROI between R1 and R2, the algorithm estimates a combination of two distances, $w_{d_{FA}} = 0.3$ and $w_{d_{ang1}} = 0.7$. As we can reason, what discriminates these two regions from the rest is their coherent orientation (45 degrees), distinct to R3, and FA, distinct to R4 and R5. The distance learning algorithm took about 1 second per example. All results were computed in a AMD Athlon 64 X2 Dual Core Processor 4800+ 2.41 GHz, with 3GB of RAM.

Figure 5.5 shows the algorithm applied to a DTI brain dataset. Two positive ROIs were selected within the *corpus callosum*. Because the random sampling of the brain selected several DTs in the gray matter, the algorithm infers d_{FA} as the most suitable measure. This results in the segmentation of the white matter. In order to improve this, a white matter masking is done by sampling of DTs with a FA threshold, i.e., 50 DTs are used as negative examples if $FA > 0.70$. Then, the algorithm estimates $w_{d_{FA}} = 0.5$ and $d_{ang1} = 0.5$ as the best discriminating combination of measures. The obtained result does not capture entirely the *corpus callosum*, as can be seen by the commissural fibers manually clustered by physicians. The result is not surprising since the defined region of interest does not represent the span of DTs orientations. The distance learning algorithm took 4 seconds to compute, with 50 negative samples and 40 positive samples.

In Figure 5.6 a positive ROI was selected within the right *cingulum*. With 30
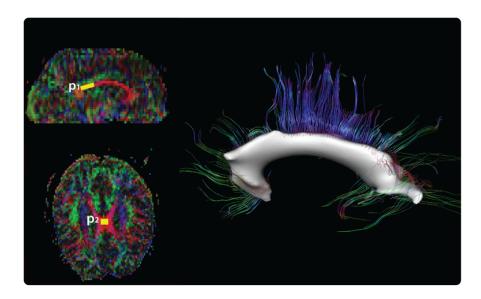
Figure 5.5: **Right:** Fusion of the segmented *corpus callosum*, in a $128 \times 128 \times 30$ DT volume, and the commissural fibers, colored using the typical RGB mapping of the main eigenvector. The estimated combination of distances is $d_{FA} = 0.5$ and $d_{ang1} = 0.5$. **Left:** P_1 and P_2 were used as positive ROIs.

random DTs taken with anisotropy $FA > 0.65$, the algorithm took 3 seconds to estimate $w_{d_{ang1}} = 1.0$ as the best measure, since the *cingulum* is a cylinder-like bundle with DTs coherently aligned.

5.5   Robustness analysis

The proposed distance learning algorithm is aimed at improving the segmentation of brain structures by semi-automatically finding the most appropriate measure for the task. Therefore it is important to evaluate the algorithm's robustness with respect to the user parameters: number of negative samples; initial ROI definition. Since the acquisition of a ground-truth for DT volumes is problematic, an objective quantitative evaluation of the segmentation quality is not possible.

For this evaluation we performed several repetitions of the algorithm with varying number of negative samples and with increasing seeding region sizes. To evaluate the consistency of the different resulting segmentations we used the Dice coefficient [22]. The Dice coefficient s is defined as the relative spatial overlap
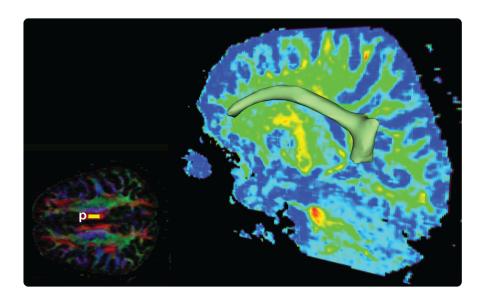
Figure 5.6: **Right:** Right *cingulum* segmented with estimated $w_{d_{ang1}} = 1.0$, in a 231 $\times$ 172 $\times$ 131 DT volume, with **p** (yellow) as positive ROI, as seen in the **Left**. The sagittal plane, on the right, shows the FA map while the plane on the left shows the RGB color coding of the main eigenvector (red: sagittal plane; green: transverse plane; blue: coronal plane).

of two volumes X and Y:

$$s = \frac{2|X \cap Y|}{|X| + |Y|} \qquad \boxed{5.20}$$

The performance of the proposed distance learning algorithm is evaluated on two different real DT volumes of the brain. The first data set has a relatively low resolution of $128 \times 128 \times 30$ (figure 5.7(a)), while the second data set has a higher resolution of $231 \times 172 \times 131$ (figure 5.7(b)). Prior to the segmentation, the DT volumes were masked by a lower FA limit of 0.6, preventing the segmentation of the gray matter.

5.5.1   Robustness to negative sample size

For different sizes of the negative DT set N, the distance learning and region growing algorithms were executed five times with increasing number of random negative samples. The overall robustness with respect to the random negative DT selection is then measured as the average of the pairwise Dice coefficients between these five segmentations.
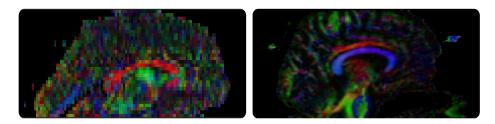
Figure 5.7: Sagittal cross-section of the (left) low and (right) high resolution DT volume used for the robustness analysis of the distance learning algorithm, using RGB color coding of the main eigenvector.
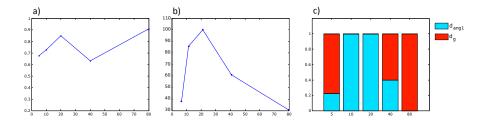
Figure 5.8: Robustness to negative sample size for low resolution *cingulum* segmentation: a) average Dice coefficient versus negative sample size; b) average segmentation size versus negative sample size; c) average weight distribution versus negative sample size.
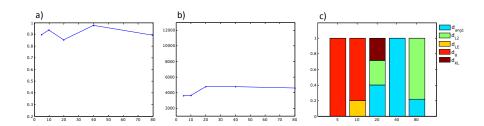
Figure 5.9: Robustness to negative sample size for high resolution *corona radiata* segmentation: a) average Dice coefficient versus negative sample size; b) average segmentation size versus negative sample size; c) average weight distribution versus negative sample size.

The segmentation of the *cingulum*, in the low resolution dataset, using a seeding region of about 20 voxels, shows a representative behaviour for the other segmentations in the same dataset. Figure 5.8(a) shows the average pairwise Dice coefficient, indicating the consistency of the segmentations with respect to the negative sample size. Two general local optima can be observed; one at 20

samples and another at 80 samples. However, figure 5.8(b) shows that the average segmentation size, using 80 negative samples approaches zero, indicating a low quality segmentation. The average weight distribution as returned by the distance learning algorithm is shown in figure 5.8(c). A general observation was that the distance learning algorithm returned one single dominant distance measure in most cases. In the case of 5 negative samples, the dominant distance was d_g, with four out of five runs of the algorithm, instead of d_{ang} with only one occurrence. The consistency of the segmentations appears to be significantly lower, when a variation in the distance measures is observed. Generally, the optimal amount of negative samples with respect to segmentation quality and consistency is 20. In this case, the dominant distance measure returned by the distance learning algorithm was d_{ang1}.

A similar behaviour can be observed in the segmentation of the *corona radiata* in the high resolution dataset, where a seeding region of about 50 voxels is used. The maximum average Dice coefficient, shown in figure 5.9(a), is achieved when using 40 negative samples. Contrary to the observations for the low resolution dataset, increasing the amount of negative samples does not effect the average segmentation size significantly as shown in figure 5.9(b). Figure 5.9(c) shows that the dominant distance weight returned by the distance learning algorithm, is d_{ang1} when using 40 negative samples.

5.5.2 Robustness to seeding region size

To illustrate the robustness of the segmentation methods with respect to the seeding region size again the *cingulum* in the low resolution dataset and the *corona radiata* for the high resolution dataset are used. All segmentations were performed using 20 negative samples. Seeding regions with increasing size were defined, keeping their location and homogeneity as constant as possible. Per seeding region, the distance learning and region growing algorithms were executed 5 times. For the low resolution *cingulum* segmentation, figure 5.10(a) shows that the average Dice coefficient appears to stabilize at a seeding region size of 20 voxels. Increasing the seeding region yields a slightly higher consistency for the one class classifier based predicate.

The segmentation sizes, shown in figure 5.10(b), yield realistic values for seeding region sizes of 20 and 40 voxels. Increasing the seeding region size to 80 voxels does not seem to lead to leaking of the segmented volume into other structures. For both seeding region sizes of 20 and 40 voxels, the dominant distance measure returned by the distance learning algorithm is d_{ang1}, as shown in figure 5.10(c).

The consistency of the *corona radiata* segmentation in the high resolution dataset also stabilizes at a seeding region of 20 voxels, as shown in figure 5.11(a). The
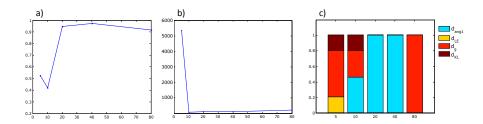
Figure 5.10: Robustness to seeding region size for low resolution *cingulum* segmentation: a) average Dice coefficient versus seeding region size; b) average segmentation size versus seeding region size; c) average weight distribution versus seeding region size.
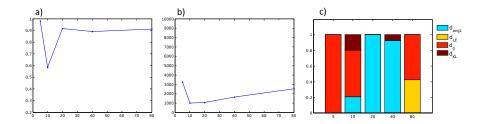
Figure 5.11: Robustness to seeding region size for high resolution *corona radiata* segmentation: a) average Dice coefficient versus seeding region size; b) average segmentation size versus seeding region size; c) average weight distribution versus seeding region size.

size of the segmentations, figure 5.11(b), again yields realistic values for seeding regions of 20 and 40 voxels. Increasing the seeding region to 80 voxels the segmentation size does not seem to lead to leaking into other structures, as expected from the FA masking. The dominant distance measure for seeding region sizes with most consistent segmentation results (20 and 40 voxels) is d_{ang1} as shown by figure 5.11(c).

Summary

Our experiments indicate that the method can be sensitive to the choice of size of the positive and negative DT samples. In our examples stable results were achieved with 20 samples for a low resolution dataset, and 40 samples for a higher resolution image.

5.6 Conclusions and future work

We proposed a novel distance learning method, based on kernel target alignment, for diffusion tensor imaging segmentation algorithms. As demonstrated, the method infers the most suitable distance(s) and parameters for the selected segmentation problem from the homogeneity/inhomogeneity characteristics of the data.

The used measures are of different nature and capture different aspects of the tensor data. Some measures isolate changes in individual degrees of freedom in the tensor data (e.g., difference in anisotropy). However, other measures, e.g., Log-Euclidean distance, d_{LE}, have no physiological significance and yield no clear intuition of distance between tensors. We present a flexible learning scheme that infers the combination of measures that better fit the data and object to be segmented. However, the resulting similarity measure will not be necessarily intuitive.

Furthermore the developed methods can be applied in other segmentation problems. For instance, Schultz [121] extended the use of structure tensors to diffusion tensor fields by combining Kindlmann's invariant gradients and rotation tangents [74, 76]. The invariant's weights used to define the distance measure are set in an ad-hoc way. Our framework could help in the definition of the weights needed to tune the segmentation, based on the specific problem at hand.

With different sizes of the positive and negative DT sample sizes, the distance learning algorithm returned inconsistent choices of measures. A probable reason is the fact that the distance learning algorithm optimizes the selected distance measure such that the distances between DTs from the positive set is minimized (intra positive class), the distance between the positive and negative DTs is maximized (inter classes), but also the distance between the negative DTs is minimal (intra negative class). This can lead to erratic results since the negative DTs are chosen randomly, and therefore can be very diverse. A possible improvement of the algorithm is to ignore the distance between DTs from set N and it will be addressed in future work.

In this chapter, we presented a proof of concept with synthetic data and real data showing the potential of the presented method. A robustness study was also performed to evaluate the quality of the results. However, doing an evaluation is a challenging problem, starting with the definition of a good ground truth.

The present algorithm could also be extended to HARDI (high angular resolution diffusion imaging) approaches to diffusion. Similarly, several distances between two spherical functions, such as DOT and Q-ball, can be defined [36, 140].

A Multi-resolution watershed-based approach for the segmentation of diffusion tensor images

"I am satisfied with the mystery of the eternity of life and with the awareness and a glimpse of the marvelous structure of the existing world, together with the devoted striving to comprehend a portion, be it ever so tiny, of the Reason that manifest."

Albert Einstein

Contents

6.1 Overview

The analysis and visualisation of diffusion tensor images is still a challenge since it is complex and exploratory in nature: we estimate tensors from the acquired images, trace fiber bundles, and analyse tracts. This quickly leads to clutter problems in visualisation but also in analysis. Furthermore, as we saw in the previous chapter 5, different tracts have different sizes, with different homogeneity characteristics.

In this chapter we therefore explore the intrinsic hierarchical nature of the brain tissue (recall chapter 2). A new framework for the multi-resolution segmentation of DTI is proposed. Based on fast and greedy watersheds operating on a multi-scale representation of a DTI image, a hierarchical depiction of a DTI image is determined conveying a global-to-local view of the fibrous structure of the analysed tissue. The multi-resolution watershed transform provides a coarse to fine partitioning of the data based on the (in)homogeneity of the gradient field. With a transversal cross scale linking of the basins (regions), a hierarchical representation is established.

This framework allows for a simple and interactive segmentation tool where different tracts can be segmented at different resolutions. Furthermore, it provides a novel hierarchical way to analyse DTI data.

We also present preliminary experimental results supporting the validity of the proposed method.

6.2 Introduction

Through chapter 2 we discussed the hierarchical nature of the brain tissue, ranging from the axons, organized in fiber bundles, and together forming coherent fiber tracts interconnecting different regions of the cortex, with different cognitive functions. This range of scales is also perceived in diffusion tensor images: the diffusion of water molecules is represented with tensors; fiber bundles are tracked from the tensor field; and bundles are grouped together into tracts. This hierarchical composition of the information drives the inspection of the data to be exploratory in essence.

Image segmentation is necessary to determine regions of interest where subsequent quantitative analysis and visualisation is performed. It provides tools to extract shape, appearance and other structural features than can then be used for the analysis of pathologies such as Parkinson's disease (Brunenberg et al. [17]) or, for instance, to identify the cognitive development of different types of population such as neonates (van Pul et al. [133]), the alterations in an aging brain (Salat et al. [119], Catani [19]).

Scalar image segmentation has been widely studied and different algorithms have been proposed through the years. However, DTI segmentation is still a challenging task. Some approaches have been proposed [118, 138, 141, 151], though, apart from the problem of defining a good distance as shown in chapter 4 and chapter 5, they often do not allow a full segmentation of the data, they segment one object at a time, have a multitude of parameters that must be set to achieve the desired result and have limited user interaction, preventing the added value of clinical users' expert knowledge in the segmentation.

In this chapter, we use a multi-scale watershed algorithm for the segmentation of DTI structures. The driving idea behind our framework is to help the user *focus* on accomplishing a given *analysis task*, by first presenting a *simplified* view of the data, while still maintaining a global one (i.e. *context*).

Given an input DTI, a *scale-space* representation is constructed [48]. The main motivation is that when increasing the scale, small details due to noise disappear, while main fibrous tissue structures (having predominant orientations) can still be reliably recovered. At each scale a watershed transform [91] is applied. Ideally, a structure would be outlined by a single region, however, in many situations, specially given the anisotropic nature of the involved tissues, the watersheds do not directly resolve to the anatomical structures. By linking several regions across scales, we infer a meaningful hierarchical representation of the data allowing novel ways to analyse and visualise diffusion tensor fields.

The main contributions of this chapter are:
- a *hierarchical representation* of the diffusion tensor field, which allows for *interactive grouping*, in an exploratory manner;

- a method for *linking* several interest regions across *scales*, such that one can infer a meaningful hierarchical representation of the data;

- a *multi-resolution segmentation* method, presented as a proof of concept of the proposed framework.

The multi-resolution watershed segmentation method is presented in section 6.4. In section 6.5, experimental results are presented supporting the validity of the methods.

6.3   Background

Different algorithms have been proposed for the segmentation of tensor fields. Zhukov et al. [151] proposed a level-set method over a scalar field derived from anisotropy measures. However this method fails to distinguish between regions with the same anisotropy but different direction.

Level-set methods using the full tensor information have been proposed by Zhizhou and Vemuri [138] and Rousson et al. [118], however, these iterative gradient descent based solutions seek a local solution and therefore are highly sensitive to initialization and parameter settings.

Watershed based methods, such as proposed by Rittner and Lotufo [91], are well known by their over-segmentation results. More recent and more efficient methods like the globally optimal graph-cuts have been applied to DTI by Weldeselassie and Hanarneh [141], however they provide a binary partition of the data, into one object and the background. Armstrong et al. [**?**] uses a livesurface approach for the segmentation of scalar volumes by establishing a 3D hierarchy of toboganned regions. This work provides an interactive segmentation tool, however, again, partitioning the data into one object and the background.

More recent work, such as Niethammer et al. [98], focus on the specific problem of segmenting a tubular structure such as the *cingulum*.

The concept of scale space has been widely studied in the image analysis field [49, 124]. An extra dimension is added to an image, representing the scale. It derives from the observation that objects are composed of different structures at different scales and therefore may appear different at different scales. The attractiveness of the method lies also in the link between scale-space theory and the biological process of vision.

Scale-space theory has been widely explored in many medical imaging problems [51]. Erik Dam and Ole Olsen's $\nabla Vision$[1] [30,31] tool, provides the multi-scale segmentation of 3D scalar volumes, such as Computed tomography (CT) images. A volume is partitioned in several regions and it is up to the user to

---

[1] http://www.itu.dk/image/appl/nablavision/index.html

sculpt his desired object by selecting the appropriate 'building blocks' in different scales. Although a somewhat cumbersome practice, this work has been used in clinical setup with promising results, given the gain in speed by orders of magnitude compared to manual segmentation.

Following the definition of a consistent multi-scale space for semi positive definite tensors by Florack and Astola [48], we extend those multi-resolution geometric studies to diffusion tensor fields. We present a multi-resolution segmentation algorithm that operates by applying the watershed transform to the different images in a generated DTI scale-space. The well-known watershed's over-segmentation, a shortcoming in most cases, is actually a core element of the presented framework.

This fast and simple partitioning of the tensor image, applied to each level in the scale-space, allows for the creation of the hierarchical representation of the data. This establishes an automated, general and interactive segmentation framework that equips the user with tools to explore and quickly segment these complex 3D tensor images.

6.4 Multi-resolution watershed segmentation
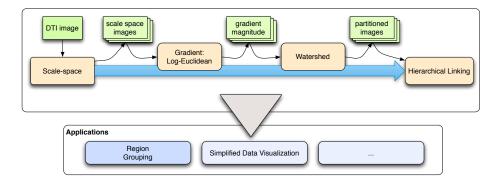
Figure 6.1: Global gist of the hierarchical segmentation.

In the following, we describe the different stages involved in the creation of a hierarchical representation of the data, see figure 6.1 for a global gist of the process. A scale-space stack of tensor field images is created by successively blurring the acquired DTI images. To each of these images, we calculate its tensor gradient magnitude, as a measure of homogeneity. The watershed transform is applied, thus obtaining a partitioning of the structures at each scale. Next, the several basins are linked to each other in a bottom-to-top manner. With this pipeline

a hierarchical representation of the data is obtained allowing further visualisation and interaction possibilities.

6.4.1   Scale-space representation of DTI

Florack and Astola [48] formulated a consistent scale-space representation for symmetric positive definite tensors. This follows the work proposed by Arsigny et al. [5] , Pennec et al. [45], and Fillard et al. [45], e.g. the so-called log-Euclidean framework.

To a tensor $\mathbf{D}$ is associated a unique logarithm map $\mathrm{Log}(\,\cdot\,)$, and its unique inverse the exponential map, $\mathrm{Exp}(\,\cdot\,)$. In practice, this framework consists in taking the matrix logarithm of tensors (by taking the scalar logarithm of $\mathbf{D}$ eigenvalues), running computations in this vector space, and mapping the result back to the tensor space with the matrix exponential.

The multi-scale representation for a tensor field f, at scale σ is achieved by the blurring operator

$$F(f,\sigma) = \mathrm{Exp}(\mathrm{Log}(f) * \phi_\sigma), \qquad \boxed{6.1}$$

where ϕ_σ is the isotropic Gaussian scale-space kernel in n dimensions, i.e.

$$\phi_\sigma(\mathbf{x}) = \frac{1}{\sqrt{2\pi\sigma^2}^n} \exp(-\frac{1}{2}\frac{\|\mathbf{x}\|^2}{\sigma^2}). \qquad \boxed{6.2}$$

Since in a multi-scale representation all scales are equivalent, a natural way to probe the archetype of DTI data is provided, resulting in a **coarse-to-fine** approach, see figure 6.2.

6.4.2   Watershed representation

The *watershed method* regards an image as a topological map, a landscape. As rain falls, water gathers in pools from the lowest points in the landscape. The landscape defines these pools, the *catchment basins*. As the water level rises, dams are built to prevent the merging of the pools. These boundaries constitute the *watersheds*. This concept was introduced by Maxwell (a Scottish theoretical physicist and mathematician) in [87]. Often the watershed provides a simple partitioning of the image based not on the original image, but based on a dissimilarity image. The gradient magnitude is a typical, simple and general measure defining region borders from the image edges. For this purpose we used the Log-Euclidean gradient.
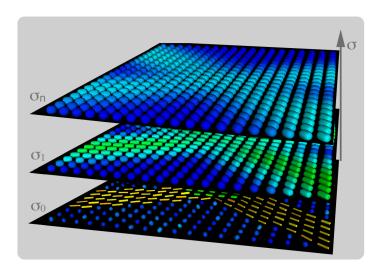
Figure 6.2: Superquadric glyphs [75] illustrating eq. (6.1) for three exponentially increasing scales σ_i. The synthetic DTI shows two fiber bundles crossing at a $65°$ angle.

Log-Euclidean gradient

At each scale, the Log-Euclidean gradient [5]

$$d_{LE}(\mathbf{A}, \mathbf{B}) = \sqrt{\operatorname{tr}((\operatorname{Log}(\mathbf{A}) - \operatorname{Log}(\mathbf{B}))^2)} \qquad \boxed{6.3}$$

is applied, hence objects are outlined with respect to the scale at which the gradient is calculated. Naturally, different object sizes require different scales.

Other dissimilarity measures for DTI watershed segmentation have been proposed such as in Rittner and Lotufo [116]. However, many different distance measures have been proposed, as mentioned in chapter 4. Which one to use is a problem *per se*. For consistency with the scale-space representation of the DTI data, we used the Log-Euclidean gradient as a common and generic dissimilarity measure. Other measures could be tested and might be more adequate for some problems, as investigated in chapter 5.

6.4.3 Cross-scale linking

Well-known scale-space theory [124] studies the progression of basins across scales. Because each basin is intrinsically related to the local minimum of the gradient magnitude, succinctly, basins can be afflicted by annihilation, creation, merge and split events. As the scale increases, the number of catchment basins decrease - they gradually merge into larger basins. These minima can be tracked by a linking process, forming a singularity string, across scales. With this process,

we combine the simplification provided by the (high level) selection scale with the fine scale basins in the (low level) representation scale (see figure 6.3). This method has been implemented, for clinical use, for scalar medical images with promising results [31].

Conceptually, in an iterative process, a region in a given scale is linked to the region at the next scale with maximum spatial overlap. From this linking process, a hierarchical representation of the data results, where each basin is linked to exactly one region at the next higher scale level, see also figure 6.5. This linking tree can be used for a "region focusing" process, where a simplified region at a selection scale is substituted by the regions at the lower representation scale. This selection scale determines the abstraction level that the user chooses to inspect the data, i.e. how simplified the data is.

| Original Tensor Field 4x downsampled | Watershed at Representation scale | Watershed at Selection scale | Selection linked to Representation scale |
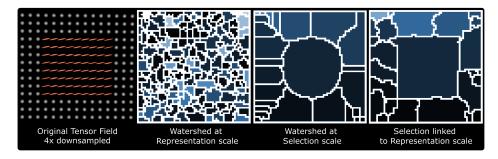
Figure 6.3: Region focusing across scales. Following the linking tree from the *selection scale* to the representation scale, fine detail is obtained; last image shows the regions at the *selection scale* represented with basins of the lower, *representation scale*.

The stack of scale-space images is produced by blurring the data-scale DTI image with increasing σ. In order to obtain a sufficiently fine space, considering its exponential nature, σ is changed per scale i

$$\sigma_i = \sigma_0 \exp(\lambda_i) \, , \lambda_i = i \frac{-\ln(\Delta_\lambda)}{3}, \qquad \boxed{6.4}$$

where parameter Δ_λ indicates the ratio at which the number of basins decrease, see figure 6.4.

6.4.4 Region grouping

With a partitioned *selection scale*, given by applying the watershed transform to the selected scale, a connected graph $G(n_i, e_j)$ is built where: each node n_i holds the Log-Euclidean mean tensor [5], representing the corresponding basin; each edge e_j holds the Log-Euclidean distance [5] between each neighbouring basin.
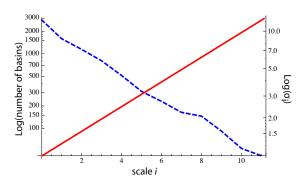
Figure 6.4: Graph plot showing the relation between increasing the scale-space parameter σ, in red, and the (decreasing) number of basins, in dashed blue, using $\Delta_\lambda = 0.5$.

The average weight of all the edges within users' selected basins (sample seeding basins) μ is taken as predicate to a simple region growing algorithm operating on the edges of graph G. A new, spatially connected, basin is added if the edge connecting it to the growing region is less than the average $\mu \times r$, where r is a user defined ratio. This algorithm quickly groups similar connected basins. These can then be 'focussed' to the lowest *representation scale*, as illustrated in figure 6.5.

6.5 Results

We start by utilizing the synthetic image described before in chapter 5, section 5.4, shown in figure 6.5. This image has five distinct regions, with different DT populations. Despite having distinct DTs, some properties are shared between regions but are different to others, e.g., R4 has the same anisotropy as R3 and R5, but different main diffusivity direction to R3. With this synthetic data we illustrate the hierarchical nature of the proposed segmentation.

The DTs in the five regions have realistic eigenvalues: $\lambda_{R1} = [3, 10, 10] \times 10^{-3} \, \text{mm}^2/\text{s}$, $\lambda_{R2} = [5, 5, 5] \times 10^{-3} \, \text{mm}^2/\text{s}$, $\lambda_{R3} = [14, 2, 2] \times 10^{-3} \, \text{mm}^2/\text{s}$, $\lambda_{R4} = [17, 3, 3] \times 10^{-3} \, \text{mm}^2/\text{s}$ and $\lambda_{R5} = [17, 3, 3] \times 10^{-3} \, \text{mm}^2/\text{s}$. In order to mimic a real DTI acquisition, noise was added to the 128×128 synthetic image. From the noiseless DTs, using the inverse Stejskal and Tanner [125] relation, the signal attenuations were obtained. To each direction value Rician noise was added with $SNR = 15.3$, and the tensors re-estimated.

Figure 6.5 shows the obtained basins for the computed scale-space. We can see the segmentation consequences of choosing appropriate selection scales and automatically tracking down objects to the lowest representation scale. As the selection scale decreases, we clearly see the hierarchical segmentation occurring.
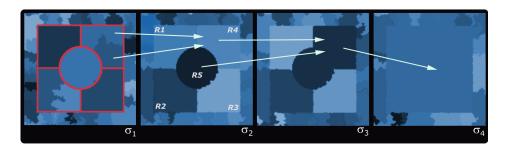
Figure 6.5: The hierarchical nature of the watershed scale-space segmentation. Different basins are linked to the ones above. Colours indicate different regions. Red outline indicates the original five regions.

At scale σ_3 we obtain regions R4 and R5 as a unique object, since they are very similar, whereas at scale σ_2 these two regions constitute different segmentation objects.

Figure 6.6 shows the results of the semi-automatic segmentation tool provided by the presented watershed based method. In a $30 \times 30 \times 20$ volume, two fiber bundles forming 'tubes' with radii of 2 voxels intersect each other. Here, the tensors, with eigenvalues $\lambda = [17, 3, 3] \times 10^{-3}$ mm^2/s and oriented tangentially to the centre line of the tube, are estimated using a mixed tensor model (chapter 3, section 3.3.1.1). Rician noise is added with $SNR = 15.3$. In this experiment, due to the use of an isotropic Gaussian scale-space kernel, no appropriate selection scale is found in order to achieve the desired segmentation: automatically distinguish the two tubes. Users' input is here necessary. An appropriate selection scale is chosen by the user, i.e. an appropriate simplification of the data is selected. Then, in a typical visualisation such as FA map, the user interactively picks some points in the structure of interest (see figure 6.6). These points select the correspondent basins, at the simplified selection scale, which are then linked down to the representation scale aggregating the corresponding voxels forming the final segmentation.

A qualitative evaluation is performed on a real $128 \times 64 \times 64$ DTI brain image. Employing an isotropic Gaussian kernel in a brain dataset, as in figure 6.7, fails to deliver an appropriate high level selection scale for the automatic segmentation of the larger structures. Brain tissue and white matter manifest itself in DTI images as anisotropic structures. In the isotropic scale-space these structures of interest successively disappear. Nevertheless, the semi-automatic segmentation still provides the desired segmentation result.

A typical user interaction follows, as illustrated in figure 6.7:

1. An appropriate selection scale is chosen, by inspection of the level of simplification provided by a scale, with a view of the corresponding basins;
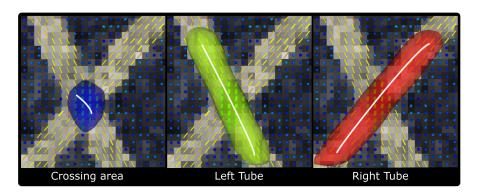
Figure 6.6: Synthetic image with two crossing tubes. User's interactive picking are represented by the white lines. Three segmentation results arise depending on the stroke selection: crossing region, tube one or tube two.

2. In a typical main eigenvector RGB colour mapping (figure 6.7(left)), by stroking some points in the image, seeding basins at the chosen scale are selected (figure 6.7(middle));

3. The region grouping algorithm collects similar connected basins, which are then 'focussed' into the representation scale.

Figure 6.7 (right) shows two tracts segmented in two different scales: the *corpus callosum* in red, visible in a higher scale σ_4; and the *cingulum* in green, visible in a lower, more detailed scale σ_2.

6.6 Conclusions

In this chapter we presented a hierarchical segmentation algorithm, using watersheds, for diffusion tensor fields. In a semi automatic manner a hierarchical representation of the data is assembled providing a new way to visualize and interact with this type of data. A multitude of possibilities arise. The automatic partitioning of the data can be used to assess statistical properties of the data. Combined with other typical uses of DTI like fiber tracking, this can be used to study the connectivity of the different regions of the brain. This framework is general and can be used for other medical purposes.

6.7 Discussion

However, we are not there yet. In the following we discuss some issues for improvement and suggestions for future work.
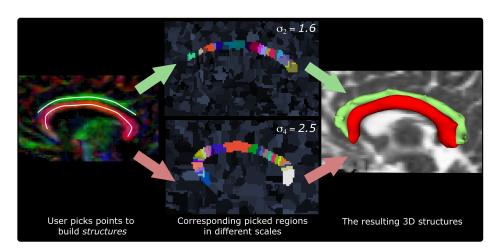
Figure 6.7: A $128 \times 64 \times 64$ DTI brain image. The user selects points in his segmentation task (Left), the corresponding basins in the respective Selection scale are highlighted (Middle) and a smooth isosurface wraps the grouped similar connected regions (Right).

The use of Log-Euclidean gradient embodies the watershed segmentation with an unclear intuition. It expresses dissimilarity between the diffusion tensors, however its connection to the underlying tissue structure is not clear. The use of other gradient alternatives, such as the tensorial morphological gradient [116] allowing the use of known distance measures (chapter 4), or Schultz's structure tensor [121] combining the intuitive Kindlmann's [76] decomposition of the tensor changes into six different gradients will be subject of future study.

Given the anisotropic nature of the brain tissue (as for muscle fibers), the use of a Gaussian isotropic kernel hinders the creation of an adequate scale space, preventing the hierarchical linking to give the desired results. However, the proposed tool is versatile enough to allow the semi-automatic segmentation of the structures of interest by combining the several building blocks. Future work will study the use of anisotropic kernels based on the diffusion tensor at each voxel as in Chung et al. [24] so to improve the inference of tissue archetype.

A more robust and anatomically significant improvement to the hierarchical linking lies on the use of tensor measures, with an underlying intuition, in addition to the spatial overlap between adjacent scales, since a tract, by definition, groups together similar DTs.

Recent research has been exploiting MRI techniques in order to obtain microstructural information of the brain tissue (Alexander et al. [4], Assaf et al. [6]) such as axon diameter and axon density, as to provide more specific biomarkers than standard indices from diffusion tensor imaging. The goal is to provide a non-invasive live *in-situ* histology, replacing tissue biopsy. The incorporation
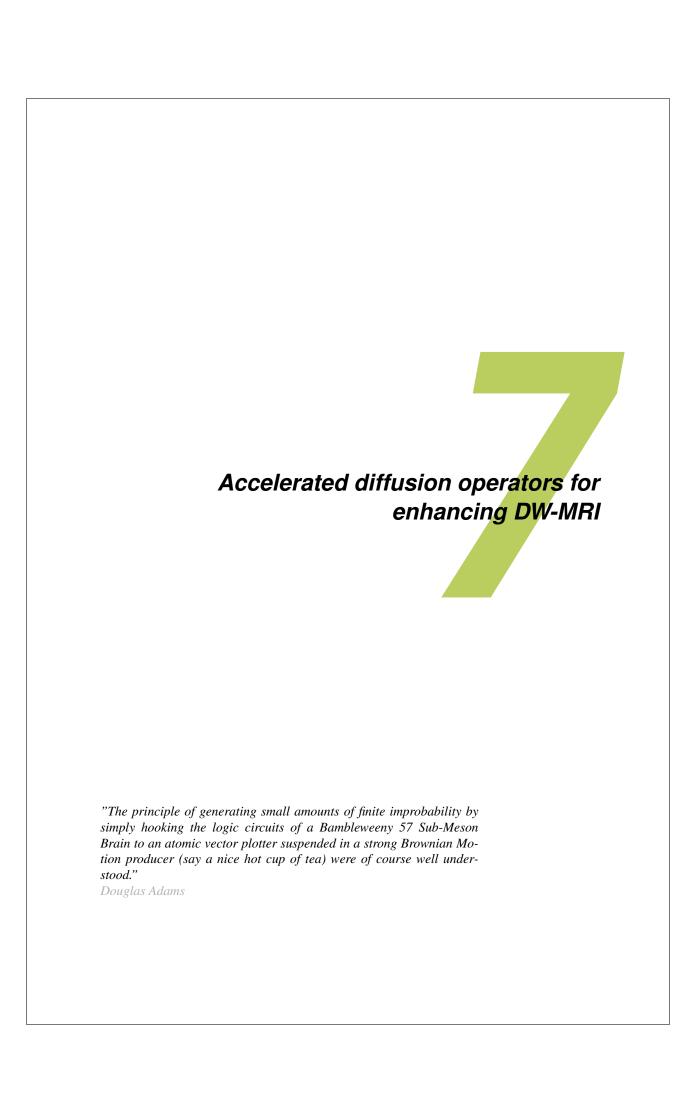
of these indices in the presented hierarchical framework would compliment the available information, specially at a lower scale.

The creation of a meaningful hierarchical representation of the data unravels new visualisation and interaction possibilities and thus novel ways to study diffusion tensor fields. One can use well-established techniques for graph visualization and interaction, to directly manipulate the data. Illustrative rendering [101] may as well be used to augment the more abstract graph visualization. Jianu et al. [65] proposed a visual exploration paradigm to navigate through the complex 3D fiber tracts using 2D linked representations. Their results suggest that this combination of traditional 3D model viewing with lower dimensional information visualisation techniques can improve the exploration of the connectivity of the brain.

The proposed hierarchical methodology is also relevant for high angular resolution diffusion imaging (HARDI) given its close analytical connection to DTI.

Artificial and real data show the potential of the presented method. However, doing a more elaborated evaluation would be necessary, and a comparison to other methods should also be performed.

110

Accelerated diffusion operators for enhancing DW-MRI

"The principle of generating small amounts of finite improbability by simply hooking the logic circuits of a Bambleweeny 57 Sub-Meson Brain to an atomic vector plotter suspended in a strong Brownian Motion producer (say a nice hot cup of tea) were of course well understood."

Douglas Adams

Contents

7.1   Overview

As we saw in chapter 2, in particular section 2.4.2, high angular resolution diffusion imaging (HARDI) is an MRI imaging technique that is better able to capture the intra-voxel diffusion pattern than its simpler predecessor diffusion tensor imaging (DTI). However, HARDI in general produces very noisy diffusion patterns and it still exhibits limitations in areas where the diffusion pattern is asymmetrical (bifurcations, splaying fibers, etc.). Enhancement and denoising of the data based on context information would help to minimize these problems. In order to achieve it, convolutions are performed in the coupled spatial and angular domain. However, these approaches have high computational complexity of an already complex HARDI data processing.

In this chapter, we present an accelerated framework for HARDI data regularization and enhancement. The convolution operators are optimized by pre-calculating the kernels, analysing kernels shape and utilizing look-up-tables. We provide an increase of speed, compared to previous brute-force approaches of simpler kernels. These methods can be used as a preprocessing for tractography and can lead to new ways for investigation of brain white matter.

7.2 Introduction

In diffusion tensor imaging (DTI), the prominent local orientation of the fiber bundles can be estimated. In DTI the local diffusivity pattern is approximated by a 2^{nd}-order diffusion tensor (DT). Although simple analysis is possible with established mathematical frameworks, these DTs fail to capture more complex fiber configurations than a single fiber bundle, such as crossings, bifurcations and splaying configurations.

To solve this problem, as we described earlier in chapter 2, section 2.4.2, new approaches based on high angular resolution diffusion imaging (HARDI) were pioneered by Tuch [130]. In HARDI more sophisticated models are employed to reconstruct more complex fiber structures and to better capture the intra-voxel diffusion pattern.

It is important to note that all of the diffusion weighted MRI modelling techniques ([?, 35, 64, 102, 104, 128, 130, 131]) model functions that reside on a sphere. For simplicity we will refer to them as spherical distribution functions (SDF). Whereas the physical meaning of these SDFs can be different (a probability density function (PDF), iso-surface of a PDF, orientation distribution function (ODF), fiber orientation distribution (FOD), etc.), in all cases they characterize the intra-voxel diffusion process, i.e. the underlying fiber distribution within a voxel. Due to the limitations in acquisition, the SDF is always antipodally symmetric and therefore can only model single fiber tracts or symmetric fiber crossing configurations. However, this can not be assumed everywhere in the white matter of the brain, especially in structures as *optic chiasm*, *the hippocampus*, *the brain stem* and others.

In HARDI, high b-values are needed to be able to capture the more complex profiles. Therefore, HARDI produces, in general, noisy diffusion patterns due to the low SNR. To overcome these limitations, post-processing of the data is crucial. As commonly done in image processing, the noise can be reduced and the data enhanced by taking into account the information in a close neighbourhood (i.e. the context).

Previous research has been done on smoothing and regularization of DTI and HARDI images [46, 47, 59, 72], however, these approaches do so by considering the spatial and orientational domains separately. In these approaches diffusion is only performed over the spherical function per voxel (i.e. the angular part). By not considering the neighbourhood information, these methods often fail at interesting locations with composite structure, since locally a peak in the profile can be interpreted as noise and therefore it can be smoothed out.

In recent work the diffusion process is done by considering the full domain, i.e. considering both spatial and orientational neighbourhood information. In Barmpoutis et al. [8] the estimated asymmetric spherical functions, called tractosemas,

are able to model local complex fiber structures using inter-voxel information. Duits and Franken [39] proposed a useful framework for the cross-preserving smoothing of HARDI images by closely modelling the stochastic processes of water molecules (i.e. diffusion) in oriented fibrous structures. These approaches increase the complexity of already complex and computationally heavy processing of HARDI data. These processes require 5D convolutions, thus a great deal of computation time.

In the work presented in this chapter, we establish a faster framework for noise removal and enhancement of HARDI datasets. We optimize the convolution operators by pre-calculating the kernels, analysing kernel's shape and accelerating convolution using the look-up-tables concept. Compared to previous brute-force approaches, we provide a significant increase of speed, enabling a contextual processing framework of HARDI data.

In Section 7.3, we start by establishing the mathematical basis on which the HARDI convolution method lives. The accelerated convolution framework is presented in Section 7.4. Subsequently, in section 7.5, we present experimental results, both in artificial and real HARDI data, supporting the validity and improvements of the method.

7.3   Background

In this section we will provide a self-contained introduction to the convolution of HARDI data over the joint domain of positions and orientations. Several kernels for these convolutions will also be exemplified as illustration of the presented method.

7.3.1   Theory

Diffusion weighted MRI modelling techniques estimate functions that reside on a sphere, the spherical distribution functions (SDF). Therefore, a HARDI image is a function not only on positions but also on orientations:

$$\mathcal{U} : \mathbb{R}^3 \rtimes S^2 \to \mathbb{R}^+ : \mathcal{U}(\mathbf{y}, \tilde{\mathbf{n}}(\tilde{\beta}, \tilde{\gamma})) \qquad \boxed{7.1}$$

Meaning that at every position $\mathbf{y} \in \mathbb{R}^3$, the probability of finding a water particle moving in a certain direction

$$\tilde{\mathbf{n}}(\tilde{\beta}, \tilde{\gamma}) = (\sin\tilde{\beta}, -\sin\tilde{\gamma} \ \cos\tilde{\beta}, \cos\tilde{\gamma}\cos\tilde{\beta})^T \in S^2, \qquad \boxed{7.2}$$

is given as a positive scalar. Here, $\tilde{\mathbf{n}}(\tilde{\beta}, \tilde{\gamma})$ is a point on the unit sphere parametrized by $\tilde{\beta} \in [-\pi, \pi)$ and $\tilde{\gamma} \in [-\frac{\pi}{2}, \frac{\pi}{2})$ (see figure 7.1).

To stress the coupling between orientation and positions we write $\mathbb{R}^3 \rtimes S^2$ rather than $\mathbb{R}^3 \times S^2$. Note that only 2 angles $(\tilde{\beta}, \tilde{\gamma})$ suffice, as rotation around the directional axis $\tilde{\alpha}$ is irrelevant.
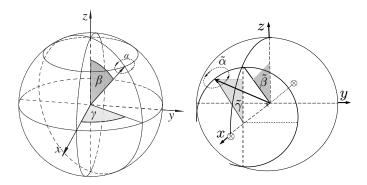
Figure 7.1: The chart on the left shows the common Euler-angle parametrization of the sphere S^2, while we use the parametrization on the right using $\tilde{\beta}$ and $\tilde{\gamma}$. Note that the first parametrization has singularities at the north and south-pole, while the second has singularities at $(\pm 1, 0, 0)$, for more convenient use in eq. 7.9.

7.3.2   Convolutions

An operator $\mathcal{U} \mapsto \Phi(\mathcal{U})$ on an SDF should be Euclidean invariant, i.e. independent on the choice of orthonormal coordinate system. In other words rotating and translating HARDI input $\mathcal{U}$ corresponds to rotating and translating the output $\Phi(\mathcal{U})$. If such operators, designed for smoothing and enhancement of HARDI data, are linear then these operators can be written as a HARDI-convolution:

$$
\begin{aligned}
&(\Phi(\mathcal{U}))(\mathbf{y}, \mathbf{n}) \\
&= \int_{\mathbb{R}^3} \int_{S^2} p(R_{\mathbf{n}'}^T(\mathbf{y} - \mathbf{y}'), R_{\mathbf{n}'}^T(\mathbf{n})) \, \mathcal{U}(\mathbf{y}', \mathbf{n}') \, d\mathbf{y}' \, d\sigma(\mathbf{n}') \\
&= \int_{\mathbb{R}^3} \int_{S^2} k(\mathbf{y}, \mathbf{n}; \mathbf{y}', \mathbf{n}') \, \mathcal{U}(\mathbf{y}', \mathbf{n}') \, d\mathbf{y}' \, d\sigma(\mathbf{n}')
\end{aligned}
\tag{7.3}
$$

where

- $\mathcal{U}$ denotes the input HARDI dataset

- $\Phi(\mathcal{U})$ denotes the output HARDI dataset (obtained by applying a convolution operation to the input with p)

- $k(\mathbf{y}, \mathbf{n}; \mathbf{y}', \mathbf{n}')$ is the full kernel in the kernel operator

- $p(\mathbf{y}, \mathbf{n})$ is the **convolution kernel** related to $k(\mathbf{y}, \mathbf{n}; \mathbf{y}', \mathbf{n}')$ by means of

$$p(\mathbf{y}, \mathbf{n}) = k(\mathbf{y}, \mathbf{n}; 0, \mathbf{e}_z), \text{ with } \mathbf{e}_z = (0, 0, 1)^T$$

  From this moment, kernels will be noted as $p(\mathbf{y}, \mathbf{n})$, i.e. the a priori probability density of finding a fiber fragment at position $\mathbf{y}$ with orientation $\mathbf{n}$, $(\mathbf{y}, \mathbf{n})$, given that there is a fiber fragment at $(\mathbf{0}, \mathbf{e}_z)$

- $R_{\mathbf{n}}$ is *any* rotation such that $R_{\mathbf{n}}\mathbf{e}_z = \mathbf{n}$. The choice of $R_{\mathbf{n}}$ does not matter as long as p has a symmetry with respect to rotations around $\mathbf{e}_z$ [39, Corr.1]

- σ denotes the surface measure on the sphere.

As mentioned previously, convolutions can operate over different domains, obviously with different outcomes. Next we consider the special cases of eq. 7.3.

Spatial domain   filtering can be applied to each of the directions without relating the directions between each other:

$$(\Phi(\mathcal{U}))(\mathbf{y}, \mathbf{n}) = \int_{\mathbb{R}^3} q(\mathbf{y} - \mathbf{y}')\,\mathcal{U}(\mathbf{y}', \mathbf{n})\,d\mathbf{y}' \qquad \boxed{7.4}$$

This relates to eq. 7.3, if p is set as a product of a spatial kernel q with a delta-spike in orientation space.

Orientational domain   filtering can be applied to each voxel independently, i.e. considering each SDF independently from each other. In this way, each voxel is smoothed locally:

$$(\Phi(\mathcal{U}))(\mathbf{y}, \mathbf{n}) = \int_{S^2} r(R_{\mathbf{n}'}^T \mathbf{n})\,\mathcal{U}(\mathbf{y}, \mathbf{n}')\,d\sigma(\mathbf{n}') \qquad \boxed{7.5}$$

Similarly, this relates to eq. 7.3 where p is a product of an angular kernel q with a delta-spike in position space.

However, appropriate treatment of crossings and bifurcations requires regularization along oriented fibers (where position and orientation are coupled) and consequently our *a priori* fiber extension probabilities $p : \mathbb{R}^3 \rtimes S^2 \to \mathbb{R}^+$ should not consist of a delta-spike in position space nor in orientation space. This means we should not restrict ourselves to eq. 7.4 or 7.5. Next we explain how to discretize full convolutions (eq. 7.3) on positions and orientations.

7.3.3 Discretization

For computational purposes, these functions are usually discretized. For the orientation domain this is done by nearly uniform sampling on the sphere using a method such as tessellation of an icosahedron (see figure 7.2).

Figure 7.2: Discrete samplings of the sphere corresponding to order 1, 2 and 3 of tessellation of an icosahedron, with correspondent $|T_1| = 12$, $|T_2| = 42$ and $|T_3| = 162$ points.

Having a discrete lattice of SDFs (the HARDI image $\mathcal{U}$), the integral over $\mathbb{R}^3$ in eq. 7.3 becomes a summation over the lattice. Since, typically, a kernel is stronger around its center (at position $\mathbf{y}$), a set of voxels P can be defined containing the lattice positions neighbour of $\mathbf{y}$. Additionally, since the SDFs are discretized over the sphere (see figure 7.2), the integral over S^2 becomes a summation over tessellation's vectors, the set T. Using these discretizations, eq. 7.3 becomes:

$$\Phi(\mathcal{U})[\mathbf{y}, \mathbf{n}_k] = \sum_{\mathbf{y}' \in P} \sum_{\mathbf{n}' \in T} q_{\mathbf{y},\mathbf{n}_k}[\mathbf{y}', \mathbf{n}'] \mathcal{U}[\mathbf{y}', \mathbf{n}'] \ \Delta \mathbf{y}' \Delta \mathbf{n}' \qquad \boxed{7.6}$$

where $\Delta \mathbf{y}'$ is the discrete volume measure and $\Delta \mathbf{n}'$ the discrete surface measure, which in the case of (nearly) uniform sampling of the sphere can reasonably be approximated by $\frac{4\pi}{|T|}$. Kernel $q_{\mathbf{y},\mathbf{n}_k}$ is the rotated and translated correlation kernel, such that it is aligned with $(\mathbf{y}, \mathbf{n}_k)$, associated with p as later explained in section 7.4.

One should notice the complexity involved in HARDI convolution. Given:

- $|P|$: number of points in kernel's lattice

- $|T|$: number of vectors in kernel's tessellation (and SDFs of the input HARDI data)

The discretized convolution expressed in eq. 7.6, per voxel of the input data, has complexity of $O(|T| \ |P| \ |T|)$. For instance, consider the convolution with a kernel discretized in a $3 \times 3 \times 3$ lattice, for 2^{nd} order tessellation ($|T| = 42$ directions). The discrete convolution in eq. 7.6, *per voxel* in the lattice of the input HARDI image, involves $42 \times 27 \times 42 = 47628$ multiplications.

7.3.4 Tractosemas

In the work of Barmpoutis et al. [8], a field of asymmetric spherical functions, called tractosemas, is extracted from a field of SDFs. The kernel that governs the smoothing process is defined as a function over space and orientation, as described before. The proposed kernel intuitively describes when a structure should be enhanced. It is constructed as a direct product of three parts involving von Mises (circular normal distribution) and Gaussian probability distributions :

$$
\begin{aligned}
k(\mathbf{y}, \mathbf{n}; \mathbf{y}', \mathbf{n}') &= k_{\text{dist}}(\|\mathbf{y} - \mathbf{y}'\|) \cdot k_{\text{orient}}(\mathbf{n} \cdot \mathbf{n}') \cdot \\
&\quad k_{\text{fiber}}\left(\tfrac{\mathbf{n}}{\|\mathbf{y}-\mathbf{y}'\|} \cdot (-(\mathbf{y} - \mathbf{y}')) \right),
\end{aligned}
\tag{7.7}
$$

where the different factors are given by

$$
k_{\text{dist}}(\|\mathbf{y} - \mathbf{y}'\|) = \frac{1}{(2\pi\sigma)^{\frac{3}{2}}} e^{-\frac{\|\mathbf{y}-\mathbf{y}'\|^2}{2\sigma^3}} ,
$$
$$
k_{\text{orient}}(\cos\phi) = k_{\text{fiber}}(\cos\phi) = \frac{\kappa e^{\kappa \cos(\phi)}}{4\pi \sinh(\kappa)} ,
$$

with $\phi \in (-\pi, \pi]$ being the angle, respectively, between the vectors $\mathbf{n}$ and $\mathbf{n}'$ and the angle between the vectors $\mathbf{n}$ and $(\mathbf{y} - \mathbf{y}')$. The two scale parameters σ and κ control kernel's sharpness. Figure 7.3 shows an example of the tractosemas kernel $p^{\sigma,\kappa} : \mathbb{R}^3 \rtimes S^2 \to \mathbb{R}^+$ given by

$$
p^{\sigma,\kappa}(\mathbf{y}, \mathbf{n}) = \frac{1}{4\pi} k_{\text{dist}}(\|\mathbf{y}\|) k_{\text{orient}}(\mathbf{e}_z \cdot \mathbf{n}) k_{\text{fiber}}(-\|\mathbf{y}\|^{-1}\mathbf{n} \cdot \mathbf{y}).
\tag{7.8}
$$

7.3.5 Diffusion Kernels

Duits et al. [39, 40] proposed a kernel based on solving the diffusion equation for HARDI images. The full derivation is out of the scope of this manuscript. This kernel represents the Brownian motion kernel, on the coupled space $\mathbb{R}^3 \rtimes S^2$ of positions and orientations. This kernel satisfies the two important requirements for a diffusion kernel:

1. **left-invariant** The kernel satisfies the right symmetry constraints, [39, Corr.1]. Thereby rotation and translation of the input $\mathcal{U}$ corresponds to rotation and translation of the output $\Phi(\mathcal{U})$.

2. **fulfill the semigroup property** When the operator is applied iteratively, the scales can be added. This allows us to simply add the kernels and process once, than to process them multiple times.
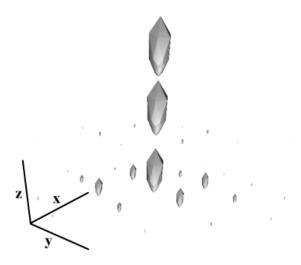
Figure 7.3: Example of the tractosemas kernel (7.8) [8]. Computed with scale parameters $\sigma = 1.3$ and $\kappa = 4$, for orientation $\mathbf{e}_z$.

The diffusion equation is solved by convolution (eq. 7.3) with the Green's function for the diffusion equation on the coupled space $\mathbb{R}^3 \rtimes S^2$ of positions and orientations. It describes Brownian motion on positions and orientations, where the angular part of a random walk prescribes the tangent vector to the spatial part of the trajectory. Next we present the close analytic approximation of the Green's function as discussed in [39]. This approximation is a product of two $2D$ kernels on the coupled space $p_{2D} : \mathbb{R}^2 \rtimes S^1 \to \mathbb{R}^+$ of 2D-positions and orientations:

$$
\begin{aligned}
p_{3D}^{D_{33},D_{44},t}&((x,y,z)^T, \tilde{\mathbf{n}}(\tilde{\beta},\tilde{\gamma})) \approx \\
&N(D_{33},D_{44},t) \cdot p_{2D}^{D_{33},D_{44},t}((z/2,x),\tilde{\beta}) \cdot \\
&p_{2D}^{D_{33},D_{44},t}((z/2,-y),\tilde{\gamma}) \,,
\end{aligned}
\tag{7.9}
$$

we recall eq. 7.1 and eq. 7.2, where $\mathbf{y} = (x,y,z)^T$, and where

$$
N(D_{33},D_{44},t) \approx \frac{8}{\sqrt{2}} \sqrt{\pi} t \sqrt{t D_{33}} \sqrt{D_{33} D_{44}}
$$

ensures that the total integral over positions and orientations is 1. The $2D$ kernel is given by:

$$p_{2D}^{D_{33},D_{44},t}(x,y,\theta) \equiv \frac{1}{32\pi t^2 c^4 D_{44} D_{33}} e^{-\frac{\sqrt{\text{EN}((x,y),\theta)}}{4c^2 t}} \qquad \boxed{7.10}$$

where we use short notation

$$EN((x,y),\theta) = \left(\frac{\theta^2}{D_{44}} + \frac{\left(\frac{\theta y}{2} + \frac{\theta/2}{\tan(\theta/2)}x\right)^2}{D_{33}}\right)^2$$
$$+ \frac{1}{D_{44}D_{33}}\left(\frac{-x\theta}{2} + \frac{\theta/2}{\tan(\theta/2)}y\right)^2$$

where one can use the estimate $\frac{\theta/2}{\tan(\theta/2)} \approx \frac{\cos(\theta/2)}{1-(\theta^2/24)}$ for $|\theta| < \frac{\pi}{10}$ to avoid numerical errors. c is a positive constant for rescaling the diffusion time t.

The diffusion parameters D_{33} and D_{44} and stopping time t allow the adaptation of the kernels to different purposes:

1. $t > 0$ determines the overall size of the kernel, i.e. how relevant the neighbourhood is

2. $D_{33} > 0$, the diffusion along the principal axis, determines how wide the kernel is

3. $D_{44} > 0$ determines the angular diffusion, so the quotient D_{44}/D_{33}, models the bending of the fibers along which diffusion takes place

7.4   Accelerated convolution

A convolution in the full HARDI domain, as addressed in section 7.3.2, is a complex task, dependent on the number of points in the kernel's lattice and number of vectors in the tessellation of the sphere, the same for both kernel and the input SDF. Applying these operations in a real dataset and for smoother (higher) orders of tessellation quickly escalates into a time consuming process.

How can this process be accelerated?

- **Pre-computing** - One immediate optimization is to pre-calculate and store the kernel for every direction $\mathbf{n}_k$, instead of calculating on-the-fly the respective kernel $q_{\mathbf{y},\mathbf{n}_k}$ per position $\mathbf{y}$ and direction $\mathbf{n}_k$. This is allowed, as the kernels are not adaptive to the data, i.e. they do not change depending on each voxel.
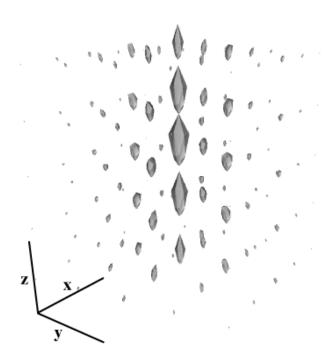
Figure 7.4: Diffusion kernel proposed in [39] computed with parameters $D_{33} = 1.0$, $D_{44} = 0.04$ and $t = 1.4$, for principal axis orientation $\mathbf{e}_z$.

- **Truncation** - As we can see in figure 7.5, these kernels typically exhibit an interesting characteristic: the probability of diffusion is larger at the locations around the starting direction e_z, and quite small at locations further from it depending on the values of the parameters D_{33} and D_{44}. We truncate the kernel such that only the meaningful directions are considered in the convolution.

Next, we explain the details of these procedures.

7.4.1 Pre-computing

Recall that in convolutions one shifts over a dummy variable $\mathbf{y}'$ whereas in correlations one shifts over the outcome variable $\mathbf{y}$. Consequently, convolutions with $k(\mathbf{x})$ are the same as correlations with $\check{k}(\mathbf{x}) = k(-\mathbf{x})$. Next we apply the same idea to convolutions on HARDI.

The *check* convolution kernel $\check{p} : \mathbb{R}^3 \rtimes S^2 \to \mathbb{R}^+$ is basically the correlation

kernel related to the convolution kernel $p : \mathbb{R}^3 \rtimes S^2 \to \mathbb{R}^+$:

$$p(\mathbf{y}, \mathbf{n}) = k(\mathbf{y}, \mathbf{n}; \mathbf{0}, \mathbf{e}_z) \text{ whereas}$$
$$\check{p}(\mathbf{y}, \mathbf{n}) = k(\mathbf{0}, \mathbf{e}_z; \mathbf{y}, \mathbf{n})$$

where we recall from eq. 7.3 that

$$k(\mathbf{y}, \mathbf{n}; \mathbf{y}', \mathbf{n}') = p(R_{\mathbf{n}'}^T(\mathbf{y} - \mathbf{y}'), R_{\mathbf{n}'}^T\mathbf{n}).$$

To align the correlation kernel with each position $\mathbf{y}$ and orientation $\mathbf{n}_k$, we define the pre-computed aligned check kernel q as:

$$q_{\mathbf{y}, \mathbf{n}_k}(\mathbf{y}', \mathbf{n}') = \check{p}(R_{\mathbf{n}_k}^T(\mathbf{y}' - \mathbf{y}), R_{\mathbf{n}_k}^T\mathbf{n}') \,,$$

which we use in our discrete convolution scheme, eq. 7.6, where we stress that

$$p(R_{\mathbf{n}'}^T(\mathbf{y} - \mathbf{y}'), R_{\mathbf{n}'}^T\mathbf{n}_k) = \check{p}(R_{\mathbf{n}_k}^{-1}(\mathbf{y}' - \mathbf{y}), R_{\mathbf{n}_k}^{-1}\mathbf{n}').$$

which explains why we must use $q_{\mathbf{y}, \mathbf{n}_k}$ rather than the original kernel p in eq. 7.6. In this step, we compute the set

$$K(\mathbf{y}, \mathbf{n}_k) = \{q_{\mathbf{y}, \mathbf{n}_k}(\mathbf{y}', \mathbf{n}') :\ \mathbf{y}' \in P,\ \mathbf{n}' \in T\} \qquad \boxed{7.11}$$

$\mathbf{n}_k \in T$, where T are the orientations in the tessellation and P is the kernel's lattice.

7.4.2 Truncation

For all positions in the pre-computed kernel set $K(\mathbf{n}_k)$ we truncate the kernel where the probability of diffusion is small enough (here "small enough" is defined by a user chosen threshold ϵ). The new truncated kernel set is then:

$$K_\epsilon(\mathbf{y}, \mathbf{n}_k) = \{(\mathbf{y}', \mathbf{n}', q_{\mathbf{y}, \mathbf{n}_k}(\mathbf{y}', \mathbf{n}')) \mid q_{\mathbf{y}, \mathbf{n}_k}(\mathbf{y}', \mathbf{n}') > \epsilon\} \qquad \boxed{7.12}$$

containing the orientations with the largest probabilities. One could simply iterate through all directions and verify the above condition 7.12. Another option would be to set all directions that do not satisfy the condition to zero, and then simply convolve with all directions. These options would then imply unnecessary iterations. To improve the truncation scheme, the probabilities are sorted, thus ensuring that only the directions corresponding to the larger probabilities are iterated. Since only a subset of all directions $\mathbf{n}' \in T$ is used, some bookkeeping is required in order to keep track of which directions should be iterated, making sure that the kernel and the input HARDI data $\mathcal{U}$ indices match.
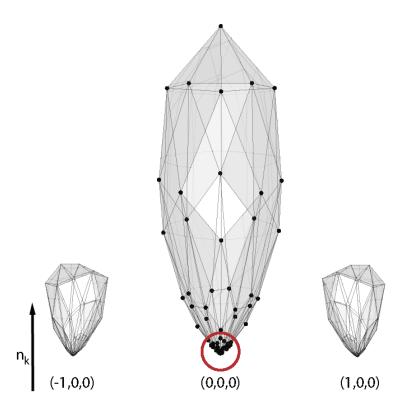
Figure 7.5: Sample discretized kernel $q_{\mathbf{y},\mathbf{n}_k}(\mathbf{y}',\mathbf{n}')$, for $\mathbf{n}_k = (0,0,1)$, $\mathbf{y}' \in P(\mathbf{0}) = \{(-1,0,0),(0,0,0),(1,0,0)\}$, where $|T| = 252$ orientations. For most of the orientations (the ones further from $\mathbf{n}_k$), the probability of diffusion is quite small. If, for example, we truncate the kernel at $\frac{1}{20}$ of the maximum value, by the red circle, 212 orientations are actually ignored in the convolution.

7.4.3 Look-up-table (LUT) convolution

Since the kernels are truncated and sorted, the convolution must now take care of matching the correct values per kernel direction to the corresponding HARDI image directions.

Figure 7.6 illustrates a simple 2D LUT convolution. Here, the kernel is discretized in $|T| = 12$ directions, and we restrict ourselves to 1 point in the spatial lattice. The top row of the figure shows $\mathbf{k}_0$ and $\mathbf{k}_1$ from the set $\mathbf{k}_i(\mathbf{y}) = K_\epsilon(\mathbf{y},\mathbf{n}_i)$, i.e. the kernels for the first two directions. The two tables hold the corresponding index tables needed for the sorted and truncated kernels. Each row in the table $(v,\mathbf{n}')$ holds the probability density value v and the respective direction $\mathbf{n}'$.

Figure 7.6's bottom row illustrates the LUT convolution of the input HARDI image $\mathbf{U}$ with the pre-computed kernel $\mathbf{k}_i(\mathbf{y})$, resulting in the output HARDI

image **O**. In the middle, we can see how eq. 7.6 is resolved. Each position **y** and direction i in the output $\mathbf{O}[\mathbf{y}, i]$ is the result of the product of the corresponding kernel $\mathbf{k}_i$ with the matching directions in the input image **U**:

$$\mathbf{O}[\mathbf{y}, i] = \sum_{\mathbf{y}' \in P(\mathbf{y})} \sum_{a=0}^{|T_\epsilon(\mathbf{d}, i)|} \mathbf{k}_i[\mathbf{d}, a].v \times \mathbf{U}[\mathbf{y}', \mathbf{k}_i[\mathbf{d}, a].n'] \qquad \boxed{7.13}$$

where $a = 0, \ldots |T_\epsilon(\mathbf{d}, i)|$ is the index of the sorted and truncated tessellation corresponding to the kernel at position $\mathbf{d} = \mathbf{y}\text{-}\mathbf{y}'$, $i = 0, \ldots |T|$, and $.v$ corresponds to the value and respective direction $.n'$. In Figure 7.6 we removed all position dependencies for clarity. It only describes for a fixed position, the gain in the angular part of the convolution ($a = 0, \ldots 4$).
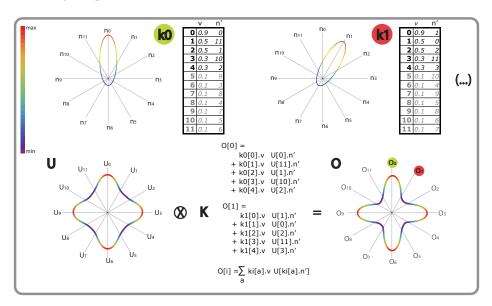
Figure 7.6: The optimized convolution illustrated. The pre-computed kernels, $\mathbf{k}_0$ and $\mathbf{k}_1$, are sorted and the pairs value/index are stored. With a threshold $t = 0.1$, only 5 out of 12 directions are used in the convolution. In the LUT convolution, each direction in the resulting image $\mathbf{O}_i$ is equal to the inner product between the corresponding kernel $\mathbf{k}_i$ and the matching values in the input image **U**.

7.5   Results

In this section we present the experiments conducted in order to analyse the performance of the proposed optimization using the methodology described in

chapter 3: a synthetic DW-MRI dataset, the FiberCup's hardware phantom [112] and a real HARDI data set from a healthy brain. In all presented experiments, Q-Balls of 4^{th} order spherical harmonics (SH) were fitted to the (simulated or acquired) signal, and the resulting SDF was sampled on a tessellated icosahedron (3^{rd} order, 162 points). The choices for SH and tessellation orders were taken since 4^{th} order of SH is the first to convey crossing information, and 3^{rd} order of tessellated icosahedron is a good balance between number of points and discretization error. Before convolution, we preprocess these SDF images by min-max normalization and squaring them. We visualize the squared output of the convolution algorithm.

For validation and illustration of the method, we generated a synthetic dataset with an underlying splaying fibers configuration. The fiber orientations follow the tangent of two ellipsoids centred in the bottom corners of the image. Using the multi-tensor model as described in chapter 3, we constructed a dataset with size 20×28, with eigenvalues for each simulated tensor to be $\lambda_i = [300, 300, 1700] \times 10^{-6}$ mm^2/s, b-value of 1000 s/mm^2 and added Rician noise with realistic SNR of 15.3. Figure 7.7(a) shows this data and the result of the convolution with the tractosemas kernel in figure 7.7(b) ($\sigma = 1, \kappa = 10$ and 3 iterations). We can observe the resulting asymmetric profile in the center region corresponding to the splaying fiber configuration.
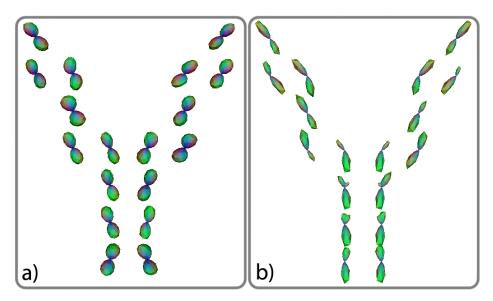
Figure 7.7: Synthetic splaying fibers example: a) Simulated data; b) The computed convolution with Barmpoutis' tractosemas.

The proposed framework was also applied to real DW-MRI datasets. For the

next experiments, the diffusion kernel was used with diffusion parameters $D_{33} = 0.4$, $D_{44} = 0.02$, $t = 1.4$. From FiberCup's data, with b-value 1500 s/mm^2 and $3 \times 3 \times 3$ mm voxel size, we estimated Q-Balls as described above. Figure 7.8 shows a region of interest (ROI) in the full dataset, where two fiber bundles cross. As we can observe, the Q-Ball model expresses a complex fiber structure in the crossing region, however, due to the low b-value, few voxels actually show the 2 expected maxima. Additionally, we can also observe the perturbation caused by noise. After convolving this dataset with the diffusion kernel, we obtain a regularized image where the crossing voxels are clearly enhanced, with evident maxima matching the underlying crossing bundles.

Applying the optimized convolution, again with the diffusion kernel, to a healthy brain from a volunteer, acquired with b-value 4000 s/mm^2, clearly illustrates the benefits of such convolution. Figure 7.9 shows a region where two major white matter structures intersect: the *corpus callosum* from the left, and the *corona radiata* from down-right. We can observe the effect of the low SNR, due to the high b-value, causing clear perturbations in the profiles, specially in the crossing voxels. After convolving this data, we obtain the expected coherency between voxels. Using the neighbourhood information allows the regularization of the data, specially in the more linear areas, and the enhancement of the crossing profiles.

7.5.1   Performance

In figure 7.10 we present a time comparison between the different convolution methods. We show the time realizations for 4 datasets (as previously described):

- **Y synthetic** - software simulated dataset (chapter 3), where the tractosemas kernel was applied

- **FiberCup** - FiberCup dataset [112], with b-value 1500 s/mm^2 and $3 \times 3 \times 3$ mm voxel size

- **Brain slate** - one coronal slate from a healthy brain's volunteer, with b-value 4000 s/mm^2

- **Brainvisa's brain** - brain dataset [111], with b-value 700 s/mm^2

Precomputing the kernels for 3^{rd} order tessellation, takes 47 seconds. This calculation, of course, is only needed once per set of parameters. To evaluate the quality of the proposed method, we quantify the difference between the image result from using the full kernel and the image result from our accelerated convolution. To quantify the differences we calculate the root mean square difference normalized by the range of the values in the accelerated convolution image (NRMSD).

Applying the proposed optimization (described in section 7.4), by truncating the kernels at $\epsilon = 0.003$ (meaning 90% of its total sum), we obtain very similar results as to using the full kernel, however *8 times faster*. Figure 7.11 shows the relation between truncation and the quality of the resulting image (in our experiments, less than 1% difference).

The used threshold was chosen by analysing visually the resulting output that differs minimally from the result using the full kernel. Since no difference can be evaluated qualitatively, no image is shown. Further work will investigate the influence of the threshold on the resulting smoothed image, but our initial results show that a substantial time improvement can be obtained at a small loss in accuracy.

7.6 Conclusions and future work

There are two key limitations inherent to DW-MRI acquisition: images can be very noisy, especially at high b-values; spherical distribution functions are symmetric, which does not always correctly express the underlying fiber structure. Processing of the data on the full domain (spatial and orientational), where contextual information plays an important role, is therefore of the utmost importance. The complexity of the involved operators is, however, a limiting factor for their use. The proposed framework allows the addition of these methods to the DW-MRI processing/visualization pipeline, with much improved computational costs. The framework's kernel independence enables the use of different kernels, for different purposes (e.g., smoothing, enhancing, completion), but still at optimized costs. Fiber tracking applications, for example, can be significantly improved with the use of a processing method such as tractosemas, resolving the problem of splaying fibers.

Further work will analyse the optimal balance between optimization (i.e. which threshold value) and accuracy. Further improvements can be achieved by making use of multiple processors or GPUs (common in today's computers) as the processing algorithm can be easily atomized to voxel level, thus becoming easily parallelized.
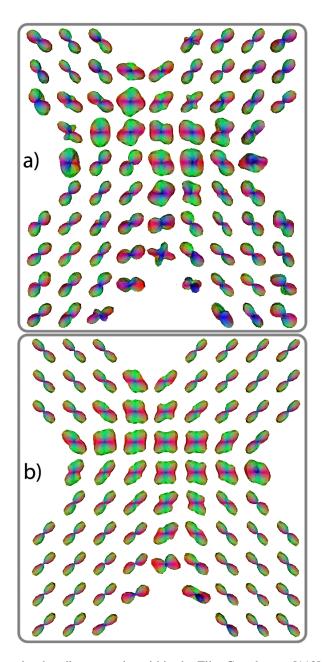
Figure 7.8: Crossing bundles example, within the FiberCup dataset [112], with b-value 1500 s/mm^2 and $3\times3\times3$ mm voxel size. a) Q-Ball's 4^{th} order of SH, sampled on a 3^{rd} order tessellation; b) After convolving with the diffusion kernel. Note the regularization of the result.
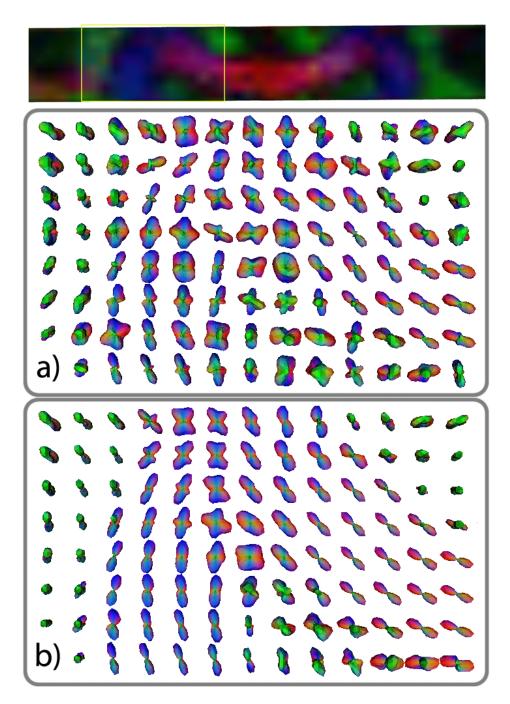
Figure 7.9: Coronal ROI (in yellow top image) of a healthy brain volunteer, acquired with b-value 4000 s/mm^2. Convolving the 4^{th} order SH Q-Ball's (a) with the diffusion kernel results in a regularized field of SDFs where the *corpus callosum* and the *corona radiata* clearly cross in the *centrum semiovale*.

Dataset	Method	Time (min)	Ratio
Y synthetic	full	2,7	
(28x20)	lut	0,23	**11,7**
Fibercup	full	17,7	
(128 x 128)	lut	2,1	**8.13**
Brain Slate	full	4,46	
(64x10)	lut	0,73	**8.12**
Brain	full	974,55	
(60x50x70)	lut	120,7	**8,07**

Figure 7.10: Table a - Time performance comparison between applying the convolution with **full** kernel or with accelerated **LUT** convolution ($\epsilon = 0.003$). All computations were conducted on a AMD Athlon X2 Dual 2.41GHz, with 3GB of RAM.
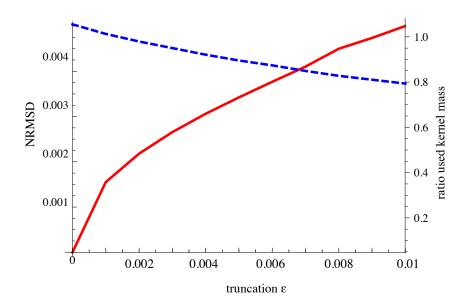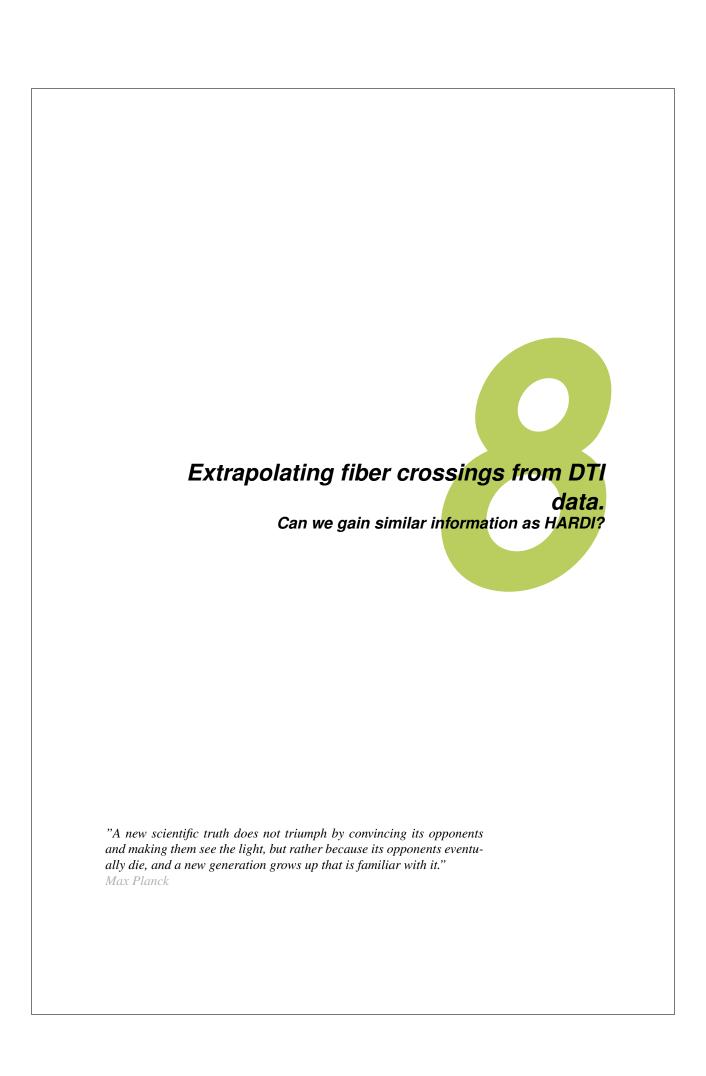
Figure 7.11: Quality comparison between applying the convolution with **full** kernel or with accelerated **LUT** convolution. In red (continuous) the NRMSD between full and accelerated convolutions. In blue (dashed) the corresponding decrease of used kernel's mass.

Extrapolating fiber crossings from DTI data.

Can we gain similar information as HARDI?

"A new scientific truth does not triumph by convincing its opponents and making them see the light, but rather because its opponents eventually die, and a new generation grows up that is familiar with it."
Max Planck

Contents

8.1 Overview

High angular resolution diffusion imaging (HARDI) has proven to better characterize complex intra-voxel structures compared to its predecessor diffusion tensor imaging (DTI). Increasing the b-value on one hand improves the angular resolution and sharpness of the profiles in the HARDI modelling techniques, but on the other hand contributes to the decrease in the SNR. Additionally, for good quality data, the total scanning time must be considerably longer than the one required for the DTI data. Therefore, the benefits from the modest acquisition costs and significantly higher signal-to-noise ratios (SNRs) of DTI make it more attractive for use in clinical setting.

In this chapter we use *contextual* information derived from DTI data to obtain similar crossing information as from the HARDI data, utilizing the accelerated framework described previously in chapter 7. We conduct a synthetic phantom study under different angles of crossing and different SNRs, and compare the extrapolated crossings with the crossings obtained from the HARDI data. We qualitatively corroborate our findings from the phantom study to real human data showing that with extrapolation of the contextual information, the obtained crossings are similar to the ones from the HARDI data, and the robustness to noise is significantly better.

8.2   Introduction

By now, we are familiar with the benefits and drawbacks of both DTI and HARDI. As we previously discussed in chapter 2 and chapter 7, there are two key limitations in HARDI: generally, it produces images with lower SNR than in DTI; and the phase of the MRI signal is commonly discarded, thus resulting in anti-podally symmetric profiles, that can only model single fiber or symmetric crossings of multiple fibers.

In the previous chapter 7, we presented an accelerated convolution framework for noise removal, regularization and enhancement of HARDI datasets. Until now, contextual processing as described above has been applied only on HARDI models, due to the natural coupling of the space of positions and orientations that describe the diffusion process.

In this chapter, we address some of the above mentioned issues. We use data from typical clinically obtained DTI acquisitions to build orientation distribution functions (ODF) that can be used for contextual processing of the data. The data initially comes with high SNR values making the local reconstruction of the ODFs reliable. The context information of well defined single direction fibers is extrapolated to areas where the fiber structure is considerably complex and therefore ill defined by the DTI model. We analyse the differences of the contextually modified ODFs comparing it to the Q-Ball reconstructions [35] without any regularization from the same data as the estimated extrapolated ODFs (E-ODFs). To be fair, we extend this comparison to the Q-Ball's "best scenario" at high b-value ($3000\,\mathrm{s/mm}^2$) and dense gradient sampling (121 number of gradients) regularized with Laplace-Beltrami (LB) smoothing as proposed by Descoteaux et al. [35]. We do quantitative analysis on synthetic crossings of two fibers at different angles and qualitative analysis on *in-vivo* data with the same acquisition as in the synthetic data. We come to a few interesting conclusions, suggesting that E-ODFs contain similar information, concerning the amount of fiber populations and their orientations, as Q-Ball's "best scenario" case. The E-ODFs could bring great improvement to the DTI data, helping to overcome the limitations in crossing regions and enabling also in these regions streamline-based tractography.

8.3   Methods

In this section we present our method for creating extrapolated ODFs (E-ODFs) from diffusion tensors (DT) estimated from our DW-MRI data. We additionally give details on the contextual image processing and perform an evaluation.

8.3.1 Creating spherical diffusion functions from diffusion tensors

As we discussed in chapter 2 of this thesis, in DTI the signal decay is assumed to be mono-exponential [125], and yields the equation:

$$S_g = S_0 \exp(-b\mathbf{g}^T\mathbf{D}\mathbf{g}) \qquad \boxed{8.1}$$

where S_g is the signal in the presence of the diffusion sensitizing gradient, and S_0 is the zero-weighted baseline signal, b is the well known b-value, $\mathbf{g}$ are the diffusion gradient unit vectors, and $\mathbf{D}$ is the 2^{nd} order symmetric, positive definite, diffusion tensor. Once the DT is calculated per voxel, the orientation distribution function (ODF) can be reconstructed and sampled on the sphere

$$ODF(\mathbf{n}) = \mathbf{n}^T\mathbf{D}\mathbf{n} \qquad \boxed{8.2}$$

where $\mathbf{n}$ is the direction vector defined by the tessellation. Figure 8.1 shows a typical linear DT and the corresponding diffusivity profile sampled on a sphere (in our case icosahedron of order 4, 642 points on a sphere). Note that this ODF, since it is derived from the DT, does not hold any crossing information and should not be confused with the apparent diffusion coefficient (ADC) whose crossing information does not necessarily coincide with the underlying fiber population as pointed out by Özarslan et al. [104].

Figure 8.1: A linear diffusion tensor (left) and the corresponding tessellated ODF (right).

From a tensor field we create an ODF field, i.e. a spherical diffusion function (SDF), in other words a HARDI-like dataset $\mathcal{U}$ defined on the coupled space of positions and orientations [39], as previously defined in equation 7.1, section 7.3.1. These DTI derived ODF images $\mathcal{U}$ can now be enhanced. Throughout this chapter we consider DTI-data as the initial condition, which means that we set $\mathcal{U}(\mathbf{y}, \mathbf{n}) = \mathbf{n}^T\mathbf{D}(\mathbf{y})\mathbf{n}$.

8.3.2 Kernels for contextual enhancing of orientation distribution functions

To extrapolate the desired crossing information from the context provided by the estimated ODF, we use Duits et al. [39, 40] diffusion kernels, presented in the previous section 7.3.5.

We now convolve this kernel with the ODF image $\mathcal{U}$, using the HARDI convolution (chapter 7) method, as expressed in equation 7.6. For our purposes, we chose the parameters for the kernel in order to give a high relevance to the diffusion along the principal axis $D_{33} = 0.6$, $D_{44} = 0.01$ and $t = 1.4$. The convolution with such a kernel will result on the extrapolation of crossing profiles where the neighbourhood information so indicates, i.e. the E-ODFs.

In order to achieve the desired results, care should be taken on the sharpness of the input image $\mathcal{U}$. Before applying the convolution, the ODFs are min-max normalized and sharpening is applied by a nonlinear transformation (i.e. power of 2) of the ODFs.

8.3.3 Data

Synthetic Data - To validate and analyse our methodology artificial datasets were generated. DT datasets were created where two fiber bundles forming "tubes" with radii of 2 voxels intersect each other. Here, the tensors, with eigenvalues $\lambda = [17, 3, 3] \times 10^{-3}\,\mathrm{mm}^2/\mathrm{s}$ and eigenvector oriented tangentially to the center line of the tube, are estimated using a mixed tensor model [2] previously discussed in chapter 3. Gaussian noise with different SNRs is added to the real and complex part of the signal reconstructed from equation 8.1. In order to evaluate the angular resolution we vary the angle between the two fiber tubes $\theta \in \{50°,$ $60°, 70°\}$. We made a choice for these angles, given that the accuracy of Q-Ball to detect multiple fiber orientations is around $60°$ [35, 113]. With these angle configurations we create two sets of data, with different acquisition parameters.

1. To evaluate the accuracy of E-ODFs we create datasets with $b = 1000\,\mathrm{s/mm}^2$ and 49 gradient directions. We add Rician noise with SNR=20, given that this is the SNR found in literature for DTI acquisitions [69,78]. From these datasets we estimate E-ODFs and Q-Balls without regularization.

2. To compare with the Qball's "best case scenario", as reported by Descoteaux et al. [35], we create datasets at $b = 3000\,\mathrm{s/mm}^2$ and 121 gradient direction. Since this kind of data is expected to have a lower SNR, we add Rician noise with SNR= 10 as calculated in our *in-vivo* data at the same b-value. We estimate Q-Balls for these datasets and regularize with Laplace-Beltrami (LB) smoothing with $\lambda = 0.006$. This choice for the regularization parameter λ was made, since it was found to be the best at $b = 3000\,\mathrm{s/mm}^2$ [35].

3. In order to evaluate the robustness to noise, we fix the angle to $\theta{=}70°$, and we vary the SNR $\in \{5, 10, 20\}$. We make the same choices for b-values and

number of gradients as previously described, and apply Laplace-Beltrami (LB) smoothing for the Q-Balls at $b = 3000 \, \mathrm{s/mm}^2$

Real Human Data - Diffusion acquisitions were performed using a twice focused spin-echo echo-planar imaging sequence on a Siemens Allegra 3T scanner, with FOV $208 \times 208 \, \mathrm{mm}$, isotropic voxels of $2 \, \mathrm{mm}$. 10 horizontal slices were positioned through the body of the *corpus callosum* and *centrum semiovale*. A uniform gradient direction scheme with 49 and 121 directions were generated with the electrostatic repulsion algorithm [67] and the diffusion-weighted volumes were interleaved with b_0 volumes every 12^{th} scanned gradient direction. Datasets were acquired at b-values of $1000 \, \mathrm{s/mm}^2$ and $3000 \, \mathrm{s/mm}^2$.

8.3.4 Analysis of synthetic data

To analyse the accuracy of the E-ODFs compared to the Q-Balls in the synthetic data sets, we calculate the angular error and standard deviation of the voxels in the crossing region. We do not expect to obtain exactly the same profile, notwithstanding it should contain the same information concerning the amount of fiber populations and their angle. To do so, we use a simple scheme for determining the error between the detected maxima, and then report the angular difference between these maxima and the simulated (true) fiber directions as explained in detail in chapter 3. We detect the maxima as the local maxima of the normalized [0,1] profiles where the function surpasses a certain threshold (here, we use 0.5). To minimize the error related to the sphere tessellation, we use 4^{th} order of tessellation of an icosahedron.

8.3.5 Analysis of human data

For qualitative analysis of the real data, we select an interesting region, the *centrum semiovale* (CS), where crossings are to be expected. This is a challenging region for DW-MRI analysis techniques, since fibers of the *corpus callosum* (CC), *corona radiata* (CR), and *superior longitudinal fasciculus* (SLF) form a three-fold crossing. A region-of-interest (ROI) was defined on a coronal slice (see figure 8.5(a)). We only do qualitative analysis for the real data, as we do not know the ground truth there.

8.4 Results

8.4.1 Phantom data results

The quantitative results of the found angular error and standard deviation of the different profiles in the crossing area from the synthetic data are presented in

Acquisition		$b=1000s/mm^2$ NG=49; SNR=20		$b=3000s/mm^2$ NG=121; SNR=10
Profiles		E-ODF	Qball	Regularized Qball
50°	l=4	x	x	x
	l=6		x	x
	l=8		x	37.6°; 19.9°
60°	l=4		x	x
	l=6	55.5°; 12°	54.3°; 15°	38.5°; 26.9°
	l=8		37.9°; 21°	33°; 27°
70°	l=4		x	14.1°; 32.1°
	l=6	12°; 5°	43°; 25.1°	5.4°; 3°
	l=8		24,4°; 7.5°	19°; 31°

Table 8.1: Results of the angular error and standard deviation from the different profiles in the crossing area of the synthetic data.

table 8.1. In the following paragraphs we relate them to some figures of interesting parameter configurations and discuss the results. In figure 8.2, we present the results of the performance of the proposed E-ODFs compared to the Q-Balls [35] for different angles of crossings, and different simulation parameters: 49 gradient directions, b-value 1000 s/mm^2 and SNR 20, (figure 8.2 middle row) ; 121 gradient directions, b-value 3000 s/mm^2, LB smoothing with $\lambda = 0.006$ (as in [35]) and SNR 10 (figure 8.2 third row).

We observe that for the angle of 50°, E-ODFs and un-regularized Q-Balls fail to find multiple maxima in the crossing areas. Only regularized Q-Ball at high $b = 3000\,\text{s/mm}^2$ and high order $\ell = 8$ detects multiple maxima. For the angle of 60° the performance of E-ODFs (angular error of 55° with standard deviation of 12°) is similar to the un-regularized Q-Ball at $b = 1000\,\text{s/mm}^2$ and truncated at order of spherical harmonics $l = 6$ (angular error of 54.3° with standard deviation of 15°). Regularized Q-Balls at $b = 3000\,\text{s/mm}^2$ outperform in this scenario. At an angle of 70°, the E-ODFs (angular error of 12° with standard deviation of 5°) outperform the best (un-regularized) Q-Ball scenario at order $\ell = 8$ (angular error of 24.4° with standard deviation of 7.5°). Only regularized Q-Ball at $\ell = 6$ outperforms E-ODFs in this scenario (see table 8.1).

The plots in figure 8.3 report the relation between the angular error and change in SNR. We observe that the E-ODFs are more stable, regardless the noise level, whereas the regularized Q-Balls improve significantly at higher SNRs. However, it is important to note that in real data at high b-value $\approx 3000\,\text{s/mm}^2$ the SNR
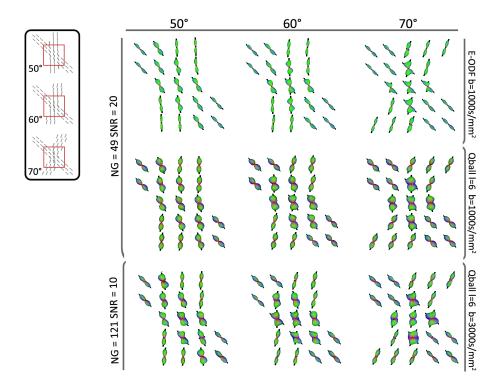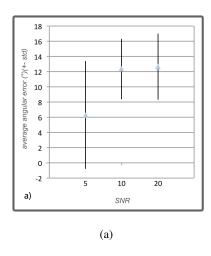
Figure 8.2: E-ODFs and Q-Balls for different angles of crossing at fixed $SNR = 20$ for $b = 1000\,\mathrm{s/mm^2}$, and SNR=10 for $b = 3000\,\mathrm{s/mm^2}$. The Q-Balls at $b = 3000\,\mathrm{s/mm^2}$ are regularized with LB smoothing with $\lambda = 0.006$. The truncation SH order $\ell = 6$.

significantly drops. Moreover, the SNR is highly dependent on the type of scanner and imaging parameters in the acquisition. In our data for $b = 3000\,\mathrm{s/mm^2}$ the SNR in the unweighed image was calculated as 10, by taking a region for computing the mean signal (from the region of interest in the *corpus callosum*) and an air region for computing the standard deviation of the noise. Figure 8.4 illustrates the previous conclusions. At higher order of truncation un-regularized Q-Ball performs much worse, giving many false positives in the linear areas where the SNR is low.

We observe that regardless the SNR, the E-ODFs preserve the coherence of the linear and crossing regions, and preserve the angular error, to almost constant value (see figure 8.3(a)). We also compared the E-ODFs, to Q-Ball's "best case scenario" with LB regularization [35]. Here, for SNR 5, the regularized Q-Ball performs worse (angular error of $14.9°$ and standard deviation $8.4°$) than the E-ODFs (angular error $6°$ with standard deviation $7°$). As noise decreases,

(a)   (b)
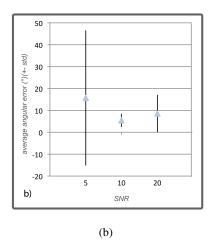
Figure 8.3: Angular error and standard deviation for angle of $70°$ (a) E-ODFs at $b = 1000 \, \mathrm{s/mm^2}$ and 49 gradient directions (b) Regularized Q-Ball with $\lambda = 0.006$, $b = 3000 \, \mathrm{s/mm^2}$ and 121 gradient directions

E-ODFs' performance is similar to the regularized Q-Balls at $b = 3000 \, \mathrm{s/mm^2}$. Regularized Q-Ball outperforms E-ODFs for SNR 20, with an angular error of $5°$ and standard deviation of $3.2°$. However, this SNR is not realistic given nowadays acquisition protocols and machinery at b-values as high as $3000 \, \mathrm{s/mm^2}$.

8.4.2   Real data results

Interestingly, even though crossing information is missing in the original DTI data, as well as in the created ODFs (as can be seen in figure 8.5(a) and figure 8.5(b)), we observe that after processing it, crossing information can be found by extrapolation (see figure 8.6(c)). The obtained information is comparable to the Q-Ball reconstructions of $l = 6$, at $b = 3000 \, \mathrm{s/mm^2}$ and 121 gradient directions and regularized with LB smoothing of $\lambda = 0.006$ (figure 8.6(b)). The crossings in the *centrum semiovale* are extrapolated and appear to correspond to the known anatomy. Additionally the spurious peaks of the data are removed, while producing a more coherent field of E-ODFs that contain peaks in the expected directions where the fibers should cross. The un-regularized Q-Balls at low b-value of $1000 \, \mathrm{s/mm^2}$ and low gradient sampling of 49 gradient directions present less obvious structures of the CC and CR, and have more chaotically oriented crossings, as observed in figure 8.6(a).

All computations were conducted in an AMD Athlon X2 Dual 2.41GHz, with
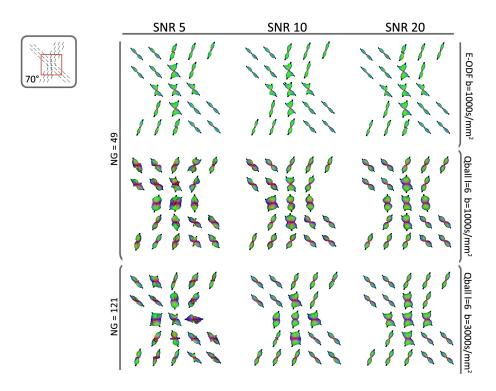
Figure 8.4: E-ODFs and Q-Balls for different SNRs, and different b-value, at fixed angle of 70°.

3GB of RAM, taking 0.5 minutes per artificial tube dataset, and about 13 minutes for the real human brain dataset for estimating the E-ODFs.

8.5   Conclusions

In this chapter we presented a new method for extrapolating crossing information using image processing of the coupled space of positions and orientations in DTI data.

We show that with typical acquisition schemes for DTI, the inferred fiber crossings are similar as 6^{th} order un-regularized Q-Ball estimated from the same data. Furthermore we compare the E-ODFs to the "best scenario" of Q-Ball at typical HARDI acquisition schemes, and we conclude that the information gain from the regularized Q-Ball is similar at low SNR, but the Q-Balls improve when increasing the SNR. However, in practice HARDI acquisitions at high b-values result in rather noisy datasets, and Q-Ball reconstructions of poor quality including LB
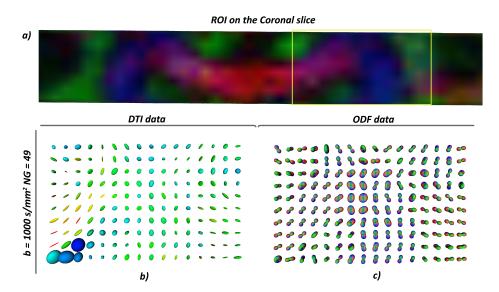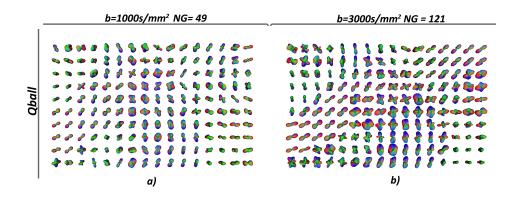
Figure 8.5: The *centrum semiovale*. Left: the original DTI data, color coded by FA. Right: the ODFs from the DTI data, RGB color coded by orientation and min-max normalized.

regularization and therefore this advantage of Q-Ball is not realistic in practice. The robustness to noise of the presented method is significantly better than from the un-regularized Q-Balls reconstructed from the same data, and comparable to the Q-Ball "best scenario". However, at very low SNRs the E-ODFs outperform both regularized and un-regularized Q-Ball.

The main contribution from this work lies on demonstrating that modest acquisitions modelled by DTI, through the use of the contextual information can result in the same information gain, with respect to number of fiber populations and their orientations, as in some of the popular HARDI reconstruction techniques that require more expensive acquisitions. The chosen kernel sets a global reasonable probabilistic model that governs how the context of a fiber fragment is taken into account. Consequently, our framework lacks adaptivity, meaning that the probabilistic model is adapted to the data. Future work will address more adaptive fiber context models to the data, such that context is only included where it is required by the data.

The proposed method has its limitations, it assumes that enough context is available for a correct extrapolation. The possible implications of this limitations for concrete brain structures should be studied. Future work should include more extensive validation to assess the exact differences between HARDI models and E-ODFs concerning acquisition parameters and anatomical areas of the brain.
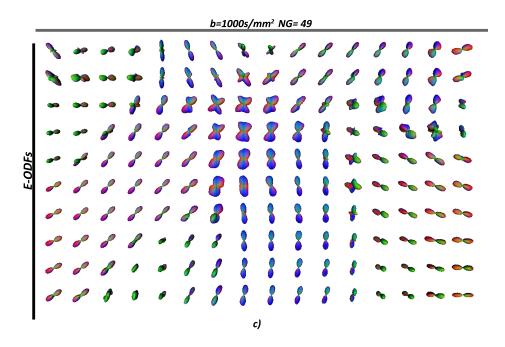
Figure 8.6: Different profiles in the *centrum semiovale* a) un-regularized Q-Ball of order 4 b) Regularized Q-Ball with $\lambda = 0.006$ of order 6 from similar region as (a) c) E-ODFs of the same region as (a)
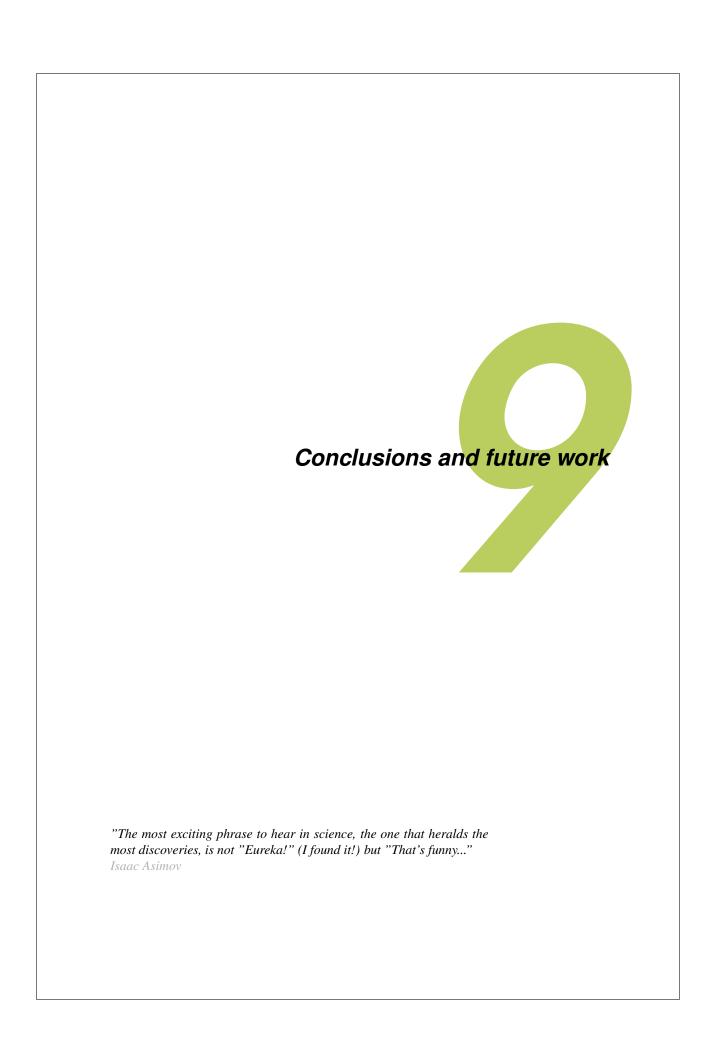
This includes synthetic data experiments with fibers of different configurations (e.g. curved bundles) and multiple crossings.

The final message from this chapter lies in the fact that contextual processing of DTI data allows overcoming one of the main drawbacks of this model. The

crossing information can be recovered, and with acquisitions that typically take 3-6 minutes and modest post-processing (13 minutes for 10 slices of a human brain on a standard PC) more accurate stream-based tractography can be applied for plain DTI data.

144

9

Conclusions and future work

"The most exciting phrase to hear in science, the one that heralds the most discoveries, is not "Eureka!" (I found it!) but "That's funny..."
Isaac Asimov

9.1 Summary of contributions

Central to this thesis is the question "how can we virtually dissect white matter structures?". Our work focussed in research and development of new post-processing and segmentation methods for the automatic identification of white matter structures of interest. These methods are based on using diffusion weighted MRI (DW-MRI) imaging techniques.

We started with a brief survey of neuroscience's history, with particular interest in neuroimaging. In chapter 2, we laid down the necessary background common to our methods. We studied the brain anatomy with focus on its tissue complexity specially the white matter. These are the basis to understand the imaging modality central to this thesis, diffusion weighted imaging.

The main contributions of the subsequent chapters can be summarized as follows:

- As a technical need for the different investigated algorithms, we studied the diverse techniques for creating artificial datasets (chapter 3), which are essential for the validation of the developed algorithms. Validation is still an issue of utmost importance in the DW-MRI field, and a crucial step towards clinical applications. This survey sheds some light over this problem. The multi-tensor model can be used for validation of DTI based techniques, since it exhibits similar behaviour as the *in-vivo* data, whereas the Söderman and Jönsson's model is more appropriate to studies with high b-values, such as in animal scanners. However, these approaches are in general too simplistic, and care should be taken when applying the methods to real data.

- Throughout the first chapters, the concept of tensor homogeneity, in terms of tensor characteristics, is explored in order to achieve new methods for segmenting, filtering and enhancing diffusion images.

  In chapter 4, we studied the diffusion tensor space, with special focus on the different measures of similarity/dissimilarity that can be used to assess the homogeneity of a tensor field. Several measures have been proposed in literature, from different natures, with different meanings, but their behaviour was not well defined and categorized. In this $4^t h$ chapter we developed an overview of the existing distance and similarity measures that can be used to discriminate diffusion tensors. We realized that depending on the application, different measures can be used.

- This lead to the development of a distance learning scheme, chapter 5, where the choice of measure (or combination of measures) is adapted to

the task at hand. Inspired by well-known methods from the fields of machine learning, pattern recognition and image processing, new segmentation techniques were investigated and explored. A semi-automatic segmentation algorithm is proposed, based on homogeneity definition and region-based growing.

- Exploiting the inherent hierarchical nature of brain's tissue, we used the well known scale-space theory to depict an equivalent hierarchical representation, starting with the measured diffusion tensors. The developed framework, discussed in chapter 6, establishes a multi-resolution segmentation tool permitting the simultaneous segmentation of different structures, of different sizes.

- In chapter 7 we explored diffusion operators with two important characteristics: they operate over the coupled space of positions and orientations; and they consider the contextual information. Although with clear benefits for processing of these types of images (denoising, smoothing, enhancement), its added complexity binds these algorithms to a theoretical research, with few practical uses, and limited clinical interest. The developed accelerated framework permits the usage of these powerful processing tools in today's typical computers, with lower time costs, and therefore of great interest for real clinical uses.

- In chapter 8, we studied the use of diffusion operators, presented in chapter 7, to overcome the limitations of DTI, namely its failure in depicting complex crossing structures (crossing, splaying, diverging). By considering the contextual information, these diffusion operators were able to extrapolate crossing information from a typical and fast DTI acquisition. The results are comparable to HARDI reconstruction techniques, whose usage is limited in a clinical setting due to its complexity, increased time costs and lower image quality.

9.2 Future work

The techniques presented in this thesis have shown promising results, although, an important next step should be taken: their true value should be fully evaluated. The main problem with validation if the lack of a "gold standard". The used artificial and phantom methods (studied in chapter 3) provided the means to perform some basic quantitative analysis. However, these datasets hold a too simplified version of the true neuronal tissue, thus they are more of a "silver standard".

Many aspects of the presented work can be further investigated and extended. In this section we depict different possibilities that arise from this research.

Our work was centred mainly around processing diffusion tensors images. The developed methods, such as distance learning, can be naturally extended into high angular resolution diffusion images, where there is also the problem of finding a good measure. A survey such as the one in chapter 4 would be of great interest, since it is not clear which measures are useful to compare the complex spherical functions on HARDI methods, such as Q-Ball or DOT. The segmentation of HARDI images could also benefit from the extension of the distance learning method, from chapter 5.

The high level information obtained from the investigated segmentation methods, i.e. white matter tracts and brain partitioning into homogeneous regions (chapter 6), can be used in new forms of visualisation and exploration of DW-MRI data. The obtained hierarchical structure can be exploited to provide new ways to explore this type of data. Navigating through the different levels in the hierarchy can give a focus-and-context paradigm, thus diminishing clutter problems. The most important data can be shown at full detail and size at the focal point, while the area around the focal point can be simplified to help making sense of the context. Illustrative rendering techniques [101] can be used here to simplify the display of the contextual fiber tracts, around a tumour area, for instance. Furthermore, the hierarchical depiction of the data, can be visualised and explored using information visualisation methodologies. With this hierarchical representation, the complex 3D structures in the brain can be abstracted into 2D depictions such as dendograms, and the different partitioned regions can be interconnected with the fiber tracking information, as proposed by Jianu et al. [65].

The creation of a meaningful hierarchical representation of the data can be used to map the global architecture of the brain, the so called connectome [57]. This structural network of the brain can be mapped at different scales: from the scale of interconnected neurons; through cortical columns linking hundreds or thousands of individual neurons; and to large brain systems responsible for specific interconnectivity patterns, as the *corticospinal tract* for example. This thesis focussed mainly on algorithms for the latter, i.e. segmenting white matter tracts, notwithstanding, the scale-space framework, explored in chapter 6 can provide valuable insight on the estimation of the full brain connectome. By virtue of the scale-space of an DW-MRI image, one can investigate the "deep structure" of the brain, i.e. the structure at all levels of resolution simultaneously. Future lines of research can investigate questions such as how fiber tractography relates to changes in resolution, which tracts remain stable across scale?

Besides knowing more about the connections within the brain, it is interesting to know more about their susceptibility to change across time and across populations. This would give more information about brain development, brain dysfunction in ageing, addiction and neurological disorders. The different processing tools developed in this thesis can improve the different stages in such

analysis: pre-processing, explorative visualisation, tract extraction, and features estimation. Additionally, the hierarchical representation can be used to build an invariant representation of the brain, and then define atlases, or statistical comparative studies.

Lastly, we would like to point out the need of knowledge transfer into and from end-users, in particular clinicians, neurologists and neurosurgeons. Both worlds, technological and clinical, each have different perspectives on the same subject: the brain. On the clinical side there are experts in acquiring and interpreting the images, i.e. begin and end of the analysis cycle, whereas in the technological side there are experts in developing tools to automate complicated and time consuming tasks, and to extract complex relations amid the data. The synergy between these two worlds can greatly increase the understanding of the human brain. The clinical world can ask the important questions, focus on real problems, while the technological world can provide the tools to foster the resolution of these problems. However, not many results are transferred back to the clinicians. In our perspective, there are two main problems. First, there is a lack of standardization technologically and acquisition wise, creating difficulties for effective collaborations between the different institutes. Second, the new technological methods should be made convenient and accessible to the clinical users, reaching out for a common language and presenting information in a meaningful way. The development of the DTITool [117] aims at exactly this point, to share our knowledge, our ideas, our methods to anyone, and in this way enable a pragmatic voice in the research process.

A next, valuable step, is wider application in a clinical setting, learning from the feedback and being confronted with further challenges.

"The limits of the possible can only be defined by going beyond them into the impossible." Arthur C. Clarke

Bibliography

[1] D. C. Alexander and G. J. Barker. Optimal imaging parameters for fiber-orientation estimation in diffusion MRI. *NeuroImage*, 27(2):357–67, 2005.

[2] D. C. Alexander, G. J. Barker, and S. R. Arridge. Detection and modeling of non-gaussian apparent diffusion coefficient profiles in human brain data. *Magnetic Resonance in Medicine*, 48(2):331–40, 2002.

[3] D. C. Alexander, J. Gee, and Bajcsy R. Similarity measures for matching diffusion tensor images. In *Proc. British Machine Vision Conference BMVC99*, 1999.

[4] D. C. Alexander, Penny L. Hubbard, Matt G. Hall, A. Moore, Elizabeth, Maurice Ptito, Geoff J. M. Parker, and Tim B. Dyrby. Orientationally invariant indices of axon diameter and density from diffusion MRI. *Neuroimage*, 52(4):1374–89, Oct 2010.

[5] Vincent Arsigny, Pierre Fillard, Xavier Pennec, and Nicholas Ayache. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic Resonance in Medicine*, 56(2):411–421, August 2006.

[6] Yaniv Assaf, Tamar Blumenfeld-Katzir, Yossi Yovel, and Peter J. Basser. Axcaliber: a method for measuring axon diameter distribution from diffusion MRI. *Magnetic Resonance in Medicine*, 59(6):1347–54, Jun 2008.

[7] Suyash P. Awate and James C. Gee. A fuzzy, nonparametric segmentation framework for DTI and MRI analysis. *Inf. Process. Med. Imaging*, 20:296–307, 2007.

[8] A. Barmpoutis, B. C. Vemuri, D. Howland, and J. R. Forder. Extracting tractosemas from a displacement probability field for tractography in DW-MRI. *In Lecture Notes in Computer Science 5241, MICCAI*, pages 9–16, 6-10 September 2008.

[9] P J Basser. Inferring microstructural features and the physiological state of tissues from diffusion-weighted images. *NMR Biomed*, 8(7-8):333–44, 1995.

[10] Peter J. Basser, James Mattiello, and Denis Lebihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66:259–267, 1994.

[11] Peter J. Basser, Sinisa Pajevic, Carlo Pierpaoli, Jeffrey Duda, and Akram Aldroubi. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44:625–632, 2000.

[12] P.J. Basser and C. Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance*, 111(3):209–219, 1996.

[13] P.G. Batchelor, M. Moakher, D. Atkinson, F. Clamante, and A. Connelly. A rigorous framework for diffusion tensor calculus. *Magnetic Resonance in Medicine*, 53:221–225, 2005.

[14] T E Behrens, H Johansen-Berg, S Jbabdi, M F Rushworth, and M W Woolrich. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage*, 34(1):144–55, 2007. 1053-8119 (Print)Journal Article.

[15] T. E. Behrens, M. W. Woolrich, M. Jenkinson, H. Johansen-Berg, R. G. Nunes, S. Clare, P. M. Matthews, J. M. Brady, and S. M. Smith. Characterization and propagation of uncertainty in diffusion-weighted mr imaging. *Magnetic Resonance in Medicine*, 50(5):1077–88, 2003.

[16] Anders Brun, Hans Knutsson, Hae-Jeong Park, Martha E Shenton, and Carl-Fredrik Westin. Clustering fiber traces using normalized cuts. *Lecture Notes in Computer Science*, 3216/2004(3216):368–375, Sep 2004.

[17] E.J.L. Brunenberg, V. Prčkovska, B. Platel, G.J. Strijkers, and B.M. ter Haar Romeny. Untangling a fiber bundle knot: Preliminary results on stn connectivity using dti and hardi on rat brains. In *Proceedings 17th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Honolulu*, page 740, 2009.

[18] Matthan W A Caan, Anne Willem de Vries, H Ganesh Khedoe, Erik M Akkerman, Lucas J van Vliet, C A Grimbergen, and Frans M Vos. Generating fiber crossing phantoms out of experimental dwis. *Med Image Comput Comput Assist Interv*, 10(Pt 1):169–76, 2007.

[19] Marco Catani. Diffusion tensor magnetic resonance imaging tractography in cognitive disorders. *Curr Opin Neurol*, 19(6):599–606, Dec 2006.

[20] Marco Catani, Robert J Howard, Sinisa Pajevic, and Derek K Jones. Virtual in vivo interactive dissection of white matter fasciculi in the human brain. *Neuroimage*, 17(1):77–94, Sep 2002.

[21] Marco Catani and Michel Thiebaut de Schotten. A diffusion tensor imaging tractography atlas for virtual in vivo dissections. *Cortex*, 44(8):1105–32, Sep 2008.

[22] Herng-Hua Chang, Audrey H Zhuang, Daniel J Valentino, and Woei-Chyn Chu. Performance measure characterization for evaluating neuroimage segmentation algorithms. *Neuroimage*, 47(1):122–35, Aug 2009.

[23] Terence C. Chua, Wei Wen, Melissa J. Slavin, and Perminder S. Sachdev. Diffusion tensor imaging in mild cognitive impairment and alzheimer's disease: a review. *Curr Opin Neurol*, 21(1):83–92, Feb 2008.

[24] Moo K. Chung, Moo K. Chung, Mariana Lazar, Mariana Lazar, Andrew L. Alexender, Andrew L. Alexander, Yuefeng Lu, Yuefeng Lu, Richard Davidson, and Richard Davidson. Probabilistic connectivity measure in diffusion tensor imaging via anisotropic kernel smoothing. Technical report, Department of Statistics, Universisty of Wisconsin-Madison, 2003.

[25] V. Comaniciu and P. Meer. Kernel-based object tracking. 2003.

[26] Luis Concha, Donald W Gross, B Matt Wheatley, and Christian Beaulieu. Diffusion tensor imaging of time-dependent axonal and myelin degradation after corpus callosotomy in epilepsy patients. *Neuroimage*, 32(3):1090–9, Sep 2006.

[27] T. E. Conturo, N. F. Lori, T. S. Cull, E. Akbudak, A. Z. Snyder, J. S. Shimony, R. C. McKinstry, H. Burton, and M. E. Raichle. Tracking neuronal fiber pathways in the living human brain. *Proc Natl Acad Sci U S A*, 96(18):10422–7, 1999.

[28] Isabelle Corouge, Sylvain Gouttard, and Guido Gerig. Towards a shape model of white matter fiber bundles using diffusion tensor MRI. In *International Symposium on Biomedical Imaging*, pages 344–347, 2004.

[29] N. Cristianini, A. Elisseeff, J. Shawe-Taylor, and J. Kandola. On kernel target alignment. 2001.

[30] Erik Dam. *Committing Medical Image Analysis*. PhD thesis, IT University of Copenhagen, 2003.

[31] Erik Dam and Mads Nielsen. Non-linear diffusion for interactive multi-scale watershed segmentation. In *MICCAI '00: Proceedings of the Third International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 216–225, London, UK, 2000. Springer-Verlag.

[32] Sarah de Rijcke. *Regarding the Brain. Practices of Objectivity in Cerebral Imaging. 17th century-present*. PhD thesis, University of Groningen, 2010.

[33] René Descartes. *Treatise of man*. 1664.

[34] Maxime Descoteaux, Elaine Angelino, Shaun Fitzgibbons, and Rachid Deriche. Apparent diffusion coefficients from high angular resolution diffusion imaging: Estimation and applications. *Magnetic Resonance in Medicine*, 56(2):395–410, 2006.

[35] Maxime Descoteaux, Elaine Angelino, Shaun Fitzgibbons, and Rachid Deriche. Regularized, fast and robust analytical q-ball imaging. *Magnetic Resonance in Medicine*, 58:497–510, 2007.

[36] Maxime Descoteaux and Rachid Deriche. Segmentation of q-ball images using statistical surface evolution. In Nicholas Ayache, Sébastien Ourselin, and Anthony J. Maeder, editors, *MICCAI (2)*, volume 4792 of *Lecture Notes in Computer Science*, pages 769–776. Springer, 2007.

[37] Maxime Descoteaux, Rachid Deriche, Alfred Anw, Thème Bio, Maxime Descoteaux, Rachid Deriche, Alfred Anw, and Projet Odyssée. Deterministic and probabilistic qball tractography: from diffusion to sharp fiber distributions. In *Tech. Rep. 6273, INRIA Sophia Antipolis*, 2007.

[38] Richard O. Duda, Peter E. Hart, and David G. Stork. *Pattern Classification (2nd Edition)*. Wiley-Interscience, November 2000.

[39] R. Duits and E. Franken. Left-invariant diffusions on $\mathbb{R}^3 \rtimes S^2$ and their application to crossing-preserving smoothing on HARDI-images. *CASA report, TU/e, nr.18*, 2009.

[40] R. Duits and E. Franken. Left-invariant diffusions on the space of positions and orientations and their application to crossing-preserving smoothing of HARDI images. *International Journal of Computer Vision 2010*, 2010. Available on http://tinyurl.com/2anfqdy.

[41] The editors of Scientific American and Antonio R. Damasio (introduction), editors. *The Scientific American book of the brain*. Scientific American, 1999.

[42] A. Einstein. Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen. *Annalen der Physik*, 322(8):549–560, 1905.

[43] A. Fick. Ueber Diffusion. *Annalen der Physik*, 170:59–86, 1855.

[44] Aaron S Field and Andrew L Alexander. Diffusion tensor imaging in cerebral tumor diagnosis and therapy. *Top Magn Reson Imaging*, 15(5):315–24, Oct 2004.

[45] Pierre Fillard, Vincent Arsigny, Xavier Pennec, and Nicholas Ayache. Clinical DT-MRI estimation, smoothing and fiber tracking with log-Euclidean metrics. *IEEE Transactions on Medical Imaging*, 26(11):1472–1482, November 2007. PMID: 18041263.

[46] L. Florack and E. Balmashnova. Decomposition of high angular resolution diffusion images into a sum of self-similar polynomials on the sphere. In Y. Bayakovsky and E. Moiseev, editors, *Proceedings of the Eighteenth International Conference on Computer Graphics and Vision, GraphiCon'2008, Moscow, Russia, June 23–27, 2008*, pages 26–31. Moscow State University. (invited paper).

[47] L. M. J. Florack. Codomain scale space and regularization for high angular resolution diffusion imaging. In S. Aja Fernández and R. de Luis Garcia, editors, *CVPR Workshop on Tensors in Image Processing and Computer Vision, Anchorage, Alaska, USA, June 24–26, 2008*. IEEE, 2008.

[48] L. M. J. Florack and L. J. Astola. A multi-resolution framework for diffusion tensor images. In S. Aja Fernández and R. de Luis Garcia, editors, *CVPR Workshop on Tensors in Image Processing and Computer Vision, Anchorage, Alaska, USA, June 24–26, 2008*. IEEE, 2008. Digital proceedings.

[49] Luc Florack and Arjan Kuijper. The topological structure of scale-space images. *J. Math. Imaging Vis.*, 12(1):65–79, 2000.

[50] L. R. Frank. Characterization of anisotropy in high angular resolution diffusion-weighted mri. *Magnetic Resonance in Medicine*, 47(6):1083–99, 2002.

[51] R.E. Fredericksen, J.M. Coggins, T.J. Cullip, and S.M. Pizer. Interactive object definition in medical images using multiscale, geometric image descriptions. In *Visualization in Biomedical Computing, 1990., Proceedings of the First Conference on*, pages 108 –114, May 1990.

[52] C. Golgi, M. Bentivoglio, and L. Swanson. On the fine structure of the pes hippocampi major (with plates XIII-XXIII). 1886. *Brain Research Bulletin*, 54(5):461–83, Mar 2001.

[53] Arnulf B. A. Graf and Silvio Borer. Normalization in support vector machines. In *Proceedings of the 23rd DAGM-Symposium on Pattern Recognition*, pages 277–282, London, UK, 2001. Springer-Verlag.

[54] Henry Gray and Charles Mayo Goss. *Anatomy of the human body*. Lea & Febiger, Philadelphia, 1973.

[55] Bernard Haasdonk and Claus Bahlmann. Learning with distance substitution kernels. In Carl Edward Rasmussen, Heinrich H. Bülthoff, Bernhard Schölkopf, and Martin A. Giese, editors, *Pattern Recognition, 26th DAGM Symposium, August 30 - September 1, 2004, Tübingen, Germany, Proceedings*, volume 3175 of *Lecture Notes in Computer Science*, pages 220–227. Springer, 2004.

[56] P. Hagmann, J.-P. Thiran, P. Vandergheynst, S. Clarke, and R. Meuli. Statistical fiber tracking on DT-MRI data as a potential tool for morphological brain studies. In *ISMRM Workshop on Diffusion MRI: Biophysical Issues (2002)*, page 240, 2002.

[57] Patric Hagmann, Leila Cammoun, Xavier Gigandet, Stephan Gerhard, P. Ellen Grant, Van Jay Wedeen, Reto Meuli, Jean-Philippe Thiran, Christopher J. Honey, and Olaf Sporns. MR connectomics: Principles and challenges. *Journal of Neuroscience Methods*, 2010.

[58] Anneriet M Heemskerk, Tuhin K Sinha, Kevin J Wilson, Zhaohua Ding, and Bruce M Damon. Quantitative assessment of dti-based muscle fiber tracking and optimal tracking parameters. *Magnetic Resonance in Medicine*, 61(2):467–72, Feb 2009.

[59] C.P. Hess, P. Mukherjee, E.T. Han, D. Xu, and D.B. Vigneron. Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic basis. *Magnetic Resonance in Medicine*, 2006.

[60] Tim Hosey, Guy Williams, and Richard Ansorge. Inference of multiple fiber orientations in high angular resolution diffusion imaging. *Magnetic Resonance in Medicine*, (54):1480–1489, 2005.

[61] Christian Igel, Tobias Glasmachers, Britta Mersch, Nico Pfeifer, and Peter Meinicke. Gradient-based optimization of kernel-target alignment for sequence kernels applied to bacterial gene start detection. *IEEE/ACM Trans. Comput. Biol. Bioinformatics*, 4(2):216–226, 2007.

[62] S. Jbabdi, P. Bellec, R. Toro, J. Daunizeau, M. Pélégrini-Issac, and H. Benali. Accurate anisotropic fast marching for diffusion-based geodesic tractography. *Journal of Biomedical Imaging*, 2008:1–12, 2008.

[63] B. Jeurissen, A. Leemans, D. K. Jones, J. D. Tournier, and J. Sijbers. Estimating the number of fiber orientations in diffusion MRI voxels: a constrained spherical deconvolution study. In *ISMRM, 18th Scientific Meeting and Exhibition*, Stockholm, Sweden, May 2010.

[64] Bing Jian and Baba C. Vemuri. Multi-fiber reconstruction from diffusion MRI using mixture of wisharts and sparse deconvolution. In *IPMI*, pages 384–395, 2007.

[65] Radu Jianu, Çagatay Demiralp, and David H. Laidlaw. Exploring 3D DTI fiber tracts with linked 2D representations. *IEEE Trans. Vis. Comput. Graph.*, 15(6):1449–1456, 2009.

[66] Lisa Jonasson, Xavier Bresson, Patric Hagmann, Olivier Cuisenaire, Reto Meuli, and Jean-Philippe Thiran. White matter fiber tract segmentation in DT-MRI using geometric flows. *Medical Image Analysis*, 9(3):223–236, June 2005.

[67] D. K. Jones, M. A. Horsfield, and A. Simmons. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magnetic Resonance in Medicine*, 42(3):515–25, 1999.

[68] D. K. Jones, A. Simmons, S. C. Williams, and M. A. Horsfield. Non-invasive assessment of axonal fiber connectivity in the human brain via diffusion tensor MRI. *Magnetic Resonance in Medicine*, 42(1):37–41, 1999.

[69] Derek K Jones. The effect of gradient sampling schemes on measures derived from diffusion tensor mri: a monte carlo study. *Magnetic Resonance in Medicine*, 51(4):807–15, Apr 2004.

[70] DK Jones, MA Horsfield, and A Simmons. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magnetic Resonance in Medicine*, 42:515–525, 1999.

[71] Richard A A Kanaan, Jin-Suh Kim, Walter E Kaufmann, Godfrey D Pearlson, Gareth J Barker, and Philip K McGuire. Diffusion tensor imaging in schizophrenia. *Biol Psychiatry*, 58(12):921–9, Dec 2005.

[72] Yunho Kim, Paul Thompson, Arthur Toga, Luminita Vese, and Liang Zhan. HARDI denoising: Variational regularization of the spherical apparent diffusion coefficient sADC. *Information Processing in Medical Imaging*, pages 515–527, 2009.

[73] G. Kindlmann. *Visualization and Analysis of Diffusion Tensor Fields*. PhD thesis, School of Computing, University of Utah, September 2004.

[74] Gordon Kindlmann, Daniel B. Ennis, Ross T. Whitaker, and Carl-Fredrik Westin. Diffusion tensor analysis with invariant gradients and rotation tangents. *IEEE Trans. Med. Imaging*, 26(11):1483–1499, November 2007.

[75] Gordon L. Kindlmann. Superquadric tensor glyphs. In Oliver Deussen, Charles D. Hansen, Daniel A. Keim, and Dietmar Saupe, editors, *VisSym*, pages 147–154. Eurographics Association, 2004.

[76] Gordon L. Kindlmann, Raúl San José Estépar, Marc Niethammer, Steven Haker, and Carl-Fredrik Westin. Geodesic-loxodromes for diffusion tensor interpolation and difference measurement. In Nicholas Ayache, Sébastien Ourselin, and Anthony J. Maeder, editors, *MICCAI (1)*, volume 4791 of *Lecture Notes in Computer Science*, pages 1–9. Springer, 2007.

[77] T Klingberg, M Hedehus, E Temple, T Salz, J D Gabrieli, M E Moseley, and R A Poldrack. Microstructure of temporo-parietal white matter as a basis for reading ability: evidence from diffusion tensor magnetic resonance imaging. *Neuron*, 25(2):493–500, Feb 2000.

[78] M Lagana, M Rovaris, A Ceccarelli, C Venturelli, S Marini, and G Baselli. DTI parameter optimisation for acquisition at 1.5T: SNR analysis and clinical application. *Comput Intell Neurosci*, NIL(NIL):254032, 2010.

[79] David H. Laidlaw, Eric T. Ahrens, David Kremers, Matthew J. Avalos, Russell E. Jacobs, and Carol Readhead. Visualizing diffusion tensor images of the mouse spinal cord. In *IEEE Visualization*, pages 127–134, 1998.

[80] D. Le Bihan and E. Breton. Imagerie de diffusion in-vivo par résonance magnétique nucléaire. *C. R. Acad. Sci.*, pages 1109–1112, 1985.

[81] D Le Bihan, E Breton, D Lallemand, P Grenier, E Cabanis, and M Laval-Jeantet. MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology*, 161(2):401–7, Nov 1986.

[82] Christophe Lenglet. *Geometric and Variational methods for Diffusion Tensor MRI processing*. PhD thesis, I.N.R.I.A., 2006.

[83] Christophe Lenglet, Rachid Deriche, and Olivier D. Faugeras. Inferring white matter geometry from diffusion tensor mri: Application to connectivity mapping. In Tomás Pajdla and Jiri Matas, editors, *ECCV (4)*, volume 3024 of *Lecture Notes in Computer Science*, pages 127–140. Springer, 2004.

[84] N F Lori, E Akbudak, J S Shimony, T S Cull, A Z Snyder, R K Guillory, and T E Conturo. Diffusion tensor fiber tracking of human brain connectivity: aquisition methods, reliability analysis and biological results. *NMR Biomed*, 15(7-8):494–515, 2002.

[85] Mahnaz Maddah, Lilla Zöllei, W Eric L Grimson, Carl-Fredrik Westin, and William M Wells. A mathematical framework for incorporating anatomical knowledge in DT-MRI analysis. *Proc IEEE Int Symp Biomed Imaging*, 4543943:105–108, 2008.

[86] J-F Mangin, C Poupon, Y Cointepas, D Rivière, D Papadopoulos-Orfanos, C A Clark, J Régis, and D Le Bihan. A framework based on spin glass models for the inference of anatomical connectivity from diffusion-weighted mr data - a technical review. *NMR Biomed*, 15(7-8):481–92, 2002.

[87] J. C. Maxwell. *On hills and dales*, volume XL. The London, Edinburgh and Dublin Philosophical Magazine and Journal of Science, 1870.

[88] M.Descoteaux. *High Angular Resolution Diffusion MRI: From Local Estimation to Segmentation and Tractography*. PhD thesis, Universite de Nice - Sophia Antipolis, 2008.

[89] Elias R. Melhem, Susumu Mori, Govind Mukundan, Michael A. Kraut, Martin G. Pomper, and Peter C.M. van Zijl. Diffusion tensor MR imaging of the brain and white matter tractography. *American Journal of Roentgenology*, 178(1):3–16, 2002.

[90] K D Merboldt, W Hänicke, and J Frahm. Diffusion imaging using stimulated echoes. *Magnetic Resonance in Medicine*, 19(2):233–9, Jun 1991.

[91] F. Meyer and S. Beucher. Morphological segmentation. *Journal of Visual Communications and Image Representation*, 1:21–46, 1990.

[92] Jyrki Lötjönen Mika Pollari, Tuomas Neuvonen. Affine registration of diffusion tensor MR images. In *MICCAI 2006*, pages 629–636. Springer-Verlag Berlin Heidelberg, 2006.

[93] A K Miller, R L Alston, and J A Corsellis. Variation with age in the volumes of grey and white matter in the cerebral hemispheres of man: measurements with an image analyser. *Neuropathol Appl Neurobiol*, 6(2):119–132, 1980 Mar-Apr.

[94] Bart Moberts, Anna Vilanova, and Jarke J. van Wijk. Evaluation of fiber clustering methods for diffusion tensor imaging. In *Proceedings of IEEE Visualization 2005*, pages 65–72, October 2005.

[95] S Mori, B J Crain, V P Chacko, and P C van Zijl. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Ann Neurol*, 45(2):265–9, Feb 1999.

[96] Susumu Mori and Peter C M van Zijl. Fiber tracking: principles and strategies - a technical review. *NMR Biomed*, 15(7-8):468–80, 2002.

[97] M E Moseley, Y Cohen, J Mintorovitch, L Chileuitt, H Shimizu, J Kucharczyk, M F Wendland, and P R Weinstein. Early detection of regional cerebral ischemia in cats: comparison of diffusion- and t2-weighted mri and spectroscopy. *Magnetic Resonance in Medicine*, 14(2):330–46, May 1990.

[98] M. Niethammer, C. Zach, J. Melonakos, and A. Tannenbaum. Near-tubular fiber bundle segmentation for diffusion weighted imaging: segmentation through frame reorientation. 03 2009.

[99] L. O'Donnell, S. Haker, and C.-F. Westin. New approaches to estimation of white matter connectivity in diffusion tensor MRI: Elliptic PDEs and geodesics in a tensor-warped space. In *Fifth International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'02)*, pages 459–466, Tokyo, Japan, 2002.

[100] Lauren J O'Donnell, Carl-Fredrik Westin, and Alexandra J Golby. Tract-based morphometry. *Med Image Comput Comput Assist Interv*, 10(Pt 2):161–8, 2007.

[101] Ron Otten, A. Vilanova, and H.M.M. van de Wettering. Illustrative white matter fiber bundles. *Computer Graphics Forum*, 29(3):1013–1022, 2010.

[102] E. Özarslan and T. H. Mareci. Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution diffusion imaging. *Magnetic Resonance in Medicine*, 50(5):955–65, 2003.

[103] E. Özarslan, B. C. Vemuri, and T. H. Mareci. Generalized scalar measures for diffusion MRI using trace, variance, and entropy. *Magnetic Resonance in Medicine*, 53(4):866–76, 2005.

[104] Evren Özarslan, Timothy M. Shepherd, Baba C. Vemuri, Stephen J. Blackband, and Thomas H. Mareci. Resolution of complex tissue microarchitecture using the diffusion orientation transform (DOT). *NeuroImage*, 36(3):1086–1103, July 2006.

[105] Geoffrey J. M. Parker and D. C. Alexander. Probabilistic anatomical connectivity derived from the microscopic persistent angular structure of cerebral tissue. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.*, 360(1457):893–902, May 2005.

[106] Geoffrey J M Parker, Claudia A M Wheeler-Kingshott, and Gareth J Barker. Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE Trans Med Imaging*, 21(5):505–12, May 2002.

[107] E. Pekalska, P. Paclik, and R.P.W. Duin. A Generalized Kernel Approach to Dissimilarity Based Classification. *Journal of Machine Learning Research*, 2:175–211, 2001.

[108] Xavier Pennec, Pierre Fillard, and Nicholas Ayache. A riemannian framework for tensor computing. *Int. J. Comput. Vision*, 66(1):41–66, 2006.

[109] C. Pierpaoli and P.J. Basser. Toward a quantitative assessment of diffusion anisotropy. *Magnetic Resonance in Medicine*, 36:893–906, 1996.

[110] C. Pierpaoli, P. Jezzard, P. J. Basser, A. Barnett, and G. Di Chiro. Diffusion tensor MR imaging of the human brain. *Radiology*, 201:637–648, Dec 1996.

[111] C. Poupon, F. Poupon, L. Allirol, and J.-F. Mangin. A database dedicated to anatomo-functional study of human brain connectivity. In *12th HBM NeuroImage*, number 646, Florence, Italie, 2006.

[112] C. Poupon, B. Rieul, I. Kezele, M. Perrin, F. Poupon, and J.-F. Mangin. New diffusion phantoms dedicated to the study and validation of HARDI models. *Magnetic Resonance in Medicine*, 60(6):1276–83, 2008.

[113] Vesna Prčkovska, Alard F. Roebroeck, W.L.P.M. Pullens, Anna Vilanova, and Bart M. ter Haar Romeny. Optimal acquisition schemes in high angular resolution diffusion weighted imaging. In *MICCAI*, volume 5242 of *Lecture Notes in Computer Science*, pages 9–17. Springer, 2008.

[114] WLPM Pullens, A Roebroeck, and R Goebel. Kissing or crossing: validation of fiber tracking using ground truth hardware phantoms. In *Proc ISMRM*, page 1479, 2007.

[115] Arish A Qazi, Alireza Radmanesh, Lauren O'Donnell, Gordon Kindlmann, Sharon Peled, Stephen Whalen, Carl-Fredrik Westin, and Alexandra J Golby. Resolving crossings in the corticospinal tract by two-tensor streamline tractography: Method and clinical assessment using fMRI.

[116] Leticia Rittner and Roberto A. Lotufo. Diffusion tensor imaging segmentation by watershed transform on tensorial morphological gradient. In *Proceedings*, Los Alamitos, Oct. 12–15, 2008 2008. IEEE Computer Society.

[117] P.R. Rodrigues, V. Prčkovska, T.H.J.M. Peeters, and B.M. ter Haar Romeny A. Vilanova Bartroli. DTITool: An insightful look into diffusion tensor images. volume EuroVis, Berlin, Germany, 2009.

[118] Mikaël Rousson, Christophe Lenglet, and Rachid Deriche. Level set and region based surface propagation for diffusion tensor MRI segmentation. In Milan Sonka, Ioannis A. Kakadiaris, and Jan Kybic, editors, *ECCV Workshops CVAMIA and MMBIA*, volume 3117 of *Lecture Notes in Computer Science*, pages 123–134. Springer, 2004.

[119] D H Salat, D S Tuch, D N Greve, A J W van der Kouwe, N D Hevelone, A K Zaleta, B R Rosen, B Fischl, S Corkin, H Diana Rosas, and A M Dale. Age-related alterations in white matter microstructure measured by diffusion tensor imaging. *Neurobiol Aging*, 26(8):1215–27, 2005.

[120] S. Schölkopf, B. Mika, S. Burges, C. Knirsch, P. Müller, K. Rätsch, and G. Smola. Input space vs. feature space in kernel-based methods. *IEEE Transactions on Neural Networks*, 10(5):1000–1017, 1999.

[121] Thomas Schultz, Bernhard Burgeth, and Joachim Weickert. Flexible segmentation and smoothing of DT-MRI fields through a customizable structure tensor. 4291:455–464, 2006.

[122] J.S. Shimony, A.Z. Snyder, N. Lori, and T.E. Conturo. Automated fuzzy clustering of neuronal pathways in diffusion tensor tracking. In *Intl. Soc. Mag. Reson. Med.*, May 2002.

[123] O. Söderman and B. Jönsson. Restricted diffusion in cylindirical geometry. *J. Magn. Reson. B*, 117(1):94–97, 1995.

[124] Jon Sporring. *Gaussian Scale-Space Theory*. Kluwer Academic Publishers, Norwell, MA, USA, 1997.

[125] E. O. Stejskal and J. E. Tanner. Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. *The Journal of Chemical Physics*, 42(1):288–292, 1965.

[126] D.M.J. Tax. One-class classification. Master's thesis, Delft University of Technology, http://ict.ewi.tudelft.nl/ davidt/thesis.pdf, June 2001.

[127] D G Taylor and M C Bushell. The spatial mapping of translational diffusion coefficients by the NMR imaging technique. *Physics in Medicine and Biology*, 30(4), 1985.

[128] J. D. Tournier, F. Calamante, D. G. Gadian, and A. Connelly. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *NeuroImage*, 23(3):1176–85, 2004.

[129] Andy Tsai, Carl-Fredrik Westin, Alfred O. Hero, and Alan S. Willsky. Fiber tract clustering on manifolds with dual rooted-graphs. In *CVPR*. IEEE Computer Society, 2007.

[130] D. S. Tuch. *Diffusion MRI of complex tissue structure*. PhD thesis, Harvard, 2002.

[131] D.S. Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52:1358–1372, 2004.

[132] P van Gelderen, M H de Vleeschouwer, D DesPres, J Pekar, P C van Zijl, and C T Moonen. Water diffusion and acute stroke. *Magnetic Resonance in Medicine*, 31(2):154–63, Feb 1994.

[133] Carola van Pul, Jan Buijs, Maurice J A Janssen, George F Roos, Marinus T Vlaardingerbroek, and Pieter F F Wijn. Selecting the best index for following the temporal evolution of apparent diffusion coefficient and diffusion anisotropy after hypoxic-ischemic white matter injury in neonates. *American Journal of Neuroradiology*, 26(3):469–81, Mar 2005.

[134] A. Vilanova, S. Zhang, G. Kindlmann, and D. Laidlaw. An introduction to visualization of diffusion tensor imaging and its applications. In J. Weickert and H. Hagen, editors, *Visualization and Processing of Tensor Fields*, Mathematics and Visualization, chapter 7, pages 121–153. Springer, 2005.

[135] von Dem and R. M. Henkelman. Orientational diffusion reflects fiber structure within a voxel. *Magnetic Resonance in Medicine*, 48(3):454–459, September 2002.

[136] S. Wakana, H. Jiang, L. M. Nagae-Poetscher, P. C. van Zijl, and S. Mori. Fiber tract-based atlas of human white matter anatomy. *Radiology*, 230(1):77–87, 2004.

[137] Setsu Wakana, Arvind Caprihan, Martina M Panzenboeck, James H Fallon, Michele Perry, Randy L Gollub, Kegang Hua, Jiangyang Zhang, Hangyi Jiang, Prachi Dubey, Ari Blitz, Peter van Zijl, and Susumu Mori. Reproducibility of quantitative tractography methods applied to cerebral white matter. *Neuroimage*, 36(3):630–44, Jul 2007.

[138] Zhizhou Wang and Baba C. Vemuri. DTI segmentation using an information theoretic tensor dissimilarity measure. *IEEE Trans. Med. Imaging*, 24(10):1267–1277, 2005.

[139] D. Wassermann, L. Bloy, E. Kanterakis, R. Verma, and R. Deriche. Unsupervised white matter fiber clustering and tract probability map generation: Applications of a Gaussian process framework for white matter fibers. *NeuroImage*, 51(1):228–241, 5 2010.

[140] Demian Wassermann, Maxime Descoteaux, and Rachid Deriche. Diffusion maps clustering for magnetic resonance Q-Ball imaging segmentation. *International Journal in Biomedical Imaging*, 2007. In press.

[141] Yonas Weldeselassie and Ghassan Hamarneh. DT-MRI segmentation using graph cuts. In *SPIE Medical Imaging*, pages 1–9, 2007.

[142] C. F. Westin, S. E. Maier, H. Mamata, A. Nabavi, F. A. Jolesz, and R. Kikinis. Processing and visualization for diffusion tensor MRI. *Med Image Anal*, 6(2):93–108, 2002.

[143] C.-F. Westin, S. Peled, H. Gudbjartsson, R. Kikinis, and F. A. Jolesz. Geometrical diffusion measures for MRI from tensor basis analysis. In *ISMRM '97*, page 1742, 1997.

[144] Dongrong Xu, Susumu Mori, Meiyappan Solaiyappan, Peter C M van Zijl, and Christos Davatzikos. A framework for callosal fiber distribution analysis. *Neuroimage*, 17(3):1131–43, Nov 2002.

[145] Santiago Ramón y Cajal. *Histology of the nervous system of man and vertebrates*. Oxford University Press, 1995.

[146] S.-S. Yoo, H. J. Park, J. S. Soul, H. Mamata, C.-F. Westin, H. Bassan, A. J. Du Plessis, R. L. Robertson, S. A. Ringer, S. E. Maier, J. J. Volpe, and G. P. Zientara. In vivo visualization of white matter fiber tracts of preterm- and term- infant brain with diffusion tensor MRI. *Investigative Radiology*, 40(2):110–115, February 2005.

[147] Fan Zhang, Casey Goodlett, Edwin Hancock, and Guido Gerig. Probabilistic fiber tracking using particle filtering. *Med Image Comput Comput Assist Interv*, 10(Pt 2):144–52, 2007.

[148] Hui Zhang, Paul A Yushkevich, D. C. Alexander, and James C. Gee. Deformable registration of diffusion tensor MR images with explicit orientation optimization. *Med. Image. Anal.*, 10(5):764–85, Oct 2006.

[149] Song Zhang and David H. Laidlaw. Hierarchical clustering of streamtubes. Technical Report CS-02-18, Brown University Computer Science Department, August 2002.

[150] L. Zhukov and A. H. Barr. Heart-muscle fiber reconstruction from diffusion tensor MRI. In *Proceedings of IEEE Visualization 2003*, pages 597–602. IEEE Computer Society, 2003.

[151] L. Zhukov, K. Museth, D. Breen, R. Whitaker, and A. Barr. Level set modeling and segmentation of DT-MRI brain data. *Journal of Electronic Imaging*, 12(1):125–133, 2003.

[152] U. Ziyan, D. Tuch, and C. Westin. Segmentation of thalamic nuclei from dti using spectral clustering. pages 807–814, 10 2006.

166

Summary

Homogeneity based segmentation and enhancement of Diffusion Tensor Images
A white matter processing framework

In diffusion weighted magnetic resonance imaging (DW-MRI) the Brownian motion of the water molecules, within biological tissue, is measured through a series of images. In diffusion tensor imaging (DTI) this diffusion is represented using tensors. DTI describes, in a non-invasive way, the local anisotropy pattern enabling the reconstruction of the nervous fibers - dubbed tractography. DW-MRI constitutes a powerful tool to analyse the structure of the white matter within a voxel, but also to investigate the anatomy of the brain and its connectivity. DW-MRI has been proved useful to characterize brain disorders, to analyse the differences on white matter and consequences in brain function. These procedures usually involve the virtual dissection of white matters tracts of interest. The manual isolation of these bundles requires a great deal of neuroanatomical knowledge and can take up to several hours of work.

This thesis focuses on the development of techniques able to automatically perform the virtual dissection of white matter structures. To segment such structures in a tensor field, the similarity of diffusion tensors must be assessed for partitioning data into regions, which are homogeneous in terms of tensor characteristics. This concept of tensor homogeneity is explored in order to achieve new methods for segmenting, filtering and enhancing diffusion images.

First, this thesis presents a novel approach to semi-automatically define the similarity measures that better suit the data. Following, a multi-resolution watershed framework is presented, where the tensor field's homogeneity is used to automatically achieve a hierarchical representation of white matter structures in the brain, allowing the simultaneous segmentation of different structures, with different sizes. The successor of DTI, high angular resolution diffusion imaging (HARDI), is also approached in this thesis. An accelerated convolution method of HARDI images is presented, where the stochastic process of water diffusion is modeled, enabling the contextual processing of diffusion images for noise re-

167

duction, regularization and enhancement of white matter structures. This thesis finishes by exploring the use of this contextual processing techniques to enhance DTI data, in order to obtain similar crossing information as from HARDI data.

These new methods are analysed on the basis of their accuracy, robustness, speed and usability - key points for their application in a clinical setting. The described methods enrich the analysis of white matter structures, fostering the exploration and understanding of the human brain.

Curriculum Vitae

Paulo Reis Rodrigues was born on 09-04-1981 in Vila Franca, Portugal. After finishing pre-university education in 1999 at Secondary School of Monserrate in Viana do Castelo, Portugal, he studied 5 year Licenciate studies (MsC equivalent) in computer science and systems engineering at the University of Minho in Braga, Portugal. In March 2006, after 2 year incursion in the IT industry, he started his PhD research in the Biomedical Image Analysis group of the department of Biomedical Engineering at Eindhoven University of Technology. His research focus was in giving insight to diffusion tensor imaging (DTI) data, by exploring and exploiting image analysis related techniques in order to improve the analysis and visualisation of this data, especially of the brain. He published several papers with new techniques for the virtual dissection of white matter structures and post processing techniques for the improvement of DTI but also of the new high angular resolution diffusion imaging technique. He developed various parts of the visualisation tool for DTI and HARDI data used within and outside the research group where his work was conducted.

Publications originated from this work

ch:publications

Book chapters

1. T. H. J. M. Peeters, P. R. Rodrigues, A. Vilanova, and B. M. Haar Romeny. *Analysis of distance/similarity measures for diffusion tensor imaging*. In Gerald Farin, Hans-Christian Hege, David Hoffman, Christopher R. Johnson, Konrad Polthier, Martin Rumpf, David Laidlaw, and Joachim Weickert, editors, Visualization and Processing of Tensor Fields, Mathematics and Visualization, pages 113136. Springer Berlin Heidelberg, 2009.

Refereed proceedings

1. V. Prčkovska, P.R. Rodrigues, R. Duits, B.M. ter Haar Romenij, A. Vilanova, *Extrapolating fiber crossings from DTI data. Can we gain similar information as HARDI?*, in MICCAI workshop on Computational Diffusion MRI, pages 26-37. Beijing, China, September 2010.

2. P.R. Rodrigues, R. Duits, B.M.ter Haar Romenij, A. Vilanova: *Accelerated diffusion operators for enhancing DW-MRI*, in Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM), pages 49-56. Leipzig, Germany, July 2010.

3. P.R. Rodrigues, A. Jalba, P. Fillard, A. Vilanova, B.M. ter Haar Romenij, *A Multi-Resolution watershed-based approach for the segmentation of Diffusion Tensor Images*, in MICCAI Workshop on Diffusion Modelling, pages 162-173. London, United Kingdom, September 2009.

4. P.R. Rodrigues, A. Vilanova, T. Twellmann, B.M. ter Haar Romenij, *Adaptive Distance Learning Scheme for Diffusion Tensor Imaging using Kernel*

Target Alignment, in MICCAI Workshop on Computational Diffusion MRI (CDMRI), pages 148-158. New York, United States, September 2008.

Conference abstracts and posters

1. V. Prčkovska, P.R. Rodrigues, R. Duits, B.M. ter Haar Romenij, A. Vilanova, *Extrapolating fiber crossings from DTI data. Can we gain similar information as HARDI?*, 3rd Dutch Biomedical Engineering (BME), Egmond aan Zee, The Netherlands, 2011.

2. P.R. Rodrigues, V. Prčkovska, W.L.P.M. Pullens, B.M. ter Haar Romenij, A. Vilanova, G.J. Strijkers, *Validating Validators: An Analysis of DWMRI Hardware and Software Phantoms*, in ISMRM; Stockholm, Sweden, May 2010.

3. P.R. Rodrigues, A. Jalba, P. Fillard, B.M. ter Haar Romenij, A. Vilanova, *A Multi-Resolution watershed-based approach for the segmentation of Diffusion Tensor Images*, in ISMRM; Stockholm, Sweden, May 2010.

4. P.R. Rodrigues, A. Jalba, P. Fillard, B.M. ter Haar Romenij, A. Vilanova, *A Multi-Resolution watershed-based approach for the segmentation of Diffusion Tensor Images*, in ISMRM Benelux Chapter; Utrecht, The Netherlands, January 2010.

5. P.R. Rodrigues, V. Prčkovska, W.L.P.M. Pullens, G.J. Strijkers, A. Vilanova, B.M. ter Haar Romeny Validating Validators: *An Analysis of DW-MRI Hardware and Software Phantoms*, in ICT Innovations 09, Macedonian Society on Information and Communication Technologies, Macedonia, October 2009.

6. P.R. Rodrigues, A. Vilanova, T. Twellmann, B.M. ter Haar Romenij, *Adaptive Distance Learning Scheme for Diffusion Tensor Imaging using Kernel Target Alignment*, in ASCI, Zeewolde, The Netherlands, June 2009.

7. P.R. Rodrigues, V. Prčkovska, T.H.J.M. Peeters, A. Vilanova, B.M. ter Haar Romenij, *DTITool: An Insightful look into Diffusion Tensor Images*, EuroVis 2009, Berlin, Germany, June 2009.

8. P.R. Rodrigues, A. Vilanova, T. Twellmann, B.M. ter Haar Romenij, *Improving Diffusion Tensor Imaging Segmentation Through an Adaptive Distance Learning Scheme*, in ISMRM 09, Honolulu, Hawaii, United States, 2009.

9. P.R. Rodrigues, A. Vilanova, T. Twellmann, B.M. ter Haar Romenij, *Adaptive Distance Learning Scheme for Diffusion Tensor Imaging using Kernel Target Alignment*, 2nd Dutch Biomedical Engineering (BME), Egmond aan Zee, The Netherlands, 2009.

10. P.R. Rodrigues, T.H.J.M Peeters, A. Vilanova, B.M. ter Haar Romenij, *Analysis of distance/similarity measures for Diffusion Tensor Imaging*, International Mathematica Symposium (IMS), Maastricht, The Netherlands, 2008.

Acknowledgements

I would like to start by exercising a bit of self indulgence, and say thank you to myself. After all, it was me that was unsatisfied with himself, somewhere in Braga, sometime during 2004, and decided that going abroad, to The Netherlands, to research, was the way to go. Enough of egocentrism! This has been quite an amazing journey. I have learned a lot, a bit in the field of diffusion, image analysis and all of that, but what really stays imprinted in my brain is much more than that. I have grown as a professional, as a scientist, but more as a person, as a man. I would like to acknowledge everybody involved in this, which basically means everybody I have met, somehow, in different levels, during this PhD.

First things first. Prof. Bart, my promotor, you have been a beacon of positivity. Always helpful, sometimes with crazy ideas and suggestions, but always interesting and useful. You are a true man of the world, and hopefully I will manage to step on, at least, one tenth of the places you have been! Anna, my supervisor, always trying to make me organized, and plan. Well, the plan got changed but I finished anyway :). And I still smile about our discussions about the colloquia! Luc, you were the connection point to Eindhoven. Through João Saraiva (I'll thank you a bit later) I've come to find The Netherlands. You have been there to help me, and always managed to put math in nice simple terms! Besides being a great scientist, you're *savoir-vivre* is an inspiration. Thorsten, it was a pleasure to work with you on my first real paper!

Then, my colleagues, actually my friends. Puuuuuttaaaa (a.k.a. Petr), as the first doctor amidst the possy, you have laid the way to a successful journey. You have enlightened me to Zouk and for that I am truly thankful. Oh, and sorry for messing up 'our' house, back in the days! Mr. Roy with whom I share so many interests. Again thank you for lending me your guitar to strum like there is no tomorrow. Ellen I didn't work with you, but wish I did (so I heard). Alex, bambino, always interesting to discuss cinema with you, and some particularities of Italian language. Ralph, always fun to find you in Portugal, in a random moment! Bram, wish I had half of your stuff, really! I've learned a lot about photography with you. Not to mention about jacuzzis and their subtleties... Remco, when I grow up into a mad scientist, I have you as my reference. It was always fun to have

your impetuous visits to the office and even more your profane exits! I really enjoyed working with you and those mathy kernels! Andrea, I will always remember your good mood and your coolness in the DTI meetings, chewing gum, gently caressing your Luna tummy. Timóteo, I wrote my first paper with you. It was one of the most boring moments of my life. We agree on this! Maybe it was its subject, maybe it was you! But hey, besides all the boringness, there are those moments where the sense of humor is just synchronized. And that should count for something, at least for spookiness. For the rest, Hans, Marcel, Markus, Margret, Frans, Mehrdad, Marieke, Fitsum, Erik, Bart, Fernando and Iñaki, it has been a pleasure. Danny, I am always curious which sordid piece of information is going to be shared! Deeply appreciated!

A special hi five (J.C. style) to the Delft possy, Jorik, Peter, Ling Xiao, Francois. And Charl, I still have opalescent dreams with that 11^{th} dimensional umbrella. It's always remarkable to talk to you.

Through these years of PhD, I've met lots of people, researchers, and scientists to be, all over different conferences, all over the world. With you I've engaged in many conversations, more or less scientific, but always delightful. Max and Pierre, I've met you in the mythical summer school of 2007. Apparently partying, dancing and drinking forested great collaborations. But, more than collaborators, I've earned friends! Thanks for being such good hosts at Neurospin, and Paris, yet if for a short visit (bloody hazardous bureaucracy!). Pierre, a special thanks for the exquisite teachings in mastering the French language. Carl, Ragini, Baba, Demian, Kathleen, Kathrine, Gordon, and Auro, we always spent more time together outside of conferences than in them, no? A special aloha for Aga, Jadrian and Erik, the Hawaii possy. Always awesome to enjoy the pleasures of the parallel life to science. Hans Knutsson, you are the coolest scientist ever! Like a rock star of science. I think we've never discussed anything scientific - just life itself. Daniel and the London crew, thanks for a warm welcome while visiting your lab, and of course for the rich discussions on diffusion and such. It's always fun to have a beer with you guys. I would also like to thank Nicolas Lori, the first Portuguese diffusion researcher I've met, in Israel, somewhere at the Dead Sea (go figure). Thanks for all your tips and advices. Leticia Rittner, from the other side of the Atlantic, with whom I had fun trying to discuss diffusion in Portuguese.

All of this started back in Portugal, with the support of a couple of teachers, especially Victor Fonte, António Ramires and João Saraiva, to whom I 'owe' finding Eindhoven and this PhD.

On a more personal and regional level, I would like to mention a couple of friends and important people in my life, from Portugal. Raul, my best friend since ever. It all started with our discussions on the meaning of life, the universe, cows and everything, on the walls of our homes, long time ago, when we were small, restless, young kids. Every time I talk to you, I get reminded why I love

Portuguese language. I love you man! (greetings to Catarina and Matias). To Carlos, and the rest of *grupinho*, Xana, Paulinha, Carlinhos, Helder and Luisa, thank you for being part of the circle! A special one to Isabel, for being here in Holland, always helpful, full of wisdom and tip$.

During a PhD, one can loose his mind, always sitting behind a computer, typing, typing, and typing the quirks and kinks of diffusion. I want to give a big thank you to all the people in the Zouk scene, especially to my teachers Adilio (Adiliozão!!!), Gilson and Dikla. Dancing Zouk is a bliss for the body and mind! On my other passion, acrobatics, I want to give a special thanks to Det and Cahit (and Minka) for being such amazing teachers, and for always having the patience for our eagerness in trying some advance move! To everybody else in the classes, I'll fly you someday!

Zoki, you are a great and unique friend. Keep on singing to life! Oh, and I can't forget your Robin Hood persona in the bostel in Amsterdam! Bubeeeeee! You are the craziest brother I never had. Respect to your muscles and perseverance. I love you (like a brother ah, no funny thoughts!). Suze, we shared some special moments while you were visiting, and I was writing this thesis ;) You are like a very cool "grown up".

I have to acknowledge two special friends, Monster^{TM} and Corona^{TM}. You have profusely fuelled this thesis! Speaking of delicacies, I also want to acknowledge ajvar for the provided feast for my gustatory senses. To Pixie and Kenzo, the greatest cats in the world. It's amazing how, in your furriness, you manage to bring peace into my life.

Para os meus pais, Fernanda e Casimiro. Um muito obrigado por serem, sim, os meus pais! Tenho estado longe, a dar noticias de vez em quando, desculpem qualquer coisinha. À minha mãe, por sempre perguntar se já comi, e o que comi. Ao meu pai por perguntar se descobri alguma coisa. :) Obrigado por estarem sempre dispostos a ajudar e me apoiarem. Amo-vos. Um beijo especial para a minha avó, Maria. Sinto tua falta. Espero que estejas orgulhosa, lá algures onde estiveres. Para a minha outra avó, Madrinha, um grande beijo, e mais um neto para ter orgulho.

vesna