

*LESS CONFUSION
IN DIFFUSION
MRI*

Colophon:

Cover: Tractography results on a 10 mm thick coronal section of the brain, showing how the algorithms developed in this thesis could contribute to less confusion in diffusion MRI. When the book is opened, the transition from a confusing geometry on the left to geometric sheet structures on the right becomes visible. Left hemisphere, from left to right: 1) Tractography on the diffusion tensor, estimated with linear least squares on a dataset with severe outliers. 2) Tractography on the diffusion tensor, estimated with REKINDLE to reduce the effect of outliers. Streamlines look more coherent, but no crossing fibers can be resolved. 3) Tractography using CSD, crossing fibers can be recognized. However, a too sharp response function is used, resulting in some spurious streamlines. Right hemisphere, from left to right: 4) Tractography using CSD and recursive calibration of the response function to reduce spurious streamlines. 5) Orientation-dependent transparency rendering reveals underlying structures. 6) Path neighborhood tractography that can be used to calculate the sheet probability index (SPI), revealing sheet structures at some locations.

ISBN: 978-94-629-5449-6

Cover design and layout: Chantal Tax and Maxime Chamberland

Printed by: Proefschriftmaken.nl || Uitgeverij BOXPress

LESS CONFUSION IN DIFFUSION MRI

MINDER CONFUSIE IN DIFFUSIE MRI

(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht
op gezag van de rector magnificus, prof.dr. G.J. van der Zwaan,
ingevolge het besluit van het college voor promoties in het openbaar te verdedigen op
donderdag 8 september 2016 des middags te 2.30 uur

door

Chantal Margaretha Wilhelmina Tax

geboren op 10 september 1988
te Weert

Promotor: Prof. dr. ir. M.A. Viergever

Copromotor: Dr. A.L.G. Leemans

The work described in this thesis is supported by a grant (No. 612.001.104) from the Physical Sciences division of the Netherlands Organization for Scientific Research (NWO), and the Marina van Damme grant. The publication of this thesis was financially supported by Philips Healthcare and Mint Labs.

CONTENTS

1 Thesis Overview.....	1
2 Introduction	7
3 The “MASSIVE” Brain Dataset: Multiple Acquisitions for Standardization of Structural Imaging Validation and Evaluation.....	51
4 REKINDLE: Robust Extraction of Kurtosis INDices with Linear Estimation	71
5 Recursive Calibration of the Fiber Response Function for Spherical Deconvolution of Diffusion MRI Data	101
6 Characterizing Diffusion MRI Profiles of Single-Fiber Populations: An Inter-Subject and Inter-Tract Analysis.....	127
7 Sheet Probability Index (SPI): Characterizing the Geometrical Organization of the White Matter with Diffusion MRI	147
8 Quantifying the Brain’s Sheet Structure with Normalized Convolution	189
9 Seeing More by Showing Less: Orientation-Dependent Transparency Rendering for Fiber Tractography Visualization.....	219
10 Discussion.....	243
References	255
Samenvatting	277
Acknowledgements.....	283
Publications.....	289
About the Author	293

1 THESIS OVERVIEW

1.1 MOTIVATION

With its unique ability to investigate tissue microstructure *in vivo*, diffusion MRI (dMRI) is one of the preferred approaches for investigating the brain's structural connectivity. Based on the fundamental principle that water molecules randomly move due to thermal energy and that this so-called Brownian motion is modulated by tissue constituents, dMRI has been widely used in clinical and biomedical research applications to infer valuable information about the underlying microstructure. The immense potential of dMRI in clinical and scientific disciplines has already been recognized since its early days. Especially in acute stroke, dMRI has been shown to be very useful for identifying ischemic regions that could not be detected with the more conventional methods (Van Gelderen et al., 1994). With the advent of diffusion tensor imaging (DTI) (Basser et al., 1994), the anisotropic behavior of the diffusion process could now also be characterized unambiguously in a three-dimensional (3D) fashion. Furthermore, DTI allowed for the virtual reconstruction of pathways – a process called tractography – of the underlying fibrous tissues such as white matter. Considered to be a quantitative imaging approach, DTI was quickly adopted in a wide range of clinical research applications (e.g. Barnea-Goraly et al., 2004; Lebel et al., 2008; Lim et al., 1999; Werring et al., 1999). However, DTI has some well-known drawbacks in that it cannot resolve complex fiber configurations such as ‘crossing fibers’ and it relies on a Gaussian approximation of the diffusion process.

Due to the ever growing interest in dMRI, the society has been continuously triggered to develop novel approaches beyond DTI to overcome its shortcomings. For example, diffusion kurtosis imaging (DKI) was proposed as an extension of DTI that can provide non-Gaussian diffusion measures (Jensen et al., 2005), and spherical deconvolution was introduced to deal with crossing fibers (Dell'Acqua et al., 2010; Tournier et al., 2007). Such new approaches have been accompanied by new developments in correction procedures for subject motion and image distortions, and in tractography, multi-subject group comparisons, and data visualization, amongst others. With the overwhelming amount of strategies currently available and still being developed, it is unfortunately not always evident to the end-users how dMRI can be optimally used to address their application. In addition, differences in processing strategies lead to ambiguities as to which conclusions can and cannot reliably be drawn from dMRI data, resulting in controversies in the field (e.g. Catani et al., 2012; Wedeen et al., 2012a; Wedeen et al., 2012b). Such issues hamper a smooth transition of dMRI processing strategies into useful tools for clinical and biomedical applications.

This thesis contributes to reducing the confusion in diffusion MRI by scrutinizing different steps of the processing pipeline. The chapters in this thesis focus on making the topics accessible for a broad audience, and propose new methodology to make more intuitive and data-driven choices in dMRI data processing and to facilitate interpretation and visualization of dMRI data. The methodology developed in this thesis is furthermore used to investigate fundamental topics such as variability in characteristics of the dMRI signal and the geometrical organization of the brain pathways. Computer simulation- and real data evaluation frameworks play an important role throughout the thesis.

1.2 ORGANIZATION OF THE THESIS

The dMRI processing pipeline is often lengthy and complex, and consists of numerous processing steps such as 1) artifact correction, 2) model selection, and 3) parameter estimation. **Chapter 2** introduces these steps and reviews the most commonly used and current state-of-the-art dMRI processing techniques with a focus on the brain. While it is by no means endeavored to cover all the literature on this topic, this chapter serves as a concise guide to *what's new* in dMRI processing and may assist in choosing the optimal dMRI processing strategy for addressing a particular application. The use of equations and technical language is kept to a minimum to make this overview more approachable. To place the work in Chapters 3 to 9 into context individually, each chapter is preceded by its own more specific introduction.

Notwithstanding recent developments described in Chapter 2, optimizing the dMRI processing pipeline is still an active area of research. In this context, a comprehensive evaluation framework for novel diffusion processing techniques and for reliable comparisons between different approaches is highly desired. **Chapter 3** describes the MASSIVE (Multiple Acquisitions for Standardization of Structural Imaging Validation and Evaluation) brain dataset containing multi-modal MR data and 8000 dMRI volumes of a single healthy subject acquired on a clinical 3 T scanner. All the datasets are designed to meet requirements for state-of-the-art processing strategies, and specifically acquired in a clinical setting to be in line with the current standards in acquisition protocols from routine examinations. As such, subsets of the MASSIVE dataset can serve as representative test beds for new developments in a wide range of dMRI data processing techniques. The dataset is used in several chapters in this thesis as an evaluation framework and is made publicly available to further facilitate dMRI processing development.

Accurate and precise parameter estimation is an ongoing field of research as it is highly challenged by noise and other types of data deterioration such as outliers. Outliers can severely bias parameter estimates and as such make subsequent analyses unreliable. For example, the added information and benefits of DKI compared to DTI are potentially nullified in a clinical setting by the higher sensitivity of DKI to outliers due to its increased model complexity and higher acquisition demands. In **Chapter 4** a robust estimation procedure is proposed coined REKINDLE (Robust Extraction of Kurtosis INDices with Linear Estimation). By means of fast linear estimation and an intuitive outlier detection criterion REKINDLE aims to identify and exclude outliers. The REKINDLE framework is by design also applicable to DTI and other dMRI models that can be linearized.

Chapter 5 focuses on the estimation of the fiber orientation density function (fODF) that can better characterize complex fiber configurations than the diffusion tensor. Spherical deconvolution (SD) based techniques are gaining popularity since they provide a balanced trade-off between acquisition demands and the reliability of the reconstructed fODF. Like any deconvolution technique SD relies on the choice of an appropriate response function (RF) that represents the diffusion-weighted signal profile of a single fiber population (SFP). An inaccurate calibration of the RF can severely impact the estimated fODF and subsequent tractography. Currently, the computation of this RF still relies on the estimation of DTI, which is inaccurate at typical acquisition settings used for SD. In this chapter, a data-driven framework is proposed that recursively excludes voxels with crossing fibers. By means of an

intuitive parameter the algorithm localizes voxels with an SFP that can be used to define the RF. In **Chapter 6**, the recursive framework proposed in Chapter 5 is used to localize SFPs in multiple subjects. While a lot of attention has been given to crossing fibers, many fundamental questions on SFPs remain still unanswered. By fitting a range of dMRI models, e.g. DKI with REKINDLE from Chapter 4, the characteristics of the signals in the identified SFP voxels can be further explored. In addition, the variability of SFP properties can be investigated across datasets (where we use the MASSIVE dataset from Chapter 1), across subjects, and across tracts. The knowledge obtained in simpler configurations like SFPs can provide important information on tissue microstructure and aid in the design of simulation experiments, amongst others.

With the advent of dMRI approaches that can model crossing fibers, the geometrical organization of pathways in the white matter can be investigated into more detail. Recently, Wedeen et al. (2012b) visually analyzed adjacency and crossing of pathways reconstructed with dMRI and tractography. The authors proposed that pathways cross each other orthogonally on surfaces somewhere along their trajectory: they form sheet structures in the brain. Furthermore, they found that no brain pathways were observed without sheet structure, and that these structures could have a relation to embryogenesis and axonal path finding. In a response, Catani et al. (2012) stated that this finding was likely an artifact introduced by the dMRI methodology used in Wedeen et al. (2012b). The dMRI technique used is less able to resolve fiber crossings at smaller angles than for example SD techniques. According to them, the low angular resolution creates a bias to orthogonal angles which makes a grid structure of interwoven sheets a very likely configuration. Wedeen et al. (2012a) rebutted these concerns and stated that there are no means known whereby this highly specific mathematical structure is created as an artifact. Importantly, Catani et al. (2012) also stated that Wedeen et al. (2012b) mainly showed qualitative results and that pathways reconstructed with dMRI and tractography cannot be equated to true axons. In **Chapter 7**, the mathematical theory behind sheet structure is recapitulated, and the condition for sheet structure and the role of orthogonal angles is discussed. The sheet probability index (SPI) is proposed as a measure to quantify the extent to which the data supports sheet structure. A method to calculate the SPI is presented that requires the reconstruction pathways with tractography. This method is extensively evaluated in simulations and real dMRI data, where we use the MASSIVE dataset and SD with the recursive framework from Chapter 5, amongst others. **Chapter 8** proposes a conceptually different method to calculate the SPI that does not rely on tractography and is less computationally intensive.

Another consequence of the ability to model crossing fibers is that the complexity of the fiber network as reconstructed with tractography increases tremendously. Many pathways interdigitate and overlap, which hampers an unequivocal 3D visualization of the network and impedes an efficient study of its organization. **Chapter 9** proposes a novel fiber tractography visualization approach that interactively and selectively adapts the transparency rendering of fiber trajectories as a function of their local or global orientation. This allows for an improved 3D visualization of the fiber network and the exploration of tissue configurations that would otherwise be largely covered by other pathways. This approach is used to efficiently extract fiber bundles and evaluated in a neurosurgical planning case, where the pathways were obtained with the technique presented in Chapter 5.

Chapter 10 discusses the findings in this thesis and sheds some light on *what's next* in dMRI processing, gratefully inspired by both formal and informal discussions with many research colleagues in the field.

1.3 SOFTWARE CONTRIBUTIONS

The algorithms developed in this thesis are or will be integrated in the software tool *ExploreDTI* (Leemans et al., 2009). This MATLAB-based software tool for processing and analysis of dMRI data is freely available upon request (<http://exploredti.com/>). The recursive calibration for SD (Chapter 5) is available in *Dipy* (Garyfallidis et al., 2014). This open source software project for dMRI analysis is written in Python and can be downloaded at <http://nipy.org/dipy/index.html>. An interactive implementation of the orientation-dependent transparency rendering of tractography pathways (Chapter 9) is integrated in the *Fibernavigator* (Chamberland et al., 2014). This open source tool written in C++ is available at <https://github.com/chamberm/fibernavigator>.

2 INTRODUCTION

Based on:

C.M.W. Tax, A. Leemans, “What’s new and what’s next in diffusion MRI processing” (to be submitted soon)

ABSTRACT

Recent advancements in diffusion MRI (dMRI) have offered exciting new avenues for investigating microstructural characteristics of fibrous tissue *in vivo*, contributing to the proliferation of dMRI studies in current literature. This growing interest for integrating dMRI in a wide range of clinical and scientific disciplines, in turn, has triggered the development of different strategies for processing and analyzing dMRI data. Since the pipeline is often lengthy and complex, and consists of numerous processing steps, it is unfortunately not always evident which particular approaches would be preferred to address a specific question. In this chapter, we review the most commonly used and current state-of-the-art dMRI processing techniques. While the majority of the presented methodology is generic, we will mainly focus here on using dMRI to study the human brain.

2.1 INTRODUCTION

Diffusion MRI (dMRI) has the unique ability to investigate tissue microstructure *in vivo*, and has been widely used in clinical and biomedical research applications to infer valuable information about the underlying microstructure and the brain's structural connectivity. dMRI is based on the fundamental principle that water molecules randomly move due to thermal energy and that this so-called Brownian motion is modulated by tissue structures. The principles of dMRI are not covered in great detail here as they already have been explained previously in numerous excellent reviews (e.g. Basser and Jones, 2002; Deprez et al., 2013; Hagmann et al., 2006; Jones and Leemans, 2011; Le Bihan et al., 2001; Le Bihan and Johansen-Berg, 2012; Mori and Zhang, 2006; Tournier et al., 2011) and reference books (e.g. Johansen-Berg and Behrens, 2009; Jones, 2010b; Mori, 2007; Van Hecke et al., 2016). For those readers that would benefit from a refresher of dMRI – and, more specifically, DTI – we present a diffusion MRI 101 in Appendix A2. While we do provide the reader with mathematical descriptions of several common concepts in the appendices, we avoid the use of equations and technical language in the body of the chapter to make the topics accessible for a broad audience.

Due to the ever growing interest in dMRI, the society has been continuously triggered to develop novel approaches to optimally process dMRI data. In this chapter, we will describe the necessary processing steps to ‘prepare’ the dMRI data for subsequent analysis and give an overview of commonly used and state-of-the-art processing methods. These processing steps include 1) checking and correcting the data for artifacts, 2) choosing an appropriate model or representation to extract useful features from the data, and 3) obtaining reliable parameter estimates for the chosen model, taking into account noise and outliers. Throughout the chapter, we take into account that the end-user often only has access to a ‘routinely-clinical’ type of acquisition: a (single-shot) echo-planar imaging (EPI) dMRI dataset acquired with multiple gradient directions and potentially multiple b -values, along with one or several anatomical (e.g. T1- and /or T2-weighted) MR images which are often acquired as part of a more general protocol.

In Section 2.2 we will discuss some common artifacts in dMRI that are often still present in the data despite regular scanner quality assurance (e.g. by imaging phantoms). Some of these artifacts are well-known and many software packages provide tools to correct for them; these corrections are more or less ‘standard’ in a dMRI pipeline. Other artifacts have largely gone unnoticed or have been ignored by the community. Section 2.3 gives a concise overview of different dMRI modelling and reconstruction techniques. There are many strategies that go beyond DTI, but the choice of model largely depends on the research question and the type of data that the end-user has access to. Section 2.4 focuses on how to reliably estimate diffusion parameters once a specific model is chosen.

Note that there is no general consensus about the order in which data processing steps are applied and that these steps are often interwoven. Some artifacts, for example, can be corrected for during parameter estimation, or dMRI models can be used to spot and simultaneously correct for artifacts. In the following, we will often display image maps derived from DTI, such as the color-coded fractional anisotropy (FA) map, to evaluate the performance of the

processing tools. As most researchers are familiar with these maps, it will facilitate the interpretation of the results.

2.2 CORRECTING FOR ARTIFACTS

One of the first steps in dMRI processing after the data has been acquired is to check image quality and correct for potential artifacts (Andersson and Skare, 2010; Jones and Cercignani, 2010; Jones and Leemans, 2011; Pierpaoli, 2010; Tax et al., 2016c; Tournier et al., 2011). Artifacts can be considered as errors in the measurements which can lead to misinterpretation of the data. In the dMRI processing pipeline, measurements from multiple acquisitions will be combined, e.g. from different diffusion images or from a T1-weighted image and a dMRI dataset. Consequently, both artifacts that occur in individual images, or that systematically occur across multiple images can have significant consequences on the reliability of subsequent analysis. For example, artifacts can cause the estimated FA values to deviate from their true values, affecting a quantitative comparison in group studies. Alternatively, errors in estimates of the diffusion tensor (DT) first eigenvector can negatively bias DTI-based tractography and might change tract geometry or cause pathways to terminate in different brain areas. This for example affects connectivity studies, or hampers correct interpretation in neurosurgical applications.

A detailed inspection of dMRI data quality is often omitted since it is generally time consuming. Given its importance, however, we will discuss in Section 2.2.1 a few straightforward and efficient tools for dMRI data quality checks which can be performed prior and after artifact correction. Sections 2.2.2 to 2.2.5 will consider various types of artifacts and their correction strategies. Note that while there are many specialized data acquisition techniques for reducing artifacts (Andersson and Skare, 2010; Peterson and Bammer, 2016; Pierpaoli, 2010; Tax et al., 2016c), we will focus here on the processing techniques. Sections 2.2.2 and 2.2.3 will discuss two artifacts that have been largely ignored in the diffusion community but which have gained attention recently: signal drift and Gibbs ringing. Whereas these artifacts are mainly acquisition-related, other artifacts are more subject-related: Section 2.2.4 discusses motion related artifacts such as bulk motion, cardiac pulsation, and signal instability between slices. Section 2.2.5 considers two sources of artifacts resulting in geometric image distortions that become more pronounced because of the EPI readout: Eddy currents (Section 2.2.5.1) and susceptibility differences in tissues (Section 2.2.5.2).

2.2.1 Data quality assessment

Visual inspection of the individual diffusion images can reveal some artifacts (e.g. instabilities between slices (Fig. 2.1a), slice dropouts (Fig. 2.1b), signal decay due to vibration (Fig. 2.2c)), but others are hard to spot in these diffusion images. An efficient strategy for highlighting artifacts while avoiding a time-consuming and detailed slice-by-slice inspection of each dMRI data set is simply combining diffusion images in various ways (Heemskerk et al., 2013; Tax et al., 2016c; Tournier et al., 2011). For instance, quickly toggling between or looping through the same slice of different diffusion images can already reveal large signal dropouts, the presence of subject motion, and the magnitude of geometric distortions.

An image of the standard deviation of the non-diffusion weighted images (DWIs) can highlight motion resulting from cardiac pulsation (Fig. 2.1d), while an image of the standard deviation of the DWIs can reveal image misalignment due to subject motion and eddy-current distortions as indicated by the bright rim in Fig. 2.1e. With further inspection of the color-coded FA map one can also recognize other artifacts. Examples include bulk subject motion which leads to misalignment between images (Fig. 2.1f), vibration (Fig. 2.1g, which is the same slice as Fig. 2.1c), eddy current distortions (Fig. 2.1h), and abnormally high and low FA values (Fig. 2.1i).

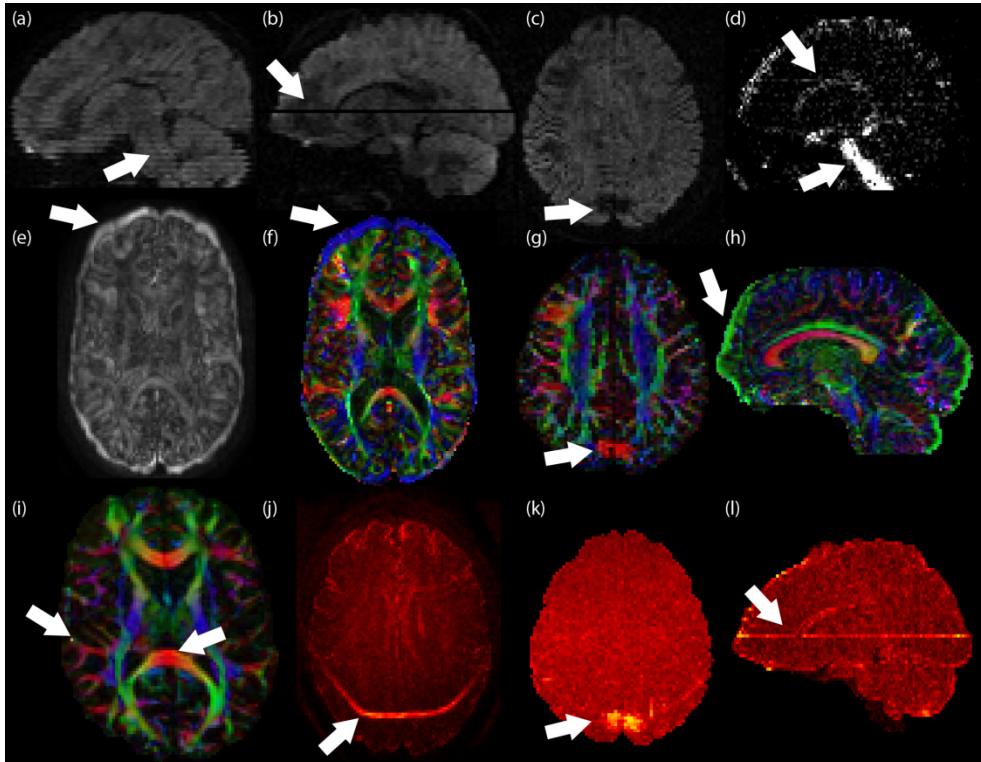


Fig. 2.1: Examples of common dMRI artifacts. Some artifacts can be recognized in the diffusion images, such as instabilities between slices (a), slice dropout (b), and signal decay due to vibrations (c). Other artifacts can be more easily spotted with ‘standard-deviation’ maps, such as cardiac pulsation around the brainstem and ventricles (d), and motion and eddy current distortions (e); FA color-coded maps ((f) motion, (g) vibration, (h) eddy current distortions, (i) artificially high/low FA values, in the corpus callosum likely due to Gibbs ringing); and residual maps ((j) chemical shift, (k) vibration, (l) slice dropout).

Residuals, which represent the difference between a measurement and the prediction of the signal according to a dMRI model (Appendix A2.3), can highlight artifacts that are not always clearly visible on individual DW images or DTI derived maps. An image of the average residual per voxel across DWIs should ideally be uniform. Large mean residuals either indicate that the chosen model is not optimal, or that artifacts are present (e.g. chemical shift artifacts (Fig. 2.1j), vibration artifacts (Fig. 2.1k, which is the same slice as Fig. 2.1c and g), or signal dropouts (Fig. 2.1l, which is the same slice as Fig. 2.1b)). Plotting the average residual per DW image across voxels can highlight problems in data along specific DW gradients. Finally, identifying and quantifying *outliers*, which are measurements that are ‘too’ distant from

expected observations, can identify data quality issues (Chang et al., 2005; Chang et al., 2012; Heemskerk et al., 2013; Pannek et al., 2012; Tax et al., 2015c, see also Section 2.4.4).

In addition to manual inspection strategies, some automated pipelines for evaluating data quality have been developed. Such pipelines can be used to facilitate quality checking in large cohort studies. Liu et al. (2010) and Oguz et al. (2014) developed an open source software tool that aims to detect and reject volumes affected by inter-slice instabilities, residual motion artifacts, and vibration artifacts according to the method in Farzinfar et al. (2013). Roalf et al. (2016) compute quality metrics such as mean and maximum amount of intensity-based outliers, mean relative subject motion, and the signal-to-noise ratio (SNR) over time to classify data as being of poor, good, or excellent quality. A poor dataset is fully excluded from subsequent analysis. Although automated pipelines aim to improve the data fidelity, they are not always sensitive to artifacts (Berl et al., 2015) and should therefore be extensively validated against visual quality assessment. Despite being labor intensive, keeping a strong (visual) connection to the data is still of paramount importance to warrant the reliability of subsequent analyses.

2.2.2 Signal drift

The combination of the EPI readout and the rapid switching of diffusion gradients puts a high work load on the MRI system, resulting in heating and temporal instability of the scanner. Consequently, a decrease in global signal intensity can often be observed over time (Vos et al., 2016b); see Fig. 2.2. This *signal drift* varies considerably in magnitude and temporal pattern when comparing scanners of different vendors. In addition, the exact underlying mechanism of signal drift is not always entirely clear. For example, heating may cause signal drift through drift in the main magnetic field (B_0), or through an altered transmission energy and flip angle over time. Although the effect of signal drift has gone largely unnoticed until recently, it can cause significant non-trivial effects on diffusion metrics and fiber tractography. Especially when the diffusion images are acquired in an ordered fashion, e.g., from high to low b -value, the systematic artificial signal decrease at the end of the acquisition can be interpreted as overestimated magnitude of diffusion. Alternatively, when all gradient directions at the end stage of the acquisition are primarily in one particular orientation, a systematic bias can occur along that orientation. Moreover, signal drift can also cause biases in acquisitions in which the diffusion directions and weighting are randomized.

During acquisition, it is possible to specifically correct for B_0 drift (e.g. Benner et al., 2006), but B_0 drift is potentially only a partial cause of the total signal drift. However, dynamically updating the center frequency might not be an option on several clinical scanners and does not fully correct for signal drift. Vos et al. (2016) therefore proposed a complementary and straightforward method to correct for signal drift during processing. They suggest to acquire the non-DW image volumes interleaved (i.e. every n^{th} image volume is a non-DW image volume). The global signal decrease can then be estimated from these non-DWIs over time within a brain mask (created for each non-DW image to take into account subject motion). From a quadratic fit on these average signal intensities, they obtain a global rescaling factor as a function of image number that can be applied to each individual image for correction of the artefactual signal decrease (Fig. 2.2).

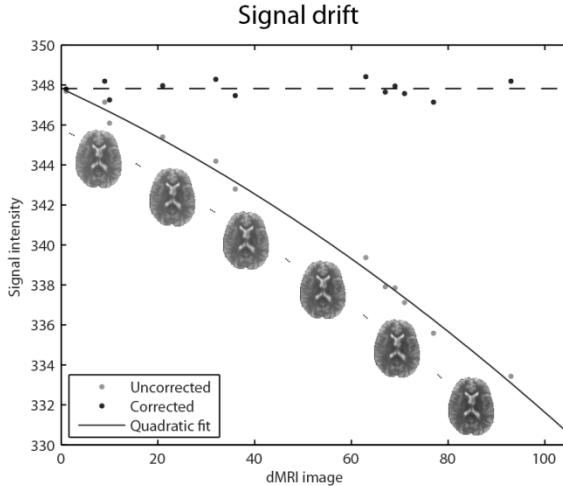


Fig. 2.2: Global signal decrease as a function of dMRI scan number. The red dots represent the average signal of the non-DW images. The blue line is a quadratic fit through these signal intensities, and the blue dots represent the corrected average signal intensities for the non-DW images (the DWIs are corrected in a similar way). In this example, there was an estimated 4.9% signal loss from the first to the last image.

2.2.3 Gibbs ringing

MR images are typically reconstructed by applying an inverse Fourier transform to a set of sampled points in finite k-space (Appendix A2.1 and A2.2). Since k-space holds the spatial frequencies of the MR image, leaving out the outer peripheral regions of k-space thus means that high frequency information is lost. Such a truncation in k-space (or multiplication with a window function) is equivalent to the convolution with a sinc function in image space. The oscillating lobes of the sinc function, in turn, result in signal oscillations or ‘under- and overshoots’ around sharp intensity edges in the image (Fig. 2.3a). Although this *Gibbs ringing* is a well-known MRI artifact, it has gone largely unnoticed in dMRI processing. After the qualitative notion that Gibbs ringing has an effect on DTI derived measures (Barker et al., 2001), its effect on dMRI has only recently been quantitatively assessed (Kellner et al., 2015; Perrone et al., 2015; Veraart et al., 2016). It was found that Gibbs ringing artifacts are more pronounced in the non-DW image because of the high intensity ‘jumps’, which can for example be observed at the interface of CSF and the corpus callosum (CC, Fig. 2.3a). The Gibbs oscillations can lead to physical implausible signals (PIS), i.e. a situation in which signals with a higher b -value are higher than signals with a lower b -value (whereas we expect that the signal decays with b -value). Therefore, a *PIS map* can potentially identify Gibbs ringing artifacts (Fig. 2.3b, Perrone et al., 2015; Tournier et al., 2011). The ringing effect is often amplified in maps of dMRI metrics (e.g. Fig. 2.1i) and can for example result in highly abnormal DTI features (Kellner et al., 2015; Perrone et al., 2015; Veraart et al., 2016).

There are several ways to correct for Gibbs ringing artifacts in MRI, including filtering to smooth out oscillations (Bakir and Reeves, 2000), piecewise reconstruction of smooth regions while preserving edges (Archibald and Gelb, 2002), extrapolating k-space data constrained by total variation (TV) regularization (Sarra, 2006), and performing a local subvoxel-shift to specifically avoid sampling of the ringing extrema (Kellner et al., 2015). In the context of dMRI, TV regularization approaches have shown to be able to alleviate the issue of Gibbs ringing (Perrone et al., 2015; Veraart et al., 2016). Here, *regularization* refers to the introduction of additional information to better be able to solve such a complex problem (see

also Section 2.4). TV methods are based on the assumption that images with spurious detail have a high total variation: the sum of the absolute gradient of the image is high. In practice, TV correction comes down to finding an image that is much similar to the original image but has less total intensity difference in neighboring voxels (i.e. the functional that is minimized includes a term based on the residuals when comparing the original and predicted image, and a TV term multiplied by a regularization parameter). This approach can be extended to capture higher order variations of the image: in this case it is not expected that the image is piecewise constant, and e.g. linear variations in intensity are ‘allowed’. This can avoid the ‘staircase effect’ that might occur in corrected images (Veraart et al., 2016). A drawback of TV methods is that the regularization parameter has to be tuned for each application individually: the optimal value is dependent on the noise level of the image and the degree of Gibbs ringing. Generally, it was found that its value should be lower for larger ringing effects, and is close to the noise standard deviation (Perrone et al., 2015; Veraart et al., 2016). Kellner et al. (2015) reported that their approach, which is based on the resampling of the image such that the sinc function is sampled at its zero crossings, has advantages over TV in that it is more robust to parameter choices and independent of the noise ratio.

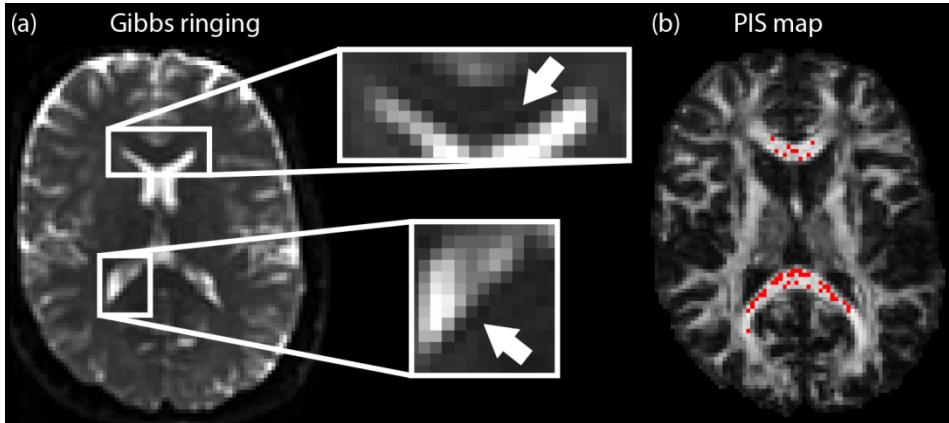


Fig. 2.3: (a) signal oscillations (or ‘under- and overshoots’) around sharp intensity edges in the non-DWI image at the interface of CSF and the CC. (b) PIS map indicating physically implausible signals where a DWI signal is larger than the non-DWI signal, overlaid on an FA image.

2.2.4 Subject motion

Because dMRI tries to infer information on the microscopic motion of particles, all other sorts of motion are important sources of artifacts and can severely impact subsequent analyses (e.g. Elhabian et al., 2014; Yendiki et al., 2013). This is one of the main reasons why single shot EPI (single readout of k-space) is still so popular: whereas multi-shot EPI (split readout of k-space) can reduce off-resonance artifacts, it elongates the scan time and further increases the risk of motion and motion-induced phase differences which can lead to considerable artifacts.

The most straightforward type of subject motion is a slow rotation or translation of the head or *slow bulk motion* (Fig. 2.4a). Lying as still as possible during the scan might be challenging for some age groups or subjects with disorders, but also for healthy subjects for example due to discomfort, sound produced by the pulse sequences, or falling asleep. Consequently, such rigid

body motion artifacts are present in almost all dMRI datasets. We often assume that subject motion does minimally occur during the acquisition of an image, and that it solely results in image misalignment. This misalignment complicates the combination of different diffusion to investigate the tissue structure at every location in the brain.

Prospective subject motion correction can be done during acquisition, but is not always feasible in a clinical setting. During processing, *registration* (i.e. realigning the images) is the method of choice to reduce the adverse effect of subject motion. Generally, each diffusion image is registered to the first non-DWI in the dataset, and misaligned images are interpolated to compute the intensities at the transformed voxel locations. An important choice in the registration process is the transformation model. For subject motion we expect that rigid transformations are sufficient (i.e. translations in x-, y-, and z-direction, and rotations along the x-, y-, and z-axis), leading to 6 parameters to estimate per image. It is essential to keep in mind that every DWI carries orientation information: the image was acquired with the diffusion weighting in a particular orientation. Therefore, when rotations are involved to align an image, the gradient direction (captured in the so-called *B-matrix* (Appendix A2.3)) should be rotated accordingly (Leemans and Jones, 2009).

Another relevant choice in dMRI registration is the similarity measure (i.e. how to define the degree of similarity between intensity patterns in images). This is tricky in dMRI: images acquired with different directions and diffusion weighting are expected to look very different. For example, non-DWIs have a much higher CSF intensity than most DWIs, and each DWI has a unique contrast according to its diffusion orientation. Therefore the *sum-of-squared-differences* (SSD), which is based on the subtraction of two images, might not be the best choice for some dMRI registration approaches. A further challenge is the registration of high *b*-value data: in images with a high *b*-value the signal loss is much higher, the *signal-to-noise ratio* (SNR) is lower, and the brain outline is not well visible (Ben-Amitay et al., 2012; Nilsson et al., 2015). This significantly impacts the accuracy of the registration. Finally, the interpolation step involves choices that may impact subsequent analyses: it is generally advised to align and interpolate all the diffusion as opposed to only the images with the most motion (Elhabian et al., 2014). Different interpolation methods can be adopted such as nearest neighbor, trilinear, and sinc interpolation. Several investigations have tried to address the challenges in the registration of diffusion to correct for subject motion. This correction is commonly integrated with registration-based approaches to correction for eddy currents among other geometric distortion artifacts, thereby reducing the amount of interpolation steps. We will therefore discuss some literature on dMRI registration in Section 2.2.5.1.

Subject motion that occurs during the generation of an image is somewhat harder to correct for in the processing stage. It can lead to a wide range of artifacts like signal dropout/attenuation, distortions, and blurring. For example, subject motion during the application of the diffusion-encoding gradient or *fast bulk motion* can lead to artificial signal attenuation within a slice, which may be wrongly interpreted as increased diffusion. Alternatively, subject motion between the acquisition of different slices can result in *inter-slice intensity instabilities*, for example when the slices are acquired in an interleaved fashion (Fig. 2.1a). Severe subject motion artifacts can be ‘corrected’ by simply rejecting corrupted measurements: this can vary from excluding the whole image based on visual inspection, to removing a single slice, to discarding voxel-wise measurements by statistical robust estimation procedures (see Section

2.4.4). It is important to keep in mind that removing measurements should not result in an imbalanced distribution of the diffusion gradients.

Cardiac pulsation and respiration (also called *physiological noise*) occurs even when the subject lies perfectly still, and can cause non-rigid displacements in the order of 1 mm. Pulsation effects are the most prominent around the lateral ventricles and brainstem (Fig. 2.1d). The combination of diffusion acquired in different stages of the cardiac cycle will result in different local deformations, and therefore misalignment of local structures in successive images. In addition, cardiac pulsation can result in artificial signal loss in certain brain regions (e.g. Nunes et al., 2005). During acquisition, cardiac gating can be used to try to avoid pulsation artifacts. In this case the images are only acquired during particular phases of the cardiac cycle, at the cost of an increased scan time. During processing, corrupted images could be discarded upon visual inspection or when information of the cardiac cycle during the scan is available. In addition, measurements severely affected by pulsation artifacts can potentially be excluded during parameter estimation by using robust parameter estimation approaches (see Section 2.4.4).

Vibration of the subject table due to switching of the diffusion gradients can lead to artificial signal decay in a particular direction, mainly in parietal-occipital regions but also sometimes in frontal regions ((Berl et al., 2015; Gallichan et al., 2010; Hiltunen et al., 2006) Fig. 2.1c). This results in regions of spurious anisotropy which can be recognized as regions of ‘artificial hue’ in color-coded FA maps (red in Fig. 2.1g, indicating artificial decay in left-right direction). Despite extensive efforts by vendors to reduce this artifact, it still affects many studies depending on the protocol, vendor, and population (Berl et al., 2015). Ways to potentially alleviate the artifact during acquisition include adding weight, slightly rotating the head (Berl et al., 2015), or increasing k-space coverage (Gallichan et al., 2010), although the latter approach did not entirely solve the problem in Berl et al. (2015). Correcting for the artifact during processing turns out to be challenging: Gallichan et al. (2010) and Scherrer and Warfield (2012) show that their (robust) estimation approaches (see also Section 2.4.4), which take into account abnormal signal attenuation inconsistent with the tensor model, can deal with the vibration artifact; whereas the robust approach tested in Berl et al. (2015) did not entirely remove the artifact.

2.2.5 Geometric distortions

Some artifacts become more pronounced because of the EPI readout and manifest predominantly in the phase encoding direction. This is caused by the way EPI uses the gradients to spatially encode the signal, i.e. the single long ‘zig-zag-like’ readout (Appendix A2.1): the traversal in k-space is very slow in the phase encoding direction compared to the frequency encoding direction. In other words, the ‘effective gradient strength’ in Hz/mm in the phase encoding direction – obtained by adding all the ‘blips’ and distributing this over the acquisition period – is much lower than in the frequency encoding direction, almost in the order of magnitude of field inhomogeneities (Andersson and Skare, 2010). Hence, when the true applied field is not the field that we think we apply (*off-resonance* field), this has an amplified effect in EPI typically along the phase encoding axis: the signal is reconstructed at the wrong location resulting in a geometrically distorted image.

Two important sources of off-resonance fields are eddy currents (Section 2.2.5.1) and susceptibility differences in tissues (Section 2.2.5.2). Other off-resonance sources, which can usually be reduced during acquisition and therefore not detailed here, are concomitant fields (higher order terms in the gradients for spatial encoding (Baron et al., 2012; Du et al., 2002)), incorrect shim, and local inhomogeneities in the field provided by the magnet (see e.g. Peterson and Bammer, 2016). Residual effects of concomitant fields can often be ‘captured’ by distortion correction during processing. Other common types of dMRI data artifacts that can be linked to or amplified by the EPI readout are Nyquist ghosting (Buonocore and Gao, 1997; Zhang and Wehrli, 2004) and chemical shift artifacts (Sarlls et al., 2011). However, these are not discussed here as these artifacts are typically corrected at the data acquisition stage.

2.2.5.1 Eddy currents

Eddy currents are important sources of off-resonance fields that particularly affect DW acquisitions because of the additional diffusion gradients. To make an image diffusion weighted the gradients have to be switched on rapidly, remain on for a reasonable time to achieve the desired diffusion weighting, and switched off rapidly. Such rapid changes of a magnetic field cause so-called *eddy currents* in conductors of the MRI, which in turn induce additional magnetic gradient fields. This effect is more problematic in dMRI than in most other acquisitions because the gradients for diffusion encoding are on for a longer time than typical gradients for spatial encoding. In the latter case, the eddy currents from switching on and off may compensate each other (Jones and Cercignani, 2010). The off-resonance fields resulting from eddy currents cause geometric distortions in the phase encoding direction when they are constant over time during the EPI readout. However, they cause image blurring when they are time-varying, but this appears to be a minor issue (Andersson and Sotiroopoulos, 2016). Finally, eddy currents can change the diffusion weighting in a non-trivial way, which affects subsequent modelling and is hard to correct for in a clinical setting. Acquisition techniques have been developed that reduce eddy currents (e.g., Reese et al., 2003) or that infer information on the eddy current off-resonance field at the cost of an increased scan time (e.g., Jezzard et al., 1998), see also Andersson and Skare (2010) for an overview). Instead, here we focus on processing techniques for the correction of the geometric distortions.

Registration-based techniques can be used to correct for eddy-current induced distortion artifacts during processing. Similar to registration-based motion correction, an appropriate transformation model is required that can describe the deformations resulting from eddy currents. Importantly, the eddy current-induced distortions vary with the diffusion-encoding direction and b -value, where high b -value DWIs are affected more and non-DWIs remain unaffected since no diffusion gradients are applied (Fig. 2.4b). Proposed approaches vary in the degrees of freedom in their transformation model and whether they correct whole-brain or slice-wise. For example, if the field resulting from eddy currents is spatially linear, the distortions are *affine*: in-plane shearing (from eddy currents along the read direction), in-plane scaling (from eddy currents along the phase encoding direction), and in-plane translation (from eddy currents along the slice selection direction) (Jezzard et al., 1998). Some software packages therefore use a whole-brain affine transformation per image to correct for both subject motion and eddy current distortions. This results in the estimation of 8 to 12 parameters per image, where the ‘simplest’ case only takes into account the distortions mentioned before (e.g. Mohammadi et al., 2010). However, the distortions can vary dependent on the slice position, for example when

residual eddy current fields from a previous slice add up to those of the current slice. This cannot be captured by a whole-brain affine transformation (Jones and Cercignani, 2010). Slice-by-slice correction for the in-plane residual distortions improves the accuracy of the registration (Mohammadi et al., 2010). Moreover, the assumption that eddy current fields are spatially linear is likely to be incorrect. Higher order terms play an important role since it cannot be expected that eddy currents only reside in the gradient coils and that gradient systems are perfectly linear (Andersson and Sotropoulos, 2016; Rohde et al., 2004). Therefore, a quadratic model with 8 or 10 parameters (Andersson and Sotropoulos, 2016; Rohde et al., 2004) or a cubic model with 16 or more parameters (Andersson and Sotropoulos, 2016; Barnett et al., 2014) can be adopted to correct for eddy currents. Simultaneous correction for subject motion (six parameters) and eddy currents assuming a quadratic model (eight parameters) thus requires 14 parameters to be estimated per image. It is important to note that the volume of a voxel might change as a result of the registration, e.g. in the case of scaling or shearing. This has the consequence that the intensity of the voxel should be modulated proportional to the volumetric change (e.g. by multiplying by the Jacobian determinant). Ignoring this step can lead to a wrong interpretation of the rate of diffusion in a given voxel and direction (Jones, 2010c; Jones and Cercignani, 2010).

Several registration-based methods have been proposed to correct for rigid subject motion and/or eddy current induced distortions. To be able to use the simple SSD similarity metric, however, it remains challenging to find a good *reference image* that is not heavily affected by eddy current distortions and has a similar contrast to DWIs. The non-DWI is free of eddy current distortions but is very different in contrast compared to the DWIs. Therefore, some studies proposed to acquire additional images with a more similar contrast, such as low b -value DWIs (e.g. $b = 300 \text{ s/mm}^2$ (Haselgrove and Moore, 1996)) or a CSF suppressed non-DWI (Bastin, 2001). As the acquisition of additional images is not always possible in a clinical setting, we focus on registration-based methods that mostly rely on the acquired set of diffusion images. These methods can be broadly divided into conventional registration methods and prediction- or extrapolation-based registration methods (Nilsson et al., 2015).

Conventional registration methods register the acquired images themselves. An approach that overcomes the problem of registering images with different contrast is the method of Rohde et al. (2004), which uses normalized *mutual information* (MI) as similarity metric (Studholme et al., 1999). Registration approaches based on MI have shown to accurately align diffusion images where all the images are generally registered to the first non-DWI, and are integrated in several diffusion MRI software packages (e.g. Leemans et al., 2009; Pierpaoli et al., 2010). Extrapolation-based methods register a predicted image from a model fit to the acquired set of diffusion images (Andersson and Skare, 2002; Bai and Alexander, 2008; Ben-Amitay et al., 2012; Nilsson et al., 2015). These methods are therefore able to use similarity metrics such as SSD. Andersson and Skare (2002) developed ‘*eddy_correct*’, which is included in FSL (Smith et al., 2004). They first fit a DT to the raw distorted data and use the tensor fit to predict DWIs with the same gradient directions as those acquired. The acquired images are then registered to the corresponding simulated images using the SSD as similarity measure, from which a new tensor fit is obtained. This process is repeated.

Registering high b -value images may require a more targeted approach due to the more significant contrast differences, considerable signal attenuation, and undefined tissue-interfaces

compared with low b -value images. The use of MI does not entirely solve these issues, and predictions of DW images from the DT are inaccurate at high b -values. Mohammadi et al. (2014) register the DWIs to a median DW image per b -value shell, and Zhuang et al. (2013) only registered high b -value DWIs that are close in orientation. In addition, several extrapolation-based methods based on other models than the DT were recently proposed (see also Section 2.3). Ben-Amitay et al. (2012) register the high b -value DW images in a multishell dMRI dataset by first rigidly registering the low b -value images to the non-DW image using MI. Subsequently, they simulate DW images at high b -value using the composite hindered and restricted model of diffusion (CHARMED (Assaf and Basser, 2005)), with the CHARMED parameters predicted from DT-derived indices obtained from the low b -value data. The acquired high b -value DWIs are then registered to the corresponding simulated DW images using MI and a global affine transformation. Nilsson et al. (2015) stated that this DT-based extrapolation may negatively impact the registration in regions of partial volume with CSF, and proposed to correct for this by separating the DT into a tissue and CSF component. In addition, they adjust for artificial diffusion anisotropy in grey matter (Nilsson et al., 2015). They used a stretched-exponential model (Bennett et al., 2003) to extrapolate the signal and found visible improvement of the registration over the CHARMED-based method.

Andersson and Sotiropoulos (2016) estimate the eddy current-induced field and predict estimates of undistorted DWIs based on *Gaussian processes*. Subsequently, they transform this prediction back to the distorted acquisition space to compare it with the acquired DW images. They thereby avoid transformation and interpolation of the original data at this stage. Based on the difference of the original and predicted image, the motion and eddy current distortion parameters are updated, and the whole process is repeated. This approach also has the option to simultaneously correct for susceptibility gradients ((Andersson et al., 2003), see also Section 2.2.5.2), and is used in the Human Connectome Project (HCP) pipeline (Glasser et al., 2013; Sotiropoulos et al., 2013). Whereas most clinical acquisitions sample one half of q -space (Appendix A2.2), the authors suggest that the performance of the method can be improved by acquiring measurements with opposing gradient sign (Bodammer et al., 2004), where the angle between the gradient axes is small. This will ‘negate’ the distortions and thus provide more information to drive the process. Alternatively, acquisition of gradient directions with opposite phase encoding directions holds valuable information on the distortions. This method (‘eddy’) is implemented in FSL and the authors report improved results over ‘eddy_correct’.

2.2.5.2 Susceptibility gradients

The *magnetic susceptibility* of a material refers to the degree of magnetization of an object when placing it in a magnetic field. Inside tissue, the magnetic field is slightly lower than the applied magnetic field because tissue is diamagnetic. This results in off-resonance fields (or macroscopic inhomogeneities) at the scale of the voxel size. These inhomogeneities are dependent on the composition and shape of the imaged body part and have a complex spatial distribution particularly at air/tissue interfaces. The distortions resulting from such a susceptibility-induced off-resonance field (also known as *susceptibility- or EPI distortions*) can be recognized as regions of signal ‘pile-up’ – where the signal of several voxels is compressed into one voxel – or signal ‘smearing’ – where the signal from one voxel is stretched over several voxels – see Fig. 2.4c white arrow. Since the susceptibility distortions only depend on the properties of the subject, all diffusion images are distorted in the same manner if one

neglects subject motion. Susceptibility distortions would therefore not cause problems for local modelling *per se*, but they do affect tractography analyses (Irfanoglu et al., 2012; Jones and Cercignani, 2010), and introduce an extra source of intersubject variability. In addition, they cause a mismatch with anatomical images (Fig. 2.4c), which hampers a more complete picture of the structural organization of the brain. Therefore, even though EPI distortion correction is not standard in many dMRI software packages, it is advised to be included in a typical dMRI processing pipeline. A few things to keep in mind when correcting for susceptibility distortions: 1) they are dominant along the phase encoding direction as discussed before, 2) resampling and interpolation of the images required for susceptibility distortion correction should ideally be combined with eddy current distortion-, motion-, and other corrections; 3) the intensity in a voxel could be modulated when volumetric changes occur during correction (Jezzard and Balaban, 1995), but this is not strictly necessary because the modulation is the same for each DW volume; and 4) this step should not be accompanied by a rotation of the B-matrix because orientation information on a voxel level is not affected by this artifact.

Several approaches exist to reduce susceptibility distortions, generally requiring the acquisition of additional data. One method is to map the field inhomogeneity at every location. This can be achieved by considering the phase of two acquisitions (commonly gradient echo non-EPI, but other choices are also possible) with different echo times: one would expect that the accumulated phase is the same for the two acquisitions if no field inhomogeneities are present. Any phase difference between the acquisitions can thus be attributed to field inhomogeneities, and from the phase difference and the echo time difference we can compute the field inhomogeneity (Jezzard and Balaban, 1995). One way to correct for susceptibility distortions is to convert such a *field map* to a *voxel displacement map* in the phase encoding direction (e.g. Pintjens et al., 2008; Reber et al., 1998), and to create the undistorted image by referring to this displacement map how far the intensity of each voxel was displaced – a process that is also called unwarping. In practice, however, this approach does not entirely solve the problem since in the case of signal pile-up there is not enough information as to how to exactly redistribute the signal back to its correct positions; it is a many-to-one mapping problem (Jones and Cercignani, 2010). It also faces other challenges (Holland et al., 2010), such as the risk of subject motion during the acquisition leading to errors in the field map, and the non-trivial process of unwrapping of the phase maps (i.e. correcting for ‘phase jumps’ by adding multiples of $\pm 2\pi$ to the phase).

Another correction strategy, which can potentially deal with the signal redistribution problem and was shown to perform better than field mapping, involves the acquisition of pairs of images with reversed and/or varying phase encoding directions (Bowtell et al., 1994; Chang and Fitzpatrick, 1992). In the case of reversed phase encoding direction (e.g. bottom-up vs. top-down k-space traversal, see Appendix A2.1) the distortion direction is also reversed: regions of signal pile-up in one image correspond to regions of signal smearing in its corresponding image. The combination of these images provides enough information on the geometrical distortions and the redistribution of signals to create an undistorted image. Such approaches, often referred to as ‘blip-up blip-down’ methods, vary in the way they combine the images with different phase encoding directions. Morgan et al. (2004) and Bowtell et al. (1994) correct the distortions along every image line in phase encoding direction, assuming that corresponding points in two reversed phase encoding images are equidistant to the ‘true’ anatomical location. This 1D correction per line, however, assumes that the subject is in exactly the same position

with an identical phase encoding direction during the reversed scans, thereby ignoring subject motion. This can result in discontinuous displacement fields, and it is sensitive to noise. Andersson et al. (2003) developed the approach TOPUP, which operates in 3D and is integrated in FSL (Smith et al., 2004) and used for the HCP data (Sotropoulos et al., 2013). The authors propose to estimate a field map and displacement field from reversed phase encoding images by modelling the image formation process of spin-echo EPI. Generally, they only use pairs of non-DW images (either one pair at the beginning of the sequence, but ideally multiple pairs throughout the acquisition) to save acquisition time. However, the entire data (DWI and non-DWI) can be used as well if it is acquired in a reversed phase encoding fashion. Importantly, this approach takes into account subject motion between two images in a pair when determining the field. From the displacement field and the reversed images they can then reconstruct least squares estimates of the undistorted images, a step that is usually integrated with the eddy current correction tools in FSL (Andersson and Skare, 2002; Andersson and Sotropoulos, 2016). However, the displacement fields in this approach are modeled with a finite set of basis functions, which limits the spatial variation (i.e. the highest variation that can be described is limited by the highest frequency in the used cosine basis (Holland et al., 2010; Ruthotto et al., 2012)). Holland et al. (2010) proposes an efficient and fast variational approach by smoothing the reversed phase encoding non-DW images so that they look rather similar, register them, and update the total deformation field according to the computed refinement. In the next iterations, the width of the smoothing kernel is decreased until there is no smoothing. This approach allows for more flexible (non-parametric) transformations. Ruthotto et al. (2012) extend this approach by introducing an additional nonlinear regularization term that guarantees diffeomorphic transformations (i.e. the mapping can be differentiated and its inverse as well, resulting in meaningful intensity modulations). However, both approaches do not incorporate the correction for subject motion and eddy currents. A critical note is that when only non-DWIs are used for correction, care must be taken that the correction is adequate in regions where the non-DWI has uniform contrast (i.e., there might not be enough local contrast information to ‘drive’ the correction process).

In a clinical routine, however, one often does not have access to more advanced acquisitions such as field maps and opposite phase encoding direction acquisitions, or there is not enough time to acquire the additional scans. In this case, registration of the dMRI data to a structural ‘undistorted’ image can be used to correct for susceptibility distortions, as such images are often acquired as part of the protocol. Registration methods have shown to perform similar (or better/worse depending on the region) to field mapping (Tao et al., 2009; Wu et al., 2008a). The transformation model for susceptibility correction typically includes non-rigid deformations only along the phase encoding orientation. Registration of dMRI data to anatomical scans is challenging because images often do not necessarily have the same spatial resolution and coverage (i.e., in T1 images, for instance, typically a larger field-of-view is taken compared to DW images). In addition, the choices of similarity measure (e.g. SSD, MI), target image (e.g. T1, T2) and source image (e.g. FA, non-DWI, mean of the DWIs) may all affect the registration performance. Different methods have been developed for registration of dMRI data to anatomical scans. Kybic et al. (2000) register the non-DWI to a T2 image using the mean-squared-difference as similarity metric and a B-spline modeled deformation field. Tao et al. (2009) developed a variational approach to register the non-DWI to the T2, where the term that forces similarity between the images depends on the derivative of the squared difference

between the images. They find that modelling the displacement field with higher-order deformations (i.e. a dense displacement field) B-splines better accounts for the distortions present in EPI. In contrast to relying on the difference between images, Wu et al. (2008a) use MI as similarity metric combined with B-splines to register the non-DWI to the T2. The MI metric allows for registration of different dMRI quantities with different modalities, such as the FA, mean of the DWIs, and ‘anisotropic power maps’ (Dell’Acqua et al., 2014) to the T1, and has been implemented in various software tools (Leemans et al., 2009; Pierpaoli et al., 2010). Glodeck et al. (2016) use MI to estimate a dense displacement field for the registration of the non-DWI to the T1, and introduce additional anisotropic regularization to make the registration process more robust.

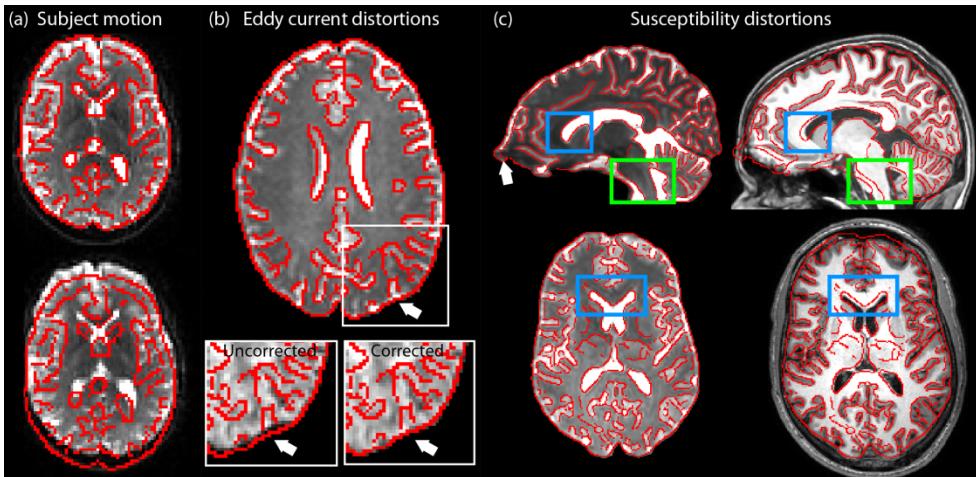


Fig. 2.4: (a) Subject motion: a slow rigid bulk motion of the head (translations and rotations) results in misalignment of images that are acquired at two different time points in the sequence. Top: first non-DWI of the acquisition, with edges indicated in red. Bottom: a non-DWI later in the acquisition, the edges of the top-image are overlaid to indicate the misalignment. (b) Geometric distortions resulting from eddy currents. Top: non-DWI which is not affected by eddy current distortions, edges are outlined in red. Bottom left: edges of non-DWI overlaid on a DWI that is affected by eddy current distortions. Bottom right: Improved alignment of the non-DWI and DWI after eddy-current distortion correction using registration. (c) Geometric distortions resulting from susceptibility gradients. Left: non-DWI with edges indicated in red. The white arrow shows an example of signal ‘smearing’. Right: edges of non-DWI overlaid on an anatomical T1 image. Green and blue squares indicate misalignments.

Finally, hybrid approaches have been proposed that combine previously discussed methods. For example, Gholipour et al. (2011) use a field map to guide the registration of DW images to a T2 image, where they used MI as similarity measure. This resulted in a better match between pairs of reversed phase encoding images and between EPI and T1 images, compared with correction based on either registration or on field maps alone. Irfanoglu et al. (2015) propose an approach called DR-BUDDI in which they combine information from blip-up blip-down non-DW images, DW images, and an anatomical undistorted T2 image to guide the registration. They use a deformation model capable to deal with very large distortions in EPI (the popular symmetric diffeomorphic time-varying velocity-based model in the ANTS software package (Avants et al., 2008)), and use an anisotropic regularization term with a cross-correlation similarity metric. In addition, they allow for incorporation of eddy current and motion correction transformations for the blip-up and blip-down dataset independently, and therefore they account for the possibility that the phase encoding direction in ‘reversed’ images is

inconsistent. The authors show that DR-BUDDI outperforms other methods (e.g. Andersson et al., 2003; Holland et al., 2010) in the presence of motion and large distortions. Even though the reversed diffusion images are averaged into a single dataset and thus provide no new diffusion information, the authors expect that the added benefit in distortion correction outweighs the increased scan time.

2.3 DATA MODELLING AND FEATURE EXTRACTION

After careful data inspection and correction for artifacts, we are now ready to capture relevant features of the measured dMRI signals. To this end, a *dMRI reconstruction technique* can be applied to the data. While many reconstruction techniques are available, DTI is still the most widely used one. It provides information on the primary direction of diffusion which is often assumed to coincide with the underlying fiber direction, and allows for the computation of the FA and MD (Appendix A2.3) which have been studied in a myriad of applications and pathologies. Although DTI is very popular, it suffers from well-known limitations. The famous ‘crossing fiber’ problem ((Jeurissen et al., 2013) the inability to reconstruct multiple fiber directions in complex configurations such as crossing, kissing, bending, fanning, or diverging fibers) has driven the development of approaches that can resolve multiple fiber directions. In addition, the low specificity of measures such as FA, which is affected by many microstructural features at the same time, has heightened the need for approaches that capture more meaningful biophysical information.

2.3.1 Diffusion MRI reconstruction techniques

To place the overwhelming amount of reconstruction techniques beyond DTI somewhat in perspective, we will first discuss several common concepts. Generally speaking, a reconstruction technique estimates parameters from the data and subsequently derives relevant features from these parameters (in some cases, the parameters directly represent the useful features). Common features that we are ultimately interested in include the local directions of underlying fibers, and measures related to the diffusion process and/or tissue microstructure in a voxel or of a distinct fiber population. In the case of DTI, for example, the parameters to be estimated are the DT components (Eq. A2.2), from which we can derive features such as FA and the primary diffusion direction (Fig. A2.5). Other common features, which are higher-dimensional, include the *ensemble average propagator* (EAP, sometimes also referred to as the spin displacement probability density function (PDF) or diffusion propagator), the *diffusion orientation density function* (dODF), and the *fiber ODF* (fODF, sometimes also referred to as the fiber orientation distribution (FOD)). The EAP provides information on the average probability of particles to have a particular displacement during a given diffusion time (Fig. 2.5a, see also Appendix A2.3). It therefore captures the complete information on the diffusion process for a given diffusion time, and depends on the tissue architecture: the displacement of spins can for example be larger in the direction of fibers that cross orthogonally (Fig. 2.5a). The dODF summarizes the information of the EAP by a radial projection on the unit sphere. Note that there are some different definitions in the literature, see for example Assemlal et al. (2011) for an overview. The dODF contains information on the probability of diffusion in a particular direction regardless of the displacement (Fig. 2.5b). The lobes of the dODF are generally broad,

and when fiber populations cross in too shallow angles they will start to overlap which will hamper the detection of crossing fibers. In contrast, the fODF directly represents the distribution of fiber directions and thus represents different information than the dODF (Fig. 2.5c). It can potentially also capture configurations like bending or fanning fibers within the broadness of the fODF lobes. While such higher order features can be visualized throughout the brain, they are often further simplified (e.g. to ‘peak directions’ or scalar measures) for subsequent analysis.

Reconstruction methods are commonly classified based on what their parameters or features represent: reconstruction methods based on *physical models* try to extract microstructurally-specific features, whereas other methods try to find coefficients of a (more abstract) *mathematical representation* of the signal¹. In the latter case, an important aim is often to accurately describe the signal (e.g. by a set of basis functions) rather than a direct interpretation of the features in terms of microstructure. Nevertheless, these representations often allow for the extraction of an EAP, dODF, or fODF, and sometimes also of microstructurally-specific features by introducing certain assumptions.

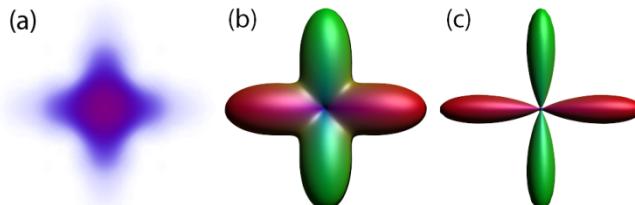


Fig. 2.5: A 90° crossing fiber configuration in a voxel. (a) The ensemble average propagator (EAP) has a value at every point in 3D displacement space. This value represents the probability density of the average relative displacement of spins in a voxel, and integrating this function over a ‘small displacement interval’ gives the probability for a spin to make that displacement during the diffusion time that was set in the acquisition. Here, red indicates a high value (the probability density near the center is high) and purple/white a low density. The crossing can be recognized by the nonzero value of the EAP in the direction of the fibers. (b) The diffusion ODF (dODF) has a value for every direction in displacement space. It is often represented by a glyph object that is deformed in each direction according to the value of the function. The dODF value in each direction reflects the proportion of spins diffusing along that direction, and thus has no information on the magnitude of displacement. The dODF is a simplified version of the EAP: it can be obtained by taking the integral of the EAP along each direction (in this process, the weighting of each point can be varied to obtain sharper dODFs (see e.g. Assemlal et al. (2011) for an overview). (c) The fiber ODF (fODF) contains information on the proportion of fibers in each direction.

Choosing an ‘optimal’ reconstruction method for a particular application is not straightforward. A systematical extensive evaluation and comparison of all reconstruction methods has not been performed thus far (although several works have contributed towards this end, e.g. Cheng et al. (2011); Daducci et al. (2014a); Ning et al. (2015); Rokem et al. (2015)). In practice, the choice of reconstruction method is often determined by multiple factors. An important factor is the aim

¹ Sometimes these are also referred to as model-based and model-free methods. However, since a set of functions to describe the signal can also be regarded as a ‘model’, we instead choose to refer to the two classes of approaches as *physical models* and *mathematical representations*.

Less Confusion in Diffusion MRI

of the study; connectivity analysis for example requires tractography and thus a method that can resolve crossing fibers, whereas a quantitative voxel-wise comparison between groups requires a method that allows for the extraction of measures indicative of the phenomenon under investigation. Another factor is the data that is available: some reconstruction methods are more demanding in that they require multi-shell or Cartesian sampled data (Appendix A2.2), which is often not available in a clinical setting. Finally, the software and support that is available determines which approaches are feasible to use since not every software package supports all reconstruction methods and most software packages require a certain level of expertise. Table 2.1 summarizes numerous current state-of-the-art reconstruction techniques, of which we will discuss several in more detail in the next sections: multi-compartment models, spherical deconvolution approaches, q -space approaches, and diffusion kurtosis imaging.

Reconstruction technique	Typical acquisition	Physical model/ Repr.	Parameters	Features	References
Diffusion tensor imaging (DTI)	Single-shell/ Multi-shell	Physical model/ Repr.*	DT components	FA, MD, AD (λ_{\parallel}), RD (λ_{\perp}), DT eigenvectors, Westin measures (C_l, C_p, C_s)	(Basser et al., 1994; Basser, 1995; Westin et al., 2002)
Diffusion kurtosis imaging (DKI)	Multi-shell	Repr.	DT-, KT components	FA, MD, AD (λ_{\parallel}), RD (λ_{\perp}), DT eigenvectors, Westin measures (C_l, C_p, C_s), MK, AK, RK, KA, dODF	(Jensen et al., 2005; Jensen and Helpern, 2010; Lazar et al., 2008; Lu et al., 2006; Neto Henriques et al., 2015; Poot et al., 2010; Tabesh et al., 2011a)
Spherical deconvolution (SD)	Single-shell/ Multi-shell	Physical model/ Repr.*	basis function coefficients/ fODF**	fODF, AFD, HMOA	(Cheng et al., 2014; Dell'Acqua et al., 2013; Jeurissen et al., 2014; Jian and Vemuri, 2007; Raffelt et al., 2012; Tournier et al., 2007)
Ball and stick model	Single-shell/ Multi-shell	Physical model	***	Diffusivities (ball and stick), vf, axis	(Behrens et al., 2003)
Composite hindered and restricted model of diffusion (CHARMED)	Multi-shell	Physical model	***	Diffusivities (cylinders and tensors), radius or distribution of radii (cylinders), vf, axis	(Assaf et al., 2008; Assaf and Basser, 2005)
White matter tract integrity (WMTI)	Multi-shell	Physical model	DT-, KT components	Diffusivities (stick and tensor), wf	(Fieremans et al., 2011)

Chapter 2: Introduction

Minimal model of white matter diffusion (MMWMD)	Multi-shell	Physical model	***	Diffusivities (tensor, cylinder), radius (cylinder), vf, axis	(Alexander et al., 2010)
Neurite orientation dispersion and density imaging (NODDI)	Single-shell/ Multi-shell	Physical model	Diffusivities (ball and sticks), concentration parameter (Watson distributed sticks), vf, orientation ***	ODI	(Zhang et al., 2012)
Bingham NODDI	Single-shell/ Multi-shell	Physical model	Diffusivities (ball and sticks), concentration parameters (Bingham distributed sticks), vf, orientation ***	Different ODIs, DA	(Tariq et al., 2016)
Ball and racket model	Single-shell	Physical model	***	Diffusivities (ball and sticks), MF (Bingham distributed sticks), vf	(Sotropoulos et al., 2012)
DIAMOND	Multi-shell and Cartesian	Physical model/ Repr*	vf, distribution parameters	dFA, cRD, cHEI, cMD	(Scherrer et al., 2015)
Q-ball imaging (QBI)	Single-shell	Repr.	dODF/ basis function coefficients **	dODF, GFA, GA	(Descoteaux et al., 2007; Hess et al., 2006; Özarslan et al., 2005; Tuch, 2004)
Constant solid angle QBI	Single-shell/ Multi-shell	Repr.	basis function coefficients	dODF, GFA, GA	(Aganj et al., 2010; Özarslan et al., 2005)
Diffusion Orientation Transform (DOT)	Single-shell/ Multi-shell	Repr.	basis function coefficients	iso-radius EAP, dODF, GFA, GA	(Canales-Rodriguez et al., 2010b; Özarslan et al., 2006)
Diffusion spectrum imaging (DSI)	Cartesian	Repr.	EAP	EAP, dODF, RTOP, MSD	(Weden et al., 2005; Wu et al., 2008b)
GQI	Single-shell/ Multi-shell/ Cartesian	Repr.	scaled dODF	scaled dODF, QA	(Yeh et al., 2010)
Radial DSI (RDSI)	Multi-shell in radial lines	Repr.	dODF	dODF, QA	(Baete et al., 2015)
Spherical polar Fourier imaging (SPFI)	Multi-shell	Repr.	basis function coefficients	EAP, dODF, RTOP	(Assemal et al., 2009; Cheng et al., 2010)

Diffusion propagator imaging (DPI)	Multi-shell	Repr.	basis function coefficients	EAP, dODF	(Descoteaux et al., 2011)
Simple harmonic oscillator based reconstruction and estimation (SHORE)	Multi-shell	Repr.	basis function coefficients	EAP, dODF, RTOP, RTAP, RTPP, MSD, NG, QIV	(Cheng et al., 2011; Fick et al., 2016; Özarslan et al., 2009)
Mean apparent propagator (MAP) MRI	Multi-shell	Repr.	basis function coefficients	EAP, dODF, RTOP, RTAP, RTPP, PA, NG, NG_{\parallel} , NG_{\perp} , QIV	(Avram et al., 2016; Fick et al., 2016; Özarslan et al., 2013)

Table 2.1: Overview of numerous current state-of-the-art dMRI reconstruction techniques, their typical acquisition strategy, and some typical features that can be derived. Repr. = representation, DT = diffusion tensor, FA = fractional anisotropy, MD = mean diffusivity, AD = axial diffusivity, RD = radial diffusivity, KT = kurtosis tensor, MK = mean kurtosis, AK = axial kurtosis, RK = radial kurtosis, KA = kurtosis anisotropy, AFD = apparent fiber density, HMOA = hindrance modulated orientational anisotropy, GFA = generalized FA, GA = generalized anisotropy, MF = mean fiber fanning extent, dODF = diffusion ODF, fODF = fiber ODF, EAP = ensemble average propagator, RTOP = return-to-origin probability, RTAP = return-to-axis probability, RTPP = return-to-plane probability, MSD = mean-squared displacement, PA = propagator anisotropy, NG = non-Gaussianity, vf = volume fraction(s), wf = water fraction, axis = primary direction(s) of diffusion, QA = quantitative anisotropy, QIV = q -space inverse variance, c = compartment, HEI = heterogeneity, ODI = orientation dispersion index, DA = dispersion anisotropy. *Depends on the interpretation of the estimated parameters and derived features. ** Depends on the implementation. *** Features and parameters are the same, and some are fixed a priori during estimation. Here, multi-shell refers to at least two non-zero b -values.

2.3.1.1 Multi-compartment models

One class of reconstruction methods models the tissue as a sum of multiple compartments. If distinct assumptions are made about the tissue structure, these compartments can directly represent tissue constituents (Stanisz et al., 1997). The hope is that such models capture the most essential information of the tissue complexity in a reliable way. Generally, each compartment is weighted by a *volume fraction* which reflects how much each compartment contributes to the dMRI signal. The exchange of water molecules between compartments during the diffusion time is hereby often neglected.

A straightforward extension of DTI is the *multi-tensor model*, which assumes that each single fiber population (SFP, i.e. configurations with pathways along one dominant orientation) in a voxel can be modeled by its own tensor (Tuch et al., 2002). In contrast to DTI, this model can thus potentially reveal crossing fiber configurations. Modeling multiple tensors, however, greatly increases the number of parameters to be estimated, and some assumptions therefore have to be made during fitting (see also Section 2.4.2). Another issue with this model is that the diffusion in an SFP is more complex than can be described by a single tensor. More specifically, the hindrance and restriction of diffusing molecules by tissue constituents makes the relation between the signal attenuation and the b -value non-monoexponential. Especially at high b -values, Eq. [A2.4] cannot describe the signal attenuation as a function of b -value sufficiently well.

Many representations and models have been proposed that focus on the characterization of signal attenuation in SFPs (Ferizi et al., 2014; Panagiotaki et al., 2012). A much investigated representation is the biexponential function, in which the signal attenuation is described by the sum of two monoexponential functions and thus two apparent diffusion coefficients (Niendorf et al., 1996), or more generally by two tensors with the same principal direction (Fieremans et

al., 2011; Panagiotaki et al., 2012). The two compartments could then be interpreted as a slow- and fast diffusing pool of water. However, assigning a true biophysical meaning to these pools (i.e. in this case the representation becomes a physical model) has been problematic: the slow- and fast water pools cannot be equated to intra- and extra-axonal components in a straightforward way (Ackerman and Neil, 2010; Le Bihan, 2010; Mulkern et al., 2009). Alternative biophysical models often explicitly model intra-axonal, extra-axonal, isotropic-restricted, and/or CSF compartments. Examples include the ball-and-stick model ((Behrens et al., 2003) intra-axonal: tensor with zero radial diffusivity ('stick'), extra-axonal: isotropic tensor), CHARMED ((Alexander, 2008; Assaf et al., 2008; Assaf and Basser, 2005) intra-axonal: cylinder or distribution of cylinders with nonzero radius, extra-axonal: tensor), and the minimal model of white matter ((Alexander et al., 2010) intra-axonal: cylinder, extra-axonal: tensor, CSF: isotropic tensor, other small compartments: isotropic tensor with zero diffusivity ('dot')).

Recently, physical models have emerged that explicitly try to model fiber incoherence or dispersion, which happens throughout the brain on a sub-voxel level. Neurite orientation dispersion and density imaging (NODDI) models a Watson distribution of intra-axonal sticks to derive an orientation dispersion index (Zhang et al., 2012), but cannot deal with crossing fiber configurations. The ball-and-racket and Bingham-NODDI models adopt a Bingham distribution, which can model more complex fanning configurations along different axes (Sotropoulos et al., 2012; Tariq et al., 2016). The Distribution of Anisotropic Microstructural Environments in Diffusion-Compartment Imaging (DIAMOND) model combines physical- and statistical modelling, and the set of tissue compartments is modeled by a sum of unimodal continuous distributions of tensors. In this way they can calculate the *microstructural heterogeneity* per compartment, which can capture dispersion effects, amongst others. In addition, this model allows one to extract compartment specific DT features such as cFA and cMD (where the c stands for compartment (Scherrer et al., 2015)).

Many of these models (with and without dispersion) are compared in the corpus callosum on different scanners with different acquisitions, in different species, and in- and ex-vivo (Ferizi et al., 2013; Ferizi et al., 2014; Ferizi et al., 2015; Panagiotaki et al., 2012). Some of these models that allow for more complex single-fiber population modelling have been applied successfully in crossing-fiber regions (e.g. CHARMED, DIAMOND).

2.3.1.2 Spherical deconvolution approaches

Multi-compartment models that explicitly model the signal in a voxel as a sum of multiple fiber populations often assume a fixed number of fiber populations. When we instead assume that there is a distribution of fiber orientations (described by the fODF) rather than a discrete number, we in fact increase the number of fiber populations to infinity and the summation becomes an integral. The DW signal can now be expressed as the spherical convolution of the fODF with a kernel or *response function* (RF) representing the DW signal of a single fiber population (Dell'Acqua et al., 2007; Jian and Vemuri, 2007; Tournier et al., 2004; Tournier et al., 2007). The fODF can then be resolved by the *spherical deconvolution* of the DW signal with the response function. This approach does thus not require a predefined number of underlying fibers, and it can often be solved efficiently without the need for computationally demanding nonlinear fitting, e.g. when representing the signal with spherical harmonics (SH)

(see Section 2.4.2.2). A typical b -value for spherical deconvolution approaches based on a single-shell acquisition is $b = 3000 \text{ s/mm}^2$ (Tournier et al., 2013b). Alternatively, spherical deconvolution can be performed on the dODF to obtain an fODF (if the response function is equal to the dODF of a single fiber population) or a sharpened dODF (Canales-Rodriguez et al., 2010a; Descoteaux et al., 2009). Spherical deconvolution on the dODF has shown to perform similar to direct deconvolution on the DW signal.

Recent work has explored whether microstructurally-specific information of distinctly different fiber populations can be derived with spherical deconvolution (Dell'Acqua et al., 2013; Raffelt et al., 2012). It was found that the amplitude of each fODF lobe contains information on the density (*apparent fiber density* AFD (Raffelt et al., 2012), the space that is occupied by fibers) and diffusion properties (radial hindrance and anisotropy) of its corresponding underlying fiber population. Dell'Acqua et al. (2013) coined this same property the hindrance modulated orientational anisotropy (HMOA). Under the assumptions that 1) the radial diffusivity of the intra-axonal compartment is zero and that 2) the signal from the extra-axonal compartment at high b -values is negligible, the fODF amplitude is approximately proportional to the AFD (Raffelt et al., 2012).

2.3.1.3 q -space approaches

q -space approaches exploit the Fourier relationship between the signal and the EAP (Appendix B2) to compute the EAP or features related to the EAP. Some approaches have focused on recovering only the angular dependence of the EAP with single-shell acquisitions. As it is not possible to perform the Fourier transform with only single-shell data, assumptions on the radial dependence of the dMRI signal (i.e. as a function of q - or b -value) are often made. *Q-ball imaging* (QBI) aims to approximate the dODF from the DW signals on a single shell. To this end, it uses the Funk-Radon transform: the estimate of the dODF in a particular direction can be obtained by integrating the signals over a circle in q -space. In other words, we take a direction in which we want to approximate the dODF, identify a plane of which this direction is the normal, and sum all the signals along the ‘equator’, which lies in such a plane. High b -values are generally required to obtain a reasonable angular contrast. Constant solid angle (CSA) QBI allows for the extraction of the mathematically correct dODF (previous QBI approaches computed a different version of the dODF) which is inherently sharper. This technique was also extended to handle multi-shell data. The *diffusion orientation transform* (DOT) allows for the computation of an iso-radius of the EAP (Özarslan et al., 2006). In other words, this feature represents the probability of a molecule displacing over a fixed distance for every direction, and is thus different from the dODF. Nevertheless, a dODF can be computed from the DOT modelling (e.g. Canales-Rodriguez et al., 2010b). This method can also be extended to multi-shell acquisitions (Özarslan et al., 2006). Several scalar indices can be extracted from dODFs. The generalized fractional anisotropy (GFA) has a value of 0 for isotropic diffusion and 1 for anisotropic diffusion similar to FA (Tuch, 2004), and can be directly computed from the SH basis. The generalized anisotropy (GA) was proposed as an alternative generalization of the FA (Özarslan et al., 2005). In addition, a quantitative anisotropy (QA) can be computed for each estimated fiber direction by subtracting the minimum ODF value from the corresponding peak dODF value and multiplying with a global scaling factor (Yeh et al., 2010).

With the development of faster acquisition techniques and improved estimation methods, it becomes feasible to sample q -space more extensively and to reconstruct the full EAP from dMRI data. Several EAP methods that do not rely on a physical model have been proposed. *Diffusion spectrum imaging* (DSI) reconstructs the EAP from a Cartesian sampled q -space (Fig. A2.4c) with very high b -values (up to $b = 17000 \text{ s/mm}^2$) by applying the discrete Fourier transform (Wedge et al., 2005). From the resulting EAP sampled on a Cartesian grid, the dODF can be computed by interpolating the values on radial lines and taking a (weighted) sum. *Radial DSI* (RDSI) computes the ODF more directly from multi-shell and radially sampled q -space using the Fourier slice theorem (Baete et al., 2015). This method can improve the angular resolution and accuracy of the dODF with fewer measurements and lower b -values than DSI.

It was proposed in recent work to fit appropriate (basis) functions to the dMRI signal that capture both its angular and radial properties. This allows for analytical computation of interesting features such as the EAP and dODF, and might facilitate robust reconstruction of the full EAP with fewer measurements (e.g. by using compressed sensing (Merlet and Deriche, 2013)). These bases include the spherical polar Fourier (SPF) basis (Assemnal et al., 2009), the solid harmonic basis used in diffusion propagator imaging (DPI) (Descoteaux et al., 2011), the simple harmonic oscillator (SHO) basis used in SHORE (Özarslan et al., 2009), and the mean apparent propagator (MAP) MRI basis which is an anisotropic extension of the SHO basis (Özarslan et al., 2013). They have in common that they are so-called *dual bases*: once the coefficients have been computed for the fit to the dMRI signals, the same coefficients can be used in combination with the inverse Fourier transform of the basis to directly compute the EAP. In particular the SHORE and MAP-MRI bases have some theoretical advantages over other bases (Cheng et al., 2011; Fick et al., 2014) and can generate robust results within a reasonable scan time (Avram et al., 2016).

Several scalar features related to the EAP have been proposed. The *return-to-origin probability* (RTOP, also called zero displacement probability) is defined as the EAP value at zero displacement (Assaf et al., 2000; Wu et al., 2008b; Wu and Alexander, 2007). If we assume that the diffusion time is sufficiently long and that the tissue consists of isolated compartments (pores) with arbitrary size, shape, and orientation, the RTOP is the reciprocal of the mean compartment volume (Özarslan et al., 2013). Since it can be defined equivalently as the integral of the whole 3D dMRI signal, it can be already computed from the raw signal without the reconstruction of an EAP. However, since the dMRI signal is truncated beyond a certain q -value, RTOP can also be computed numerically from the EAP (e.g. in the case of DSI) or analytically from the basis function estimation (e.g. in the case of SPF imaging, SHORE, and MAP-MRI). Özarslan et al. (2013) and Fick et al. (2016) give these analytical relationships and their interpretation, as well as relationships to calculate other features from the basis function coefficients like *return-to-axis probability* (RTAP), *return-to-plane probability* (RTPP), *mean-squared displacement* (MSD), and *q-space inverse variance* (QIV). If the compartments are highly anisotropic, like in white matter, it makes sense to study the zero displacement probability in certain directions. The RTAP computes this property along the directions in which the diffusion is the most restricted. In case of a single coherent fiber population, for example, the RTAP can be calculated by integrating the dMRI signal on the plane that has the fiber direction as its normal. If the intra-axonal compartment is represented by restricted cylinders, the RTAP is the reciprocal of the mean cross-sectional area of the axons. The RTPP can in this case be computed by taking the integral of the dMRI signal along the fiber direction.

The RTPP is the reciprocal of the length of the cylinder, but as the condition of long diffusion time is hard to fulfill, its value is close to that of unrestricted diffusion. The MSD can be computed by integrating the EAP weighted by the squared displacement, and describes the average displacement within a voxel. It can be related to the MD by the Einstein diffusion equation (Wu et al., 2008b). The QIV can be interpreted as the inverse of the ‘variance’ of q (the signal as function of q is not a probability density function, hence the quotation marks (Hosseini et al., 2013)). It is equivalent to the MSD in the case of Gaussian diffusion but deviates from the MSD in the case of non-Gaussian diffusion: the contrast between different tissue types has shown to be higher in QIV maps than in MSD maps (Hosseini et al., 2013; Wu et al., 2008b). Finally, the *propagator anisotropy* (PA) quantifies the dissimilarity of the EAP with its isotropic part, and the non-Gaussianity (NG) quantifies the dissimilarity of the EAP with its Gaussian part (i.e. the DT). NG can be computed in axial (NG_{\parallel}) and in radial direction (NG_{\perp}) (Özarslan et al., 2013).

2.3.1.4 Diffusion kurtosis imaging

In the previous paragraphs we have discussed some methods that try to capture deviations from Gaussian behavior. One approach that is increasingly being used, and therefore separately discussed here, is *diffusion kurtosis imaging* (DKI (Jensen et al., 2005; Jensen and Helpern, 2010)). *Kurtosis* is a dimensionless descriptor of the shape of a probability distribution, measuring its ‘tailedness’. The kurtosis of a Gaussian distribution is 3, and the *excess kurtosis* is defined as the kurtosis minus 3 to compare a distribution with the Gaussian distribution. The excess kurtosis (hereafter shortened to kurtosis) therefore quantifies the ‘degree of non-Gaussianity’ of a distribution: a positive kurtosis indicates a higher ‘peakedness’ of the distribution compared to a Gaussian, whereas a negative kurtosis indicates a distribution that is less ‘peaked’. For diffusion, the kurtosis would only become negative in the case of fully restricted spherical pores. In dMRI of biological tissue, which consists of multiple compartments with hindered and restricted diffusion, it is therefore generally assumed that the kurtosis is *positive* (De Santis et al., 2012b; Jensen et al., 2005).

The diffusion kurtosis can be derived from dMRI data by using the Fourier relationship between the signal and the EAP. This form possesses a useful mathematical property, the so-called *cumulant expansion* (Kiselev, 2010). This means that we can write the *natural logarithm* of the signal as a power series of the q -value: the n^{th} term in this sum is proportional to the n^{th} power of q times the n^{th} *cumulant*. The cumulants are quantities that can be computed from a probability distribution (‘features’ of the distribution) and that are often easier to interpret than moments of a probability distribution. To give an example: the first moment and cumulant are the same and simply represent *the mean of the distribution*, but the second moment depends on both the mean and the width of the distribution. We therefore often work with the second cumulant instead, better known as the variance. In dMRI, the variance is directly related to the ADC according to the Einstein equation. The fourth cumulant is related to the kurtosis: the kurtosis is defined as the fourth cumulant divided by the square of the second cumulant. In dMRI the underlying probability is unknown but, it is possible to try to estimate the cumulants as useful descriptors of the propagator without reconstructing the full EAP itself. This is in fact what DTI and DKI both pursue: DTI estimates the second cumulant (ADC) and DKI estimates both the second and the fourth cumulant (ADC and *apparent kurtosis*).

Similar to DTI, the directional dependence of the kurtosis is better represented by a tensor. The *kurtosis tensor* (KT) is a symmetric fourth order tensor with 15 unique parameters. To reconstruct the DT and KT with DKI, 21 parameters have to be estimated in total (22 if an estimate of the non-DW signal is also included). Estimating the DT in the DKI framework has its advantages: compared to ‘conventional’ DTI, the DT estimates are more accurate (Veraart et al., 2011a). Several scalar features can be computed from the KT that have shown to provide useful information in addition to DTI (Jensen et al., 2005). The *mean kurtosis* (MK) is the average apparent kurtosis over the sphere. The *axial kurtosis* (AK) is the apparent kurtosis along the principle diffusion direction as derived from the DT, whereas the *radial kurtosis* (RK) is the average apparent kurtosis over the plane spanned by the second and third eigenvector of the DT. The *kurtosis anisotropy* (KA) is defined as the standard deviation of the apparent kurtosis, in analogy with the FA. Since the computation of these features is more challenging than in DTI, multiple implementations exist for several of these features (e.g. Jensen and Helpern, 2010; Lu et al., 2006; Poot et al., 2010; Tabesh et al., 2011a). A drawback of KT features is that they are hard to interpret in terms of tissue microstructure. To provide a meaningful interpretation, these metrics are often considered in the context of physical models (De Santis et al., 2012b; Fieremans et al., 2011; Hui et al., 2015; Jensen and Helpern, 2010).

The acquisition for DKI is more demanding than for DTI: a multi-shell scheme must consist of at least three b -values (of which one can be zero). There is a theoretical upper bound to the highest b -value of around 2500 to 3000 s/mm^2 in the human brain (Jensen and Helpern, 2010; Lazar et al., 2008). To facilitate the acquisition and processing of DKI in a clinical setting, methods for fast acquisition and computation of MK were recently proposed (Hansen et al., 2013; Hansen et al., 2015).

2.4 PARAMETER ESTIMATION, REGULARIZATION, AND DENOISING

Once we have agreed on a dMRI model or representation (in the following the term ‘model’ will refer to either of these), we are left with the important challenge of how to obtain precise and accurate estimates of the parameters from the acquired dMRI images. Here, *precision* refers to estimates being consistent when measurements are repeated many times, and *accuracy* to estimates having no systematic error. Solving such an *inverse problem* is, unfortunately, challenging because 1) a model is an approximation and never exactly fits the measurements, 2) measurements always contain measurement noise; 3) if an exact solution does exist it is not necessarily unique, that is, multiple (physically plausible) solutions can potentially result in the same predicted signals; and 4) the process of finding an inverse solution can be unstable: a small perturbation in the measurement can in some problems generate large changes in the estimates, this is called an ill-conditioned problem. Computations in our application domain are often *ill-posed*. *Regularization*, which refers to the introduction of additional information to solve an ill-posed problem or to prevent overfitting of the noise, can help to obtain more reliable and stable results.

In Section 2.4.1 we discuss the phenomenon of noise and the type of noise that is typically present in dMRI images. In Section 2.4.2 we will review parameter estimation techniques that take noise into account. We give an overview of parameter estimation for several different dMRI reconstruction techniques. Section 2.4.3 discusses denoising techniques that can be

applied prior to dMRI reconstruction, and contextual processing techniques that can be applied after dMRI reconstruction to reduce the effect of noise and to emphasize coherent structures. In Section 2.4.4 we will discuss measurement errors beyond noise that are often present in dMRI data: outliers.

2.4.1 Noise

While we have discussed several sources of data quality deterioration in Section 2.2, we did not discuss one important source that affects all acquisitions: *noise*. Noise introduces a random scatter in the data causing the measured value to deviate from the true value. This can cause parameter estimates to be biased (i.e. there is a systematic deviation) and/or physically implausible. Here we will focus on *thermal noise*, which arises from random fluctuations in the receive coil electronics and the subject to be imaged. The noise in the complex image domain (Fig. A2.2a) is Gaussian for each receiving coil under certain assumptions (Aja-Fernández and Tristán-Vega, 2013; Henkelman, 1985). If we furthermore assume that a single coil fully samples k-space, the complex signal in image space is the sum of the true signal (i.e. the signal not contaminated by noise) and a complex additive white Gaussian noise term with zero mean and variance σ^2 ($N(0, \sigma^2)$). Hence, the measured signal has a mean that is equal to the true signal, and a variance σ^2 .

The magnitude image derived from the data in the complex image domain has a PDF corresponding to a *Rician distribution* with scale parameter σ (Fig. A2.2b, this is the image we usually deal with). This distribution is more complicated than a Gaussian distribution: it has a mean that is not equal to the true signal value but also depends on σ , and a variance that in turn depends on the true signal value. Assuming a Gaussian distribution in the estimation process will thus bias the parameter estimates. Fortunately, in the case of high SNR (i.e. when the true signal is much larger than σ), the distribution of the derived magnitude signal converges to a Gaussian distribution. The mean of this Rician distribution depends on the true signal value and an ‘offset’ related to σ (Gudbjartsson and Patz, 1995) Fig. 2.6a, right), and is sometimes referred to as an *offset-Gaussian function*. In the image background, which consists of air and where the true signal is thus zero, the Rician PDF reduces to that of a *Rayleigh distribution* (Fig. 2.6a, left). For nonzero σ the mean of this distribution is also nonzero, which gives rise to a minimum signal measurable even if the true signal is zero (also called the *rectified noise floor* (Jones and Basser, 2004)). At low SNR this significantly affects DWI measurements and has a deleterious effect in Gaussian estimation procedures, see also Section 2.4.2.

The Rician distribution is valid in the case of a single coil. However, the signal is typically acquired with multiple coils, and the composite magnitude signal is reconstructed by combining the signals from each coil. If the composite signal is calculated by taking the sum-of-squares of the signals, there is no correlation between the coils, and the variance of the noise is the same in each coil, the composite magnitude signal follows a *noncentral chi distribution* (Aja-Fernández and Tristán-Vega, 2013; Constantinides et al., 1997). This distribution is a more general case of the Rician distribution, and reduces to the Rician distribution for a single coil. If the acquired signals are correlated between coils, the number of ‘effective coils’ decreases and the noncentral chi distribution is an approximation of the real noise distribution. While several approaches can take the Rician distribution into account, not many approaches can be generalized to the noncentral chi distribution.

The Rician and noncentral chi distribution assumptions are valid when k-space is fully sampled. Parallel imaging approaches are frequently used to shorten the acquisition time by acquiring only a subsampled version of k-space and combining the information from different coils. Such techniques can change the distribution of the noise dependent on the reconstruction process of the image. Examples of parallel imaging techniques include generalized autocalibrated partially parallel acquisitions (GRAPPA (Griswold et al., 2002)), and sensitivity encoding (SENSE (Pruessmann et al., 1999)). In the case of GRAPPA, the noncentral-chi model does not hold but can be used as an approximation for the magnitude-signal distribution. In the case of SENSE, the magnitude-signal is Rician distributed, just as in single-coil acquisitions (Aja-Fernández and Tristán-Vega, 2013).

In addition to acquisition with multiple coils and parallel imaging approaches, other acquisition and processing factors can influence the noise distribution. Importantly, some pre-processing steps that we discussed in Section 2.2 to correct for artifacts can also alter the noise distribution in a non-trivial way. For example, motion-, eddy current-, and susceptibility distortion correction require the interpolation of the noisy data. The combined effect of various artifact correction techniques on the noise distribution has not been extensively studied thus far.

Having knowledge on the noise variance is useful in parameter estimation, denoising techniques, outlier detection, and tuning regularization parameters, and there are several methods that try to estimate it from the data (Aja-Fernandez et al., 2009). For example, the noise can be estimated from background areas of the image (e.g. Chang et al., 2005; Henkelman, 1985; Koay et al., 2009b; Sijbers et al., 2007), or from the image object (e.g. Chang et al., 2012; Coupé et al., 2010; Sijbers et al., 1998; Veraart et al., 2013a). The correlation between coils and the use of parallel imaging may result in signal distributions that are spatially varying. In this case there is thus not a single SNR or noise standard deviation that can be estimated for the image, but a standard deviation for every location. The spatial dependence of the noise distribution can be visualized in so-called *noise maps*. Several methods have been proposed to estimate spatially varying noise maps from the data (e.g. Aja-Fernandez et al., 2015; Glenn et al., 2015; Landman et al., 2009; Tabelow et al., 2015; Veraart et al., 2013a; Veraart et al., 2015). Alternatively, noise maps can be obtained as part of the dMRI acquisition by switching of the RF power and gradients but keeping other acquisition parameters the same (e.g. Froeling et al., 2016).

2.4.2 Dealing with noise during the estimation process

Ideally, more measurements are performed than parameters to be estimated so that the system of equations is *overdetermined*. During parameter estimation we can then take into account the presence of noise in the measured data. Parameter estimation in dMRI is an active area of research: many approaches exist that vary in their underlying assumptions, incorporated prior information, and means of optimization (i.e. the method used to find a ‘best guess’) (Koay, 2010). For example, some approaches explicitly take into account the Rician noise model and the noise floor associated with it. In addition, estimation approaches are mostly targeted towards a particular reconstruction approach. As such, it is difficult to unify the subject of estimation in dMRI, and we will concisely review common estimation techniques for some dMRI reconstruction methods in this section.

Less Confusion in Diffusion MRI

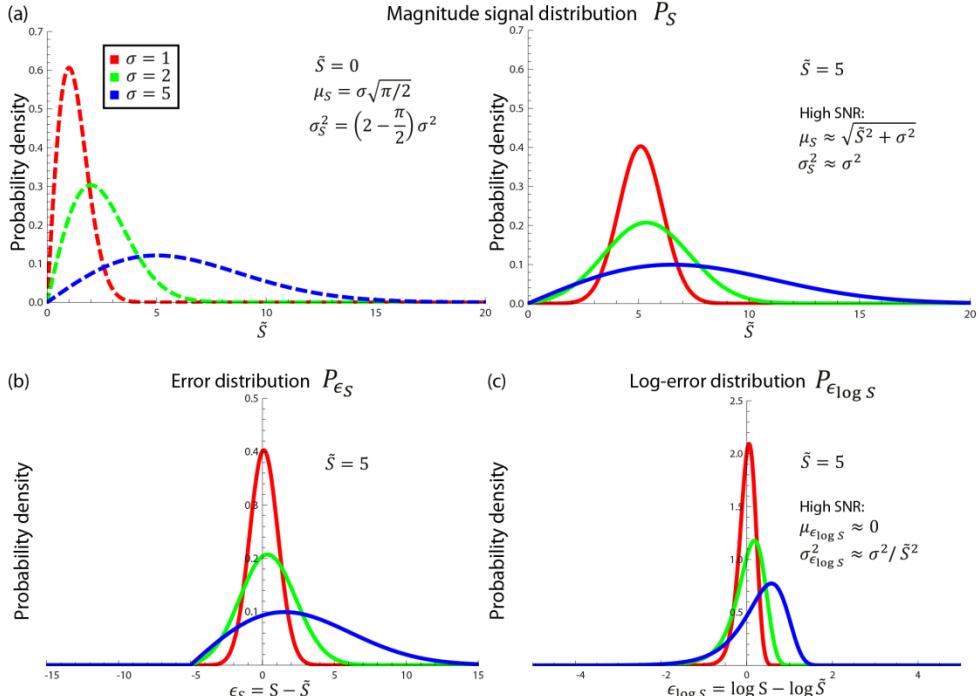


Fig. 2.6: (a) The measured magnitude signal (here denoted by S) follows a Rician distribution with parameters \tilde{S} indicating the true signal and σ indicating the standard deviation of the Normally distributed complex signal. When the true signal is zero, the distribution reduces to a Raleigh distribution (left). In the case of high SNR ($\tilde{S} \gg \sigma$) the distribution starts to approximate a Gaussian distribution. (b) The error distribution for a signal with $\tilde{S} = 5$ for different σ , for higher SNR this starts to look like a Gaussian. (c) The error distribution for a log-transformed signal with $\tilde{S} = 5$ for different σ , for SNR > 2 this starts to approximate a Gaussian.

2.4.2.1 DTI and DKI

Parameter estimation in DTI has been intensively investigated since its proposal, but even for this relatively ‘straightforward’ case estimation remains challenging. For example, noise can cause the eigenvalues to become negative, which is physically implausible (note that the DT should be *positive definite*). The estimation approaches in DTI can be subdivided into 1) least squares, 2) maximum likelihood, and 3) Bayesian methods.

Least squares (LS) approaches are particularly appealing because they have been extensively studied and can be fast (Koay, 2010; Tristan-Vega et al., 2012). They try to find a ‘best guess’ by minimizing an *objective function* based on the *sum of squared residuals*, where a *residual* represents the difference between the measurement and the model prediction. Here, we will discuss nonlinear LS, linear LS, and weighted linear LS. Since the DT has a nonlinear (exponential) relationship with the signal (Eqs. A2.4 and A2.5) the estimation process can be written as a nonlinear regression function, and *nonlinear LS* (NLS) can be employed to estimate the DT. The nonlinear regression problem can be solved with off-the-shelf optimization algorithms such as the Levenberg–Marquardt optimization strategy. However, NLS estimation faces some challenges. First, NLS has to find an optimum in a ‘nonlinear parameter landscape’ by gradually moving around and thus takes a long time. In addition, there exists a risk of getting

stuck in local optima; we therefore might not be entirely sure if the estimated parameters are also the ‘best guess’ globally. Finally, when writing the problem as a nonlinear regression function including an error term, it is assumed that this error term has zero mean and a given variance. It can be shown that in the case of Rician distributed measurements this is generally not true (Fig. 2.6b), and as a result the parameter estimates will be *biased* (i.e. they will have a systematic error). This problem becomes more severe for lower SNR. It can be ameliorated (but not entirely solved), for example, by applying NLS where the model prediction is replaced by its approximate expectation under a Rician PDF (Jones and Basser, 2004). More specifically, the objective function is not directly the sum of squared residuals, but an offset is added to the model prediction before calculating the residuals. This is sometimes referred to as the *offset-Gaussian objective function*. Omitting this bias correction can result, for example, in an underestimation of the ADC.

In addition to NLS, *linear LS* (LLS) can be adopted by linearizing Eq. A2.4 (i.e. taking the log-transform of the measured data), which makes it much easier to implement and faster than NLS. Another advantage is that the error terms in the linear regression equation turn out to be normally distributed for each measurement when the SNR of the DWIs is at least 2 (Fig. 2.6c, (Koay, 2010; Salvador et al., 2005)). However, the variance is not equal for all error terms in this case: as a result of the log-transform it now depends on the magnitude of the true underlying signal. This phenomenon is called *heteroscedasticity* (homoscedasticity thus refers to equal variance across measurements). This means that LLS is not optimal in that it assigns an equal weight to each residual; we should instead assign a higher weight to those measurements with a lower variance. This can be achieved with a *weighted LLS* (WLLS) fit, where the weight of each measurement is equal to the square of the true underlying signal. WLLS estimation theoretically yields the best linear unbiased estimator (BLUE) with the highest precision within its class. However, here lies another problem: we only have the measured signals which are affected by noise and thus not optimal to determine the weight. Instead, an iterative WLLS (IWLLS) can be used: predictions of the true signals are generated from an initial (W)LLS fit, and these are in turn used as weights for the next iteration. In practice, IWLLS shows a high performance in terms of accuracy and precision and may even be preferred over NLS (Veraart et al., 2013b).

Constrained versions of NLS, LLS, and WLLS exist to prevent that the DT eigenvalues become negative. In DTI this can be achieved quite elegantly: instead of the representation of the DT in its elements (Eq. A2.6), we can simply write it differently. The *Cholesky representation*, for example, implicitly constrains the eigenvalues to be positive (in fact, the well-known *eigenvalue decomposition* (Eq. A2.6) is yet another way to represent the DT). Instead of estimating the tensor elements directly, we can then estimate elements of this other representation using any LS approach and compute the tensor elements from them (Koay, 2010).

Maximum likelihood (ML) estimation maximizes the probability of measuring the observed data given a set of parameters and an assumed noise model. Here, we can thus explicitly incorporate the Rician distribution (ML reduces to LS when the noise model is Gaussian). Also in this framework the DT can be constrained to be positive definite by representing it in a different way (e.g. as a Log-Euclidean metric (Fillard et al., 2007)). Alternatively, when representing the DT using eigenvalue decomposition, the likelihood function can be multiplied

by a *prior* on the DT eigenvalues and the non-DWI signal to ensure positivity (Andersson, 2008). In the latter case, we have entered the area of *Bayesian* statistics: *maximum a posteriori* (MAP) estimation maximizes the probability of the parameters given the data according to Bayes rule (which includes the term for ML estimation, a prior term, and a ‘data likelihood’ term that is usually constant). Including a Rician noise model using this framework was found to result in less biased estimates, but in a poorer precision of the estimates compared to a Gaussian noise model for very low SNR (Andersson, 2008). The authors suggest that in group studies the high precision may be more important, and Gaussian noise models could remain to be used.

In addition to introducing priors to limit the possible range of the estimates, *spatial regularization* in the context of DTI can be performed. The purpose of spatial regularization is to suppress the influence of noise by, for instance, using neighborhood information, and to obtain a more ‘smooth’ appearance of the tensor field and derived features. Spatial priors can be introduced during the estimation process itself, which turns this into MAP estimation. For example, Fillard et al. (2007) include an anisotropic prior that preserves edges while smoothing homogeneous regions in an ML framework, and Liu et al. (2013) combine smoothness constraints with an NLS term. Both works enforce positive definiteness of the DT and show to be able to effectively deal with noise.

Parameter estimation in DKI bears a resemblance to estimation in DTI since it can be written as a similar linear or nonlinear regression equation. For example, LS approaches have also been adopted for DKI estimation (e.g. Lu et al., 2006; Veraart et al., 2013b), where IWLLS shows high performance characteristics. In the context of DKI, however, the occurrence of physically implausible parameter estimates as a result of noise is an even bigger problem than in DTI. Specifically, the constraints that the kurtosis along any direction is likely positive in biological tissue and has a theoretical upper bound are often violated. Tabesh et al. (2011a) proposed an LLS approach with linear constraints only along the discrete imaging gradient directions, which can be solved by a computationally heavy algorithm or a faster heuristic algorithm that approximates the optimal solution. Ghosh et al. (2014) evaluated a WLS approach in which the DT and KT were reparameterized (Cholesky and ternary quartic, respectively) to fulfill the positivity constraints along all directions and thus not only the imaging directions. The maximum kurtosis constraint was applied explicitly along the imaging directions by using efficient sequential quadratic programming. Glenn et al. (2015) used WLLS with noise bias correction to account for bias in DKI estimates derived from measurements with low SNR. In addition to LS approaches, constrained ML approaches have been developed to properly account for the Rician noise, thereby avoiding overestimation of the kurtosis values (Ghosh et al., 2014; Veraart et al., 2011b).

2.4.2.2 Spherical deconvolution approaches

SD approaches rely on a response function as input. Choosing an incorrect response function can result for example in spurious lobes and angular deviations (Parker et al., 2013a). Even though being an over-simplification, it is generally assumed that there is a single response function throughout the whole brain. Several methods have been proposed to determine a suitable white matter RF. For example, it can be modeled by a tensor with predefined axial and radial diffusivities (Dell’Acqua et al., 2007; Tournier et al., 2004). However, at the *b*-values

typically used for SD the mono-exponential assumption underlying the DT is not valid anymore. Such a DT-predicted single fiber response does thus not include the restriction effects that typically can be measured at this high b -value. Instead, Tournier et al. (2007) proposed to estimate the response function from the data by selecting the 300 highest FA voxels, or only voxels with an FA exceeding a predefined FA threshold and averaging their signals. This method is still based on a DT fit to the data and it is not always straightforward which FA threshold to use, e.g. in ex-vivo-, spinal cord, or muscle dMRI data. Alternatively, a region-of-interest can be drawn in single fiber regions such as the corpus callosum to estimate the single fiber response function. However, this method introduces a certain level of user-dependence. Tax et al. (2014c) proposed to recursively calibrate the response function from the data by iteratively removing voxels that exhibit more than one fODF peak with SD. This method aims to find single-fiber-population voxels using a more intuitive criterion than FA: a voxel is removed from calibration when the second peak exceeds a predefined *peak ratio threshold* compared to the first peak. The white matter response function is suboptimal in regions of partial volume effect with CSF and grey matter, which may result in an increase in the detection of false peaks and a decreased precision of the detected fiber orientations (Roine et al., 2014). Therefore, the RF can be modified e.g. based on tissue fractions estimated from anatomical data (Jeurissen et al., 2014; Roine et al., 2015).

There are multiple implementations of the SD approach of which the majority can be cast in a unified deconvolution framework (Jian and Vemuri, 2007). In these approaches, the fODF is generally expressed as a linear combination of basis functions, where each basis function is multiplied by a *weight*. This allows the convolution problem to be written as a linear problem: the signal vector is a matrix multiplication of a coefficient vector with a design matrix which can analytically or numerically be computed from the response function and the basis functions. Estimation of the fODF then comes down to estimating the basis function coefficients. Regarding the choice of basis functions, SD methods can be classified into two groups (Cheng et al., 2014): 1) SD based on a continuous representation such as the SH basis (Tournier et al., 2004; Tournier et al., 2007), and 2) SD based on a discrete representation – a mixture of rotated versions of the response in discrete directions (Dell’Acqua et al., 2007; Dell’Acqua et al., 2010; Ramirez-Manzanares et al., 2007). The latter representation holds a close connection to multi-compartment models in which each compartment has a different orientation. In this case, the coefficients are the weights for each compartment. The drawback of this representation is that the angular resolution is limited by the amount of directions on the sphere along which the response function is rotated. Other examples of basis functions include higher order tensors (Feng et al., 2015; Welselassie et al., 2012), and Wisharts (Jian et al., 2007).

Solving the linear SD problem faces some challenges. For example, if a weight has to be estimated for every compartment, then the number of parameters easily exceeds the number of measurements (the system is under-determined). In this case there exists no unique LS solution and the problem is ill-posed. In addition, the problem is often ill-conditioned: noise has a tremendous influence on the stability of SD approaches and can cause spurious lobes and physically implausible negative weights. Regularization and constraints are therefore commonly employed based on the assumptions that 1) only a few weights are non-zero – a property called *sparsity* –, and 2) the fODF is a non-negative function.

A popular SD approach is constrained spherical deconvolution (CSD), in which the fODF is expressed in SH basis functions (Tournier et al., 2007). This method iteratively obtains an improved fODF estimate employing Tikhonov regularization, which has an analytical solution and is in fact a *regularized LLS* estimation. More specifically, the objective function consists of the sum of squared residuals and a sum of the squared negative and small weights below a certain threshold. The l_2 -norm regularization employed in CSD essentially drives negative and small values to zero while promoting large values, but it does neither guarantee positivity nor sparsity. This approach can be further extended with a correction for the Rician bias and spatial regularization (e.g. Tournier et al., 2013a).

Another commonly used SD approach is Richardson-Lucy (RL) SD in which the solution can be constrained to be positive (Dell'Acqua et al., 2007). This method uses a set of diffusion tensors with predefined diffusivities oriented along many directions on the sphere as basis. The weight of each tensor is estimated in an iterative procedure. A damped version of this algorithm (dRL) can reduce the influence of isotropic partial volume effects from CSF or grey matter without adapting the white matter response function (Dell'Acqua et al., 2010). dRL is less sensitive to the choice of response function and causes less spurious peaks than CSD, but has a lower ability to resolve crossing fibers that have a low anisotropy (Parker et al., 2013a). The ability of solving fiber crossings was shown to be improved by including spatial regularization and Rician/noncentral-chi noise models (Canales-Rodriguez et al., 2015; Liu et al., 2015).

Reconstructing a sparse fODF from a limited amount of measurements can also be obtained by using concepts from the theory of *compressed sensing* (Donoho, 2006) and convex optimization. Minimizing the l_0 -norm – equaling the number of nonzero entries – theoretically results in the sparsest solution. In practice, however, l_0 -minimization problems are difficult to solve and it is often proposed to minimize the l_1 -norm instead – which is the sum of the absolute coefficients – or a weighted version of the l_1 -norm, or a combination of the l_1 -norm and l_2 -norm (Daducci et al., 2014b; Feng et al., 2015; Jian and Vemuri, 2007; Landman et al., 2012; Ramirez-Manzanares et al., 2007). These sparse regularizations can be accompanied by spatial regularization to promote spatial coherence of, for instance, fiber directions or weights (Auria et al., 2015; Ramirez-Manzanares et al., 2007; Ye et al., 2016).

Non-negativity constraints are not always enforced in the previously discussed works. To this end, (Jian and Vemuri, 2007) propose to solve a non-negative LS problem using quadratic programming. Alternatively, Cheng et al. (2014) guarantee non-negativity over the whole sphere instead of discretized points on the sphere by representing the fODF differently: the square root of the fODF is written as a linear combination of SH basis functions.

2.4.2.3 Multi-compartment models

Non-linear estimation techniques are commonly adopted for parameter estimation of multi-compartment models. Accurate and precise estimation with such routines is highly challenged by the large amount of parameters to be estimated. Therefore, simplifying assumptions are commonly made to reduce the amount of parameters and increase the precision of the estimates, e.g. equal diffusivities in different compartments (e.g. Zhang et al., 2012). Ferizi et al., 2014; Ferizi et al. (2015) and Panagiotaki et al. (2012) use a Levenberg-Marquardt algorithm for NLS fitting in single-fiber population voxels with an offset-Gaussian objective function (Jones and Basser, 2004). Alexander (2008), Alexander et al. (2010), Tariq et al. (2016), and Zhang et al.

(2012) use ML estimation with a Gauss-Newton optimization technique, sometimes followed by a Markov chain Monte Carlo (MCMC) procedure (Alexander et al., 2010). These estimation methods are computationally expensive and time-consuming, particularly in large cohort studies.

To address the computational issue in multi-compartment model fitting, Daducci et al. (2015a) propose to decouple the estimation of the number and orientation of distinct fiber populations, and the assessment of microstructural properties per fiber population. When an estimate of the underlying fiber population direction(s) is obtained (e.g. from DTI or SD), they can express the multi-compartment model for each population as a linear system. This is based on a similar principle often used in SD optimization as discussed in the previous paragraph: the signal vector is a matrix multiplication of a coefficient vector with a design matrix or *dictionary*. In this case, the dictionary can be built from the expected signal attenuations at the measured q -space points for a range of parameter settings for each compartment. For example, in the case of NODDI, the dictionary is constructed from the signal attenuations corresponding to a range of intra-cellular volume fractions and orientation dispersions. This framework drastically accelerates the fitting and has shown to produce estimates that are in agreement with estimates from conventional fitting procedures.

Some recent works aim to improve or facilitate the estimation of microstructural parameters by first fitting a more general representation. For example, Fieremans et al. (2011); Hui et al. (2015) derive biophysical parameters by first fitting the DKI model. The purpose is two-fold: it facilitates a clearer interpretation of the KT in terms of microstructure, and the multi-compartment model estimation can benefit from advanced estimation methods that are already available for DKI. However, such approaches can only capture the information that is available in the DT and KT, and is not sensitive to microstructural features that affect higher order properties of the signal. Fieremans et al. (2011) derive analytical relationships between KT features and parameters of a physical model for a single fiber direction within a voxel. Hui et al. (2015) developed a framework that can be used with more general models (e.g. crossing fibers, grey matter), but that requires non-linear optimization of the parameters. Alternatively, Fick et al. (2016) first fit a regularized version of the MAP method (Özarslan et al., 2013) and then extrapolate the signal as a preprocessing step for the fitting of multi-compartment models. They show a reduced variance of the estimates of AxCaliber (Assaf et al., 2008) and NODDI (Zhang et al., 2012) after this preprocessing step.

2.4.2.4 q -space approaches

In Q-ball imaging, integration over a circle in q -space can be done by using interpolation (Tuch, 2004) or analytically by using SH (Descoteaux et al., 2007; Hess et al., 2006). In DSI, the truncation of q -space beyond a certain value can lead to Gibbs ringing artifacts in the EAP when taking the discrete Fourier transform. Therefore, the q -space data is typically first multiplied by a filter (e.g. Hanning filter) to enforce a smooth attenuation of the signal with q -value (Wedeen et al., 2005). The signal is assumed to be zero beyond a certain q -value and is ‘zero-padded’ before taking the Fourier transform. Only the real and non-negative parts of the resulting EAP are maintained. To compute the dODF, the EAP is typically radially integrated from zero up to a certain maximum radius. In practice, however, there is no consensus as to what the interval of integration should be to avoid propagation of ringing and other artifacts

into the dODF. Paquette et al. (2014) propose to integrate from a non-zero radius to a maximum radius dependent on the SNR. Tian et al. (2016) propose to use the unfiltered q -space data and truncate the EAP beyond its ‘extent’, which was approximated by the mean displacement distance (as upper bound the ADC parallel to the corpus callosum was taken).

Methods that are based on basis functions to reconstruct the EAP require a reliable estimation of the coefficients. The fitting of such bases is often sensitive to noise, and several regularization methods have been proposed. These methods generally enforce the reconstructed signal to be smooth (angularly and/or radially). Examples of regularizations include Laplace-Beltrami regularization (solid harmonic basis (Descoteaux et al., 2011)) combined with a radial low-pass filter (SPF basis (Assemal et al., 2009)), and Laplacian regularization (SPF basis (Caruyer and Deriche, 2012), SHORE and MAP-MRI bases (Fick et al., 2016)). Fick et al. (2016) use Laplacian regularization in an LS framework to fit the MAP basis, and show that their method outperforms a previously proposed method for MAP MRI fitting (Özarslan et al., 2013) and Laplacian regularization of the SPF basis (Caruyer and Deriche, 2012) in terms of signal fitting and reconstruction of the EAP and dODF.

2.4.3 Dealing with noise prior to or after the estimation process

The effect of noise can also be reduced prior to parameter estimation by noise bias correction and *denoising*. The goal is to estimate or reconstruct the ‘true signal intensity image’ given the noisy image. Many approaches have been developed for denoising of MRI images in general (e.g., Aja-Fernandez et al. 2008; Manjon et al., 2010), or specifically for denoising in dMRI (e.g. Bao et al., 2013; Becker et al., 2014; Brion et al., 2013; Gramfort et al., 2014; Lam et al., 2014; Manjon et al., 2013; St-Jean et al., 2016; Tristan-Vega and Aja-Fernandez, 2010). Some approaches estimate the true signal intensity and the noise variance simultaneously (e.g. Koay and Basser, 2006; Sijbers and den Dekker, 2004), while other approaches use noise estimates as input. Gudbjartsson and Patz (1995) propose to correct for the Rician noise bias by using the relationship between the mean magnitude signal and the true signal for high SNR (Fig. 2.6a right), and apply this also to low SNR data. Koay and Basser (2006) derive the analytically exact correction, and show that the correction in Gudbjartsson and Patz (1995) is a special case for high SNR. However, the corrected data are not Gaussian distributed after application of these methods. Koay et al. (2009a) aim to solve the noise-induced bias by transforming non-central chi distributed signals to Gaussian-distributed signals. *Nonlocal means filters* (Buades et al., 2005) restore the intensity in a voxel by taking a weighted average of voxels in the image. These voxels do not necessarily have to be close in position, hence the term nonlocal. The weights of the voxels depend on how similar these voxels are to the target voxel. Manjon et al. (2010) extended this approach to deal with spatially varying Rician noise. Such approaches may increase the reliability of dMRI parameter estimates (e.g. Zhou et al., 2015).

Denoising approaches specific to dMRI typically take advantage of the multiple acquired images in a dMRI dataset: they depict the same structure with different diffusion-weighting settings. Becker et al. (2014) proposed an approach that combines information from multiple shells to denoise dMRI data by using weighted means of neighboring points in imaging and q -space. By combining information from the ‘position-orientation’ space, that is, the voxel position and the diffusion sensitizing orientation, the algorithm benefits from the high-SNR and low-orientation contrast in the position space and low-SNR high-orientation contrast in the

orientation space. St-Jean et al. (2016) decompose the position-orientation space into ‘patches’ to capture the local spatial and angular structure of the signal. In this way, they can account for spatially varying noise. To denoise the image they find a local sparse representation by learning a dictionary instead of predefining it.

Alternatively, *contextual processing* can be performed as a post-processing step after estimation to reduce the noise and enhance the structures present in dMRI data (e.g. Coulon et al., 2004; Westin et al., 2006). For example, dODFs can be sharpened by using the orientational context at each location (Descoteaux et al., 2009; Florack, 2008). Other approaches work on the joint space of positions and orientations to enhance coherent ODF lobes and suppress incoherent structures (Barmoutis et al., 2008; Dela Haije et al., 2014; Duits and Franken, 2011; Portegies et al., 2015; Reisert and Skibbe, 2016). These techniques have shown to be promising in clinical applications (Prckovska et al., 2015; Tax et al., 2014b).

2.4.4 Outliers

In many practical situations, the data is not affected by thermal noise alone. Measurement artifacts that are not or insufficiently corrected for in the preprocessing stage can result in signal values that are very unlikely or even physically impossible, also called *outliers*. Examples include outliers resulting from vibration, cardiac pulsation, and system related artifacts such as temporal scanner instabilities, spike noise, and signal dropouts. Alternatively, outliers can result from insufficient correction for chemical shift, breathing and head motion, and eddy current/susceptibility distortions. Note that thermal noise can also result in outliers, and that it is difficult to disentangle these from outliers resulting from measurement artifacts. Some artifacts such as physiological noise are likely to result in signal dropouts (Chang et al., 2012), while other artifacts may lead to a signal increase. Outliers can significantly bias the estimates, and it is therefore important to minimize the effect of outliers before subsequent analysis.

Robust methods aim at reducing the effect of outliers during model estimation. Various robust estimation procedures have been developed specifically for dMRI (Chang et al., 2005; Chang et al., 2012; Collier et al., 2015; Mangin et al., 2002; Pannek et al., 2012; Parker et al., 2013b; Scherrer and Warfield, 2012; Tax et al., 2015c; Tobisch et al., 2016; Zhou et al., 2011; Zwiers, 2010). The Robust Estimation of Tensors by Outlier Rejection (RESTORE) algorithm has been widely used in DTI, and reduces the weight of outliers in an iterative procedure (Chang et al., 2005). An adapted version of this algorithm, called informed RESTORE or iRESTORE, mostly targets signal dropouts and constraints the amount of data points that can be excluded as outliers (Chang et al., 2012). Both RESTORE and iRESTORE are based on a NLS in each iteration, and are thus computationally expensive. Other approaches have used iterative reweighted LLS (IRLLS) fitting to speed up the process, without significant reduction of accuracy and precision of the estimates (Collier et al., 2015; Tax et al., 2015c). The Robust Extraction of Kurtosis INDices with Linear Estimation (REKINDLE) approach is designed for robust estimation in DTI, DKI, and other models that can be linearized, and takes into account the heteroscedasticity of the data after linearization. Other methods use higher order models or representations to reduce the effect of outliers (Pannek et al., 2012; Tobisch et al., 2016). In general, care should be taken that rejecting outliers does not result in an ill-defined problem and that the fit is not biased by rejections of points clustered along certain directions (Chen et al., 2015a). In such cases, results should be interpreted with extreme caution.

2.5 APPENDICES

2.5.1 Appendix A2: Diffusion MRI 101

2.5.1.1 Appendix A2.1 Basics of dMRI acquisition

The most common acquisition technique in dMRI is a single-shot *pulsed gradient spin-echo* (PGSE) blipped *echo-planar imaging* (EPI) sequence. An EPI sequence can generate an entire 2D image in a single excitation (hence the term single shot). Whereas an EPI image can also be acquired with multiple shots, single-shot EPI is still popular because it is fast and less sensitive to subject motion. The PGSE part specifically sensitizes the acquisition to diffusion. A typical (simplified) sequence diagram is shown in Fig. A2.1a, and we will briefly go through the different steps in the following paragraphs.

The 90° RF pulse will bring the net magnetization vector in the transverse plane in a given slice, with all the spins being perfectly in phase. From there, the transversal magnetization component starts precessing around the z-axis with the Larmor frequency. We can write the transverse magnetization as a complex number that varies over time, with the x-component (sinusoidal) the real and the y-component (cosinusoidal) the imaginary part. However, due to the interactions of spins with surrounding spins (related to the T2) and *microscopic inhomogeneities* (related to T2'), each individual spin will precess with a slightly different frequency (*dephasing*), causing the magnitude of the transverse magnetization to decay exponentially according to the T2* relaxation time (combination of T2 and T2').

Then, a diffusion gradient is switched on in a particular unit *gradient direction* \mathbf{g} by combining the spatial encoding gradients (in this example, the diffusion gradient is in the readout direction). The diffusion gradient has a *strength* G and a *duration* δ , and these acquisition settings are often captured with the so-called *q-value*: $q = \gamma\delta G/2\pi$ or a vector $\mathbf{q} = q\mathbf{g}$, where γ is a constant defining the gyromagnetic ratio. The typical maximum gradient strength of a clinical system is around 40 - 80 mT/m . The gradient causes the total magnetic field to vary spatially: for a very short time, spins precess a bit slower or faster dependent on their position. At the end of the gradient pulse, the spins got a position-dependent phase shift (they ‘dephased’, Fig. A2.1b, where we ignore relaxation). After the application of the diffusion gradient, all spins start to precess again with their ‘normal’ frequency.

At time point TE/2 (with TE the echo time) a 180° pulse is applied, ‘flipping’ the phase of the spins. This pulse is part of the spin-echo sequence and has now started the correction for the effects contributing to the T2' relaxation. This pulse will yield an *echo* at time TE, with the signal decay now solely originating from T2 decay if no diffusion occurs. Note that it also inverts the sign of the phase shift caused by the diffusion gradient.

At *diffusion time* Δ after the start of the first diffusion gradient, a second diffusion gradient with the same duration δ and strength G is applied which now tries to undo the phase shift caused by the first gradient. This is the essence of the diffusion encoding: when the molecules remain at the same location the effects of both pulses indeed cancel out, but when these molecules have diffused the effect of the first pulse cannot be undone. The *net dephasing* of a spin is proportional to the displacement in the direction of the diffusion gradient, G , and δ . The displacement, in turn, depends on the diffusion time. This dephasing of spins will result in a

signal drop in addition to the T2 decay at time point TE. The sensitivity to diffusion of the sequence is often quantified with the *b-value* $b = \gamma^2 G^2 \delta^2 (\Delta - \delta/3) = (2\pi)^2 q^2 (\Delta - \delta/3)$. Analogously to q , one can define a *b*-vector $\mathbf{b} = b\mathbf{g}$. We come back to the *b*- and *q*-value in Appendix A2.2.

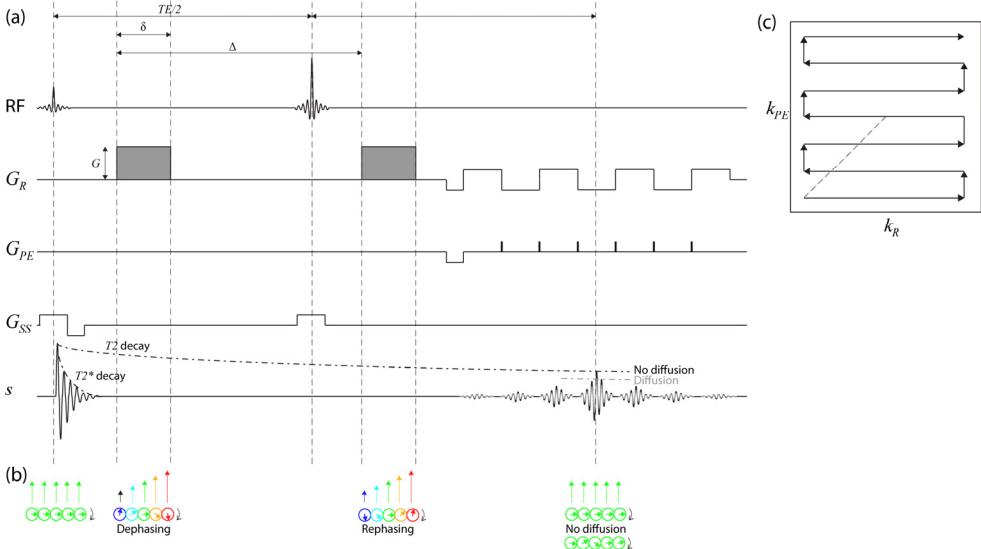


Fig. A2.1: (a) Sequence diagram of a single-shot (SS) pulsed gradient spin-echo (PGSE) echo-planar imaging (EPI) sequence. RF = radio frequency pulses, G denotes gradient magnitude in readout (G_R), phase encoding (G_{PE}), and slice selection (G_{SS}) direction, and s represents the real or imaginary signal. (b) Illustration of the dephasing and rephasing due to the diffusion gradients, and the loss of coherency when diffusion occurs. Note that we ignore relaxation here. The colored arrows denote the local magnetic field magnitude. (c) Zig-zag like sampling of k-space, corresponding to the sequence in (a).

The last pulses are part of the *EPI readout*. In this part, *k-space* is filled by sampling the ‘train’ of echoes. Linear gradients are used to spatially encode the signal: they cause signals from different positions to be emitted at different frequencies. This linear relationship between position and frequency allows for an elegant way to reconstruct the image from the signal through the inverse Fourier transform. Note that field deviations from the intended linear gradient will ‘disturb’ this relationship and result in wrong location information. For example, the diffusion- and spatial encoding is obtained by the same gradients (Fig. A2.1a), and overlap in these gradients can lead to distortions. The combination of readout gradients and phase encoding ‘blips’ results in a zig-zag sampling of k-space (Fig. A2.1c, in this case a ‘blip up’ sequence). A quadrature detector, consisting of two receiver coils, detects both the real and imaginary component of the signal and stores this in 2D k-space.

2.5.1.2 Appendix A2.2 k-space, q -space, and *b*-space

The complex *k-space* holds the spatial frequencies of an image (Fig. A2.2a), and a complex 2D image can be computed by taking the inverse Fourier transform (Fig. A2.2b). From these images, the scanner typically computes the magnitude image by adding the squared imaginary and real image signal and taking the square root, and outputs this to the user. The phase image is mostly not provided automatically.

Less Confusion in Diffusion MRI

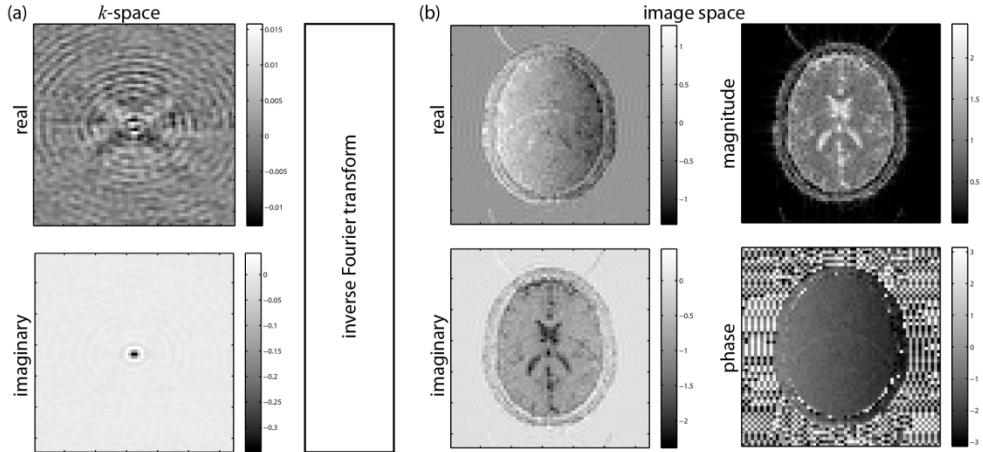


Fig. A2.2: (a) Complex k -space. (b) Complex image space, the scanner generally only outputs the magnitude image. EPI images simulated with JEMRIS (Stocker et al., 2010).

A *dmRI dataset* consists of several magnitude images with different gradient directions \mathbf{g} and/or one or more diffusion weightings (b -values or q -values). We will call the diffusion images with non-zero diffusion weighted images (DWIs), and images with zero diffusion weighting, the *non-DWIs*. Fig. A2.3a shows DW images with the same b -value but different directions: regions of low signal intensity indicate a high rate of diffusion and vice versa. Fig. A2.3b shows diffusion images with different b -values: the signal and therefore SNR decreases with b -value. High b -value data provides additional information on the tissue microstructure (e.g. on restricted diffusion effects by certain tissue constituents) and a better contrast-to-noise to detect crossing fibers. The low SNR, however, often has to be compensated by lowering the resolution or acquiring more data.

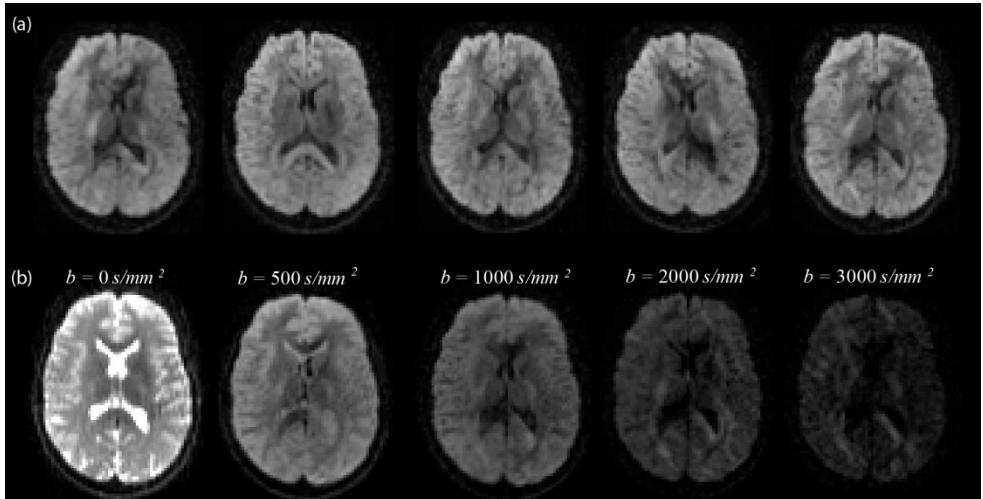


Fig. A2.3: DW images with the same b -value (1000 s/mm^2) but different gradient directions. (b) diffusion images (1 non-DW and 4 DW images) with different b -values.

dMRI sampling schemes can be represented in *q-space* or *b-space*, where each point (having location $q\mathbf{g}$ or $b\mathbf{g}$, respectively) corresponds to a diffusion volume (Fig. A2.4). Typical schemes include single-shell, multi-shell, or Cartesian sampling. Single-shell schemes capture angular information on the diffusion process, whereas multi-shell and Cartesian schemes also capture radial information. Fig. A2.4 shows sampling schemes that cover both the upper and lower half of *q*-space. In practice, however, shell schemes often only sample one half of the hemisphere, since the diffusion information captured in images with \mathbf{g} and $-\mathbf{g}$ is the same for the same *b*-value (although the distortions are different, see Section 2.2.5). Cartesian schemes are often sampled within a given sphere, leaving out the edges in Fig. A2.4c.

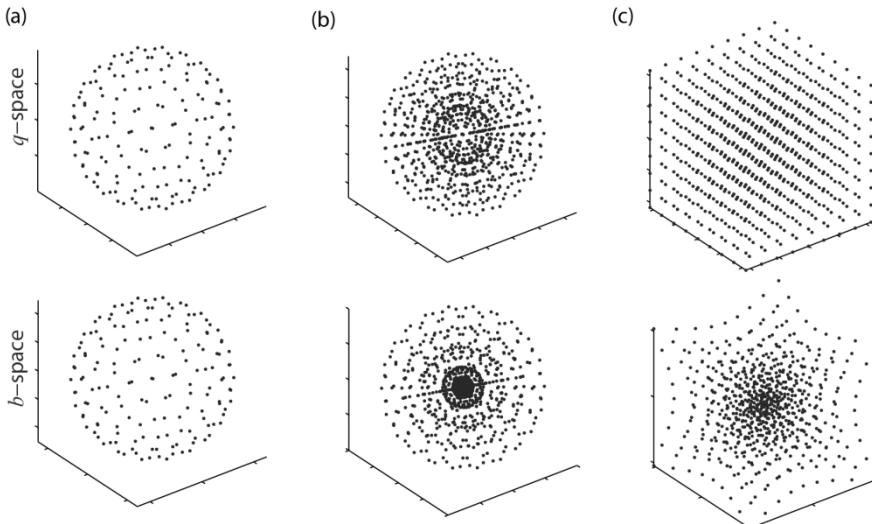


Fig. A2.4: dMRI sampling schemes represented in *q*-space and *b*-space. (a) Single-shell, (b) Multi-shell, (c) Cartesian in *q*-space.

2.5.1.3 Appendix A2.3: From diffusion weighted imaging to diffusion tensor imaging (DTI)

The diffusion images in Fig. A2.3 show a signal decay that is related to diffusion but not yet quantifies the diffusion process as such. Therefore, we need an approach that gives a relationship between the signal decay and the amount of diffusion. Assuming that the diffusion gradient pulses in Fig. A2.1 are infinitely short (such that we can ignore diffusion during the gradient pulses), the signal decay (i.e. the DW signal divided by the non-DW signal) as a function of \mathbf{q} has a Fourier relationship (already the second time that this comes in handy!) with the so-called *ensemble average propagator* (EAP). The EAP describes the density of the average relative displacements of spins in a voxel, i.e. integrating it over a small ‘displacement interval’ gives the probability for a spin to make that displacement during a given diffusion time Δ . Hence if we would measure signal decays for a dense sampling of *q*-space (e.g. such as in Fig. A2.4c), we could reconstruct the whole EAP for a given Δ by taking the Fourier transform.

If we assume that the diffusion process is Gaussian and the same in all directions (i.e. *isotropic*), this relationship becomes very simple and we do not need many measurements to

characterize the diffusion: we can derive one *apparent diffusion coefficient* (ADC) per voxel from one non-DW signal ($S(0)$) and one DW signal ($S(b)$):

$$S(b) = S(0)e^{-b \text{ ADC}} \Leftrightarrow \text{ADC} = \frac{-\ln(\frac{S(b)}{S(0)})}{b}. \quad [\text{A2.1}]$$

The term apparent here refers to the fact that the diffusion coefficient is a measured one obtained with certain acquisition parameters: the ADC is generally lower than the theoretical diffusion coefficient of water because water molecules are hindered and restricted by tissue constituents.

However, the isotropy-assumption is often too simplistic: the diffusion process is known to be very *anisotropic* in white matter. We could already see this in Fig. A2.3a: in some voxels the signal decay depends on the gradient direction. In *diffusion weighted imaging* this issue is usually addressed by measuring the signal in three orthogonal gradient directions and calculating a mean ADC. When we relax the isotropy-constraint but still assume Gaussian diffusion, the process can be described by an apparent *diffusion tensor* (DT) \mathbf{D} :

$$\mathbf{D} = \begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{xy} & D_{yy} & D_{yz} \\ D_{xz} & D_{yz} & D_{zz} \end{pmatrix}. \quad [\text{A2.2}]$$

The DT is symmetric and therefore only has six unique parameters. The diagonal values correspond to the diffusivities along the axes of the coordinate system that we have chosen, and the off-diagonal elements correspond to correlations between diffusivities along those axes. The ADC in a given unit direction $\mathbf{g} = (g_x \ g_y \ g_z)^T$ is given by

$$\mathbf{g}^T \mathbf{D} \mathbf{g} = g_x^2 D_{xx} + g_y^2 D_{yy} + g_z^2 D_{zz} + 2g_x g_y D_{xy} + 2g_x g_z D_{xz} + 2g_y g_z D_{yz}. \quad [\text{A2.3}]$$

By acquiring one non-DW image and at least six DW images, the components of the DT can be calculated (i.e. we have six equations and six unknowns):

$$S(b, \mathbf{g}) = S(0)e^{-b \mathbf{g}^T \mathbf{D} \mathbf{g}} \Leftrightarrow \mathbf{g}^T \mathbf{D} \mathbf{g} = \frac{-\ln(\frac{S(b, \mathbf{g})}{S(0)})}{b}. \quad [\text{A2.4}]$$

In practice, more than six DW images and one non-DW image are acquired which makes the system overdetermined and the DT estimates more reliable. Under the assumption of Gaussian diffusion, DW images with one b -value and different directions should be sufficient to determine \mathbf{D} and a single-shell sampling scheme is therefore often adopted. *Parameter estimation* techniques are then used to find an estimate $\hat{\mathbf{D}}$ of \mathbf{D} that ‘fits’ all the data the best (see also Section 2.4). In other words, when plugging ‘guesses’ $\hat{\mathbf{D}}$ into Eq. [A2.4] we can compute what the corresponding predicted values for the signal $\hat{S}(b, \mathbf{g})$ would be. These values should ideally lie as close as possible to the measured values. For parameter estimation we often compose the DT components and $S(0)$ in a vector $\boldsymbol{\beta} = (\log(S(0)) \ D_{xx} \ D_{xy} \ D_{xz} \ D_{yy} \ D_{yz})$, so that we can write efficiently:

$$\mathbf{z} = e^{B\beta}, \quad \text{with}$$

$$\mathbf{B} = \begin{pmatrix} 1 & b_1 g_{1x}^2 & 2b_1 g_{1x} g_{1y} & 2b_1 g_{1x} g_{1z} & b_1 g_{1y}^2 & 2b_1 g_{1y} g_{1z} & b_1 g_{1z}^2 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & b_n g_{nx}^2 & 2b_n g_{nx} g_{ny} & 2b_n g_{nx} g_{nz} & b_n g_{ny}^2 & 2b_n g_{ny} g_{nz} & b_n g_{nz}^2 \end{pmatrix}, \quad [\text{A2.5}]$$

and $\mathbf{z} = (S(b_1, \mathbf{g}_1), \dots, S(b_n, \mathbf{g}_n))^T$ the n measured signal intensities. \mathbf{B} is also sometimes referred to as the *B-matrix*. Note that in this case we thus also obtain an estimate for the non-DW signal $S(0)$. An alternative formulation is also possible in which β only contains the DT elements, but this only allows the use of one b -value in addition to the non-DW image (i.e. a single shell acquisition (Kingsley, 2006)).

Often, we also mention the simulation of signals (e.g. Section 2.2.1): we then just assume that we know the properties of the diffusion process and then generate the signals that correspond to this process. A very common way to do this is to use Eq. A2.4: we choose a particular \mathbf{D} , set $S(0)$ to a constant value, and generate (noise-free) signals $S(b, \mathbf{g})$ for a given sampling scheme (e.g. a single shell sampling scheme consisting of 32 directions \mathbf{g} on a hemisphere where each measurement has $b = 1000 \text{ s/mm}^2$).

Even though the DT elements capture all the information of the DT, they do not provide much intuition as to how the diffusion process is ‘oriented’. In other words, the x , y , or z axes are not necessarily chosen along the major direction of diffusion. The major axis of diffusion can be derived by simply representing the DT in a different way, i.e. with the *eigen decomposition* of the DT. \mathbf{D} can be rewritten then as follows:

$$\mathbf{D} = \begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{xy} & D_{yy} & D_{yz} \\ D_{xz} & D_{yz} & D_{zz} \end{pmatrix} = \mathbf{E} \Lambda \mathbf{E}^{-1} =$$

$$\begin{pmatrix} e_{1x} & e_{2x} & e_{3x} \\ e_{1y} & e_{2y} & e_{3y} \\ e_{1z} & e_{2z} & e_{3z} \end{pmatrix} \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix} \begin{pmatrix} e_{1x} & e_{1y} & e_{1z} \\ e_{2x} & e_{2y} & e_{2z} \\ e_{3x} & e_{3y} & e_{3z} \end{pmatrix}. \quad [\text{A2.6}]$$

Here, $\lambda_1 \geq \lambda_2 \geq \lambda_3$ are the eigenvalues (where λ_1 is often also called the *axial diffusivity* (AD), and the mean of λ_2 and λ_3 the *radial diffusivity* (RD)), and \mathbf{e}_1 , \mathbf{e}_2 , and \mathbf{e}_3 the eigenvectors (where \mathbf{e}_1 is often called the *primary diffusion direction*), see Fig. A2.5a. From the eigenvalues, common features such as the *fractional anisotropy* (FA), and the *mean diffusivity* (MD) can be computed (Figs. A2.5b and c):

$$FA = \sqrt{\frac{3}{2} \frac{\sqrt{(\lambda_1 - \bar{\lambda})^2 + (\lambda_2 - \bar{\lambda})^2 + (\lambda_3 - \bar{\lambda})^2}}{\sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}}},$$

$$MD = \bar{\lambda} = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3}. \quad [\text{A2.7}]$$

For simulations, we often want to predefined a DT with a given FA and MD along a given axis \mathbf{e}_1 (\mathbf{e}_2 , and \mathbf{e}_3 can then be chosen perpendicular to this). If we assume $\lambda_2 = \lambda_3$, we can

Less Confusion in Diffusion MRI

compute the corresponding eigenvalues (Jones and Basser, 2004; Leemans et al., 2005) and generate \mathbf{D} and S with Eq. [A2.6] and [A2.4], respectively:

$$\begin{aligned}\lambda_1 &= MD \left(1 + \frac{2FA}{\sqrt{3-2FA^2}} \right), \\ \lambda_2 = \lambda_3 &= MD \left(1 - \frac{FA}{\sqrt{3-2FA^2}} \right).\end{aligned}\quad [\text{A2.8}]$$

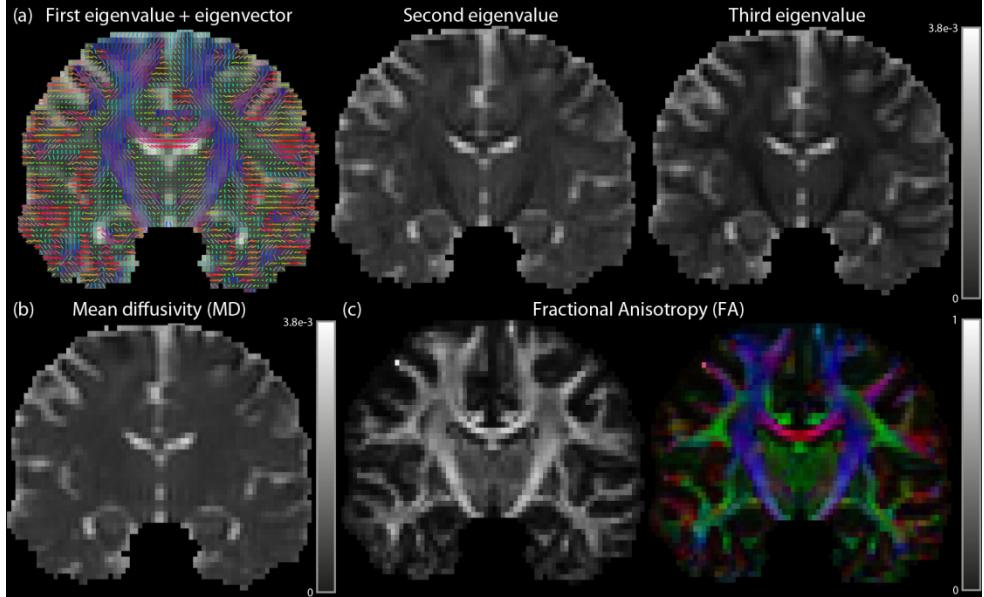


Fig. A2.5: The three eigenvalues λ_1 , λ_2 and λ_3 (units in mm^2s^{-1}) are shown in (a). Eigenvector e_1 is overlaid on λ_1 and colored according to the commonly used color scheme: left-right in red, inferior-superior in blue, and anterior-posterior in green. (b) Mean diffusivity (MD) map. (c) Fractional anisotropy (FA) map (left), and the same map color-coded according to the first eigenvector (right).

3 THE “MASSIVE” BRAIN DATASET: MULTIPLE ACQUISITIONS FOR STANDARDIZATION OF STRUCTURAL IMAGING VALIDATION AND EVALUATION

Based on:

M. Froeling, C.M.W. Tax, S.B. Vos, P.R. Luijten, A. Leemans, “The “MASSIVE” brain dataset: Multiple Acquisitions for Standardization of Structural Imaging Validation and Evaluation”, Magnetic Resonance in Medicine, (in press).

ABSTRACT

In this chapter we present the MASSIVE (Multiple Acquisitions for Standardization of Structural Imaging Validation and Evaluation) brain dataset of a single healthy subject, which is intended to facilitate diffusion MRI (dMRI) modelling and methodology development. MRI data of one healthy subject (female, 25 y) were acquired on a clinical 3 T system (Philips Achieva) with an 8-channel head coil. In total, the subject was scanned on 18 different occasions with a total acquisition time of 22.5 hours. The dMRI data were acquired with an isotropic resolution of 2.5 mm^3 and distributed over five shells with b -values up to 4000 s/mm^2 and two Cartesian grids with b -values up to 9000 s/mm^2 . The final dataset consists of 8000 dMRI volumes, corresponding B0 field maps and noise maps for subsets of the dMRI scans, and ten 3D FLAIR, T1-, and T2-weighted scans. The average signal-to-noise-ratio (SNR) of the non-diffusion-weighted images was roughly 35. This unique set of in vivo MRI data will provide a robust framework to evaluate novel diffusion processing techniques and to reliably compare different approaches for diffusion modelling. The MASSIVE dataset is made publically available (both unprocessed and processed) on www.massive-data.org.

3.1 INTRODUCTION

Diffusion magnetic resonance imaging (dMRI) is used in a wide range of clinical and scientific disciplines for its ability to infer information about tissue architecture and microstructure *in vivo* (Assaf and Pasternak, 2008; Basser and Jones, 2002; Sundgren et al., 2004; Tournier et al., 2011). Investigating brain tissue characteristics with dMRI, however, remains challenging mainly due to the presence of numerous artifacts during data acquisition and the high complexity of the diffusion-weighted (DW) signal for modelling purposes. Consequently, multiple processing steps have to be performed to be able to extract meaningful and reliable features from dMRI data. A variety of correction strategies have been developed that address data imperfections (e.g., eddy current induced distortions, susceptibility based deformations caused by magnetic field inhomogeneities, noise and physiological artifacts (Tax et al., 2016c; Tournier et al., 2011)) to minimize the presence of confounds that could convolute data interpretation. In addition, a wide range of dMRI models and reconstruction methods have been proposed that aim to extract tissue characteristics in a reliable way, including diffusion tensor imaging (DTI) (Basser et al., 1994), diffusion kurtosis imaging (DKI) (Jensen et al., 2005), diffusion spectrum imaging (DSI) (Weden et al., 2005), Q-ball imaging (QBI) (Tuch, 2004), spherical deconvolution (SD) (Dell'Acqua et al., 2007; Dell'Acqua et al., 2010; Tournier et al., 2007), CHARMED (Assaf and Basser, 2005), and NODDI (Zhang et al., 2012).

Notwithstanding recent developments in dMRI artifact correction and modelling, optimizing the dMRI processing pipeline is still an active area of research. For example, there is currently no consensus on the optimal dMRI model or representation to characterize the DW signal, or on the optimal sampling scheme for dMRI reconstruction techniques that use 'multi-shell' acquisitions (i.e., gradient directions distributed over multiple *b*-value shells). Moreover, these two issues may even depend on each other, i.e., optimal sampling might be different for different models. In parallel, many new artifact correction strategies are currently being developed (Perrone et al., 2015; Vos et al., 2016b). In this context, a comprehensive evaluation framework for such novel diffusion processing techniques and for a reliable comparison between different approaches is highly desired.

Unbiased and reliable evaluations are, however, generally hampered by the lack of a genuine gold-standard, and there is an urgent need for a reliable framework that can facilitate the development of dMRI methodology. Hardware phantoms and simulations are very valuable because of their known ground truth, but are often too simplistic (Bach et al., 2014) or can be biased towards a specific model (Close et al., 2009; Leemans et al., 2005; Neher et al., 2014). Real data acquisitions can, despite their unknown ground-truth, serve as valuable references to complement phantom and simulated data in the validation and evaluation of new processing strategies. For this purpose, a comprehensive dMRI dataset is required.

Data repositories and databases are becoming more readily available (Eickhoff et al., 2016), greatly facilitating the development of dMRI methods. Many of these repositories contain cross-sectional and/or longitudinal data allowing for research on normal brain development and function (Evans, 2006; Hodge et al., 2016; Karayannidis et al., 2016; Van Essen et al., 2013; Walker et al., 2016). Other databases aim to give insight into brain anatomy and resolve complex neuronal microarchitecture, either by deriving templates or atlases from data of a single subject or multiple subjects (Hsu et al., 2015; Mori et al., 2005; Varentsova et al., 2014;

Chapter 3: The “MASSIVE” Brain Dataset:
Multiple Acquisitions for Standardization of Structural Imaging Validation and Evaluation

Zhang et al., 2011), or using high resolution post mortem data (Calabrese et al., 2015a; Calabrese et al., 2015c; Dyrby et al., 2011). An example of a cross-sectional repository is the Human Connectome Project (HCP) database, which will contain dMRI data (among others) of 1200 subjects acquired with maximum gradient strengths surpassing that of clinical scanners (Hodge et al., 2016; Van Essen et al., 2013). Although this database contains data acquired with multiple b -values and more diffusion gradient directions than most acquisitions it is still limited in its sampling of q -space (270 directions on three shells). More densely sampled q -space data (512 directions) with higher b -values ($b = 10.000 \text{ s/mm}^2$) for a broader age range are also available (Fan et al., 2016). These datasets are acquired with innovative gradient systems and therefore not comparable to those typically acquired in a clinical setting in terms of resolution and SNR, among others. Another dMRI database is provided as part of the ‘MyConnectome Project’ (Poldrack et al., 2015). This database contains 19 (15 usable) repeated scans of a single subject acquired over the course of 18 months with the purpose of specifically investigating the dynamics of brain function, and the scans are thus identical in terms of acquisition parameters (i.e., b -values and gradient orientations). Despite the availability of many excellent high quality repositories, for validation of processing methods and algorithms (Tax et al., 2014c; Vos et al., 2013) in a clinical setting, typically, synthetic phantoms (Fillard et al., 2011; Neher et al., 2014; Poupon et al., 2008) or small clinical datasets (Pujol et al., 2015) are still used.

In this chapter we present the MASSIVE (Multiple Acquisitions for Standardization of Structural Imaging Validation and Evaluation) brain dataset containing multi-modal MR data and 8000 dMRI volumes of a single healthy subject acquired on a clinical 3 T scanner. All the datasets were specifically acquired in a clinical setting, i.e., using single-shot echo-planar imaging (EPI), ‘conventional’ gradient strengths and hardware, no dedicated head fixation or advanced high density receive coil, to be in line with the current standards in acquisition protocols from routine examinations. As such, subsets of the MASSIVE dataset are comparable to data acquired in clinical studies, and can serve as representative test beds for new developments in a wide range of dMRI data correction strategies, image processing techniques, and microstructural modelling approaches. The MASSIVE dataset consists of 8000 dMRI volumes with b -values up to 9000 s/mm^2 , sampled in configurations of five shells and two Cartesian grids. Data was acquired with echo-planar imaging (EPI) phase-encoding in both anterior-posterior (AP) and posterior-anterior (PA) directions, and with gradient directions both in positive and negative z-direction resulting in 2000 scans for each combination. In addition, the dataset contains B_0 field maps, noise maps, and ten 3D fluid-attenuated inversion recovery (FLAIR), T1-, and T2-weighted datasets, which often play an important role in dMRI processing and analysis methods (e.g., Huang et al., 2008; Irfanoglu et al., 2012; Smith et al., 2013) and can also be used independently for test-retest experiments and methodological evaluations and comparisons (e.g., De Bresser et al., 2011; Duning et al., 2005; Maclarens et al., 2014). The MASSIVE dataset, which was first presented at the 22nd Scientific Annual Meeting of the ISMRM (Froeling et al., 2014), is made publicly available on www.massive-data.org.

3.2 METHODS

3.2.1 Data acquisition

All the MRI data of the healthy subject (female, 25 y) were acquired on a clinical 3 T system (Philips Achieva) with an 8-channel head coil. The subject gave informed consent to participate in this study under a protocol approved by the University Medical Center Utrecht ethics board. In total, the subject was scanned on 18 different occasions (total acquisition time: 22.5 h). A schematic overview of the protocol for a single session is shown in Fig. 3.1. Each of the 18 scan sessions consisted of four dMRI acquisition blocks of 15 minutes in which a unique subset of the 8000 DW volumes was acquired with B_0 -maps being acquired before and after each of these four dMRI acquisition blocks (Andersson et al., 2003; Jezzard and Balaban, 1995; Ruthotto et al., 2012). Additionally, noise maps were obtained at the end of each dMRI acquisition block by switching off the RF pulses and imaging gradients. Finally, ten 3D FLAIR, T1-, and T2-weighted datasets were acquired in five of the eighteen sessions. In these sessions, two FLAIR, T1-, and T2-weighted datasets were acquired with a two hour pause in between. Positioning of the head in the coil and planning the field-of-view in each session was done manually, which resulted in small offsets in rotation and translation. The coronal positioning was intentionally varied between sessions to minimize the systematic effects of ghosting artifacts. The design of the diffusion and anatomical MRI acquisitions will be outlined in more detail in the next paragraphs (further acquisition details can be found in Table 3.1).

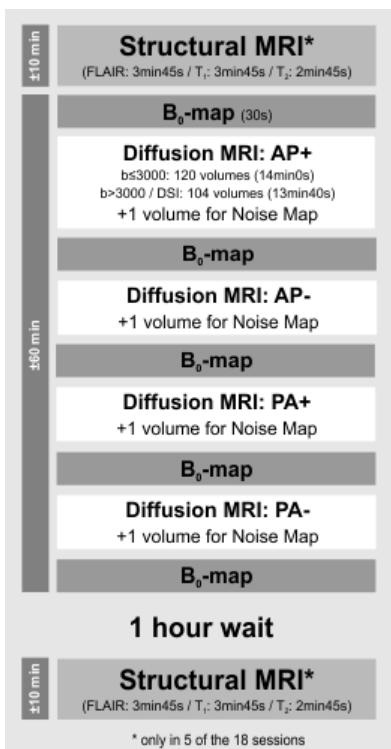


Fig. 3.1: Schematic overview of the MRI acquisition protocol of a single session.

Chapter 3: The “MASSIVE” Brain Dataset:
Multiple Acquisitions for Standardization of Structural Imaging Validation and Evaluation

	FLAIR	T1	T2	DWI	B_0
Sequence	3D-IR-TSE	3D-TFE	3D-TSE	2D-SE-EPI	Dual echo 2D-FFE
FOV [mm^3/mm^2] (AP/RL/IS)	240×180×140	240×180×140	240×180×140	240×240	240×240
Acquisition matrix	240×90×140	240×90×140	240×90×140	96×96	96×96
Reconstruction matrix	240×180×140	240×180×140	240×180×140	96×96	96×96
Slice thickness [mm]	-	-	-	2.5	2.5
Voxel size [mm^3]	1×1×1	1×1×1	1×1×1	2.5×2.5×2.5	2.5×2.5×2.5
Slices	-	-	-	56	56
SENSE: AP/RL:	2/2	2/2	2/2	2.5 (AP)	-
Flip angle [°]	90	8	90	90	20
Partial Fourier	-	-	-	1 / 0.77 *	
TSE/TFE/EPI factor	182	122	124	36 (24 ms) / 31 (25 ms) *	-
Startup echo's	6	4	6	-	-
TE [ms]	Effective: 300 Equivalent: 128	1.25	Effective: 213 Equivalent: 92	100	First: 2 Second: 4
TR [ms]	4800	8000	2500	7000 / 7500 *	322
TI [ms]	1650	-	-	-	-
Fat suppression	SPIR	none	SPIR	SPIR	none
BW frequency [Hz]	1111	191	1111	3035/1991 *	2804
BW EPI [Hz]	-	-	-	50/30 *	-
NSA	2	1	2	1	1
$\Delta/\delta/\zeta$ [ms]	-	-	-	51.6 / 32.8 / 0.9	-
G_{\max} [mT/m]	-	-	-	61.7	-
G_D [mT/m] (Diffusion-weighting gradient strength along one gradient axis)	-	-	-	12.6 (b500) / 17.8 (b1000) / 25.2 (b2000) / 30.9 (b3000) / 35.6 (b4000) / 27.0 (DSI8) / 30.9 (DSI9)	-
Scan time [min:s]	3:45	3:46	2:47	14:08 / 13:10 *	0:32

* Settings for session 1-8 and ($b \leq 3000 s/mm^2$) and for session 9 to 18 ($b > 3000 s/mm^2$ and DSI grids) respectively.

Table 3.1: Overview of the acquisition parameters. (Abbreviations: IR – inversion recovery, TSE – turbo spin echo, TFE – turbo field echo, SE – spin echo, EPI – echo planar imaging, FFE – fast field echo, FOV – field of view, AP – anterior posterior, RL – right left, IS – inferior superior, SENSE – sensitivity encoding, SPIR – spectral pre-saturation by inversion recovery, BW – bandwidth, NSA – number of signal averages, G – gradient strength, DSI – diffusion spectrum imaging)

3.2.1.1 Diffusion MRI

The MASSIVE dataset comprises 8000 unique DW volumes, subdivided into four ‘sets’ with both positive and negative gradient directions, and with both AP and PA phase encoding directions (in the following referred to as AP+, AP-, PA+ and PA-). The acquisition of each set of 2000 DW volumes was divided in 18 sessions (see examples shown in Fig. 3.2) of which eight sessions contained 120 dMRI volumes and 10 sessions contained 104 dMRI volumes. The ordering of the diffusion gradient orientations and *b*-values was randomized throughout every set to prevent an acquisition bias across sessions (Fig. 3.2d-f). The exact scan order can be found in the lookup table which can be downloaded from the website www.massive-data.org.

The gradient directions are distributed over five shells and two Cartesian grids. The diffusion gradient orientations on the shells were generated using the approach described in Caruyer et al. (2013). In short, this approach uses static repulsion of particles to homogenize the gradient orientations on half a sphere (Caruyer et al., 2013; Cook et al., 2007; Jones et al., 1999). In this work, the solution was obtained by an iterative solver where in each iteration the particles repel each other and are subsequently back-projected onto the unit sphere until convergence is reached (see Supplementary Fig. S3.1). The software tool that was developed for computing these gradient orientations is also made available on the MASSIVE website.

The five shells consisted of 125, 250, 250, 250 and 300 gradient orientations on the half sphere with a *b*-value of 500, 1000, 2000, 3000 and 4000 s/mm², respectively. The two Cartesian grids were evenly spaced in half a cube, one with an even ($8^3/2 = 256$) and one with an odd ($(9^3+1)/2 = 365$) number of samples. The maximum *b*-values along the axes were 2296 s/mm² ($q = 0.038 \mu\text{m}^{-1}$) for the even grid and 3000 s/mm² ($q = 0.043 \mu\text{m}^{-1}$) for the uneven grid. The maximum *b*-values for the corner points were 6890 s/mm² ($q = 0.066 \mu\text{m}^{-1}$) for the even grid and 9000 s/mm² ($q = 0.075 \mu\text{m}^{-1}$) for the uneven grid. For the diffusion data with $b > 3000 \text{ s/mm}^2$ and for the Cartesian grids the EPI bandwidth was reduced from 50 to 30 Hz (keeping the TE, δ and Δ constant) to maximize the SNR for the high *b*-values and to be able to acquire the data in a feasible scan time given the system’s duty cycle limitations. Additionally, 204 $b = 0 \text{ s/mm}^2$ images were acquired, resulting in the 2000 dMRI volumes per set with approximately a 1:9 ratio between the non-DW and the DW volumes. These non-DW volumes were randomly interleaved throughout each dMRI acquisition to avoid any measurement bias and to allow for signal drift correction (Vos et al., 2016b; see also section 3.2.2).

3.2.1.2 Anatomical MRI

Ten anatomical MRI datasets (T1- and T2-weighted, and FLAIR) were acquired as they often support dMRI processing and analysis methods (facial features were removed for anonymization). For instance, T1-weighted data can be used for segmentation of grey and white matter regions, which can be used to improve fiber tractography (Kleinnijenhuis et al., 2012; Smith et al., 2013). Similarly, as T2-weighted data provide a comparable contrast with the non-DW data, these can be used to correct for susceptibility induced distortions (Huang et al., 2008; Irfanoglu et al., 2012). The FLAIR data, which has a similar contrast to the T2-weighted data but with suppression of the signals originating from the cerebrospinal fluid, may be useful to investigate the contributions of partial volume effects (Concha et al., 2005; Metzler-Baddeley et al., 2012b). Details of the acquisition protocols for these anatomical MRI data are included in Table 3.1.

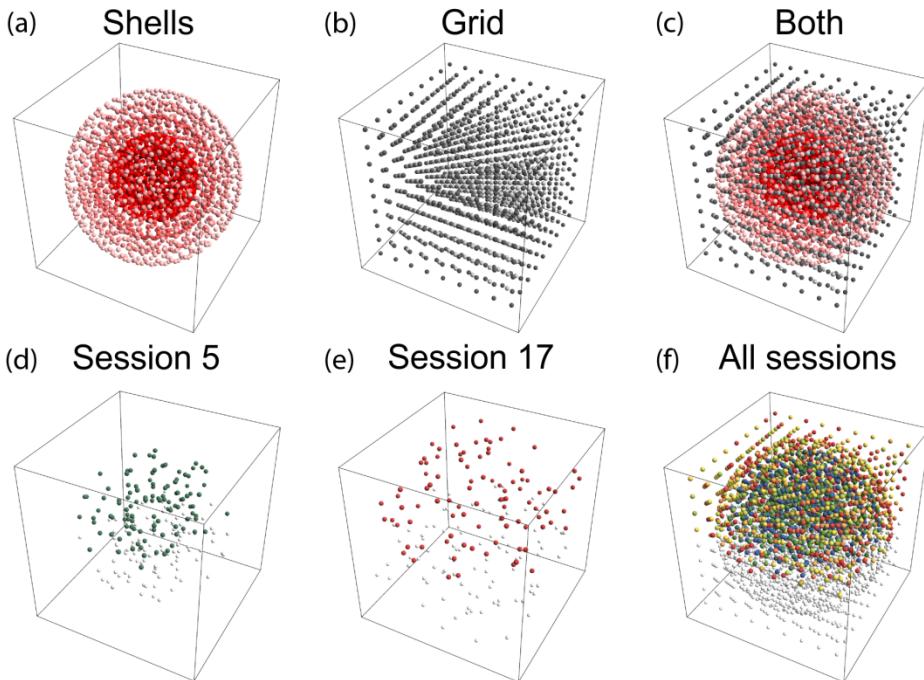


Fig. 3.2: Schematic representation of the encoding schemes in q -space. (a) The 5 shells ($b = 500, 1000, 2000, 3000$ and 4000 s/mm^2). (b) The two Cartesian DSI grids (DSI8: light grey, DSI9 dark grey). (c) The five shells and the two Cartesian grids combined. (d-e) Examples of the random selection of gradients for each session. The gradients defined along the positive direction are shown in color, whereas the mirrored gradient directions are shown in light grey. All the gradients are color-coded for their session. (f) The 18 sessions combined fill up half of the q -space (2000 orientations) and by mirroring the gradients the q -space is fully filled.

3.2.2 Data processing

In addition to the raw data, we provide further information on the acquired data, such as SNR estimates and the ‘true’ applied B-matrix for the raw data as derived from the scanner. In addition, we make available three processed datasets: 1) only the intensity-normalized data; 2) both intensity-normalized and signal-drift corrected data; and 3) data that has been intensity-normalized and corrected for signal-drift, subject motion, eddy current distortions, and EPI deformations. Note that for each of these results conventional processing tools were used as described in the following subsections.

3.2.2.1 Signal-to-noise-ratio (SNR) estimation

Having knowledge of the image SNR is often important in dMRI modelling and processing (Anderson, 2001; Farrell et al., 2007; Jones, 2003; Jones, 2004; Landman et al., 2007a; Lu et al., 2006; Tournier et al., 2011; Veraart et al., 2013a). The SNR can be quantified in numerous ways, however, and may not be the same across different methods (Dietrich et al., 2007; Firbank et al., 1999; Kaufman et al., 1989; Polders et al., 2011; Reeder et al., 2005). Therefore, noise maps were obtained after each acquisition of a 15-minute diffusion block, by switching

off the gradients and RF power but leaving the acquisition channels open. The acquisition parameters of the noise map were equal to the dMRI acquisition parameters and, as such, an accurate noise measurement is obtained. To calculate a rough approximation for the global SNR, the average whole brain signal within a mask was computed for every DW volume and divided by the noise standard deviation within the same brain mask as derived from the noise map.

3.2.2.2 Intensity normalization and signal drift correction

Since the data were acquired in different sessions, intensity normalization between sessions is needed to be able to combine the DW volumes into a single data set (Jeurissen et al., 2013). To this end, the first volume of each 15-minute dMRI acquisition block, which was always a non-DW volume, was used to normalize the signal intensities of all the other volumes within that acquisition block.

In dMRI, the heavy duty cycle of the EPI-readout and the diffusion gradients can lead to temporal instability of the scanner. This instability typically causes a decrease in global signal intensity of the DW images over time, as explained in detail in Vos et al. (2016b). To correct for this so-called signal drift, the non-DW volumes, which were randomly interleaved throughout each dMRI acquisition block, were identified. Subsequently, the signal drift that occurred during this 15 minute time-window was characterized by a quadratic fit of the mean signal of the $b = 0 \text{ s/mm}^2$ volumes as a function of the scanned volume (Vos et al., 2016b), i.e.

$$S(n|b = 0 \text{ s/mm}^2) = d_1 \cdot n^2 + d_2 \cdot n + S_0, \quad [3.1]$$

where S is the normalized measured signal, n the ordering number of the acquired volume, S_0 the signal offset at $n = 0$, and d_1 and d_2 describe the quadratic and linear signal drift per volume, respectively. The corrected signal of the n^{th} volume, $S_{\text{cor}}(n)$, is then given by:

$$S_{\text{cor}}(n) = \alpha \frac{S(n)}{d_1 \cdot n^2 + d_2 \cdot n + S_0}, \quad [3.2]$$

where the factor α is an arbitrary chosen signal scaling factor.

3.2.2.3 True B-matrix calculation

In addition to the DW pulsed field gradients, imaging gradients can also contribute to the diffusion-weighting. Therefore, the actual B-matrix was calculated in addition to the prescribed B-matrix (Mattiello et al., 1997; Mattiello et al., 1994). The amount of diffusion weighting $\mathbf{b}_{i,j}$ along the coordinate axes $i = \{x, y, z\}$ and $j = \{x, y, z\}$ can be expressed as the time (t) integral over the echo time (TE) of the zeroth-order ($n = 0$) moments (\mathbf{M}_n) of the gradients, i.e.,

$$\mathbf{b}_{i,j} = \gamma^2 \int_0^{TE} \mathbf{M}(t)_{n=0,i} \mathbf{M}(t)_{n=0,j} dt, \quad [3.3]$$

with

$$M(t)_n = \int_0^t t'^n G(t) dt', \quad [3.4]$$

where γ is the gyromagnetic ratio and $G(t)$ the gradient wave form.

3.2.2.4 Correcting for subject motion, eddy current distortions, and EPI deformations

dMRI acquisitions suffer from subject motion and eddy current induced distortions within an acquisition session (Jones and Cercignani, 2010; Tournier et al., 2011). In this study, the dataset was scanned in multiple sessions which caused an additional source of misalignment. As a result, the final gradient distribution will slightly differ from the applied one, because the B-matrix needs to be rotated when correcting for subject motion (Leemans and Jones, 2009). For each session, the dMRI data was registered using *ExploreDTI* (Leemans et al., 2009) using an affine method with 12 degrees of freedom to also correct for eddy current induced distortions in the same step (Rohde et al., 2004). The first $b = 0 \text{ s/mm}^2$ image of each acquisition was chosen as a reference image. To correct for EPI distortions and subject motion between the different session, all data was transformed to a common T1-weighted anatomical target dataset using a rigid-registration for rough alignment, followed by a non-rigid b-spline registration to correct for susceptibility induced deformations (Irfanoglu et al., 2011; Irfanoglu et al., 2012; Vos et al., 2016a; Wu et al., 2008a). Here, only non-rigid deformations along the phase-encoding axis of the dMRI data were allowed, as this is the axis along which susceptibility distortions occur. Note that the transformations from these two last steps are combined with the previous eddy current correction procedure to ensure that only one interpolation step is needed minimizing unwanted smoothing effects due to resampling. All other anatomical datasets (FLAIR, T1-, and T2-weighted) were also transferred to the same T1-weighted target dataset using rigid registration.

3.3 RESULTS

Fig. 3.3 shows representative images of the acquired data. The top two rows (Fig. 3.3a-h) show images from the different shells and Cartesian grids. The FLAIR, T1-, and T2-weighted anatomical scans are shown in Fig. 3.3i-k and a B_0 map is shown in Fig. 3.3l.

3.3.1 Signal-to-noise-ratio (SNR)

The noise as derived from the noise map is Rayleigh distributed (see Fig. 3.4a). The SNR estimates are presented in Fig. 3.4b. The average SNR was between 35 and 40 for the non-DW images and was in the same range across all sessions. The SNR of the data in sessions 9 to 18 ($b = 4000 \text{ s/mm}^2$ shell and both DSI grids) was around 15 percent higher than for the shells with $b \leq 3000 \text{ s/mm}^2$, which was to be expected with the lower bandwidth. The estimated SNR values per acquired volume can be found in the lookup table which can be downloaded from the website www.massive-data.org.

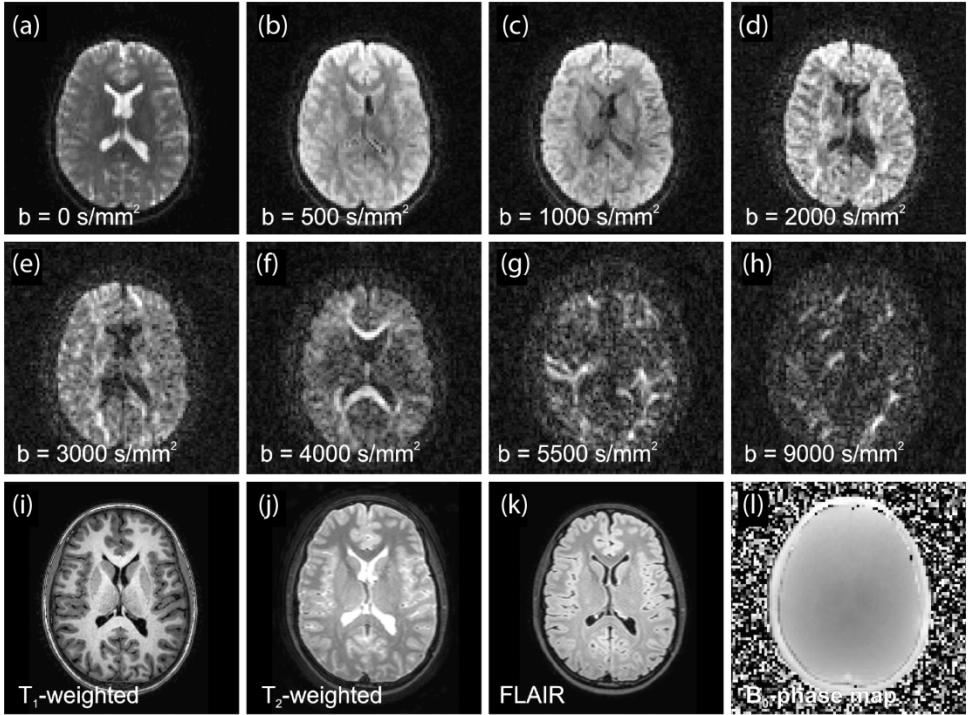


Fig. 3.3: Representative images of the acquired data (intensity windowing is adjusted for each image). Diffusion-weighted images with an applied b -value of: (a) 0 s/mm^2 , (b) 500 s/mm^2 , (c) 1000 s/mm^2 , (d) 2000 s/mm^2 , (e) 3000 s/mm^2 , (f) 4000 s/mm^2 , (g) 5500 s/mm^2 , (h) 9000 s/mm^2 . Anatomical images: (i) T₁-weighted image, (j) T₂-weighted image, and (k) FLAIR. (l) B₀ phase map.

3.3.2 Intensity normalization and signal drift correction

In Fig. 3.5, the mean signal of the non-DW volumes is shown as a function of the measurement number (red markers) for 24 randomly chosen acquisition blocks. The quadratic fit used to correct for signal drift is shown in black and the signal-drift corrected data are shown in blue. The mean signal drift during a 15 minute acquisition block was 9.3% with a standard deviation of 3.7% (range: 3.0 to 18.8%) with respect to the initial volume. The mean signal for all the 8000 acquired volumes is shown in Figs. 3.6a and 3.6b. By comparing Fig. 3.6a with Fig. 3.6b, one can appreciate that the mean signal of the normalized and drift-corrected volumes across all sessions is more constant now.

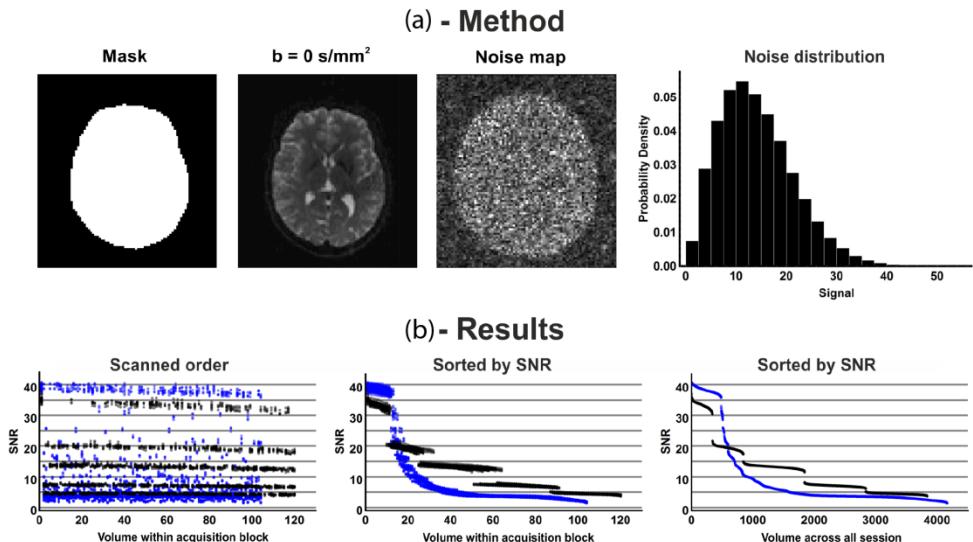


Fig. 3.4: Method and result used for SNR calculation. (a) Schematic representation of the method used for SNR calculation. Using a whole-brain mask the average signal was estimated from the diffusion weighted volume. The noise standard deviation was estimated using the acquired noise map in which the noise is Rayleigh distributed as shown in the probability density histogram. (b) Estimated SNR using the acquired noise map. Black markers are for session 1 to 8 ($b \leq 3000 \text{ s/mm}^2$) and blue markers are for session 9 to 18 ($b > 3000 \text{ s/mm}^2$ and DSI grids). The left image shows the SNR per session in the scanned order, the middle image shows the SNR per session sorted for the SNR, and the right image shows the SNR for all the volumes. The SNR levels of the different shells can clearly be identified.

3.3.3 True B-matrix calculation

Since the slice selection gradients are always in the orientation of the z-axis, the actual b -value for the non-DW images, using Eqs. [3.3] and [3.4], is 0.25 s/mm^2 instead of 0 s/mm^2 . The average value and the range of the actual b -values of the five shells are listed in Table 3.2. For the DW images the actual b -value differs up to 2% from the applied b -value. The relative difference is largest for lower b -values. The effect of these small differences is shown in Figs. 3.6c and 3.6d, where the intensity-normalized mean signal (with and without signal-drift correction) is plotted as a function of both the applied and the actual b -value. The imaging gradients not only affect the b -value, but also the gradient direction. However, the median change in orientation due to the slice selection gradients across all gradients was found to be only 0.004° . The applied and actual B-matrix values can be found in the lookup table which can be downloaded from the website www.massive-data.org.

Less Confusion in Diffusion MRI

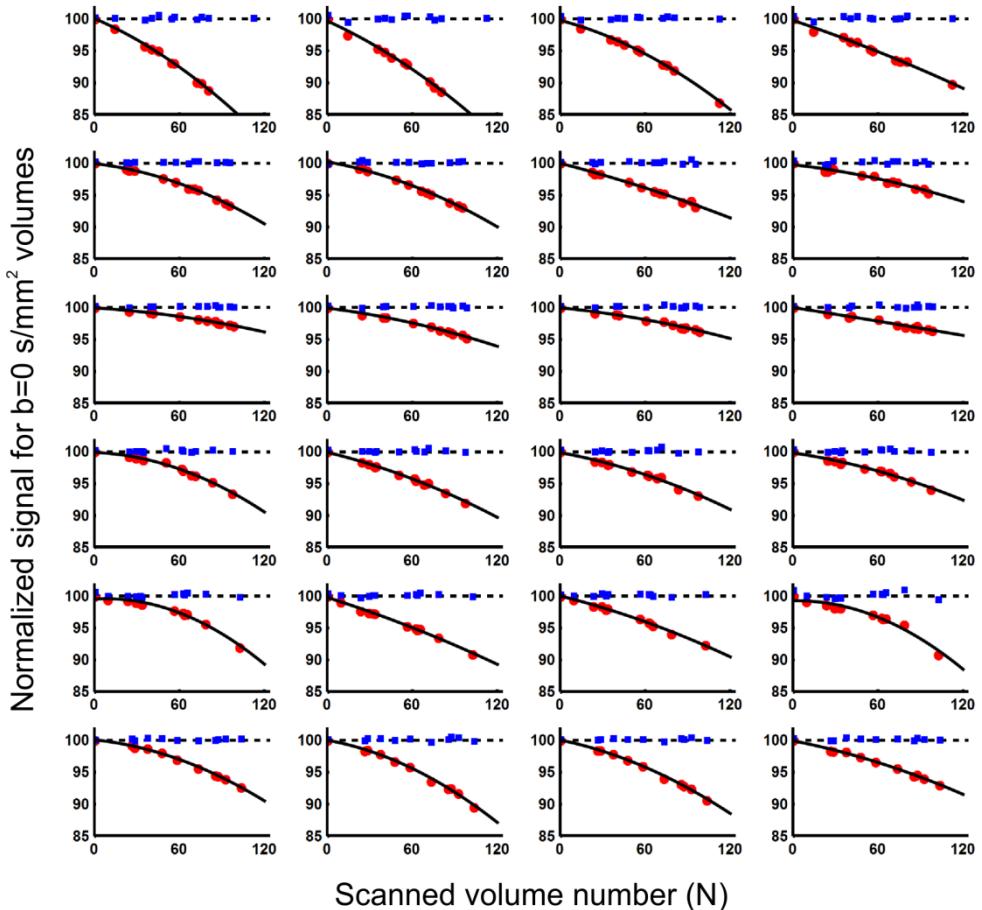


Fig. 3.5: Original (red) and signal-drift corrected (blue) average signal of the non-diffusion-weighted volumes for 24 acquisition blocks of 15 minutes. The black line describes the quadratic signal drift (see Eq. [3.1]) and the dashed line is the theoretically constant signal of 100%.

Applied b -value in units s/mm^2	Actual b -value (mean \pm SD) in units s/mm^2	Actual b -value range in units s/mm^2	Percent deviation from applied b -value
500	500.3 ± 5.2	$491.3 - 509.2$	± 1.8
1000	1000.3 ± 7.3	$987.6 - 1013.0$	± 1.3
2000	2000.3 ± 10.4	$1982.3 - 2018.2$	± 0.9
3000	3000.3 ± 12.7	$2978.3 - 3022.3$	± 0.7
4000	4000.2 ± 15.1	$3974.1 - 4026.4$	± 0.7

Table 3.2: Mean values and range of the actual b -values for the 5 acquired shells.

Chapter 3: The “MASSIVE” Brain Dataset:
Multiple Acquisitions for Standardization of Structural Imaging Validation and Evaluation

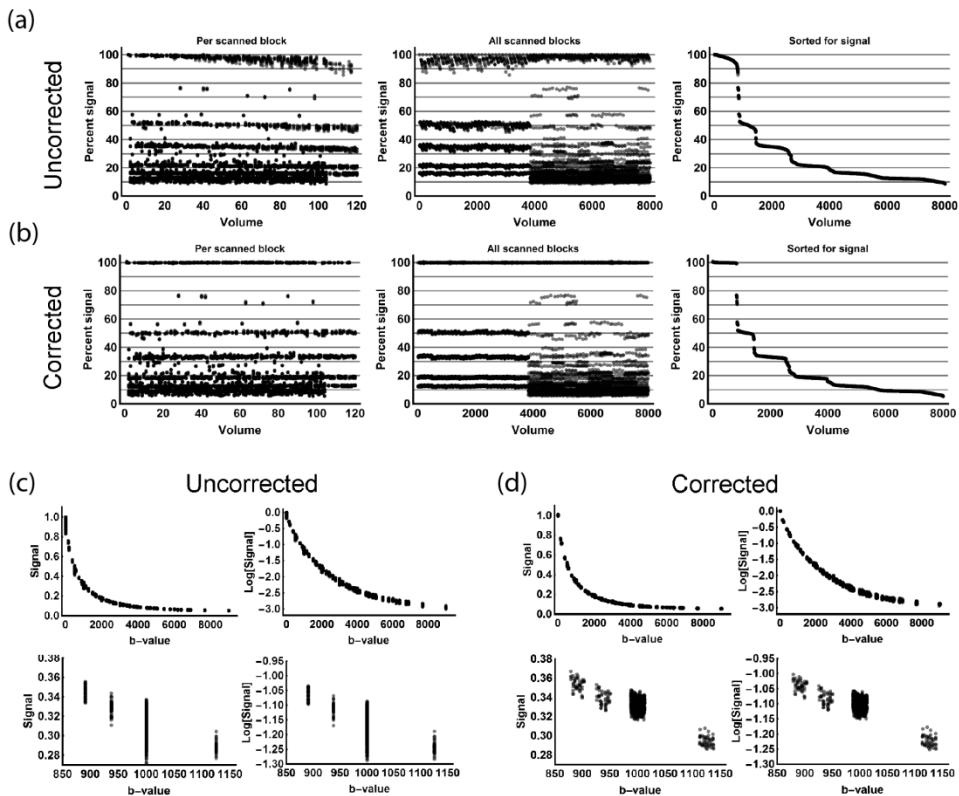


Fig. 3.6: (a-b) The mean signal per acquired volume of the drift-uncorrected volumes (top row - a) signal-drift corrected volumes (bottom row - b). The left column shows the signal per session in the scanned order, the middle column shows the mean signal for all sessions in the scanned order and the right column shows the signal for all the volumes sorted for the signal values. (c-d) The intensity-normalized mean signal per acquired volume as a function of the b -value for all 8000 volumes (top row) and a zoomed b -value range (bottom row). The signal is plotted using a linear and a logarithmic scaling of the y-axes revealing the obvious non-exponential decay of the diffusion weighted signal. c) The normalized mean signal (but not corrected for signal drift) as a function of the predefined b -value. d) The intensity-normalized and signal-drift corrected mean signal as function of the actual b -value. The bottom row clearly shows how the ‘effective’ b -value causes a spread of the signal over the b -value axes, that correctly follows the signal decay.

3.3.4 Correcting for subject motion, eddy current distortions, and EPI deformations

The average maximal translations over all individual sessions were $0.3 \pm 3.3 \text{ mm}$ (range -4.7 to 4.6 mm), $0.4 \pm 1.2 \text{ mm}$ (range -2.5 to 2.8 mm), and $-0.8 \pm 1.3 \text{ mm}$ (range -3.6 to 2.3 mm) for the coronal, sagittals and axial directions, respectively. The average maximal rotations over all individual session were $-0.6 \pm 1.0^\circ$ (range -2.5 to 2.9°), $0.2 \pm 1.2^\circ$ (range -2.9 to 2.6°), and $-0.2 \pm 1.1^\circ$ (range -2.3 to 2.4°) along the coronal (roll), sagittal (pitch), and axial (yaw) axes, respectively. The maximal range of the rotations between and within the sessions was only $\pm 5^\circ$ with respect to the mean. The difference between the imposed gradient distribution and the gradient distribution after B-matrix correction for subject motion is shown for a subset of the data ($b = 1000 \text{ s/mm}^2$ and $b = 3000 \text{ s/mm}^2$ volumes) in Supplementary Fig. S3.2.

Fig. 3.7 shows the differences in image distortions between AP and PA phase encoding directions, on the one hand, and positive and negative diffusion weighting gradient directions, on the other hand (AP+, AP-, PA+, PA-). The differences in the distortions between the AP and PA phase-encoding directions can be appreciated most in the regions of the eyes, the temporal lobes, and the cerebellum. The difference in distortions between positive and negative gradient directions is reflected by the opposite eddy current distortions (e.g., compressions vs. stretches along the phase-encoding orientation). The color-coded FA map (Fig. 3.8a) and the T1-weighted data (Fig. 3.8b) in anatomical space with a $1 \times 1 \times 1 \text{ mm}^3$ voxel size are shown in Fig. 3.8. The alignment of both datasets after motion correction, eddy current distortion correction, and EPI distortion correction by registration to the T1-weighted image is illustrated in Fig. 3.8c. The registered data in the anatomical space is available with a $2.5 \times 2.5 \times 2.5 \text{ mm}^3$ (website) and a $1 \times 1 \times 1 \text{ mm}^3$ voxel size (upon request, due to limited online storage capacity).

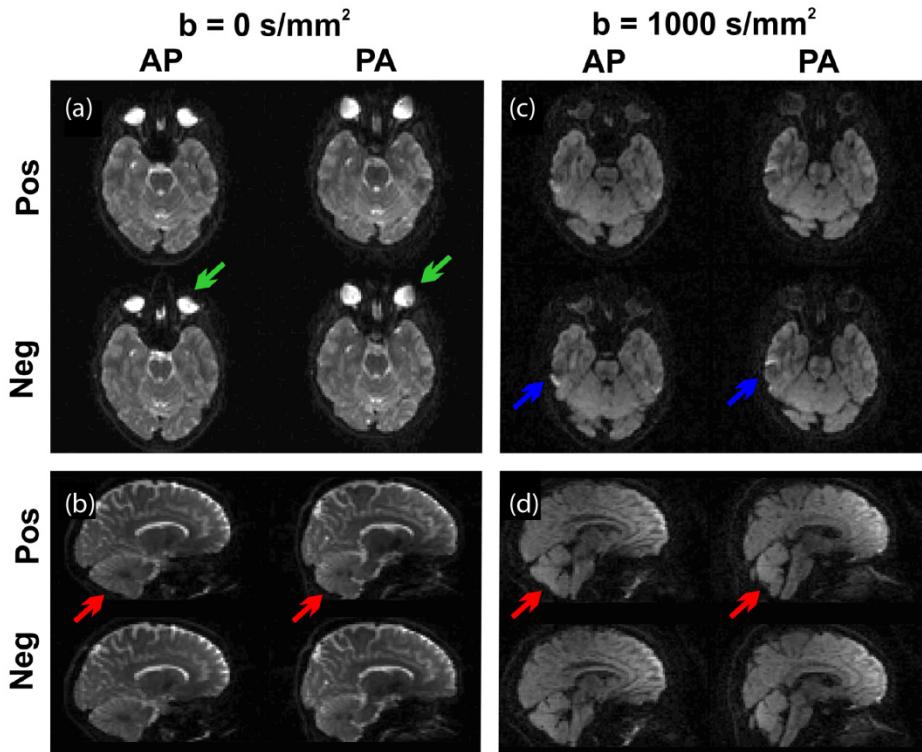


Fig. 3.7: Representative images for the different phase encoding (AP+, AP-, PA+, PA-) directions. Axial (a and c) and sagittal (b and d) cross-sectional images of a non-diffusion-weighted (a and b, $b = 0 \text{ s/mm}^2$) and a diffusion-weighted (c and d, $b = 1000 \text{ s/mm}^2$) volume with the different phase encoding directions and gradient signs. The differences in the distortions between the AP and PA phase encoding directions can be appreciated most in the eyes (green arrows), temporal region (blue arrows), and the cerebellum (red arrows).

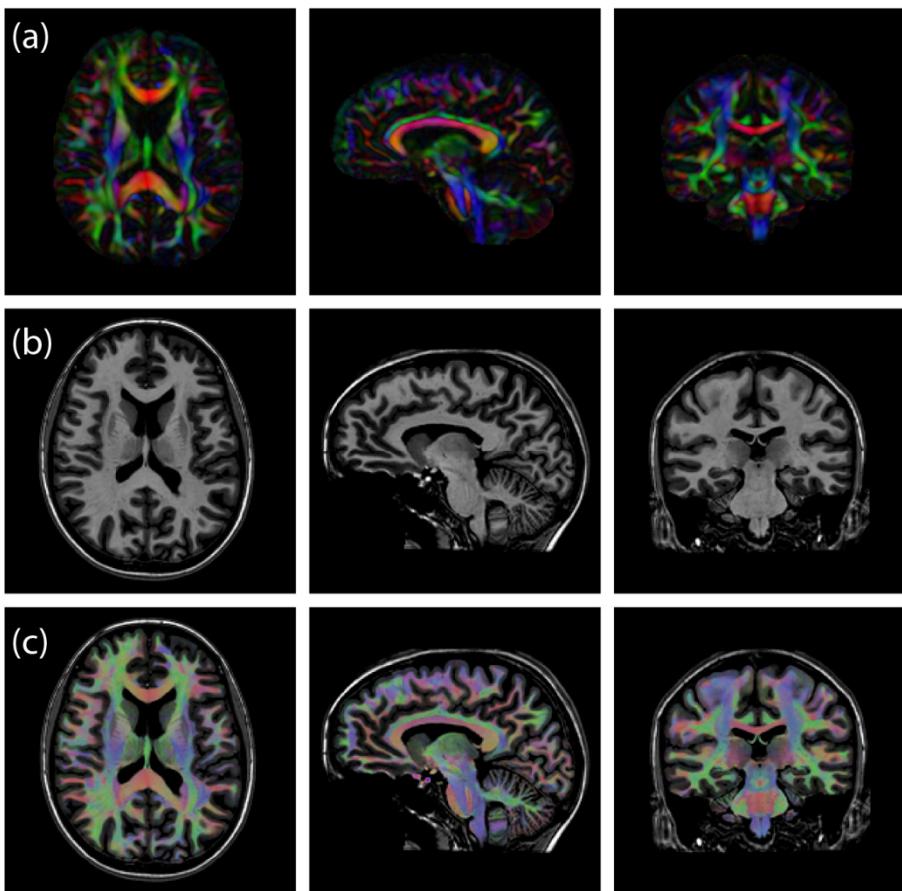


Fig. 3.8: Example DTI reconstruction of the MASSIVE data. (a) Color coded FA maps after applying all the correction procedures. (b) T1 weighted data, which were used here to correct for EPI deformations. (c) Color coded FA maps fused with T1 weighted data to appreciate the quality of the processing.

3.4 DISCUSSION

In this work, we have presented the MASSIVE brain dataset, which contains 8000 *in vivo* dMRI volumes of a healthy subject. Currently, the raw, intensity-normalized, signal-drift corrected, and subject motion / eddy current distortion / EPI distortion corrected dMRI data can be downloaded from www.massive-data.org. All the B_0 field maps, noise maps, and the volumetric FLAIR, T1-, and T2-weighted datasets are also made available.

We have established a platform through the MASSIVE website to share improvements of specific processing steps and updates of the processed data. Such methodological developments encompass novel subject motion and distortion correction methods, new microstructural modelling approaches, etc. Similarly, we anticipate that segmentations of the anatomical images

(e.g., T1- and T2-weighted images) using common brain atlases and processing tools will also become available.

The purpose of MASSIVE is to serve as an extensive dataset to compare, evaluate, and validate existing or novel diffusion MRI methods, such as preprocessing steps, signal modelling, tissue characterization, and analysis strategies. While existing brain dMRI databases can provide data from many subjects with only moderate coverage of q -space, none of these provide dMRI data from a single subject with as dense a q -space sampling as provided in the MASSIVE database. MASSIVE is unique in consisting of 8000 DW volumes that are sampled on shells as well as two Cartesian grids. The data was acquired on a standard clinical system using coil and acquisition settings that are commonly available. This makes the data quality of each individual dMRI volume comparable to data typically acquired in most clinical studies, which means that subset of the data, e.g. containing 100 dMRI volumes with $b = 1000 \text{ s/mm}^2$, would closely resemble a clinical acquisition. As such, methods and models derived from this database can easily be transferred to other clinical and pre-clinical research workflows.

Examples of research questions that have already been addressed using the MASSIVE brain database include characterizing signal drift in dMRI acquisitions (Vos et al., 2016b), investigating the *in vivo* trade-off between accuracy and precision of multi-fiber methods with respect to b -value and number of gradient orientations (Vos et al., 2014), studying the difference between interpolation methods for transforming dMRI data between grids and shells (Tax et al., 2014e), and characterizing single fiber population signal profiles using a wide range of reconstruction strategies (Tax et al., 2015b). Furthermore, the MASSIVE dataset could be useful in investigating the optimality of different EPI distortion correction techniques, e.g. using field maps, registration to an anatomical image, or using opposite phase encoding images. Commonly used dMRI acquisition protocols sample either the upper or lower hemisphere in q -space. However, eddy current correction techniques might benefit from sampling on the whole sphere (Andersson and Sotropoulos, 2016). In addition to correcting for image distortions, correction for subject motion remains an active field of research in which new methods (e.g. targeted to high b -values) are constantly being developed and evaluated (Ben-Amitay et al., 2012; Huizinga et al., 2016). Therefore, the optimal acquisition scheme to adequately correct for artifacts and motion remains an open question, and the MASSIVE dataset could be subsampled to investigate such issues. In addition, whereas there is a reasonable consensus of an ‘ideal’ single shell acquisition for diffusion tensor imaging (Jones and Leemans, 2011), optimal acquisition strategies for many of the other diffusion models are still being investigated and new models are constantly under development (Caruyer et al., 2013; De Santis et al., 2014; Hoy et al., 2015; Jelescu et al., 2016b; Jensen et al., 2016; Kuo et al., 2008; Zhang et al., 2012). The MASSIVE brain database, and in particular the unique dMRI dataset, which – to the best of our knowledge – represents the largest *in vivo* dMRI dataset of a single subject to date, will avert the need to continuously reacquire optimized data and boost new developments in diffusion modelling and processing.

3.5 CONCLUSION

We have presented the MASSIVE (Multiple Acquisitions for Standardization of Structural Imaging Validation and Evaluation) brain dataset, consisting of an unprecedented set of 8000

Chapter 3: The “MASSIVE” Brain Dataset:
Multiple Acquisitions for Standardization of Structural Imaging Validation and Evaluation

DW volumes of a single human subject. This unique set of in vivo MRI data will provide a robust framework to evaluate novel diffusion processing techniques and to reliably compare different approaches for diffusion modelling. All data is made publicly available on www.massive-data.org.

3.6 SUPPLEMENTARY MATERIAL

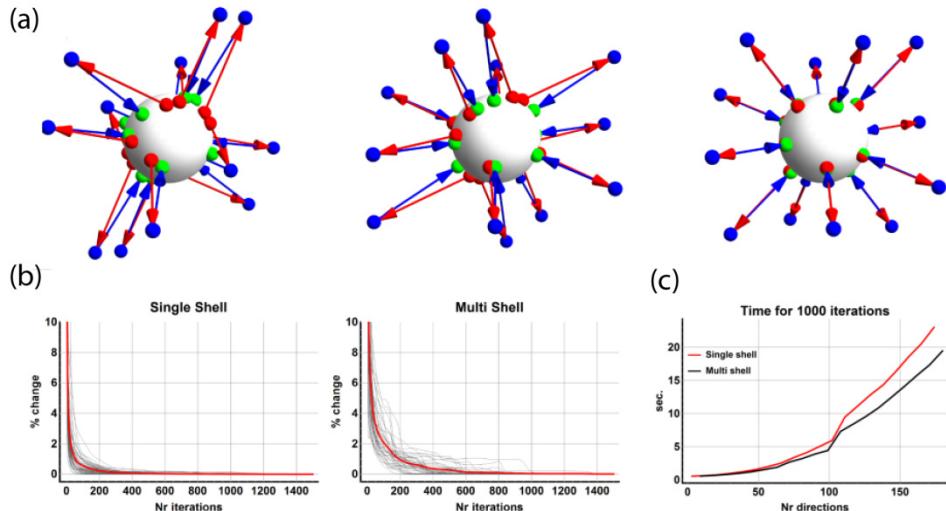


Fig. S3.1: (a) Iterative method of optimization, first all particles are repulsed (red) after which they are projected back on the unit sphere (blue). The resulting particles (green) are then used for the next iteration. (b) The percent change in entropy of the system over 1500 iterations for single- and multi-shell optimizations. (c) Time it takes to perform 1000 iterations (Intel Core i5-2520M CPU, 2.5GHz) for shell ranging from 3 to 180 gradient orientations.

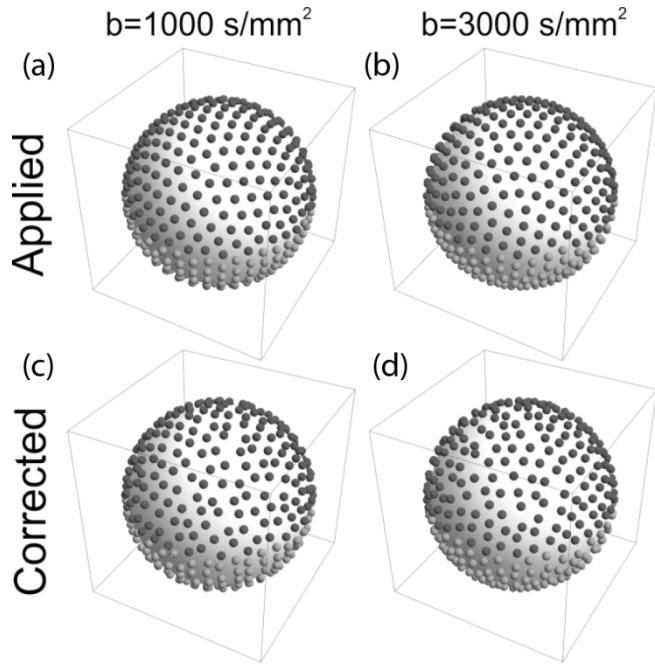


Fig. S3.2: Effect of subject motion on the imposed gradient orientations. The 500 gradient direction for $b = 1000 \text{ s/mm}^2$ and $b = 3000 \text{ s/mm}^2$ as they were defined (a and b) and after registration with B-matrix correction (Leemans and Jones, 2009) (c and d). The condition number of the $b = 1000 \text{ s/mm}^2$ and $b = 3000 \text{ s/mm}^2$ shells (1.5813 for both) changed to 1.5826 and 1.5861, respectively.

4 REKINDLE: ROBUST EXTRACTION OF KURTOSIS INDICES WITH LINEAR ESTIMATION

Based on:

C.M.W. Tax, W.M. Otte, M.A. Viergever, R.M. Dijkhuizen, A. Leemans, “REKINDLE: Robust Extraction of Kurtosis INDices with Linear Estimation”, Magnetic Resonance in Medicine, 2015, nr. 2, vol. 73, pp. 794-808

ABSTRACT

Recent literature shows that diffusion tensor properties can be estimated more accurately with diffusion kurtosis imaging (DKI) than with diffusion tensor imaging (DTI). Furthermore, the additional non-Gaussian diffusion features from DKI can be sensitive markers for tissue characterization. Despite these benefits, DKI is more susceptible to data artifacts than DTI due to its increased model complexity, higher acquisition demands, and longer scanning times. To increase the reliability of diffusion tensor and kurtosis estimates, we propose a robust estimation procedure for DKI. We have developed a robust and linear estimation framework, coined REKINDLE (Robust Extraction of Kurtosis INDices with Linear Estimation), consisting of an iteratively reweighted linear least squares approach. Simulations are performed, in which REKINDLE is evaluated and compared with the widely used RESTORE (Robust EStimation of Tensors by Outlier REjection) method. Simulations demonstrate that in the presence of outliers, REKINDLE can estimate diffusion and kurtosis indices reliably and with a ten-fold reduction in computation time compared to RESTORE. We have presented and evaluated REKINDLE, a linear and robust estimation framework for DKI. While REKINDLE has been developed for DKI, it is by design also applicable to DTI and other diffusion models that can be linearized.

4.1 INTRODUCTION

In diffusion MRI, microstructural properties of tissue are revealed by probing the local diffusion of water molecules, which can give useful insights into brain disorders (Carballido et al., 2012; Langen et al., 2012; Reijmer et al., 2012; Verhoeven et al., 2012; Wang et al., 2012). For clinical applications, it is challenging to get an accurate and complete picture of diffusion characteristics within a limited acquisition time. Diffusion tensor imaging (DTI) (Basser et al., 1994) is the most commonly used technique to extract microstructural features, such as fractional anisotropy (FA) and mean diffusivity (MD), from a set of diffusion weighted images (DWIs). To this end, a diffusion tensor (DT) with six parameters is fitted to the data.

In addition to the metrics obtained with DTI, diffusion kurtosis imaging (DKI) (Jensen et al., 2005) can provide non-Gaussian diffusion measures by means of the kurtosis tensor (KT). DKI has shown to be more sensitive to tissue microstructural changes in both normal and pathological neural tissue (e.g., Cheung et al., 2009), and kurtosis indices such as mean kurtosis (MK) (Poot et al., 2010) and other derived quantities (Hansen et al., 2013) can provide additional information in clinical applications (Gooijers et al., 2014; Van Cauter et al., 2012). Furthermore, DKI is able to estimate DTI measures such as FA more accurately (Veraart et al., 2011a). In a clinical setting, however, these benefits are often nullified by numerous acquisition artifacts, such as signal dropouts due to subject motion or cardiac pulsation (Heemskerk et al., 2013; Jones and Cercignani, 2010). Owing to its increased model complexity, higher acquisition demands, and longer scanning times, DKI is also more sensitive to data outliers than DTI. Though it would be best to avoid signal perturbations during acquisition, this is not always possible and outliers are often only detected during image processing. At this stage, tensor estimation procedures that can deal with such signal perturbations are crucial to ensure the reliability of subsequent tractography- or voxel-based analyses.

There exist various approaches to estimate the DT and KT from the diffusion signal, like maximum likelihood, least squares, Bayesian, and robust approaches. These methods differ in parameterization, a priori information needed, and the incorporation of constraints, which makes tensor estimation an important ongoing field of research (Koay, 2010). Least-squares (LS) estimation approaches are generally adopted for tensor estimation (Koay et al., 2006). Unlike more advanced diffusion parameter estimators, e.g. the maximum likelihood (Veraart et al., 2011b) and maximum a posteriori (Andersson, 2008) estimators, LS methods do not depend on knowledge of the full data distribution. LS approaches are, however, based on certain assumptions, such as a normally distributed and additive error term. These assumptions are often invalidated by the occurrence of outliers and, consequently, parameter estimates become unreliable.

Robust estimation procedures aim at reducing the influence of outliers on the eventual model estimates (Chang et al., 2005; Chang et al., 2012; Mangin et al., 2002; Pannek et al., 2012; Parker et al., 2013b; Zhou et al., 2011; Zwiers, 2010). The popular Robust Estimation of Tensors by Outlier Rejection (RESTORE) method improves DT estimation in the presence of artifacts and has been widely used in DTI (Chang et al., 2005). This method iteratively reweights each data point in nonlinear LS (NLS) regression to reduce the impact of outliers on the fit. Subsequently, potential outliers are identified based on the residuals and excluded from estimation.

In this chapter, we rekindle the RESTORE robust estimation procedure in the context of DKI. In addition to the extension of the conventional RESTORE approach to DKI, we propose an adapted linearized framework, coined REKINDLE (Robust Extraction of Kurtosis INDices with Linear Estimation). Linearized approaches have several strengths compared with nonlinear strategies: They are fast, not prone to getting stuck in local optima, and linear LS estimators are easy to compute (Veraart et al., 2013b). Owing to linearization, however, the standard deviation (SD) of the log transformed signal varies as a function of the signal intensity (Salvador et al., 2005). REKINDLE takes into account this heteroscedasticity in robust estimation (Carroll and Ruppert, 1982; Giltinan et al., 1986; Maronna et al., 2006). Furthermore, REKINDLE does not require a predefined estimate of the signal SD to identify potential outliers. Many (robust) estimation methods require such an estimate and assume it to be constant throughout the volume. However, the assumption of spatially constant signal variance is invalid (Veraart et al., 2013a), which may complicate outlier identification. REKINDLE, instead, uses an intuitive outlier detection criterion that allows the user to adjust for the desired sensitivity and specificity. Our results on simulations and real data demonstrate that REKINDLE is able to identify and exclude outliers and reduces computational cost drastically without any significant reduction in accuracy. Preliminary results of this work on REKINDLE have been presented at the 2013 ISMRM meeting in Salt Lake City, USA (Tax et al., 2013).

4.2 METHODS

REKINDLE uses iteratively reweighted LS regression (IRLS) to identify outliers. These outliers are subsequently excluded and a final fit is performed on the remaining data points. Before explaining the framework in detail, we first recapitulate the relation of the DT and KT with the signal attenuation and the corresponding regression analysis, both nonlinear and linear. We then elaborate on two types of estimators: LS estimators and robust estimators. Based on this, the REKINDLE framework is explained.

4.2.1 Estimation framework

4.2.1.1 The DKI model

In DKI, the DT and KT can be computed from their relation with the signal attenuation as follows:

$$S(b, \mathbf{g}) = S(0) \exp \left(-b \mathbf{g}^T \mathbf{D} \mathbf{g} + \frac{1}{6} b^2 \left(\frac{1}{3} \sum_{i=1}^3 D_{ii} \right)^2 \sum_{i=1}^3 \sum_{j=1}^3 \sum_{k=1}^3 \sum_{l=1}^3 g_i g_j g_k g_l W_{ijkl} \right), \quad [4.1]$$

with $S(b, \mathbf{g})$ the noise-free DW signal along gradient direction $\mathbf{g} = (g_1, g_2, g_3)$, b the diffusion weighting, $S(0)$ the non-DW signal, and D_{ij} and W_{ijkl} the DT and KT elements, respectively.

4.2.1.2 Regression analysis

Since the DTI and DKI models are nonlinear, solving the NLS problem yields estimates of the diffusion parameters (Jones and Basser, 2004). Eq. [4.1] can be written as nonlinear regression function

$$\mathbf{z} = \exp(\mathbf{X}\boldsymbol{\beta}) + \boldsymbol{\epsilon}, \quad [4.2]$$

where \mathbf{z} is a $n \times 1$ column vector (n equals the number of measurements) with the observed signal intensities $\mathbf{z} = (\tilde{S}(b_1, \mathbf{g}_1), \dots, \tilde{S}(b_n, \mathbf{g}_n))^T$, and the independent variable \mathbf{X} a ‘B-matrix’ of size $n \times 22$ that combines both DT and KT multiplication components:

$$\mathbf{X} = (1 \ -\mathbf{b}_{DT} \ \mathbf{b}_{KT}), \quad [4.3]$$

with

$$\mathbf{b}_{DT} = \begin{pmatrix} b_1 g_{11}^2 & 2b_1 g_{11}g_{12} & 2b_1 g_{11}g_{13} & b_1 g_{12}^2 & 2b_1 g_{12}g_{13} & b_1 g_{13}^2 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ b_n g_{n1}^2 & 2b_n g_{n1}g_{n2} & 2b_n g_{n1}g_{n3} & b_n g_{n2}^2 & 2b_n g_{n2}g_{n3} & b_n g_{n3}^2 \end{pmatrix}, \quad [4.4]$$

and

$$\begin{aligned} \mathbf{b}_{KT} = & \frac{1}{54} \begin{pmatrix} b_1^2 g_{11}^4 & b_1^2 g_{12}^4 & b_1^2 g_{13}^4 & 4b_1^2 g_{11}^3 g_{12} & 4b_1^2 g_{11}^3 g_{13} & 4b_1^2 g_{12}^3 g_{11} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ b_n^2 g_{n1}^4 & b_n^2 g_{n2}^4 & b_n^2 g_{n3}^4 & 4b_n^2 g_{n1}^3 g_{n2} & 4b_n^2 g_{n1}^3 g_{n3} & 4b_n^2 g_{n2}^3 g_{n1} \\ 4b_1^2 g_{12}^3 g_{13} & 4b_1^2 g_{13}^3 g_{11} & 4b_1^2 g_{13}^3 g_{12} & 6b_1^2 g_{11}^2 g_{12}^2 & 6b_1^2 g_{11}^2 g_{13}^2 & 6b_1^2 g_{12}^2 g_{13}^2 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 4b_n^2 g_{n2}^3 g_{n3} & 4b_n^2 g_{n3}^3 g_{n1} & 4b_n^2 g_{n3}^3 g_{n2} & 6b_n^2 g_{n1}^2 g_{n2}^2 & 6b_n^2 g_{n1}^2 g_{n3}^2 & 6b_n^2 g_{n2}^2 g_{n3}^2 \\ 12b_1^2 g_{12} g_{13} g_{11}^2 & 12b_1^2 g_{11} g_{13} g_{12}^2 & 12b_1^2 g_{11} g_{12} g_{13}^2 \\ \vdots & \vdots & \vdots \\ 12b_n^2 g_{n2} g_{n3} g_{n1}^2 & 12b_n^2 g_{n1} g_{n3} g_{n2}^2 & 12b_n^2 g_{n1} g_{n2} g_{n3}^2 \end{pmatrix}. \end{aligned} \quad [4.5]$$

In Eq. [4.2], $\boldsymbol{\beta}$ represents the parameters to be estimated, i.e. the DT and KT components, which can be written in a 22×1 matrix

$$\begin{aligned} \boldsymbol{\beta} = & (\log(S_0) \ D_{11} \ D_{12} \ D_{13} \ D_{22} \ D_{23} \ D_{33} \ tr(\mathbf{D})^2 W_{1111} \ tr(\mathbf{D})^2 W_{2222} \\ & tr(\mathbf{D})^2 W_{3333} \ tr(\mathbf{D})^2 W_{1112} \ tr(\mathbf{D})^2 W_{1113} \ tr(\mathbf{D})^2 W_{2221} \ tr(\mathbf{D})^2 W_{2223} \\ & tr(\mathbf{D})^2 W_{3331} \ tr(\mathbf{D})^2 W_{3332} \ tr(\mathbf{D})^2 W_{1122} \ tr(\mathbf{D})^2 W_{1133} \ tr(\mathbf{D})^2 W_{2233} \\ & tr(\mathbf{D})^2 W_{2311} \ tr(\mathbf{D})^2 W_{1322} \ tr(\mathbf{D})^2 W_{1233})^T, \end{aligned} \quad [4.6]$$

where $tr(\mathbf{D}) = (D_{11} + D_{22} + D_{33})$. The term $\boldsymbol{\epsilon}$ in Eq. [4.2] is a $n \times 1$ vector of random errors with $E(\epsilon_i) = 0$ and $Var(\epsilon_i) = \sigma^2$.

Both DTI and DKI diffusion models are frequently linearized by the natural log transformation. The log-transformed signal intensity $\mathbf{y} = (\log(\tilde{S}(b_1, \mathbf{g}_1)), \dots, \log(\tilde{S}(b_n, \mathbf{g}_n)))^T = \log(\mathbf{z})$ is a heteroscedastic random variable with the variance being dependent on the true signal intensity. The nonlinear model in Eq. [4.2] can be written as:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\sigma} \circ \boldsymbol{\epsilon}, \quad [4.7]$$

with \circ denoting the entrywise product. Under the condition of high signal-to-noise ratio (SNR), the vector $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)^T$ includes scaling constants that express the heteroscedasticity of the linear model (Appendix A4) (Carroll and Ruppert, 1982; Giltinan et al., 1986; Salvador et al., 2005). The random error term $\sigma_i \epsilon_i$ has variance $\sigma_i^2 \sigma^2$ (scale equivariance: $Var(\sigma_i \epsilon_i) = \sigma_i^2 Var(\epsilon_i) = \sigma_i^2 \sigma^2$). When no replicated measurements are available to estimate σ_i , one often assumes a model $\sigma_i = h(\boldsymbol{\lambda}, \mathbf{x}_i^T \boldsymbol{\beta})$, where σ_i thus depends on the values $\mathbf{x}_i^T \boldsymbol{\beta}$ (with \mathbf{x}_i the i^{th} row of \mathbf{X}) and a vector parameter $\boldsymbol{\lambda}$. For example, one could model the error scale as Bickel (1978); Maronna et al. (2006):

$$h(\boldsymbol{\lambda}, \mathbf{x}_i^T \boldsymbol{\beta}) = \exp(\lambda_1 + \lambda_2 \mathbf{x}_i^T \boldsymbol{\beta}). \quad [4.8]$$

Under the condition that the DW data are Rician distributed and the SNR of the DW signals is greater than 2, the linear estimator is unbiased. Furthermore, Salvador et al. (2005) showed that

$$\text{Var}(\boldsymbol{\sigma} \circ \boldsymbol{\epsilon}) \propto (S^{-2}(b_1, \mathbf{g}_1), \dots, S^{-2}(b_n, \mathbf{g}_n))^T, \quad [4.9]$$

and thus $\sigma_i \propto S^{-1}(b_i, \mathbf{g}_i) = \exp(-\mathbf{x}_i^T \boldsymbol{\beta})$ (see also Appendix A4), which corresponds to taking λ_1 as some constant and $\lambda_2 = -1$ in Eq. [4.8]. If $\boldsymbol{\beta}$ is known, then the transformed variables

$$y_i^* = \frac{y_i}{\exp(-\mathbf{x}_i^T \boldsymbol{\beta})}, \quad \mathbf{x}_i^* = \frac{\mathbf{x}_i}{\exp(-\mathbf{x}_i^T \boldsymbol{\beta})} \quad [4.10]$$

would follow the homoscedastic regression model.

4.2.1.3 Least squares estimators

The linear LS (LLS) approach minimizes the sum of squared residuals $\sum_{i=1}^n e_i^2$, with $e_i = y_i - \hat{y}_i$, where y_i represents the measured and \hat{y}_i the predicted signal (Montgomery and Runger, 2010). The parameter estimates can then be obtained by solving

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}. \quad [4.11]$$

To account for the differences in variance in the log-transformed signals, one should instead minimize the weighted residuals $\sum_{i=1}^n W_{ii} e_i^2$ with $W_{ii} = \sigma_i^{-2} \propto S^2(b_i, \mathbf{g}_i)$ (corresponding to Eq. [4.9]) in weighted LLS (WLLS), obtaining:

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}^T \mathbf{W} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{W} \mathbf{y}, \quad [4.12]$$

and thus $\mathbf{W} = \text{diag}(S^2(b_1, \mathbf{g}_1), \dots, S^2(b_n, \mathbf{g}_n))$ (Salvador et al., 2005; Veraart et al., 2013b).

4.2.1.4 Robust estimators

Instead of minimizing $\sum_{i=1}^n e_i^2$, robust estimators try to reduce the effect of outliers by minimizing the sum of another symmetric, positive-definite function of the residuals $\rho(e_i)$ (Meer et al., 1991). The minimization problem is converted into an IRLS problem, in which the weights are recalculated iteratively and depend on $\rho(e_i)$ (Holland and Welsch, 1977; Street et al., 1988) (see Appendix B4).

Mangin et al. (2002) propose to use the Geman-McClure M-estimator $\rho(d_i) = d_i^2/2/(d_i^2 + 1)$, where $d_i = e_i/\sigma$ is the standardized residual, to reduce the influence of outliers in DTI estimation (Mangin et al., 2002). Compared with $\rho(e_i) = e_i^2$, which is driven by large residuals, this function cancels out the influence of residuals far beyond σ . The corresponding Geman-McClure weighting function, used for iterative reweighting, is (Appendix B4):

$$w(d_i) = \frac{1}{(d_i^2+1)^2}. \quad [4.13]$$

Chang et al. (2005) use this approach only to identify potential outliers, and subsequently exclude them in the final fitting procedure (Chang et al., 2005). Their RESTORE approach uses the nonlinear relationship of Eq. [4.2] and performs a final nonlinear fit with constant weights.

4.2.1.5 REKINDLE

Similarly to the DTI based RESTORE (Chang et al., 2005) algorithm, we propose REKINDLE as robust linearized DKI estimation framework (Fig. 4.1). Since the log-transformed signal intensity is a heteroscedastic variable, we have to take this into account when detecting outliers. It is not trivial, however, how to combine the heteroscedasticity correction with the Geman McClure weights (Chang et al., 2005), since the measured signal intensities are influenced by noise and outliers. Therefore, the framework consists of a multistep procedure based on a method described in Maronna et al. (2006) to detect and exclude outliers (grey box in Fig. 4.1). More background on the motivation of the multistep nature of our method can be found in the Supplementary Information.

Initialized by an LLS estimate (step 1), a robust estimate for homoscedastic regression is computed in step 2), assuming that the σ_i are equal for all measurements at this stage. The variables y_i and x_i are subsequently transformed to y_i^* and x_i^* in step 3), and robust estimation for homoscedastic regression based on the transformed variables (which are now more likely to be ‘truly homoscedastic’) is performed in step 5) (Giltinan et al., 1986; Maronna et al., 2006). Robust homoscedastic regression was done using IRLS, constituting the ‘inner loops’ in steps 2) and 5) of Fig. 4.1 (Chang et al., 2005; Holland and Welsch, 1977; Maronna et al., 2006). The framework was implemented in MATLAB (MathWorks, Natick, Massachusetts, U.S.A.):

- 1) Perform an initial LLS fit (Eq. [4.11]). Note that we do not check for goodness of fit here, which prevents the need for estimating the signal SD a priori.
- 2) Compute a robust estimate for homoscedastic regression using IRLS by doing the following steps for each voxel:
 - a) Calculate the residuals $e = y - \hat{\mathbf{X}}\hat{\beta}$ in the linear domain.

- b) Obtain an estimate of the dispersion of the residuals σ by calculating the median absolute deviation (MAD) (Chang et al., 2005; Maronna et al., 2006):

$$\hat{\sigma} = 1.4826 \cdot \text{MAD}(\mathbf{e}) = 1.4826 \cdot \text{Median}(|\mathbf{e} - \text{Median}(\mathbf{e})|). \quad [4.14]$$

- c) Recompute the weights according to Eq. [4.13].
- d) Perform a WLLS fit with these new weights (Eq. [4.12]).
- e) Check if the estimated parameters $\widehat{\boldsymbol{\beta}}$ have not changed significantly with respect to the previous iteration in this inner loop $j-1$ (check if all $|\widehat{\boldsymbol{\beta}}_j - \widehat{\boldsymbol{\beta}}_{j-1}| \leq c \cdot \max(|\widehat{\boldsymbol{\beta}}_j|, |\widehat{\boldsymbol{\beta}}_{j-1}|)$, where we typically choose $c = 10^{-3}$), or if the maximum number of iterations n has been reached (set to 5). In the very first iteration, $\widehat{\boldsymbol{\beta}}_j$ is compared to the initial LLS estimate from step 1). If convergence has been reached, continue with step 3). Otherwise, return to step a).
- 3) Transform the variables according to Eq. [4.10] using the estimated parameters $\widehat{\boldsymbol{\beta}}$ from step 2).
- 4) Perform an initial LLS fit on the transformed parameters \mathbf{X}^* and \mathbf{y}^* .
- 5) Compute a robust estimate for homoscedastic regression based on the transformed variables using IRLS with $\mathbf{e}^* = \mathbf{y}^* - \mathbf{X}^* \widehat{\boldsymbol{\beta}}$, similar to steps 2a) through 2e). To check for convergence in the first iteration, $\widehat{\boldsymbol{\beta}}_j$ is compared to the initial LLS estimate from step 4).
- 6) Steps 2) to 5) can be iterated (the ‘outer loop’ in Fig. 4.1) until convergence is reached. Here, the estimated parameters $\widehat{\boldsymbol{\beta}}_k$ are compared to the estimates of the previous iteration in this outer loop $\widehat{\boldsymbol{\beta}}_{k-1}$ (maximum amount of iterations m set to 20). In the very first iteration, $\widehat{\boldsymbol{\beta}}_k$ is compared to the initial LLS estimates from step 1). If convergence has been reached, continue with step 7). Otherwise, return to step 2). The final improved fit $\widehat{\boldsymbol{\beta}}$ from step 5) is then used to recalculate the residuals and weights in step 2).
- 7) Based on the residuals $\mathbf{e}^* = \mathbf{y}^* - \mathbf{X}^* \widehat{\boldsymbol{\beta}}$, accept points for final fit when $|e_i^*| < \kappa \cdot 1.4826 \cdot \text{MAD}(\mathbf{e}^*)$. Here, κ is a trade-off parameter that balances the specificity and sensitivity of the outlier rejection step. When residuals lie outside this interval, we exclude them as outliers. Note that our outlier detection criterion does not involve prior knowledge on the SD of the data.
- 8) Do the final fit on \mathbf{X} and \mathbf{y} , only including the points selected in step 7). Here, we use the multistep WLLS approach in which the weights are initialized by preliminary estimates of $S^2(b, \mathbf{g})$, as predicted by the LLS estimate of $\boldsymbol{\beta}$ and the weights are updated according to $\mathbf{w} = \exp(2\mathbf{X}\widehat{\boldsymbol{\beta}})$ (Veraart et al., 2013b) (checking for convergence is done as described in step 2e). This approach, which we will refer to as iteratively WLLS (IWLLS) in analogy to Veraart et al. (2013), was shown to have high accuracy and precision and can even be preferred over nonlinear estimation methods as it also circumvents the need to estimate the noise level (Veraart et al., 2013b).

In addition, we implemented a RESTORE-based extension using nonlinear fitting throughout IRLS, which we will refer to as RESTORE-DKI. Here, the residuals were observed in the nonlinear domain: $\mathbf{e} = \mathbf{z} - \exp(\mathbf{X}\widehat{\boldsymbol{\beta}})$, and the weights and outlier detection were based on these

residuals. For nonlinear fitting the Levenberg-Marquardt algorithm is used, initialized by the LLS estimates.

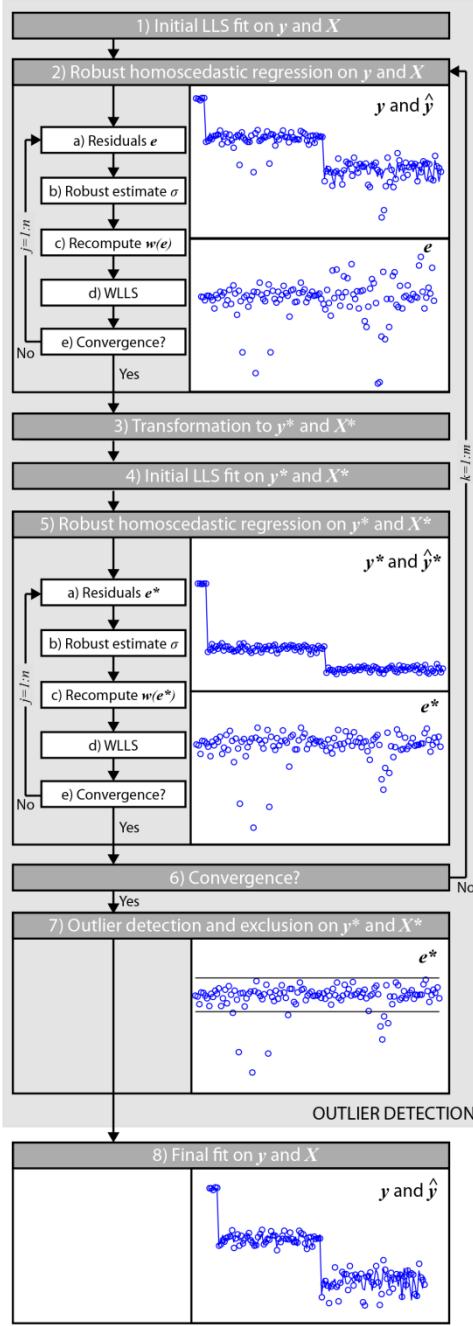


Fig. 4.1: Flow diagram of the REKINDE algorithm. 1) Initial LLS estimate. 2) Homoscedastic regression is performed on the log-transformed signal that includes outliers. Visually, the variance of the residuals $e = y - \hat{y}$ is higher for the on average lower DW signals of the high b-shell. Here, σ represents the standard deviation of the error term and $w(e)$ the weights computed based on the residuals. After transformation of the variables y and X to y^* and X^* (step 3), the residuals appear to have similar variance, independent of signal intensity (step 5). 6) This process is iterated until convergence is reached. 7) Outliers are rejected based on the transformed variables. 8) Fit is performed on the remaining data points.

4.2.2 Simulation experiments

We performed Monte Carlo simulations (10 000 trials) to evaluate the performance of outlier detection, and to investigate the influence of artifacts on the estimated DT and KT and their derived metrics. The DT and KT from a voxel of the corpus callosum (CC) and a deep grey matter (GM) nucleus were extracted from the real human dataset (Fig. 4.2, see Section Real data experiments). In the CC voxel, the FA, MD, MK, kurtosis anisotropy (KA), and radial kurtosis (RK) (Poot et al., 2010) were 0.78, $0.90 \cdot 10^{-3} \text{mm}^2/\text{s}$, 1.1, 0.47, and 2.0 respectively, and in the GM voxel, these values were 0.14, $0.89 \cdot 10^{-3} \text{mm}^2/\text{s}$, 0.60, 0.086, and 0.55 respectively. Two shells of diffusion signals were simulated based on these DT and KT data with b -values 1200 and 2500 s/mm^2 (60 isotropically distributed directions on the hemisphere for each shell), together with five $b = 0$ images. To show the general applicability of our framework, we have additionally simulated a set consisting of 15 directions per shell and b -values ranging from 500 to 2500 s/mm^2 in increments of 500 s/mm^2 , along with five $b = 0$ images (the results for these simulations can be found in supplementary Figs. S4.4 and S4.5). Data points were corrupted by either decreasing, or increasing, or randomly decreasing and increasing the signal intensity with 50% ($b = 0$ points excluded). Noise was added with SNR defined on the $b = 0$ image, using the Rician data distribution.

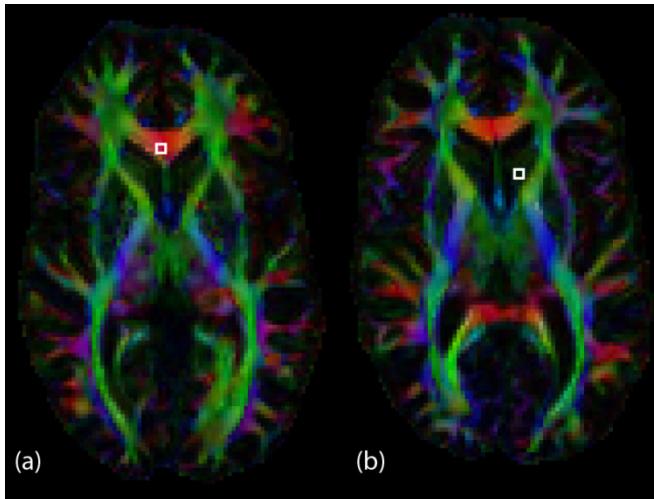


Fig. 4.2: Location of the CC voxel (a) and GM voxel (b) from which the DT and KT were selected for further use in simulations.

4.2.2.1 Simulation experiment 1

The sensitivity and specificity of the REKINDLE outlier detection step were evaluated for the CC voxel as a function of the parameter κ for varying percentages of outliers. Here, sensitivity is defined as the number of true detected outliers (true positives) divided by the total simulated amount of outliers (true positives + false negatives), and specificity is defined as the number of true non-outliers (true negatives) divided by the total amount of non-outliers (true negatives + false positives). We added noise with an SNR of 35 (see simulation experiment 2 for the dependence on SNR).

4.2.2.2 Simulation experiment 2

The sensitivity and specificity for the CC voxel were evaluated for varying SNR as a function of κ . 6.67% of the data points was corrupted (for the dependence on percentage of outliers see simulation experiment 1).

4.2.2.3 Simulation experiment 3

The median FA, MD, MK, KA and RK computed from the corrupted DT and KT tensors obtained with REKINDLE were evaluated as a function of κ for varying percentages of outliers (SNR = 35).

4.2.2.4 Simulation experiment 4

We evaluated the distributions of FA, MD, MK, KA, and RK computed from the corrupted DT and KT tensors obtained with IWLLS, REKINDLE, and RESTORE-DKI approaches. For the robust approaches, κ was set to 3. We varied the number of corrupted points (signal decrease), and added noise with an SNR of 35.

4.2.3 Real data experiment

4.2.3.1 Data acquisition

A healthy volunteer (male, 25 y) was scanned on a 3.0 T Philips Achieva MR scanner (Philips, Best, NL) equipped with an 8-channel receiver head coil. The subject gave informed consent to participate in this study under a protocol approved by the University Medical Center Utrecht ethics board.

DW images were acquired using a single-shot spin-echo echo planar imaging (EPI) sequence with an acquisition matrix of 112×112 , field of view of $224 \times 224 \text{ mm}^2$, and 70 slices with thickness 2 mm resulting in isotropic voxels of $2 \times 2 \times 2 \text{ mm}^3$. Other parameters were SENSE acceleration factor of 2 and TE/TR = $73/6718 \text{ ms}$. Diffusion sensitizing gradients were applied in 60 directions uniformly distributed over the hemisphere (Jones et al., 1999) with a b -value of 1200 and 2500 s/mm^2 (total of 120 DWIs, number of signal averages (NSA) = 1) along with one $b = 0$ (NSA = 6, magnitude averaged on the scanner) image.

The DW scans were corrected for subject motion and eddy current induced geometric distortions including the required B-matrix adjustments (Leemans and Jones, 2009). Processing and visualization was performed with *ExploreDTI* (Leemans et al., 2009).

A ground-truth DWI dataset was created by anisotropic smoothing of the DW images (Full width at half maximum of 6 mm) before using the IWLLS tensor estimation to obtain the DT and KT values. From the DT and KT tensors, the DW signals $S(b, \mathbf{g})$ were recalculated.

4.2.3.2 Artifact simulation

An interleaved motion artifact was simulated by decreasing the signal intensities in 10% of the DW images with 50%. Such an artifact is common when a subject moves during EPI, in which the even/uneven slices are collected sequentially. Rician distributed noise was added with SNR

35. The DT and KT were estimated from this corrupted dataset using IWLLS, REKINDLE, and RESTORE-DKI.

4.3 RESULTS

4.3.1 *Simulation experiments*

4.3.1.1 Simulation experiment 1

In Fig. 4.3, the sensitivity and specificity of the outlier detection step using REKINDLE is shown in the CC voxel as a function of $\kappa = [6, 4, 3, 2, 1, 0.5, 0.25]$. The colored lines represent different percentages of simulated outliers. In general, at a lower value of κ , more and more points are excluded, leading to an increase of true positives and therefore sensitivity. However, when one is too strict in setting the value of κ , an increasing number of points is excluded wrongly (more false positives and decreasing specificity). With an increasing percentage of outliers, specificity and sensitivity both slightly drop. The outlier detection performs best in terms of sensitivity and specificity in the case of signal decrease. Supplementary Fig. S4.1 shows the same results for the GM voxel. Fig. S4.4 shows results for the CC voxel, but now for the 15 direction, five shell protocol.

4.3.1.2 Simulation experiment 2

In Fig. 4.4, the sensitivity and specificity of the REKINDLE outlier detection step is shown for the CC voxel, as a function of $\kappa = [6, 4, 3, 2, 1, 0.5, 0.25]$ for different SNR values (colored lines). When the SNR drops, it becomes harder to detect outliers which results in lower specificity and sensitivity. The highest sensitivity and specificity can be obtained in the case of signal decrease only. Supplementary Fig. S4.2 shows the same results for the GM voxel.

4.3.1.3 Simulation experiment 3

In Fig. 4.5, the median MD, FA, MK, KA, and RK of the CC voxel are shown as a function of percentage corrupted points, for different $\kappa = [6, 3, 1, 0.25]$ (colored lines). The dashed lines represent the simulated values of the diffusion measures in absence of corrupted points. As the percentage of outliers increases, the median value starts to deviate from the true values. Computed values generally show the smallest deviation when $\kappa = 3$. For a signal decrease, median DT indices show virtually no error up to 15% simulated outliers. In case of both decrease and increase, the deviation of the median value is negligible up to 25% simulated outliers. Median KT indices are in general less robust to outliers for either signal decrease or increase. The errors appear to be the smallest in case of outliers defined by randomly simulated signal increase/decrease. Supplementary Fig. S4.3 shows the same results for the GM voxel.

Less Confusion in Diffusion MRI

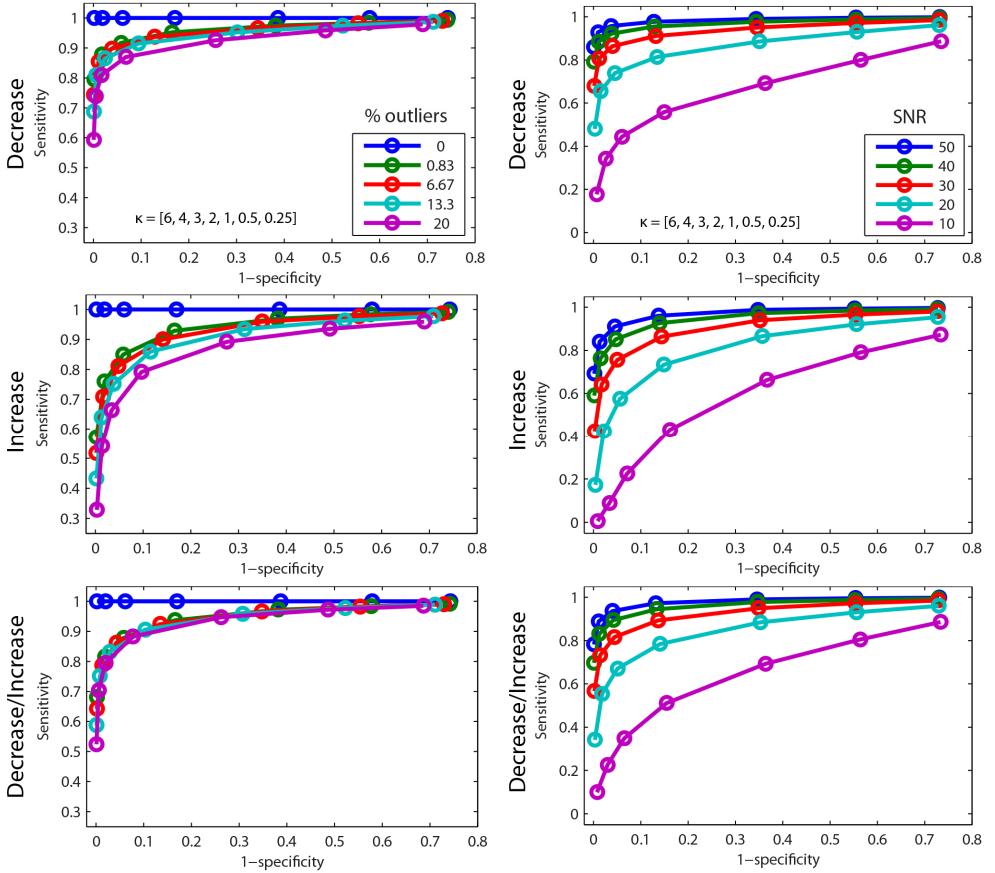


Fig. 4.3: Sensitivity and specificity of the outlier detection step using REKINDLE in the CC voxel, as a function of $\kappa = [6, 4, 3, 2, 1, 0.5, 0.25]$. The colored lines represent different percentages of simulated outliers ($\text{SNR} = 35$). We simulated signal decrease (top), increase (middle) and both decrease and increase randomly (bottom).

4.3.1.4 Simulation experiment 4

Figs. 4.6 and 4.7 show distributions of FA, MD, MK, KA, and RK for the CC and GM voxels, respectively. The distributions were calculated with a varying number of corrupted data points (represented by the colored lines), using IWLLS, REKINDLE, and RESTORE-DKI fitting approaches. For the IWLLS approach, all DT and KT estimated indices are significantly biased in the presence of outliers. Both the REKINDLE and RESTORE-DKI approach are more robust in the presence of outliers, given that the number of corrupted points is not too high. For both the CC and the GM voxel, the curves for MD and FA are almost superimposed. The distributions of the KT indices look noisier and less symmetric than the DT indices. For the CC voxel (Fig. 4.6), the curves for KA and RK almost overlap up to 13.3% outliers. The distribution of MK appears to be most sensitive to outliers. For the GM voxel (Fig. 4.7), the curves for all kurtosis indices almost overlap up to 6.67% outliers. The kurtosis indices start to

Fig. 4.4: Sensitivity and specificity of the outlier detection step using REKINDLE in the CC voxel, as a function of $\kappa = [6, 4, 3, 2, 1, 0.5, 0.25]$. The colored lines represent different SNR values (% outliers = 6.67%). We simulated signal decrease (top), increase (middle) and both decrease and increase randomly (bottom).

get biased when the percentage of outliers is larger than 13.3%, this effect is more pronounced for the GM voxel then for the CC voxel.

The distributions resulting from the robust approaches are broader with slightly lower number of occurrences at peak locations than for the IWLLS case. The specificity of outlier detection for the CC and GM voxels in the case of no simulated outliers is around 94% for $\kappa = 3$ and SNR of 35 (Figs. 4.3 and S4.1). Consequently, there are a few false positives in the case of no simulated outliers, resulting in a slightly lower accuracy. Notice that this effect is more prominent in RESTORE-DKI than in REKINDLE. Fig S4.5 shows the distributions resulting from estimation with REKINDLE for the 15 direction, five shell protocol.

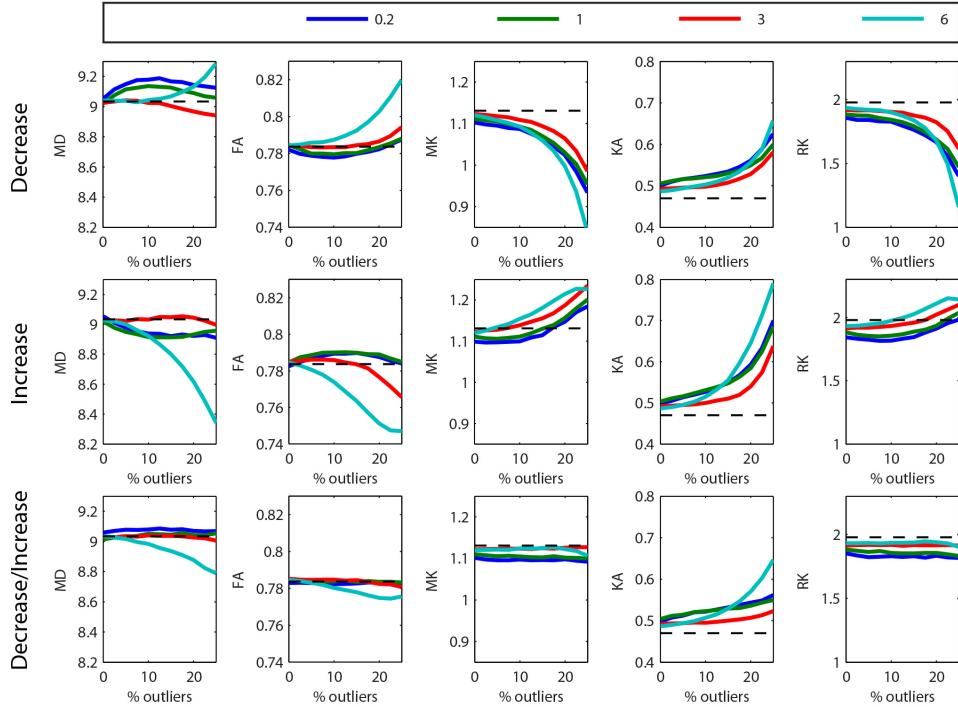


Fig. 4.5: Median MD, FA, MK, KA, and RK as a function of percentage corrupted points in the CC voxel (SNR = 35), for different κ (colored lines). The expected values of the median in absence of corrupted points are indicated by the dashed line. We simulated signal decrease (top), increase (middle) and both decrease and increase randomly (bottom). MD is given in units $10^{-4} \text{ mm}^2 \text{s}^{-1}$.

4.3.2 Real data experiment

Fig. 4.8a displays a sagittal slice of the simulated direction-encoded color FA (DEC) map (top) and the interleaved artifact in the DW images with low (middle) and high (bottom) b-value. In Fig. 4.8b, the motion artifact is clearly visible for the FA, KA and RK maps using IWLLS (similar results were obtained for the other kurtosis/diffusion maps). Both RESTORE-DKI and REKINDLE are able to deal with the artifact by excluding the signal perturbations as outliers.

Less Confusion in Diffusion MRI

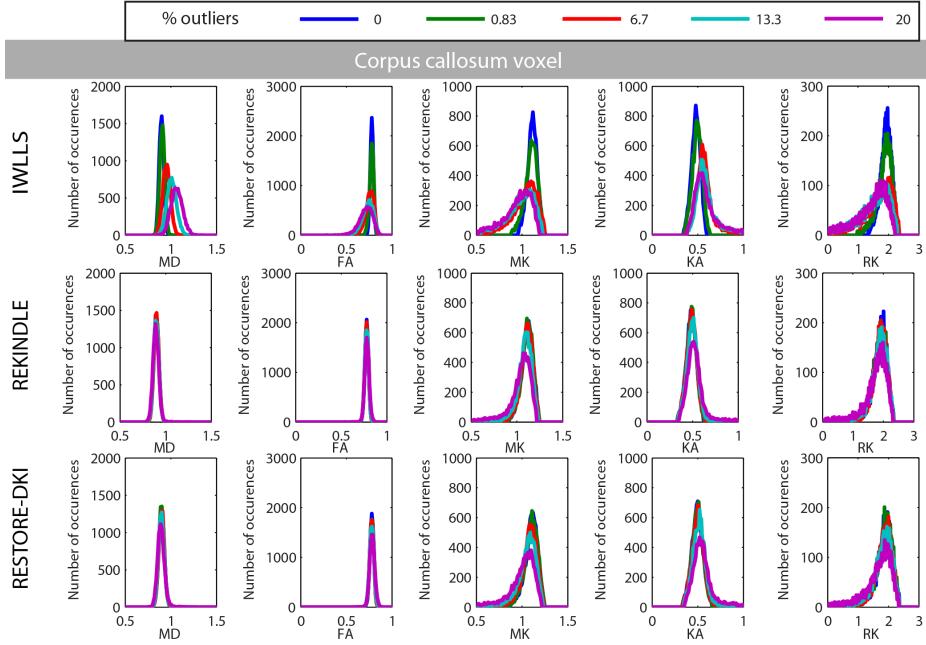


Fig. 4.6: Distributions of FA, MD, MK, KA, and RK for a varying number of corrupted data points (represented by the colored lines), using IWLLS (top), REKINDLE (middle), and RESTORE-DKI (bottom) fitting approaches for the CC voxel. MD is given in $10^{-3} \text{ mm}^2\text{s}^{-1}$.

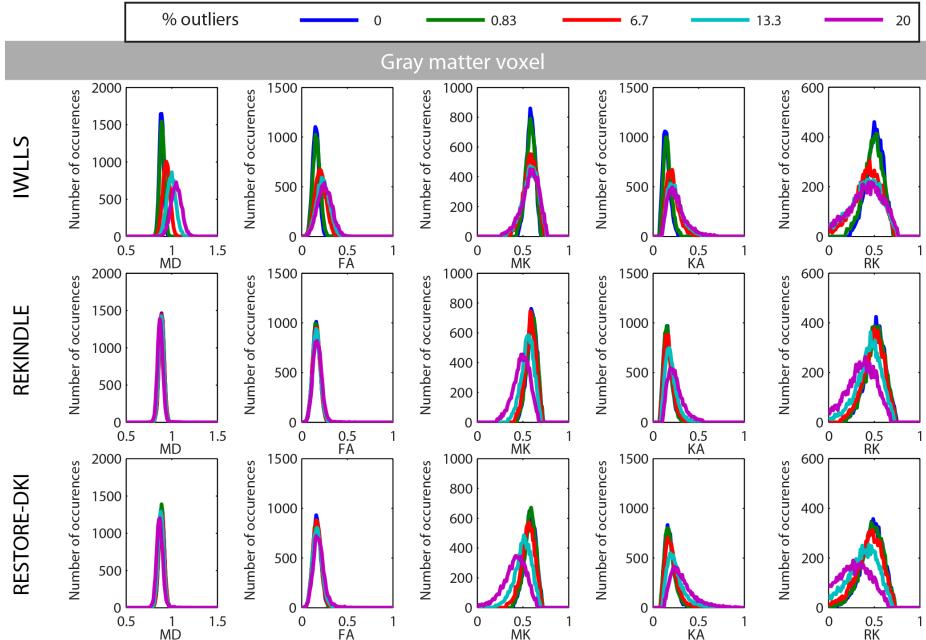


Fig. 4.7: Distributions of FA, MD, MK, KA, and RK for a varying number of corrupted data points (represented by the colored lines), using IWLLS (top), REKINDLE (middle), and RESTORE-DKI (bottom) fitting approaches for the GM voxel. MD is given in $10^{-3} \text{ mm}^2\text{s}^{-1}$.

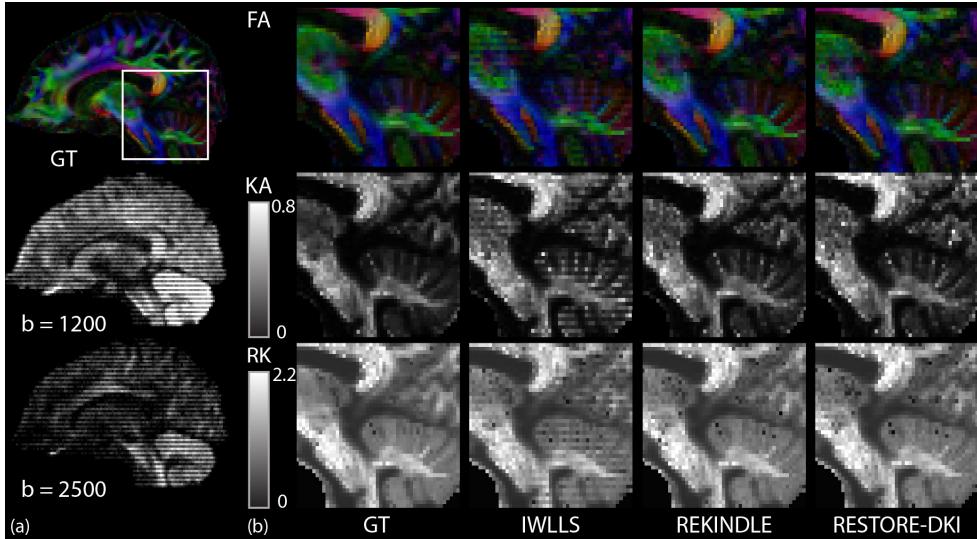


Fig. 4.8: (a) Direction-encoded color FA map (top) and the simulated interleaved artifact in the DW images with low (middle) and high (bottom) b -value. (b) The artifact is clearly visible in the FA, KA and RK maps for the IWLLS approach. Both the REKINDLE and RESTORE-DKI approach are robust towards outliers.

4.4 DISCUSSION

In a clinical setting, artifacts in DW-MRI are common and may affect DKI measures even more than DTI metrics due to prolonged scanning times and higher b -values. Although DKI will clearly benefit from robust estimation procedures, thorough investigations in this context have not been reported previously. REKINDLE is a linearized estimation framework that is robust to outliers. REKINDLE is also applicable to other linear diffusion models (spherical harmonic fitting) or models that can be linearized (DTI). The performance of our estimation framework will depend on the choice of model, for example the tensor model is expected to perform worse in terms of outlier rejection at higher b -values (Pannek et al., 2012). We will discuss the benefits and challenges of linearization, as well as the modifications with respect to the conventional RESTORE.

4.4.1 Benefits of linearization

The DTI and DKI models have in common that they can be written in the same form (Eq. [4.2]), which allows for direct extension of RESTORE to DKI. The long computations times, increased model complexity of DKI, and the potential strengths of linear estimators (Salvador et al., 2005; Veraart et al., 2013b) motivated us to develop a linearized version of the algorithm. REKINDLE is, compared with RESTORE-DKI, roughly ten times faster, and will thus be a valuable method if results have to be displayed shortly after acquisition, such as in a clinical setting. RESTORE-DKI uses NLS fitting throughout the pipeline, which is an iterative and computationally expensive procedure. In every iteration of the IRLS this NLS fitting is repeated, resulting in long computation times. It might even be beneficial, instead, to directly

solve the minimization problem associated with the robust estimator, that is, directly solve $\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta}} \left\{ \sum_{i=1}^n \rho \left(\frac{e_i(\boldsymbol{\beta})}{\sigma} \right) \right\}$.

NLS fitting approaches are susceptible to local minima, leading to inaccurate parameter estimates. This problem may occur even more often in DKI because of the increased model complexity and lower effective SNR in the high b -value shells. Chang et al. (2005) reported that they sometimes encountered such adverse effects in the periphery of the brain in DTI (Chang et al., 2005).

For the final fit, we use the IWLLS estimation as it may be preferred over NLS, depending on the number of DW signals and the parameters to be estimated (Landman et al., 2007b; Veraart et al., 2013b). Linear estimators are unbiased under certain conditions and have potentially high accuracy and precision. The precision of WLLS estimators is increased by accounting for the heteroscedasticity after the log-transformation, preferably by updating the weights iteratively (Veraart et al., 2013b).

4.4.2 Challenges of linearization

It is well known that the variance of the log-transformed signal depends on the signal intensity. Fig. 4.1 step 2) shows that the dispersion of the residuals increases with decreasing signal intensity (in general, the average intensity of the high b -shell is lower and the dispersion is larger). This has to be taken into account when detecting the outliers (see also the Supplementary Information). If we would have considered the data to fulfill homoscedasticity, and thus stop after step 2) in our pipeline, the robust estimate of σ for outlier rejection as in Eq. [4.14] would be some average of the standard deviations of each residual. Consequently, the probability of wrongly excluding a measurement is larger for low signal intensities. The specificity of outlier detection in that case (at 0% simulated outliers) would be lower for comparable κ (around 82% for a GM voxel, see Supplementary Information). In the current implementation, the assumption of homoscedasticity is only used to obtain a first robust fit in step 2) to transform the variables in step 3). The final detection of outliers is performed on these transformed variables in step 6) after robust fitting in step 5). For robust fitting we used IRLS. For the interested reader, we would like to refer to other implementations of robust regression in heteroscedastic linear models (Carroll and Ruppert, 1982; Giltinan et al., 1986; Hekimoglu and Berber, 2003).

Transformation of the variables in this work was based on the theoretical relation between σ_i and the signal intensity $S(b_i, \mathbf{g}_i)$. This relation is valid on the condition that the SNR is high enough and that the data are Rician distributed. Owing to, inter alia, parallel imaging techniques, these conditions are not always fulfilled (Aja-Fernandez and Tristan-Vega, 2012). Alternatively, one could perform a robust fit to estimate the parameter λ in the model for $\sigma_i = h(\lambda, \mathbf{x}_i^T \boldsymbol{\beta})$ (Maronna et al., 2006). For computational reasons we did not implement this approach here.

4.4.3 Other modifications

4.4.3.1 Weight function

We used a slightly different weight function than RESTORE. We re-derived this function from the Geman-McClure M-estimator, which resulted in an additional square in the denominator (see Appendix B4). Although the weight function in Fig. A4.1 is slightly sharper than without square, this did not yield significant differences in the results. Though we have chosen to use this estimator to compare our work to previous work (Chang et al., 2005; Mangin et al., 2002), there are many robust estimators in the literature – each providing its own trade-off between sensitivity and specificity of outlier detection.

4.4.3.2 Outlier rejection criterion

Robust estimation procedures such as RESTORE require knowledge of the signal variance. Correct estimation of this parameter is critical to their performance. There are various noise estimation methods (Aja-Fernandez et al., 2009), which use either background areas (Chang et al., 2005; Henkelman, 1985; Koay et al., 2009b; Sijbers et al., 2007), or the image object itself (Chang et al., 2012; Coupé et al., 2010; Koay and Basser, 2006; Landman et al., 2009; Manjón et al., 2010; Maximov et al., 2012; Sijbers et al., 1998; Veraart et al., 2013a). In addition to the potential inconsistency in noise variance estimates using different methods, noise estimation is often complicated by suppression of the background by the scanner, spatial variance of noise, movement of the subject, user subjectivity, parallel imaging techniques, and preprocessing techniques. The REKINDLE framework uses an outlier detection criterion that does not depend on the a priori determined signal variance. Instead, it uses the MAD as robust dispersion estimate of the random error in homoscedastic regression of the transformed parameters \mathbf{y}^* and \mathbf{X}^* . This means that the threshold value for a residual to be classified as outlier differs for each individual voxel and is based only on the residuals in that voxel, potentially allowing spatial variation in signal variance since the criterion is not dependent on one a priori determined value for the signal SD. However, if a voxel has too many outliers, the MAD will eventually be biased (it can handle 50% incorrect observations before giving an incorrect result). Other dispersion measures exist (e.g. Rousseeuw and Croux, 1993) that are potentially more robust, efficient, or capable of handling non-symmetric distributions. Here, we have chosen the MAD for its ease of computation, and the fact that we expect the error term to approximate a Gaussian distribution at sufficiently high SNR values (Salvador et al., 2005).

4.4.3.3 Goodness of fit criterion

We do not perform a goodness-of-fit criterion prior to robust estimation for the reason that we prefer not to be biased to any a priori (ad hoc) criterion that potentially relies on invalid assumptions. The consequence is that all voxels are evaluated with the robust estimation framework, which increases computation time.

4.4.3.4 User parameter κ

The parameter κ makes a trade-off between sensitivity and specificity. Lowering the value of κ will result in more points being excluded, which will lead to an increase of true positives and therefore sensitivity. However, if the value of κ is set too low, an increasing number of points is

excluded incorrectly (decreasing specificity). In our experiments, we found that a value of $\kappa = 3$ is a good trade-off between sensitivity and specificity both in WM and GM for the two shell protocol (Figs. 4.3-4.5 and S4.1-S4.3). REKINDLE is then robust to signal perturbations in up to 10% of the DWIs (Figs. 4.5 and S4.3). In real data, REKINDLE was able to deal with the interleaving artifact by excluding the signal perturbations as outliers with $\kappa = 3$ (Fig. 4.8).

4.4.4 Limitations and potential future improvements

The ROC curves show that the specificity of outlier detection is not 100% at $\kappa = 3$, even when we simulate 0% outliers (Figs. 4.3, S4.1, S4.4). One should note, however, that outliers may be a result of variability in the measurement or an experimental error, and that these are often mixed (Chang et al., 2005). In our algorithm we define an ‘outlier’ as a data point having a residual larger than $\kappa \cdot 1.4826 \cdot \text{MAD}$, but we do not distinguish between ‘outliers due to measurement variability’ (which we indicate here with outliers resulting from adding Rician noise), and ‘outliers due to experimental error’ (which are our simulated outliers). We only know the simulated outliers, whereas outliers that are caused by adding Rician noise may be counted as false positives in some cases. The effect of no 100% specificity at 0% outliers may thus be less pronounced than the ROC curves may suggest, but it is present: in fact, Figs. 4.6 and 4.7 show that the distributions resulting from REKINDLE get slightly broader compared with the IWLLS distributions at 0% outliers, likely owing to the exclusion of good data points. If, for some reason, the user does not expect many outliers, the convergence criterion can be adapted or the value of κ can be adjusted to balance between sensitivity and specificity.

The ability of REKINDLE to identify outliers depends on the type of outlier (signal decrease/increase), the percentage of outliers in the data (i.e., the ‘oversampling’ of good points) and the acquisition scheme, but also on SNR and the underlying tissue configuration. The method shows the least bias in DT and KT estimates with increasing percentage of outliers in case of combined artificial signal increase/decrease (Figs. 4.5 and S4.3). For a signal dropout in DW signals in a CC region, median DT indices show virtually no error up to 15% simulated outliers. Median KT indices are in general less robust to outliers (Fig. 4.5). The percentage of outliers that can be tolerated is dependent on the gradient sampling scheme (Chang et al., 2005), as the exclusion of too many points in a particular direction may result in an imbalanced B-matrix. Figs. S4.4 and S4.5 show results for the protocol of 15 directions and b -values ranging from 500 to 2500 s/mm^2 in increments of 500 s/mm^2 , along with five $b = 0$ images. For a low percentage of outliers (< 13.3%), REKINDLE is still robust towards outliers when looking at the indices (Fig. S4.5), but with a slightly lower sensitivity (Fig. S4.4) compared to the 60 direction, two shell protocol.

The exclusion of too many points may result in an ill-defined B-matrix and could occur in acquisitions with only few DW gradient directions. In such cases, however, one could be less strict in defining the κ parameter (i.e., increase the value of κ), resulting in fitting on more data points (with the risk of including outliers). Recently, Chang et al. (2012) added constraints to their RESTORE algorithm, setting criteria for the exclusion of data points (Chang et al., 2012). These constraints could also be added to the REKINDLE framework in future work, which would be particularly relevant for low-redundancy DW datasets. Furthermore, Chang et al. also proposed an approach called informed RESTORE (iRESTORE) that uses the fact that physiological noise is more likely to result in signal dropouts (Chang et al., 2012). At the end of

each iteration, they exclude the data point with the maximum negative residual, which would correspond to the largest signal dropout from what is expected based on the fit. A signal dropout, however, can theoretically also result in positive residuals, for example when the fit is severely affected by this outlier. To keep our algorithm general to any type of outlier (decrease/increase), we did not include this approach. Cardiac pulsation and patient movement, for example, can result in signal dropout, while EPI distortions and eddy currents may lead to either signal dropout (stretch) or signal increase (pile up). Insufficient fat suppression can lead to a ghost resulting in a local signal increase. Residual errors in motion/eddy current/EPI distortion correction in post-processing can lead to artificial signal decreases or increases. Furthermore, there are system related artifacts, such as spike noise, temporal scanner instabilities, and thermal noise that can also result in signal increasing and decreasing outliers. As outlier detection and exclusion is performed voxel-wise, REKINDLE will be mainly beneficial in artifacts that cause a decrease or increase of signal which is inconsistent over the diffusion weighting volumes, e.g. subject motion, cardiac pulsation, eddy current distortions and noise. In future work, REKINDLE could be extended with the ideas of iRESTORE.

Noise and artifacts may cause the tensor estimates and derived indices to be biologically or physically implausible. The signal $S(b, \mathbf{g})$ should be a monotonically decreasing function of b -value in biological tissues, and directional diffusivities and kurtoses should typically be positive and within a certain range (Jensen et al., 2005). In our real data set these implausible signals occur for example in the cerebellum, causing the ‘black spots’ in the RK map. These are indications of implausible signals rather than disease. To be able to deal with these implausible signals, constrained counterparts of LS approaches have been developed in DTI and DKI estimation (Koay et al., 2006; Tabesh et al., 2011a). Such a constraint approach could be incorporated in REKINDLE, for example in the final fit.

4.4.5 Robust estimation in population analysis

Artifacts that cause outliers will likely have an effect on the outcome of a population analysis, manifesting itself as increased population variability or bias in diffusion metrics. Though Habib et al. (2010) suggest that DWI artifacts from cardiac pulsation has negligible effects in DTI group studies, Walker et al. (2011) and Pannek et al. (2012) did observe changes in population mean and variability when outliers were not removed by means of a robust procedure, deeming artifact removal crucial in population studies. The results in population studies can be biased by artifacts owing to increased involuntary head motion or altered cardiac pulsatile motion in patient populations compared to healthy controls (Pannek et al., 2012). Outlier rejection probability (ORP) maps indicate regions that are often affected by outliers, such as regions prone to susceptibility artifacts and cardiac pulsation, and can be used to attribute potential population differences to outliers rather than to the disease investigated. These regions are, however, variable within the population, and robust procedures can help to reduce spurious findings due to artifacts in population analysis. Szczerpankiewicz et al. (2013) suggest that between-individual variability dominates strongly over noise-induced variability in DKI group studies, but the statistical power is dependent on location and may be confounded by artifacts. Our simulation and real data experiments confirm that there is a bias in DT and KT indices in the presence of outliers when using non-robust approaches (Figs. 4.6 and 4.7). Though Fig. S4.3 suggests that REKINDLE potentially introduces a bias in the GM voxel, Fig. S4.6 shows

that other (non-robust) methods also introduce a bias for this GM voxel, even at 0% outliers. REKINDLE is the most robust approach when outliers are present. In future work, an additional validation study in a population can be done to evaluate the performance of REKINDLE in group studies.

4.5 CONCLUSION

DKI may provide new avenues for more complete tissue characterization within clinically feasible scanning times, but suffers from an increased sensitivity to artifacts due to increased model complexity and higher acquisition demands. We have investigated the performance of RESTORE-DKI and the linearized version REKINDLE in the presence of outliers. REKINDLE can estimate the diffusion and kurtosis tensors robustly at low computational cost, and will therefore be valuable for robust DKI tensor estimation in clinical applications. While REKINDLE has been developed for DKI, it is by design also applicable to DTI and other diffusion models that can be linearized.

ACKNOWLEDGEMENTS

The authors thank Jelle Veraart for valuable discussions on (robust) estimators and Sjoerd Vos for critically reviewing parts of the manuscript.

4.6 APPENDICES

4.6.1 Appendix A4: Modeling heteroscedasticity after linearization with scaling constants σ_i

After applying the log-transform to Eq. [4.2], one obtains:

$$\begin{aligned} \log(z_i) &= \log(\exp(\mathbf{x}_i^T \boldsymbol{\beta}) + \epsilon_i) \\ &= \log(\exp(\mathbf{x}_i^T \boldsymbol{\beta})(1 + \exp(-\mathbf{x}_i^T \boldsymbol{\beta})\epsilon_i)) \\ &= \log(\exp(\mathbf{x}_i^T \boldsymbol{\beta})) + \log(1 + \exp(-\mathbf{x}_i^T \boldsymbol{\beta})\epsilon_i) \\ &= \mathbf{x}_i^T \boldsymbol{\beta} + \log(1 + \exp(-\mathbf{x}_i^T \boldsymbol{\beta})\epsilon_i) \\ &= \mathbf{x}_i^T \boldsymbol{\beta} + \sigma_i \epsilon_i. \end{aligned} \quad [\text{A4.1}]$$

Hence in general, the expression for σ_i is

$$\sigma_i = \frac{\log(1 + \exp(-\mathbf{x}_i^T \boldsymbol{\beta})\epsilon_i)}{\epsilon_i}. \quad [\text{A4.2}]$$

Under the condition of high SNR ($\exp(-\mathbf{x}_i^T \boldsymbol{\beta})\epsilon_i \ll 1$) and by expanding the logarithm ($\log(1 + x) \approx x$) we have

$$\sigma_i \approx \exp(-\mathbf{x}_i^T \boldsymbol{\beta}). \quad [\text{A4.3}]$$

4.6.2 Appendix B4: Robust estimators and IRLS

Regression M-estimates $\hat{\boldsymbol{\beta}}$ are defined as solutions of the minimization problem

$$\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta}} \left\{ \sum_{i=1}^n \rho \left(\frac{e_i(\boldsymbol{\beta})}{\sigma} \right) \right\}, \quad [\text{A4.4}]$$

where $\rho(x)$ is a symmetric, positive definite function having a unique minimum at $x = 0$, and σ is a scale parameter. When differentiating with respect to $\boldsymbol{\beta}$ one finds the analog to the normal equations:

$$\sum_{i=1}^n \psi \left(\frac{e_i(\hat{\boldsymbol{\beta}})}{\sigma} \right) \mathbf{x}_i = 0. \quad [\text{A4.5}]$$

Here, $\psi(x) = \frac{d\rho}{dx}$ is the influence function, measuring the influence of a data point on the value of the parameter estimate. By defining a weight function $w(x) = \psi(x)/x$, Eq. [A4.5] becomes

$$\sum_{i=1}^n w \left(\frac{e_i(\hat{\boldsymbol{\beta}})}{\sigma} \right) \mathbf{x}_i e_i(\hat{\boldsymbol{\beta}}) = 0. \quad [\text{A4.6}]$$

This is equivalent to the minimization problem

$$\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta}} \left\{ \sum_{i=1}^n w \left(\frac{e_i(\hat{\boldsymbol{\beta}})}{\sigma} \right) e_i^2(\hat{\boldsymbol{\beta}}) \right\}, \quad [\text{A4.7}]$$

which can be solved by IRLS (Hogg, 1979; Holland and Welsch, 1977; Street et al., 1988). For LS regression one defines $\rho(x) = x^2/2$, $\psi(x) = x$, and $w(x) = 1$, attributing equal weights to all residuals, see Fig. A4.1a. The scale σ will only be reflected as a constant factor in Eqs. [A4.4] to [A4.7] in this case, and therefore it is sufficient to minimize $\rho(e_i^2)$.

The Geman-McClure M-estimator

$$\rho(x) = \frac{x^2/2}{x^2+1} \quad [\text{A4.8}]$$

flattens out at large x , see Fig. A4.1b. The influence and weight functions are, respectively:

$$\psi(x) = \frac{x}{(x^2+1)^2}, \quad [\text{A4.9}]$$

and

$$w(x) = \frac{1}{(x^2+1)^2}. \quad [\text{A4.10}]$$

See Fig. A4.1b for their shapes for different values of σ .

Less Confusion in Diffusion MRI

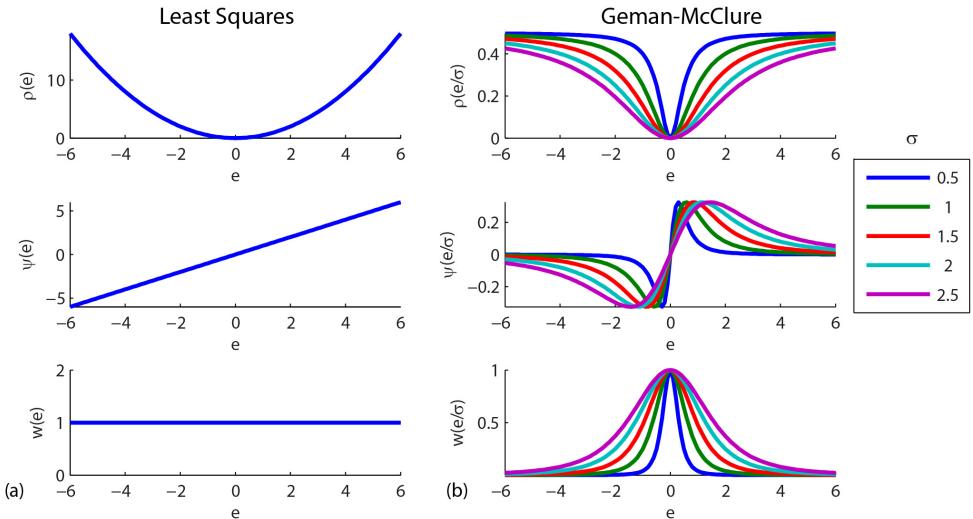


Fig. A4.1: (a) Least squares estimator and its influence and weight function. (b) Geman-McClure estimator and its influence and weight function for different values of σ .

4.7 SUPPLEMENTARY MATERIAL

4.7.1 Supplementary figures

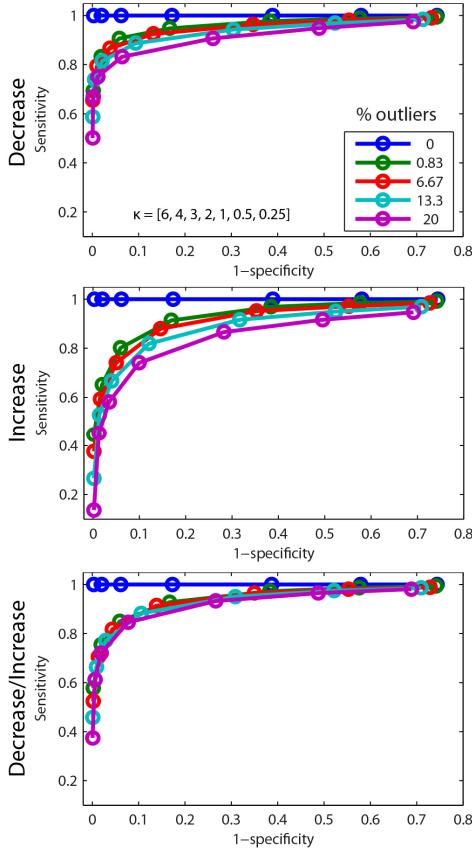


Fig. S4.1: Sensitivity and specificity of the outlier detection step using REKINDLE in the GM voxel, as a function of $\kappa = [6, 4, 3, 2, 1, 0.5, 0.25]$. The colored lines represent different percentages of simulated outliers (SNR = 35). We simulated signal decrease (top), increase (middle) and both decrease and increase randomly (bottom).

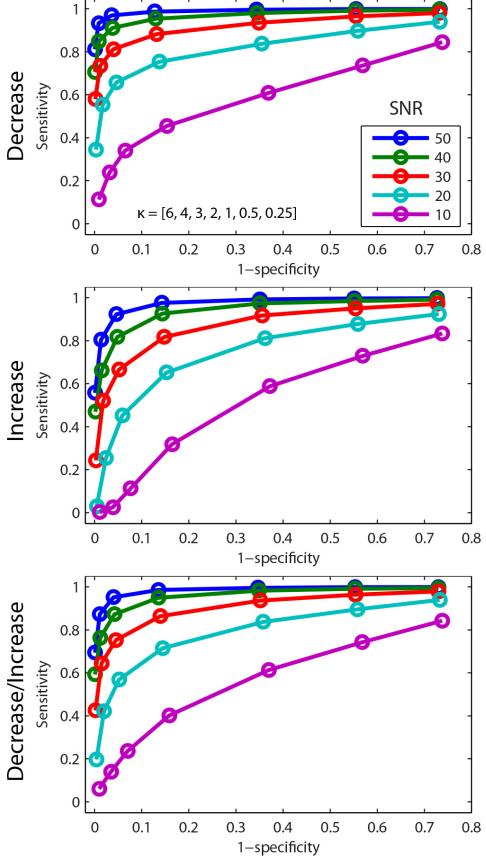


Fig. S4.2: Sensitivity and specificity of the outlier detection step using REKINDLE in the GM voxel, as a function of $\kappa = [6, 4, 3, 2, 1, 0.5, 0.25]$. The colored lines represent different SNR values (% outliers = 6.67%). We simulated signal decrease (top), increase (middle) and both decrease and increase randomly (bottom).

Less Confusion in Diffusion MRI

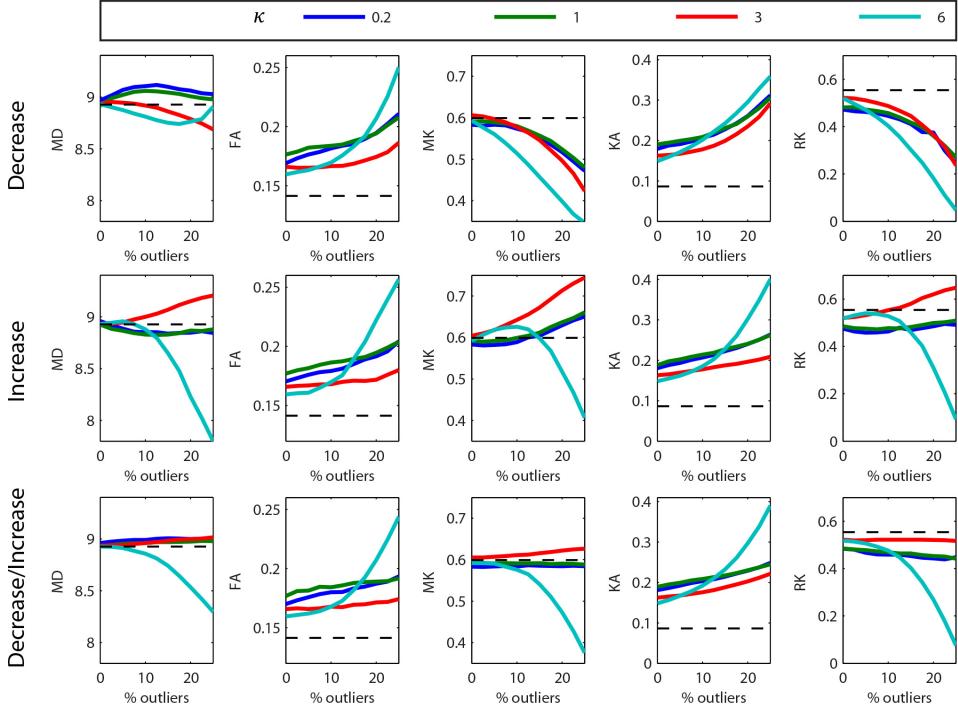


Fig. S4.3: Median MD, FA, MK, KA, and RK as a function of percentage corrupted points in the GM voxel (SNR = 35), for different κ (colored lines). The expected values of the median in absence of corrupted points are indicated by the dashed line. We simulated signal decrease (top), increase (middle) and both decrease and increase randomly (bottom). MD is given in units $10^{-4} \text{ mm}^2 \text{s}^{-1}$.

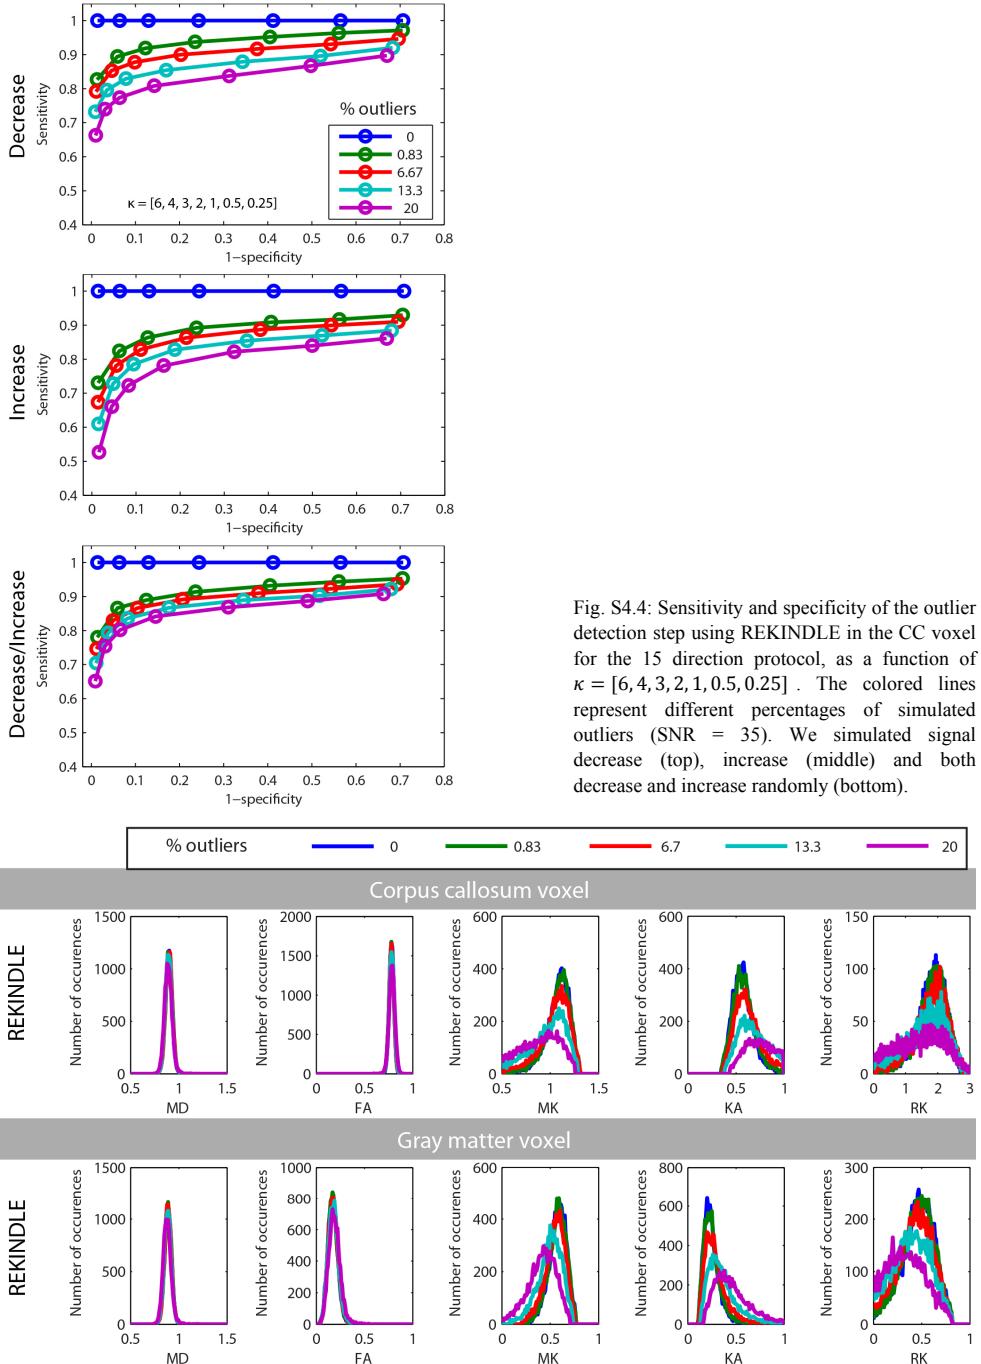


Fig. S4.4: Sensitivity and specificity of the outlier detection step using REKINDLE in the CC voxel for the 15 direction protocol, as a function of $\kappa = [6, 4, 3, 2, 1, 0.5, 0.25]$. The colored lines represent different percentages of simulated outliers (SNR = 35). We simulated signal decrease (top), increase (middle) and both decrease and increase randomly (bottom).

Fig. S4.5: (a) Distributions of FA, MD, MK, KA, and RK for a varying number of corrupted data points (represented by the colored lines), using REKINDLE for the CC voxel for the 15 direction, five shell protocol ($\kappa = 4$). (b) Same distributions for the GM voxel. MD is given in $10^{-3} \text{ mm}^2 \text{s}^{-1}$.

Less Confusion in Diffusion MRI

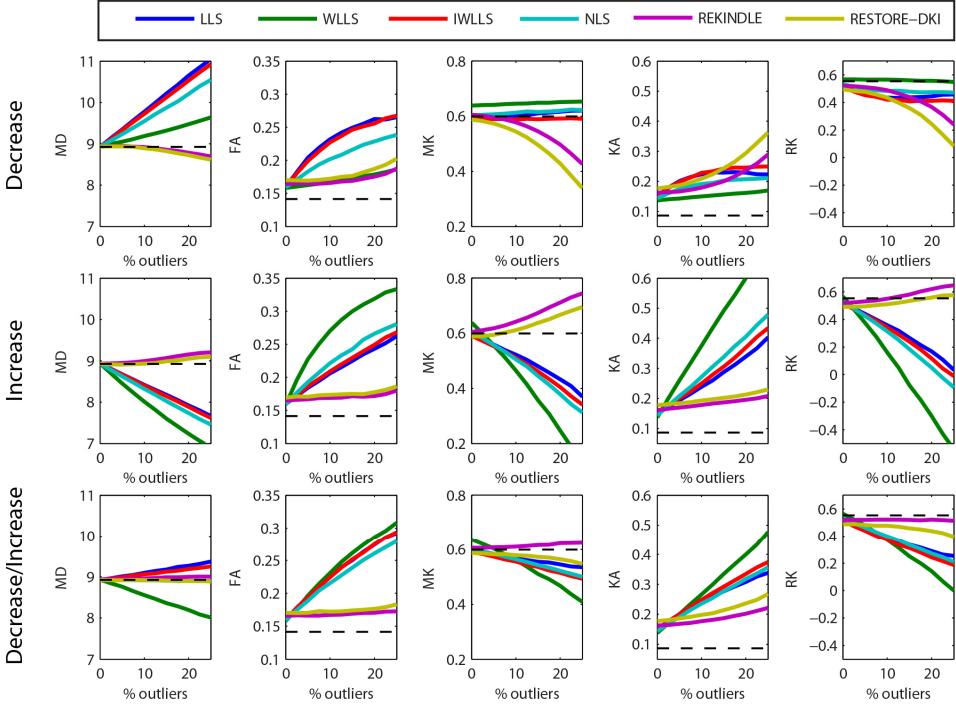


Fig. S4.6: (a) Median MD, FA, MK, KA, and RK as a function of percentage corrupted points in the GM voxel for the 60 directions, two shell protocol (SNR = 35), for different estimation methods (colored lines). The expected values of the median in absence of corrupted points are indicated by the dashed line. We simulated signal decrease (top), increase (middle) and both decrease and increase randomly (bottom). In case of a signal decrease, WLLS seems to be fairly robust, but this is due to the inherent low weights outliers get in case of signal decrease. REKINDLE is more robust than RESTORE-DKI. DT indices can be estimated robustly up to 15 – 20% outliers, and KT indices up to 10% outliers. All methods show a bias in FA and KA, even at 0% outliers. MD is given in units $10^{-4} \text{ mm}^2 \text{ s}^{-1}$.

4.7.2 Supplementary information

To see why it would make a difference that so many fits have to be performed in different domains and why one cannot stop and define outliers on the basis of the first step, we have compared the performance of the single step approach (where we stop after step 2) and do outlier rejection) and our proposed multistep approach based on Maronna et al. (2006) in Figs. S4.7 to S4.8 below.

Since the log-transformed signal intensity is a heteroscedastic variable, we have to take this into account when detecting the outliers. From Salvador et al. (2005), we know that for a lower signal intensity, the variance increases inversely proportional to the true signal intensity. Empirically, this can be seen from the residual plots after initial LLS estimation, where the high b -value signals appear to have greater variance (the outliers are encircled) as in Fig. S4.7a:

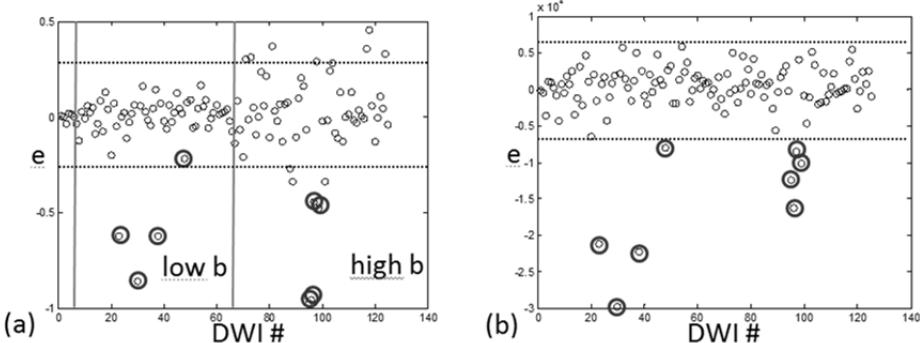


Fig. S4.7: (a) errors and outlier rejection without transformation (based on \mathbf{y} and \mathbf{X}) The horizontal dashed lines indicate $\kappa \cdot 1.4826 \cdot MAD$ with $\kappa = 3$, all points outside the line would be excluded as outliers. Encircled points are true outliers. Vertical lines show separation between low b -shell and high b -shell. (b) errors and outlier rejection with the additional transformation (based on \mathbf{y}^* and \mathbf{X}^*).

If we would just calculate the MAD as a robust measure of signal variance, this would be some averaged value of the different (signal dependent) variances. This possibly results in more false positives, more likely in the high b -shell (see dashed black lines). Therefore, heteroscedasticity should be taken into account before outlier exclusion. It is not recommended, however, to use the signal intensities directly to transform the data (or use them somehow as weights in the regression, but it is unclear how to combine this with the Geman McClure weights in a single iterative loop (Chang et al., 2005)) to correct for heteroscedasticity, since they are influenced by noise (Veraart et al., 2013b) and outliers. Instead, we want to transform both \mathbf{y} and \mathbf{X} based on a robust fit for the scaling constants σ_i . This robust fit is obtained in step 2) via the first iterative loop where we assume the signal to be homoscedastic. We thus estimate a single robust estimate of scale (the MAD) to incorporate in the Geman McClure weights Eq. [4.13], and update the weights iteratively. In the case of heteroscedastic variables, the regression equation is:

$$\begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix} = \begin{pmatrix} x_{11} & x_{12} & \cdots & x_{1p} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \cdots & x_{np} \end{pmatrix} \begin{pmatrix} \beta_1 \\ \vdots \\ \beta_p \end{pmatrix} + \begin{pmatrix} \sigma_1 \\ \vdots \\ \sigma_n \end{pmatrix} \circ \begin{pmatrix} \epsilon_1 \\ \vdots \\ \epsilon_n \end{pmatrix},$$

which gives

$$y_i = x_{i1}\beta_1 + \cdots + x_{ip}\beta_p + \sigma_i\epsilon_i.$$

Dividing both sides by σ_i gives

$$y_i/\sigma_i = x_{i1}\beta_1/\sigma_i + \cdots + x_{ip}\beta_p/\sigma_i + \epsilon_i = \left(\frac{x_i}{\sigma_i}\right)^T \beta + \epsilon_i.$$

This is why we need to transform both \mathbf{X} and \mathbf{y} . After transformation (step 3)), the variable is closer to homoscedastic, and we do another homoscedastic robust fit via the second iterative loop in step 5) (now, the scale can be represented by a single robust estimate such as the MAD). The overarching iterative loop (from step 6) to step 2)) improves the transformation in step 3).

Fig. S4.8a shows the ROC curve when we stop after step 2) and do outlier rejection, while in (b) the ROC curve obtained with our multistep method is shown (Fig. S4.1, GM voxel,

Less Confusion in Diffusion MRI

increase). While sensitivity remains approximately the same for a given specificity < 0.8 in both cases, one can see that if we require higher specificity (less false positives), this comes with a price of rapidly dropping sensitivity when we do not do the transformation and stop after step 2) (see (a)).

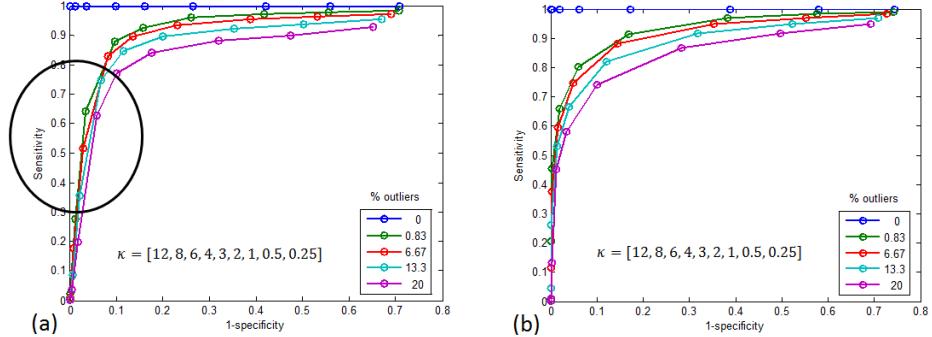


Fig S4.8: ROC curves as a function of κ for a signal decrease (a) without transformation (based on \mathbf{y} and \mathbf{X}). (b) with transformation (based on \mathbf{y}^* and \mathbf{X}^*).

Fig. S4.9a Shows the distributions of the estimated indices if we stop after step 2) and do outlier rejection, while (b) shows the REKINDLE distributions of Fig. 4.6b (GM voxel, decrease). The distributions are significantly lower and slightly broader in (a) compared to (b), probably caused by the greater amount of false positives when not correcting for heteroscedasticity. This difference is the most prominent for the kurtosis measures.

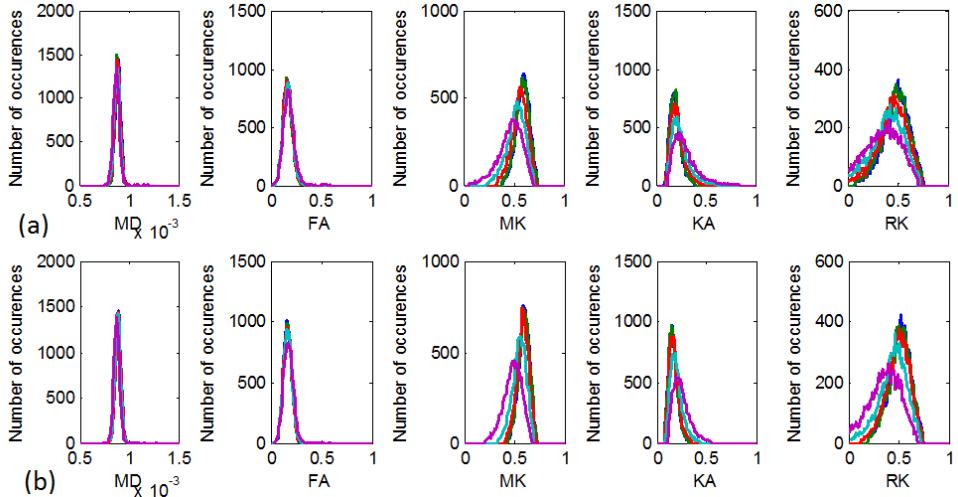


Fig S4.9: Diffusion and kurtosis indices distributions for a signal decrease (a) without transformation (based on \mathbf{y} and \mathbf{X}). (b) with transformation (based on \mathbf{y}^* and \mathbf{X}^*).

5 RECURSIVE CALIBRATION OF THE FIBER RESPONSE FUNCTION FOR SPHERICAL DECONVOLUTION OF DIFFUSION MRI DATA

Based on:

C.M.W. Tax, B. Jeurissen, S.B. Vos, M.A. Viergever, A. Leemans, 2014. "Recursive calibration of the fiber response function for spherical deconvolution of diffusion MRI data", NeuroImage vol. 86, pp. 67-80

ABSTRACT

There is accumulating evidence that at current acquisition resolutions for diffusion-weighted (DW) MRI, the vast majority of white matter voxels contains ‘crossing fibers’, referring to complex fiber configurations in which multiple and distinctly differently oriented fiber populations exist. Spherical deconvolution based techniques are appealing to characterize this DW intra-voxel signal heterogeneity, as they provide a balanced trade-off between constraints on the required hardware performance and acquisition time on the one hand, and the reliability of the reconstructed fiber orientation distribution function (fODF) on the other hand. Recent findings, however, suggest that an inaccurate calibration of the response function (RF), which represents the DW signal profile of a single fiber orientation, can lead to the detection of spurious fODF peaks which, in turn, can have a severe impact on tractography results. Currently, the computation of this RF is either model-based or estimated from selected voxels that have a fractional anisotropy (FA) value above a predefined threshold. For both approaches, however, there are user-defined settings that affect the RF and, consequently, fODF estimation and tractography. Moreover, these settings still rely on the second-rank diffusion tensor, which may not be the appropriate model, especially at high b -values. In this chapter, we circumvent these issues for RF calibration by excluding ‘crossing fibers’ voxels in a recursive framework. Our approach is evaluated with simulations and applied to in vivo and ex vivo data sets with different acquisition settings. The results demonstrate that with the proposed method the RF can be calibrated in a robust and automated way without needing to define ad-hoc FA threshold settings. Our framework facilitates the use of spherical deconvolution approaches in data sets in which it is not straightforward to define RF settings a priori.

5.1 INTRODUCTION

Diffusion weighted MRI (DW-MRI) allows for non-invasive characterization of brain microstructure and the directional organization of neural fiber tissue *in vivo* (Basser et al., 1994; Le Bihan et al., 1986; Le Bihan and Johansen-Berg, 2012; Moseley et al., 1990). The signal that is measured with DW-MRI is related to the amount of water diffusion and depends on several properties of the underlying tissue. Descriptive models that relate this signal to diffusion in fibrous tissue can provide markers of microstructural tissue organization. The most commonly used model to obtain such biomarkers is the second-rank diffusion tensor (Basser et al., 1994), which describes the signal adequately in case of aligned axons that hinder the diffusion. Measures derived from diffusion tensor imaging (DTI), like mean diffusivity (MD) and fractional anisotropy (FA), can be modulated by microstructural tissue components, such as myelination and cell membrane density (Beaulieu, 2002), and have already shown their value in a wide range of applications (Carballido et al., 2012; De Groot et al., 2009; Langen et al., 2012; Lebel et al., 2008; Scholz et al., 2009; Verhoeven et al., 2012; Wang et al., 2012).

It is well-known that DTI is unable to model multiple fiber populations within one voxel (Alexander et al., 2001; Frank, 2002; Tournier et al., 2011). The amount of white matter (WM) voxels that exhibits such a configuration with multiple fiber populations has recently been estimated to be around 90% (Jeurissen et al., 2013), which makes this a significant problem for the validity of DTI and, consequently, DTI based tractography approaches. Several approaches have been developed to address this issue of multiple fiber directions within a voxel. Techniques such as Q-ball imaging (Descoteaux et al., 2007; Hess et al., 2006; Tuch, 2004) and diffusion spectrum imaging (Wedeen et al., 2005) can reconstruct the diffusion orientation distribution function (dODF) that is related to water diffusion. Another category for dealing with ‘crossing fibers’ consists of the so-called spherical deconvolution (SD) techniques, which attempt to infer information about the fiber orientations themselves (Anderson, 2005; Dell’Acqua et al., 2007; Descoteaux et al., 2009; Tournier et al., 2004). Constrained SD (CSD), one of the more popular SD approaches (Emsell et al., 2013; McGrath et al., 2013; Metzler-Baddeley et al., 2012a; Reijmer et al., 2012; Reijmer et al., 2013a; Reijmer et al., 2013b; Vos et al., 2012), can reconstruct the fiber orientation distribution function (fODF) in a robust way with scanning protocols that can be used in clinical practice (Tournier et al., 2007). The diffusion signal is then modeled by the convolution of the fODF with a kernel that represents the DW signal corresponding to a single fiber orientation. The fODF is then deconvolved from the DW signal (hence, spherical *deconvolution*) using this kernel, coined the response function (RF).

There are several ways to construct the RF. For instance, it can be modeled by an axially symmetric tensor with a given FA and MD (Anderson, 2005; Descoteaux et al., 2009; Tournier et al., 2004), providing a good fit to the DW signal for a single fiber orientation at low b -values. At high b -values, however, the DW signal deviates significantly from its Gaussian distribution, even in single fiber populations (Clark and Le, 2000), due to restricted diffusion components (Assaf et al., 2004). In other work, the RF is estimated (also referred to as the ‘RF calibration’ step) from the DW signals in voxels with FA values that are typically larger than a user-defined threshold (Jeurissen et al., 2011; Jeurissen et al., 2013; Tournier et al., 2007). The reliability of the FA, however, is generally poor for high b -value and low-SNR DW-MRI data (Jones and

Basser, 2004). In addition, high FA values can still be present in regions with ‘crossing fibers’ (e.g., in case of two fiber populations, of which one has a small volume fraction (VF)), violating the assumption that the corresponding DW signals represent the diffusion profile of a single fiber population. Irrespective of these concerns, prior knowledge to calibrate the RF is required to perform the SD step for estimating the fODF. Fig. 5.1 illustrates an example of the difference in resulting fODF peak orientations for two different FA thresholds for RF calibration. In this example, there is an angle difference between the main peak orientations. In addition, the fODF with FA threshold of 0.8 exhibits two peaks, whereas with an FA threshold of 0.7, there is only peak orientation. Although the sensitivity of such user-defined settings is typically considered to be insignificant, recent research by Parker et al. (2013) clearly demonstrates that inaccurate characterization of the RF can produce spurious peaks, angular deviations of fODF peaks, and poor estimations of volume fractions of the various fiber populations.

In this chapter, we present a recursive framework to calibrate the RF which avoids the need of any user-defined settings related to DTI. After a coarse initialization of the RF, CSD is performed to identify and, subsequently, exclude voxels that contain multiple fODF peaks. From the remaining voxels the RF is updated and used for the next iteration. This procedure is repeated until only voxels with one fODF peak are left, thus, assuring that the final RF reflects the DW signal for a single fiber orientation. This framework is evaluated with simulated data to investigate its convergence properties and its robustness with respect to data noise. We also applied the proposed RF calibration approach to (i) conventional *in vivo* human brain data; (ii) *ex vivo* human brain data (Miller et al., 2011); and (iii) *ex vivo* data from the cervical spinal cord of the vervet monkey (Lundell et al., 2011), demonstrating the general applicability of our framework.

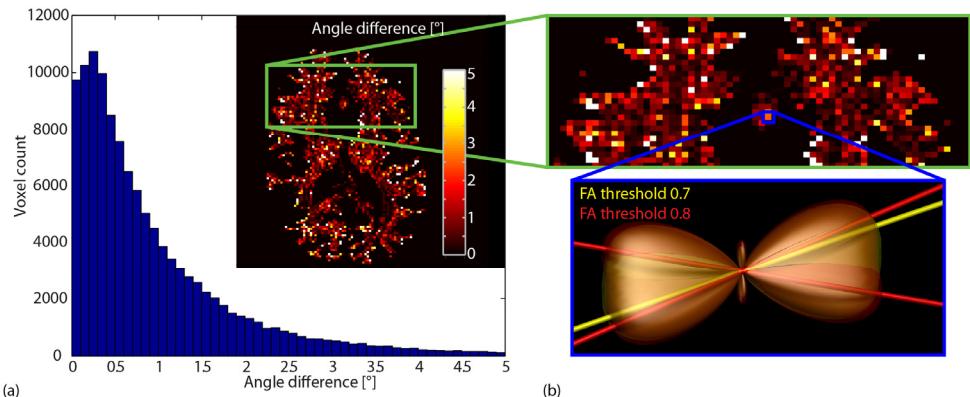


Fig. 5.1: (a) Histogram of the angles between the dominant peaks of the fiber orientation distribution obtained with constrained spherical deconvolution and with FA thresholds of 0.7 and 0.8 for estimating the response function. The spatial distribution of this angular difference, illustrated for an axial slice, is shown enlarged in the top image in (b). The bottom image in (b) represents the fODF and its peak orientations for both FA thresholds in a single voxel in the corpus callosum. In this example, both the main peak orientations and the number of peaks differ between the two threshold settings.

5.2 METHODS

5.2.1 fODF estimation using CSD

SD techniques use high angular resolution diffusion imaging (HARDI) sampling schemes to enable an accurate angular characterization of the diffusion signal (Tuch et al., 2002). Typically, a single shell of the q -space is sampled at a higher b -value compared to DTI. The diffusion signal on this shell can then be modeled as the convolution of the RF with an fODF. If the RF is known, the fODF can be found by the deconvolution of the diffusion signals with this RF, a processing step that is generally ill-posed (Tournier et al., 2004). By using regularization, Tournier et al. further improved the SD approach, which is then referred to as *constrained* SD, drastically reducing its sensitivity to noise (Tournier et al., 2007). A maximum harmonic degree of 8 (i.e., $l_{max} = 8$) is generally recommended for the spherical harmonics (SH) representations when performing CSD (Tournier et al., 2009). After a succinct description of the conventional way to obtain the RF (see Section 5.2.1.1), we outline our proposed recursive framework for RF calibration in Section 5.2.1.2.

5.2.1.1 Conventional RF calibration

In summary, voxels are identified that satisfy a predefined criterion, typically $FA > 0.7$, assuming that the corresponding DW signals represent diffusion profiles of a single fiber population. The DW signals in these selected voxels are then reoriented such that the first eigenvector of the diffusion tensor becomes aligned along a particular axis (usually the z-axis). Subsequently, the SH representations of the adjusted DW signals are averaged across these voxels to obtain a reliable estimate for the RF. During this step, only the SH coefficients with $m = 0$ are estimated (where m is the order of the spherical harmonic function), thereby constraining the RF to be an axially symmetric function (Tournier et al., 2007). In the following, we will refer to this conventional way of RF calibration as the ‘FA method’, given that the shape of the RF will be mainly determined by the FA threshold that is chosen.

5.2.1.2 Recursive RF calibration

Although there are slight variations in the conventional methodology to calibrate the RF, most approaches to date still rely on properties of the diffusion tensor (i.e., first eigenvector and FA). With the FA being an unreliable measure to differentiate between single and multiple fiber populations, it is desirable to avoid the diffusion tensor model for RF calibration altogether. Therefore, in this work, we propose a recursive framework to compute the RF, consisting of the following steps (for a schematic overview, see Fig. 5.2):

- 1) Initialize the RF with a diffusion profile shape that is significantly less sharp than that of the expected (i.e., ‘true’) RF. With some abuse of terminology, this initial RF will be referred to as a *fat* RF. In practice, such a coarse initialization of the RF can be easily achieved by taking DW signal profiles that corresponds to a near-isotropic diffusion tensor ($FA = 0.05$). The only requirement at this stage is that this fat RF should still exhibit a dominant orientation, as CSD will produce degenerate fODFs otherwise.

- 2) For voxels that were selected from the previous iteration, calculate the fODF using CSD with the RF from the previous iteration. Obviously, if this is the first iteration, then the fat RF initialization from the previous step will be used and the selected voxels could be the entire brain parenchyma or could be constrained to a WM mask (e.g., derived from a T1 image) to shorten computation time.
- 3) Calculate the direction and magnitude of the fODF peaks using a Newton optimization algorithm (Jeurissen et al., 2011).
- 4) Select the voxels for which the second largest fODF peak is absent or sufficiently small (e.g., below the noise level) compared to the largest fODF peak. In other words, if the ratio of the second largest fODF peak magnitude to the largest fODF peak magnitude (in the following, this will be called the peak ratio – PR) drops below a specific PR threshold, the voxel will be assumed to contain only one genuine fODF peak and, therefore, will be included for RF estimation in the next iteration. Note that this PR threshold represents an elegant user-defined parameter for balancing between sensitivity and specificity given a specific noise level. For instance, if the PR threshold is chosen too small, noisy fODF peaks will be detected as being sufficiently large. As a result, voxels are excluded erroneously for computation of the RF in the next iteration (false-negatives). If, on the other hand, the PR threshold is set too high, voxels with multiple fODF peaks will be included in the next iteration (false-positives).
- 5) Reorient the DW signals in the voxels obtained in step 4) according to their main fiber direction, that is, the largest fODF peak.
- 6) Calculate the new RF from the SH representations of the reoriented DW signals obtained from step 5) and constrain it to be axially symmetric. Continue with step 2) unless one of the following convergence criteria is met: (i) the maximum number of iterations is reached (set to 20, although the algorithm typically converged in nine iterations) or (ii) the RF shape did not change significantly compared to the previous iteration (e.g., less than 1% size difference for each of the SH coefficients).

Note that by using an RF that is fatter than the true RF when performing CSD, the number of fODF peaks will be underestimated, providing an overestimation of the number of voxels containing only one dominant fiber orientation. As such, for a given PR threshold, convergence of the RF computation is always guaranteed. In the remainder of this manuscript, we will refer to our proposed approach as the ‘recursive method’.

5.2.2 Simulations

Noise free DW data were generated assuming axially symmetric diffusion tensor profiles for each fiber population ($MD = 0.7 \cdot 10^{-3} \text{ mm}^2/\text{s}$) (Leemans et al., 2005), using a typical 60 direction sampling scheme (Jones et al., 1999) and b -value of 2500 s/mm^2 . Three types of simulation data sets were generated to investigate (i) the occurrence of spurious peaks in single fiber simulations and the influence that PR threshold settings would have, (ii) the influence of fiber configuration on convergence of the recursive method, and (iii) the effect of the PR/FA threshold and SNR on the performance of the recursive/FA method. Details of these simulations are provided in the following sections.

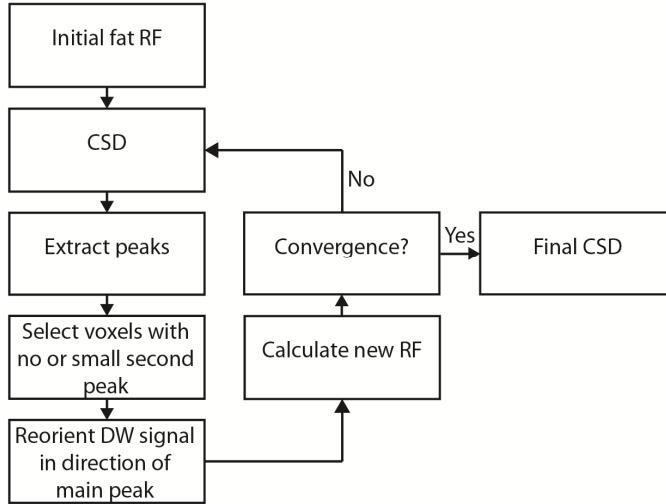


Fig. 5.2: Schematic overview of the recursive framework for RF estimation.

5.2.2.1 Simulation data set (i)

In recent work by Parker et al. (2013), simulations have shown that spurious peaks may also occur for single fiber populations. Such simulations can be used to investigate the optimal PR threshold required in step 4) of the recursive RF calibration, to avoid wrong exclusion of true one-fiber voxels. In line with the parameters defined in Parker et al. (2013), we constructed the following simulation data set. A single fiber population aligned with the z-axis is defined, with varying FA values ranging from 0.1 to 0.9 (with steps of 0.1). Noisy data sets (using the Rician data distribution) are generated with SNR of 10, 15, 22, and 30 (defined on the non-DW signal), with 500 noisy realizations each. The resulting signal was deconvolved with an RF, also simulated as a tensor with chosen FA (same ranges). Peaks were extracted and the magnitude of the second peak was calculated as a fraction of the first peak's magnitude. The purpose of this simulated data set is to investigate the relative magnitude of spurious secondary peaks with respect to the main fODF peak for these parameter settings. Ideally, only the voxels containing a single fiber population should be selected for RF calibration.

5.2.2.2 Simulation data set (ii)

The second simulation set contains only one- and two-fiber populations, constructed from two prolate tensors with fixed FA value of 0.8. The ratio of voxels containing one-fiber versus two-fiber populations was set to 1:9, given that approximately 90% of the WM voxels contain ‘complex fiber configurations’ for common voxel sizes (Jeurissen et al., 2013). VF and inter-fiber angle were the same for all the two-fiber voxels. Rician distributed noise was added to result in $\text{SNR} = 22$ for the $b = 0 \text{ s/mm}^2$ signal, corresponding to the SNR of the in-vivo data (see Section 5.2.3.1). This simulation set was used to investigate convergence of our recursive RF calibration method. We created multiple data sets where we either varied VF or inter-fiber angle for all two-fiber populations. Ideally, the RF should converge to the DW profile of the predefined one-fiber population that was characterized by $\text{FA} = 0.8$. In other words, we expect that the final ‘FA of the RF’, estimated by fitting the diffusion tensor to the resulting RF, will approximate the FA value of the simulated single-fiber population.

5.2.2.3 Simulation data set (iii)

To construct the DW signals for the third simulation data set, we incorporate realistic distributions of FA values for a single fiber population, VFs, and inter-fiber angles derived from *in vivo* human brain data as described in Jeurissen et al. (2013). Specifically, one-, two- and three-fiber populations were constructed with the ratio 10:45:45. The VFs of the two-fiber populations, i.e., VF_1 and VF_2 , were characterized by a Normal distribution (denoted as $N(\mu, \sigma)$ with μ the mean and σ the standard deviation), with VF_1 sampled from $N(0.5, 0.15)$ and with $\text{VF}_2 = 1 - \text{VF}_1$. The VFs of the three-fiber populations are sampled from the standard uniform distribution (range between 0 and 1) and subsequently normalized. Inter-fiber angles were sampled from $N(90, 30)$ (in degrees). The FA values for the single fiber populations were sampled from either $N(0.6, 0.05)$ – then referred to as the ‘ $\text{FA}_s = 0.6$ ’ simulation data set – or $N(0.8, 0.05)$ – then called the ‘ $\text{FA}_s = 0.8$ ’ simulation data set. The purpose of this simulation data set is to investigate differences in consistency of the resulting fODF peaks with the true simulated peaks, when performing CSD with the RF either calibrated from the recursive or FA method. An fODF peak is termed ‘true positive’ when it has an angle difference smaller than 5° with a true simulated fiber. We applied an fODF threshold of 0.1 normalized to the magnitude of the first peak, and set the algorithm to detect up to 4 fODF peaks. False positives fODF peaks were all remaining peaks that did not correspond to a true simulated peak. False negative peaks were, in turn, defined as all remaining true simulated fiber populations.

5.2.3 Real data

The majority of our experiments on real data were performed on an *in vivo* human brain data set $b = 2500 \text{ s/mm}^2$ typically used for CSD ($l_{\max} = 8$). With this data set we compared the recursive method with the FA method and extended findings from simulations to real data to provide guidelines for practical use. To demonstrate the general applicability of our framework, we also applied it to *ex vivo* data sets with different acquisition settings.

5.2.3.1 In vivo human brain data

5.2.3.1.1 Data acquisition

A healthy volunteer (male, 25 y) was scanned on a 3.0 T Philips Achieva MR scanner (Philips, Best, NL) equipped with an 8-channel receiver head coil. The subject gave informed consent to participate in this study under a protocol approved by the University Medical Center Utrecht ethics board.

First, DW images were acquired using a single-shot spin-echo echo planar imaging sequence with the following settings: acquisition matrix of 112×112 , FOV of $224 \times 224 \text{ mm}^2$, and 70 slices with thickness 2 mm resulting in isotropic voxels of $2 \times 2 \times 2 \text{ mm}^3$. Other parameters were SENSE acceleration factor of 2 and TE/TR = $73/6718 \text{ ms}$. Diffusion sensitizing gradients were applied in 60 directions uniformly distributed over the hemisphere (Jones et al., 1999) with a b -value of 2500 s/mm^2 (NSA = 1) along with one $b = 0 \text{ s/mm}^2$ (NSA = 6) image. Total scan time was approximately 12 minutes. The SNR of the data set is 22, as estimated in WM regions of the $b = 0$ image (Sijbers and den Dekker, 2004).

Next, a T1 image was acquired to obtain a WM mask, used in step 2) of our framework to shorten computation time. A sagittal 3D turbo field echo sequence was used with an FOV (AP \times LR \times IS) of $240 \times 160 \times 240 \text{ mm}^3$ and imaging matrix $240 \times 160 \times 240$, reconstructed to $256 \times 160 \times 256$ matrix size with $0.94 \times 1 \times 0.94 \text{ mm}^3$ voxel size. A SENSE acceleration factor of 1.6 was used in both the AP and IS directions. A flip angle of 8° was used with TE/TR/TI = $3.5/7.6/991 \text{ ms}$, for a total acquisition time of 200 s.

5.2.3.1.2 Image processing

The DW scans were corrected for subject motion and eddy current induced geometric distortions including the required B-matrix adjustments (Leemans and Jones, 2009). Processing and visualization was performed with *ExploreDTI* (Leemans et al., 2009).

WM voxels from the DW images were extracted using a T1 based WM mask as described previously (Jeurissen et al., 2013). In summary, a tissue probability map was estimated from the T1 image using the unified segmentation tool from SPM (Ashburner and Friston, 2005). Next, the T1 image was registered to the FA image using 3D nonrigid b-spline-based registration using Mattes mutual information (Klein et al., 2010). WM voxels in the DW image were then identified by warping the tissue probability map to the DW volume, and select voxels with WM probability higher than 95%. Finally, voxels with $\text{MD} > \text{median}(\text{MD}) + 1.5 \cdot \text{IQR}(\text{MD})$ were removed due to avoid remaining partial voluming with CSF (where IQR is the interquartile range over the mask).

5.2.3.1.3 Fiber tractography

We performed fODF streamline tractography (Jeurissen et al., 2011) to evaluate the influence of RF calibration with both methods on tract propagation. We set an fODF and angle threshold of 0.1 and 35° , respectively, and a step size of 1 mm. The corpus callosum (CC, commissural fibers) was reconstructed, using the guidelines in Wakana et al. (2007) for defining the regions of interest.

5.2.3.2 Ex vivo data

One human brain and one monkey spinal cord data set were used for analysis. The human brain data set was acquired in collaboration with K. Miller and consisted of 53 DWI images (NSA = 3) with $b = 4500 \text{ s/mm}^2$, along with one $b = 0$ image with $1.2 \times 1.2 \times 1.2 \text{ mm}^3$ voxel size. TE/TR = $20/35 \text{ ms}$ with a flip angle of 35° was used. The monkey spinal cord data set consisted of 105 DWI images ($b = 4525 \text{ s/mm}^2$) and 25 non-DWI images, with axial in-plane resolution of $188 \times 188 \mu\text{m}^2$, and a slice thickness of $375 \mu\text{m}$ (Lundell et al., 2011). For more detailed acquisition parameters and image processing steps of the spinal cord and human brain data sets, the reader is referred to Lundell et al. (2011) and Miller et al. (2011), respectively.

The RF was calibrated with the recursive method and subsequently used to perform CSD. Appropriate selection of FA thresholds is unintuitive for these ex vivo data sets due to the intrinsically lower FA values. As such, we compared the recursive method to the FA method using two conventional approaches: 1) Define the FA threshold based on the method from Tournier, 2004, in which the voxels with the 300 highest values of FA are selected to calculate the response function; 2) Set the FA threshold as the average of FA values in a region known to

have single fiber populations. For the latter approach, we have drawn ROIs in the corpus callosum of the ex vivo human brain data set, and in the WM located superior in the spinal cord of the ex vivo monkey data set.

For the monkey spinal cord data we used $l_{max} = 8$; for the human brain data $l_{max} = 6$ was used. In theory, 53 DW images would allow us to fit SH coefficients up to degree 8. However, as the gradient orientations were not evenly distributed on the sphere, we could only capture the angular frequencies up to degree 6 in the human brain data set.

On the human brain data set, we performed streamline tractography with fODF threshold 0.01 and step size to 0.5 mm (which differ from the in vivo human brain data set due to the higher spatial resolution and generally smaller fODF peaks). Fiber tracts of the CC, CST, and superior longitudinal fasciculus (SLF) were reconstructed at the area where they are expected to intersect each other. For the monkey spinal cord data, we were specifically interested in detecting multiple fiber populations which can be used to study commissural fibers and complex features of the grey matter and white matter collaterals (Lundell et al., 2011).

5.3 RESULTS

5.3.1 Simulations

In this section, the results of the experiments on simulation data sets (i), (ii), and (iii) are presented in Sections 5.3.1.1, 5.3.1.2, and 5.3.1.3, respectively.

5.3.1.1 Spurious peaks in single fiber populations

To illustrate the problem of spurious peaks, Fig. 5.3a shows the resulting fODF when a simulated single fiber population signal was deconvolved with the wrong (sharper) RF, in the case of infinite SNR (Parker et al., 2013). These spurious peaks result in erroneous tractography results, if they become too large. The median ratio of the second (spurious) peak with respect to the first peak and its upper and lower quartile ranges are shown in Fig. 5.3c in the ‘ideal’ case, i.e., when we would have chosen exactly the right RF for deconvolution of the simulated signal. This corresponds to the diagonals of the color graphs in Fig. 5.3b, which display the median ratio by a color (blue indicating a ratio of 0 and red represents a ratio of 1) for different combinations of simulated signal FA and RF FA ranging from 0.1 to 0.9, and for different SNR values. This can be used to make an informed decision regarding the PR threshold, where the ‘ideal’ case in Fig. 5.3c can be used as a lower bound to avoid incorrect exclusion of one-fiber populations due to spurious peaks. In the case of SNR = 22, for example, a PR threshold of 0.05 would be a reasonable choice, since 75% of the spurious peaks have a relative magnitude below this threshold and, therefore, can be considered ‘sufficiently small’. One can be even stricter in this threshold setting, as the RF calibration is initiated with a fat RF. Consequently, the magnitude of such spurious peaks will be even smaller.

5.3.1.2 Convergence of the recursive method

Fig. 5.4 shows the FA of the RF as function of iteration for different VF values (a) and angles (b) of the included two-fiber populations, setting a PR threshold of 0.01. For all VFs (at angle 90°), the recursive method correctly includes single fiber voxels for RF calibration, as the FA of the RF converges to 0.8 after approximately 2 – 4 iterations (Fig. 5.4a). Fig. 5.4b illustrates that the RF converges to the correct one for inter-fiber angles 60° and 90° in 2 – 9 iterations. For inter-fiber angles of 45° and smaller, one can see that the RF converges to FA values lower than 0.8 (i.e., the calibrated RF is fatter than the true RF). This means that in the two-fiber voxels, with angles of 45° and smaller between the two fiber orientations, only one peak could be detected. Consequently, these voxels were not correctly excluded from RF calibration.

5.3.1.3 Comparison of angular consistency between recursive and FA method

Figs. 5.5ac and 5.5bd display the results for the simulation sets with FA_s = 0.8 and FA_s = 0.6, respectively. In Figs. 5.5a and 5.5b, we plotted the sensitivity (ratio of true positive fODF peaks to the sum of true positive and false negative peaks) against the number of false positives as function of the threshold value (circled points). The colored lines represent different SNR values. The top plots in figures (a) and (b) show results for the FA method with FA threshold ranging from 0.1 to 0.9 in steps of 0.1 (in Fig. 5.5b FA ranged from 0.1 to 0.8, since there were no voxels with FA higher than 0.9 for the different SNR levels). The bottom plots show the results for the recursive method (with PR threshold values = [0.5, 0.2, 0.1, 0.05, 0.02, 0.01, 0.005]).

From Fig. 5.5a, one can see that for high SNR (15 to 30), the change in sensitivity and number of false positives for the recursive method is almost independent of the PR threshold, despite a 100-fold increase (from 0.005 to 0.5). For the FA method, on the other hand, this is not the case. One can notice that for the FA method at high SNR, the amount of false positives is high at a low FA threshold, drops down, and subsequently increases again with increasing FA threshold. When using a fatter RF (and thus low FA threshold), angular deviation of fODF peaks with true peaks will be larger, which results in more false positives. An RF that is too sharp will result in spurious peaks, which will be counted as false positives. The graphs of the FA method and recursive method overlap for the same SNR, which implies that the same sensitivity and specificity can be reached with both methods. The recursive method, however, yields optimal sensitivity and specificity values that are less dependent on the PR threshold at high SNR.

In Fig. 5.5b, there is an overall decrease of sensitivity and increase of false positives for this data set with lower FA for a single fiber population. One can again notice the higher robustness of the recursive method towards the PR threshold. The main result here is that the PR thresholds that yield the highest sensitivity and specificity in the recursive method are independent of the simulated fiber FA of the underlying data set, when comparing Fig. 5.5a to Fig. 5.5b. For instance, a peak threshold of 0.01 or 0.005 for SNR 22 would give a good trade-off between sensitivity and number of false positives in both the FA_s = 0.8 and the FA_s = 0.6 data sets, whereas one should use different FA thresholds (0.7 and 0.5, respectively) to obtain the same sensitivity and number of false positives with the FA method.

Chapter 5: Recursive Calibration of the Fiber Response Function for Spherical Deconvolution of Diffusion MRI Data

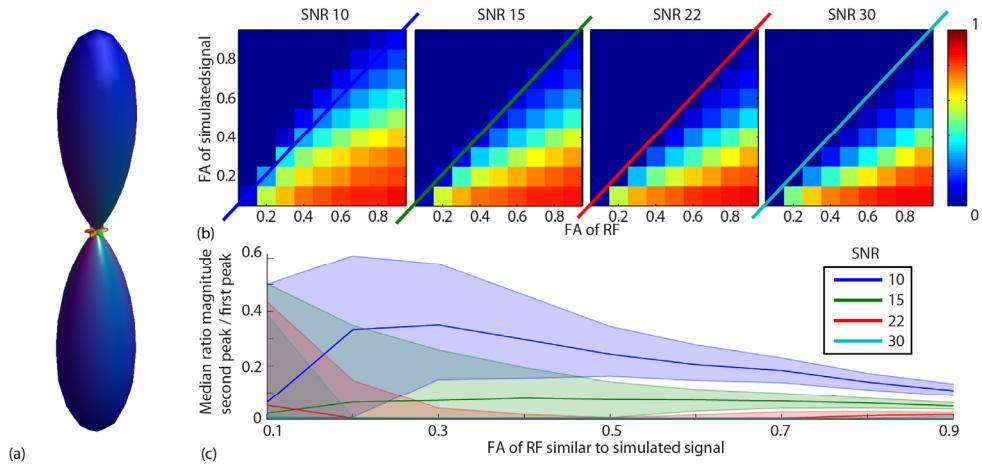


Fig. 5.3: Results of simulation set (i). (a) Illustration of spurious peaks occurring when the RF is not properly chosen. Here, the resulting fODF is shown for a simulated signal with $FA = 0.5$ (infinite SNR) and an RF with $FA = 0.8$. (b) CSD of a single fiber population simulated as a tensor with a particular FA (vertical axis) with an RF (also a tensor with chosen FA, horizontal axis) gives spurious peaks. Here the median relative magnitude of the second (spurious) peak compared to the first (true) peak is shown in color for all combinations of simulated signal- and RF-FA (Parker et al., 2013a), for $SNR = [10, 15, 22, 30]$. (c) For the ideal case, where the RF has exactly the same shape as the simulated signal (same FA), the median relative magnitude of the spurious peak is visualized together with Q_1 and Q_3 for different FA. This corresponds to the diagonals of (b).

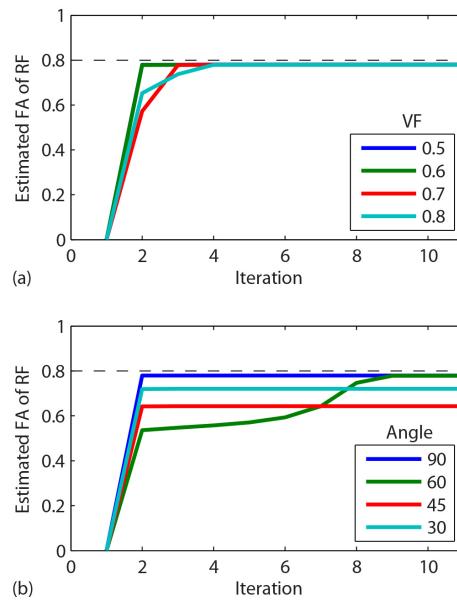


Fig. 5.4: Results of simulation set (ii). FA of the RF over iteration for (a) different volume fractions with a 90° angle, and (b) different angles with equal volume fractions for the included two-fiber voxels. Dashed line represents the simulated fiber population. We set PR threshold = 0.01. For the data sets that include only single fibers and crossings of 45° or 30° , one can see that the two-fiber voxels cannot be distinguished from one-fiber voxels and, therefore, will also be selected for RF estimation, which results in a fatter RF and thus lower FA of the RF.

Figs. 5.5c and 5.5d show histograms of the angle between the main fODF peak and its closest true simulated peak at an SNR of 22 for the FA_s = 0.8 and the FA_s = 0.6 data sets, respectively (top rows: FA method; bottom rows: recursive method). It can be observed that for the FA method, an FA threshold of 0.8 and 0.6 would result in the smallest median angle deviation for the FA_s = 0.8 and FA_s = 0.6 data sets, respectively. By contrast, a peak threshold of 0.005 yields the smallest median angle deviation for both simulation data sets (FA_s = 0.8 and FA_s = 0.6) using the recursive method.

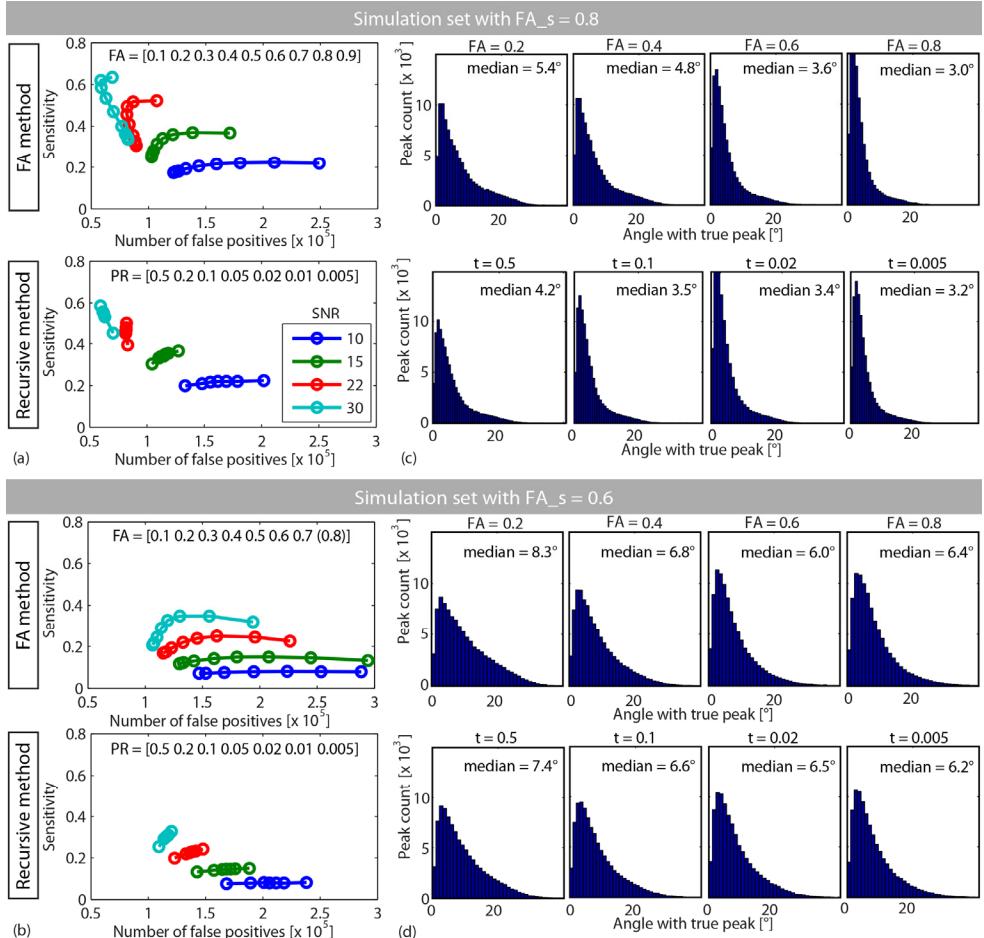


Fig. 5.5: Results of simulation set (iii). We used mean FA_s = 0.8 for the distribution of a single fiber population for Figs. (a) and (c), and FA_s = 0.6 for Figs. (b) and (d), respectively. (a) Top: Sensitivity of the FA threshold method against the number of false positives for FA threshold [0.1 - 0.9] in steps of 0.1, for different SNR values. The point with the lowest sensitivity corresponds to FA = 0.1. Bottom: Same graph for the recursive method with PR threshold values = [0.5, 0.2, 0.1, 0.05, 0.02, 0.01, 0.005], the point with the lowest sensitivity corresponds to PR threshold = 0.5. (b) The recursive and FA methods for RF calibration are used on the data set with mean FA_s = 0.6 for a single fiber population. For SNR 10, there were no voxels with FA > 0.8. (c) and (d) Histograms of the angle between the main fODF peak and its closest true simulated peak at SNR 22 for data sets FA_s = 0.8 and FA_s = 0.6, respectively, and for the FA (top row) and recursive (bottom row) methods, separately.

5.3.2 Real data

5.3.2.1 In vivo human brain data

In Fig. 5.6a, the method is illustrated on the in vivo human data for a PR threshold of 0.01. The graph represents the number of voxels included for RF calibration as function of iteration, convergence is reached after ten iterations. The resulting fODFs in the centrum semiovale are visualized for iteration one, two, and ten. We start off with a fat RF and all WM voxels derived from the T_1 , and the resulting fODFs show secondary peaks in voxels with crossing fiber populations of approximately equal volume fractions. These voxels will be discarded from RF calculation in the first iteration. For the second iteration, we end up with an RF that is flatter, and deconvolution yields secondary and tertiary peaks in voxels with multiple fiber populations of unequal VF. fODF peaks of single fiber populations become larger. After ten iterations, the algorithm has converged and the number of voxels used for RF calculation remains unchanged. Fig. 5.6b top shows which voxels are selected in each iteration (up to 10 iterations). The brightest voxels are used for final RF calculation, and comprise a range of FA values (0.3 – 1) with a mean of 0.67, as can be seen from the histogram in Fig. 5.6b bottom. Lower FA voxels are thus also selected, while some high FA voxels are discarded as can be seen by comparing the FA maps of Fig. 5.6b top and 5.6c.

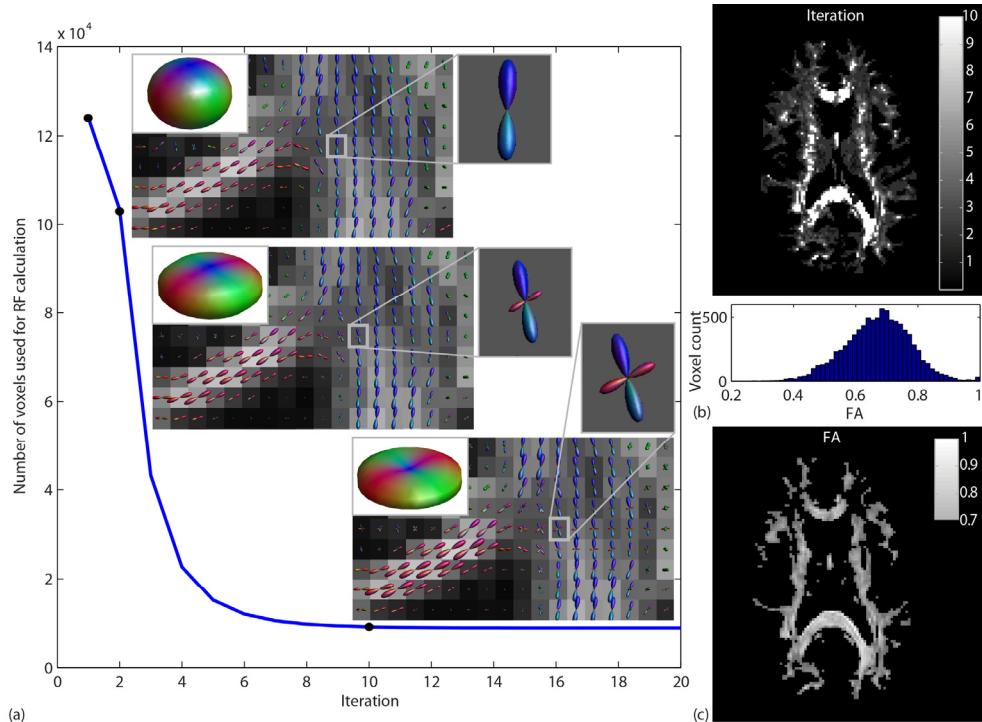


Fig. 5.6: Illustration of the method on real human data, using PR threshold = 0.01. (a) Number of voxels used for RF calibration, together with the used RF and the resulting fODFs, as function of iteration. (b) Top: Voxels selected in different iterations, visualized up to iteration 10. The brightest voxels are used for final RF calculation. Bottom: Histogram of the FA of the voxels used for final RF calculation. (c) FA > 0.7 map.

The recursive and FA method are compared in Fig. 5.7. The figure shows the angle differences between the main peak orientations, both in histograms (whole brain) and in color plots (ROI as indicated in Fig. 5.7a). Fig. 5.7b shows the angle difference when comparing two different thresholds (0.01 and 0.05) within the recursive method. Fig. 5.7c shows the same figure for two different thresholds (0.7 and 0.8) within the FA method. Note that a five-fold increase of the peak threshold (from 0.01 to 0.05) yields a smaller median angle difference (median = 0.5°) than the median angle difference when comparing FA threshold 0.7 to 0.8 using the FA method (median = 0.7°, see Fig. 5.7c). This supports simulation results indicating increased robustness of the recursive method towards its PR threshold value. Figs. 5.7d-g compare different thresholds between methods. One can see that between the methods, fODF peaks obtained with PR threshold = 0.01 and FA threshold = 0.7 are the most comparable for this particular data set (median = 0.5° in Fig. 5.7d).

Fig. 5.8 shows the results of fODF streamline tractography of the CC, for both the FA method (FA threshold = [0.7, 0.8]) and the recursive method (PR threshold = [0.01, 0.05]). The difference between both methods can be appreciated in the region of the centrum semiovale, where the lateral projections of the CC are delineated more completely using the recursive method.

5.3.2.2 Ex vivo data

Fig. 5.9 shows the results for the ex vivo human brain data set. In 5.9a, the fODFs reconstructed with the recursive method are shown in a region as indicated in Fig. 5.9b. Figs. 5.9c and 5.9e show the same region for the FA method, using the two different approaches: 1) The 300 highest FA voxels and 2) the average FA value in a superior region of the WM. The FA thresholds for approach 1) and 2) were 0.67 and 0.48, respectively. Although differences in the fODFs are very subtle between Figs. 5.9a and 5.9c, the fODFs clearly differ in magnitude and direction compared to 5.9e. Differences become more prominent for the tractography results. We were able to track the CST, CC and SLF (see Fig. 5.9b for the recursive method, and 5.9d and 5.9f for the FA method, showing a coronal view left and a sagittal view right), despite the major challenges in these crossing fiber regions that are known to exist for ex vivo data (Miller et al., 2011). Using the recursive method, the lateral projections of the CC could be reconstructed with greater detail when comparing Fig. 5.9b left with Figs. 5.9d and 5.9f. Although the SLF pathways extend more anteriorly using the FA method, there are trajectories that deflect into other regions (see white arrows in right image in 5.9d and 5.9f) or there seem to be parts that are missing (see yellow arrow in right image in 5.9d and 5.9f).

In Fig. 5.10, results for the ex vivo monkey spinal cord are displayed. For visualization purposes, we performed DTI streamline tractography (Basser et al., 2000; Lundell et al., 2011) as shown in Fig. 5.10a in a dorsal coronal view, with most fibers oriented inferiorly – superiorly. Figs. 5.10b and 5.10c show two transversal slices of interest for both the recursive and the FA method, one at the location of the cervical enlargement (bottom three) and one more superiorly (top three). For each slice, the recursive method is compared to the FA method where the FA threshold was determined using the two different approaches. The FA thresholds for approach 1) and 2) were 0.82 and 0.59, respectively. Fig. 5.10b shows the inferiorly-superiorly oriented largest fODF peaks, whereas Fig. 5.10c shows both the second and third fODF peaks. The radial organization in the WM and commissural connections in the grey

Chapter 5: Recursive Calibration of the Fiber Response Function for Spherical Deconvolution of Diffusion MRI Data

matter can be appreciated in both slices for both the recursive and the FA method. In analogy to Lundell et al. (2011), we found that more crossing fibers could be found when moving closer to the cervical enlargement. Note the difference in crossing fibers for the FA method between the two different thresholds. It can also be seen in both slices that the FA method with threshold 0.82 reveals more crossing fibers than both the FA method with the lower threshold and the recursive method.

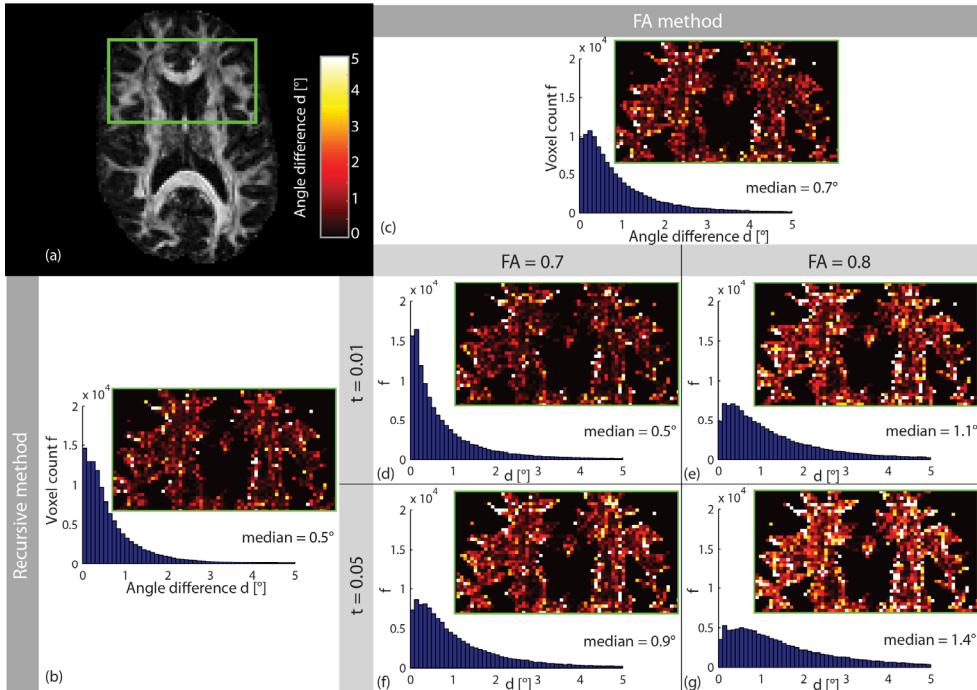


Fig. 5.7: Histograms of the angle difference d between the main peak orientations. ROIs in Figs. (b-g) are chosen as indicated in (a). (b) Comparison of main peak orientations for different thresholds using the recursive method. (c) Same as (b) for the FA method. (d-g) Comparison of both methods, different thresholds.

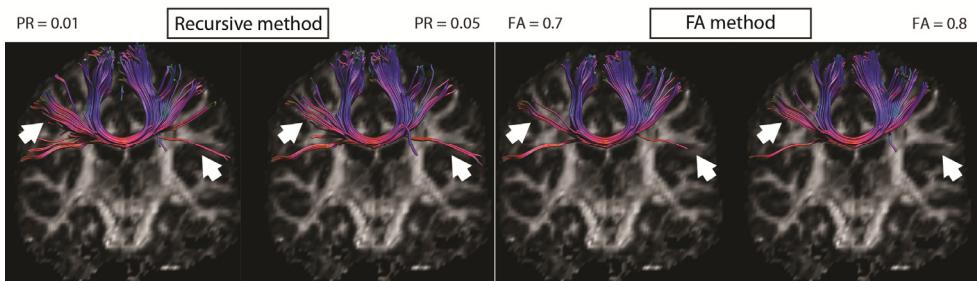


Fig. 5.8: Tractography results of CC using RF estimation with both methods, for different FA and PR thresholds respectively. The recursive method reveals the lateral projections of the Corpus Callosum better (arrows).

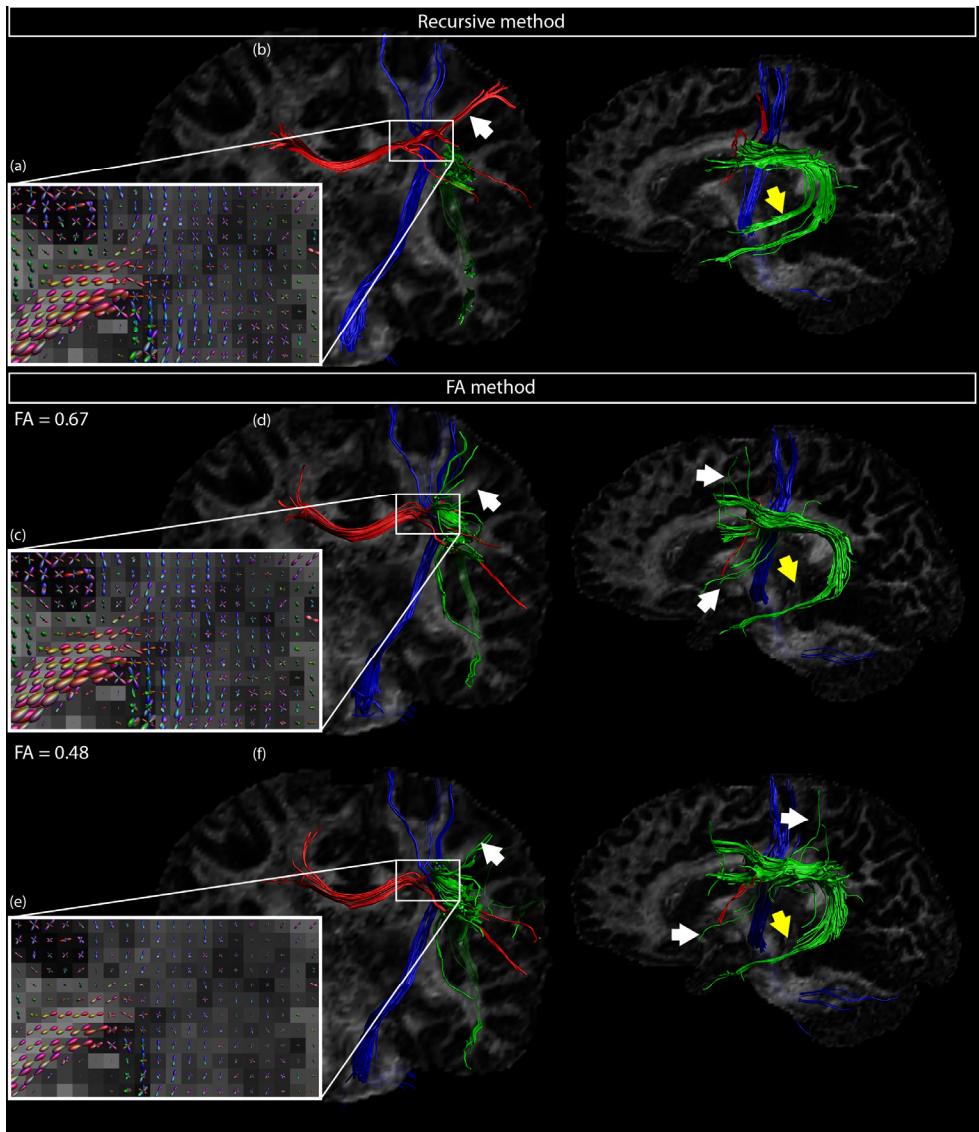


Fig. 5.9: Results of CSD with the RF resulting from the recursive method (PR threshold = 0.01) compared to the FA method (two different threshold approaches) for the ex vivo human brain data. (a, c, e) fODFs visualized in the centrum semiovale, where the CC, CST and SLF pathways intersect each other. fODFs in (a) and (c) look similar, but in (e) they differ a lot in direction and magnitude. (b, d, f) Streamline tractography of the CC (red), CST (blue) and SLF (green), in a coronal (left) and sagittal (right) view. White arrows in the coronal view indicate difference in lateral projections of the CC. In the sagittal image, white arrows show false positive pathways of the SLF, whereas yellow arrows indicate missing areas.

Chapter 5: Recursive Calibration of the Fiber Response Function for Spherical Deconvolution of Diffusion MRI Data

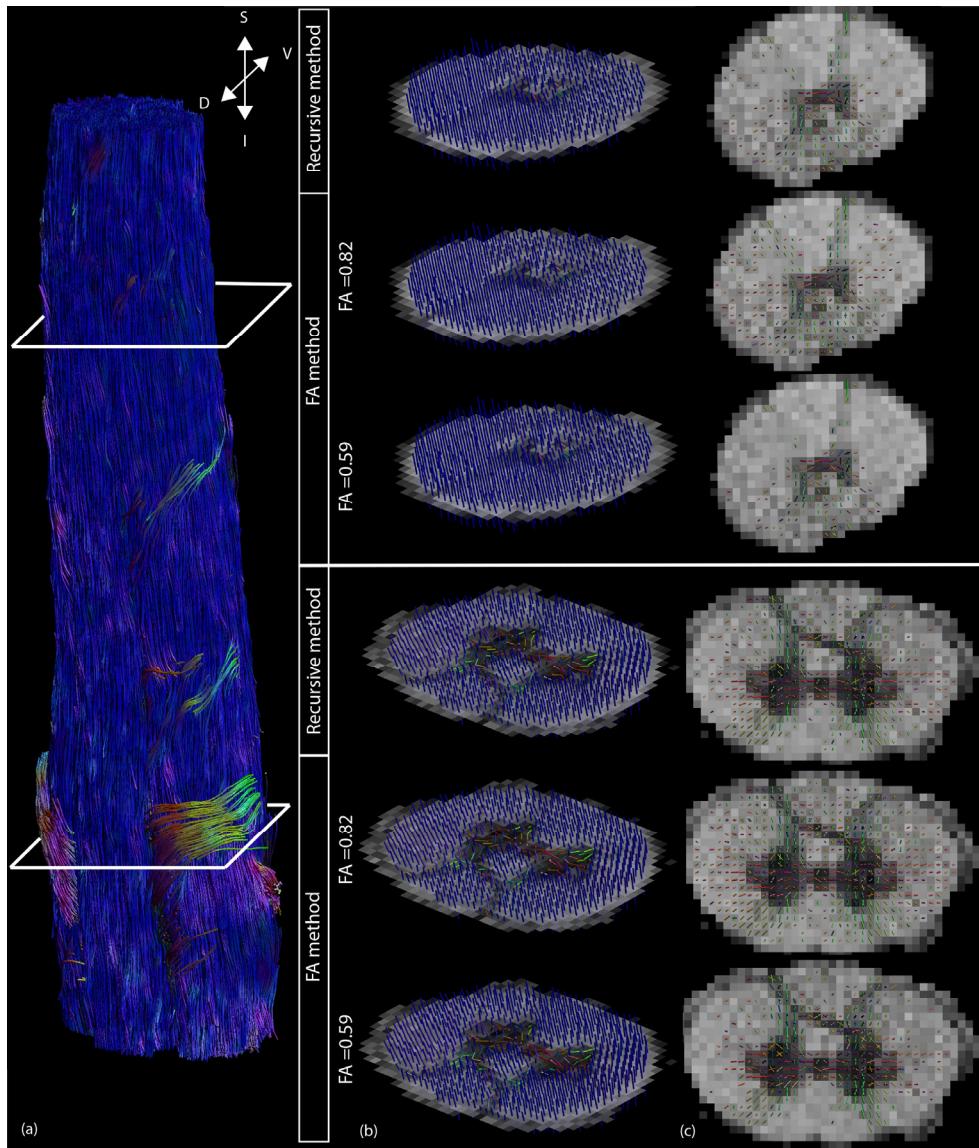


Fig. 5.10: (a) Tractography of the ex vivo monkey spinal cord to highlight the location of the two transversal slices shown in (b) and (c). Results of CSD are visualized by means of the first peak (b) and the second and third peaks (c). Results are shown for both the recursive (PR threshold = 0.01) and FA method (with two different FA threshold approaches). In the middle row, the result when 300 highest voxels are chosen determined the FA threshold, whereas in the bottom row, the FA threshold was based on the ROI.

5.4 DISCUSSION

Spherical deconvolution based techniques are appealing to tackle the crossing fiber problem, as they provide a balanced trade-off between constraints on the required hardware performance and acquisition time on the one hand, and the reliability of the reconstructed fODF on the other hand. The underlying assumption for most SD techniques is that there is a single appropriate RF throughout the whole brain, which is not known a priori. Although it is questionable whether there actually exists one RF for the whole brain (and not a distribution of one-fiber RFs instead), it provides computational advantages to assume so.

Recent work has extensively investigated the effect of RF miscalibration on the resulting fODFs (Parker et al., 2013a). Besides angular deviations, an overestimation in RF anisotropy will lead to predictable spurious fODF peaks in true single fiber populations due to truncation of the SH series. Their work states that a more informed calibration might lead to vast improvements in results, for example in cases of WM degradation and other fibrous tissue, where it is difficult to make an assumption about the RF a priori.

The RF could be modeled by axial symmetric tensors on a voxel-by-voxel basis (Anderson, 2005), but it is well known that the DW signal deviates significantly from this model at high b -values, even in single fiber populations. The FA method therefore determines the RF from the signal itself, by averaging the signal in high FA voxels. We will discuss the drawbacks of this method in the following paragraphs, which the recursive method tries to address.

5.4.1 Dependence on the diffusion tensor model

Although CSD itself is not based on the tensor model, the calibration step for the FA method still is. It uses the FA to select voxels for RF calibration, and the first eigenvector of the diffusion tensor to rotate the DW signal before averaging. Besides the violation of the assumption of Gaussianity at high b -values, the accuracy of FA and eigenvector estimates in this low SNR regime is poor. The choice of diffusion tensor estimation procedure (e.g., ordinary-, weighted-, and non-linear-least squares) may play a role, especially if gross artifacts are present (as data outliers may produce artificially high FA values), which will contaminate the RF estimation. Noteworthy is that the recursive framework presented in this work is completely independent of DTI based measures.

5.4.2 Dependence on underlying data properties

For the FA method, it is difficult to make an assumption about which FA to choose for RF calibration, especially in case of white matter degradation, neuro-development, or when other types of fiber tissue are investigated. Due to unknown anisotropy of underlying tissue, the results with the FA method are very dependent on the predefined FA threshold setting. In our framework, we can make a more informed decision by setting the PR threshold, and the method furthermore shows to be robust towards this threshold.

By setting the PR threshold, we impose the relative magnitude of the secondary fODF peak compared to the first fODF peak that would still be allowed for a voxel being classified as one-fiber voxel. For true single fiber populations, we approximately know what spurious peak-

magnitudes to expect at given SNR (see Fig. 5.3), and adapt the PR threshold accordingly. A lower bound PR threshold may be set at which spurious peaks in one fiber populations can still be distinguished from true multiple-fiber voxels. For reasonable SNR levels (15 to 30), this minimal threshold is almost independent of the FA of underlying tissue (Figs. 5.3 and 5.5). For SNR 22, for example, a PR threshold of 0.01 would correctly select one-fiber voxels for RF estimation in case of a ‘perfectly’ chosen RF. In practice, however, we start off with a fat RF and spurious peaks will be smaller, so one could even be very strict in this threshold setting (for example, setting it to 0.005, see Fig. 5.5).

We have shown that our recursive framework for RF estimation is more robust than the FA method towards differences in anisotropy of underlying tissues for high SNR (see Fig. 5.5). For both the $FA_s = 0.6$ and $FA_s = 0.8$ simulation data sets, the same PR threshold of 0.005 yields the minimal difference in angle between the main fODF peaks and the true simulated peaks. This is supported by our evaluations on real data. The ex vivo human brain data has intrinsically lower FA values than the in vivo data (0.32 in the CC according to previous work by Miller et al. (2011)), and the ex vivo spinal cord data has a mean FA of approximately 0.53 in the WM as estimated by Lundell et al. (2011). Despite the difference in underlying diffusion properties between these data sets, we were able to generate plausible results with the recursive RF calibration using the same PR threshold of 0.01. As demonstrated in Figs. 5.9 and 5.10, such uniformity could not be realized by a single threshold setting for the FA. The results for the FA method show that selecting an FA threshold with two plausible approaches already gives very different results. Selecting the 300 highest FA voxels is still quite ad-hoc (Why 300, and how will this choice affect the results?) and selecting an ROI to determine the FA threshold can be very rater-dependent.

In the ex vivo human brain data, tractography on the fODFs obtained with the FA method resulted in missing and deflecting tracts. If the FA threshold is taken too low, the fODFs are smaller in magnitude, the number of false negative peaks is larger, and the fODFs have a lower angular resolution. In the ex vivo monkey spinal cord data, a radial organization of the WM and commissural connections in the grey matter were found with both RF calibration methods. The highest number of fODF peaks could be found using the FA method with the highest FA threshold. Note, however, that one always needs to make a trade-off between sensitivity and specificity, which relies heavily on the choice of FA threshold. Though the sensitivity with this FA threshold appears to be higher compared to the recursive method in this case, the number of false positives (spurious or deflecting peaks) increases with a higher FA threshold. We thus argue that when no reliable assumption can be done on the underlying FA of the tissue, the recursive method should be the method of choice.

For the FA method, sensitivity and number of false positive fODF peaks change significantly between different FA thresholds, even at the high SNR of 30 (see Fig. 5.5) and despite the relatively small change in FA threshold. When comparing this to the recursive method, sensitivity and number of false positives do not change a lot for a relatively large change in PR threshold. More specifically, at $SNR = 22$, sensitivity and number of false positives increase by 0.01 and 637, respectively, when decreasing PR threshold from 0.05 to 0.01, which is a five-fold decrease. This statement is also supported by experiments on human data, where the angle difference between the main fODF peaks is smaller for the recursive method when changing the

PR threshold from 0.05 to 0.01 than for the FA method when changing the FA threshold from 0.7 to 0.8. For low SNR data, which is known to be difficult for SD methods (Fillard et al., 2011; Tournier et al., 2013a), one should make an informed decision balancing between decreased specificity and increased number of false positives.

5.4.3 Convergence

Since the RF calibration in our proposed method is recursive, it is important to investigate its convergence properties. As shown by our simulation experiment in Fig. 5.4, the algorithm converges towards the predefined profile of a single fiber population after roughly nine iterations. For all VFs with fiber orientations crossing at an angle of 90°, the FA of the RF converges to 0.8 (see Fig. 5.4a). For equal VFs and inter-fiber angles of 60° and 90°, the RF converges also correctly, but for inter-fiber angles of 45° and smaller, one can see that the two-fiber voxels were not correctly excluded from RF calibration, resulting in an FA value of the RF lower than 0.8 (i.e., a fatter RF than the true one – see Fig. 5.4b).

The recursive method assumes that fODFs with a single significant peak are unambiguously associated with a single fiber orientation. However, it is well-known that for some configurations, like fanning and bending fibers as well as crossing fibers with low VF or small separation angle, only a single FOD peak can be detected, irrespective of the underlying true response function profile. As a result, inclusion of voxels with such configurations will affect the RF calibration in the sense that it will become fatter than the true RF, ultimately leading to a lower angular resolution (Parker et al., 2013). This effect becomes more apparent when relatively few ‘true single fiber populations’ are present in the data. Furthermore, Tournier et al. (2007) state that, under specific simulation conditions (b -value of 3000 s/mm², SNR 20, 60 directions and exact RF), CSD cannot resolve fiber populations separated by angles smaller than 40° in a reliable way. Regarding the simulations in Fig. 5.4b, given our simulation settings (with a lower b -value of 2500 s/mm², non-exact RF, and inclusion of 90% two-fiber voxels as described in Section 5.2.2.2), it is to be expected that angles below 45° cannot be distinguished correctly from a single fiber population.

For the analyses on the in/ex vivo human brain data and the ex vivo spinal cord data, the recursive method converged after eight to ten iterations. For the human brain data, processing time for the RF calibration is in the order of five minutes.

5.4.4 Voxels used for RF calibration

For in vivo human data, the voxels that are eventually selected for RF calculation comprise FA values ranging from roughly 0.3 to 0.9 (see Fig. 5.6); hence, the final RF is not simply the average of high FA voxels. The FA values are approximately normally distributed with mean 0.67 and standard deviation 0.10. This indicates that there are single fiber populations in voxels with FA values smaller than 0.7 (which was confirmed by visual inspection); voxels that would otherwise have been discarded from RF calibration. As the initial distribution of all WM voxels has mean FA 0.37 (± 0.17) and is slightly positively skewed, a large amount of low FA voxels are discarded from RF calculation. One has to keep in mind, however, that the reliability of FA estimates is low at high b -values. In the WM, 25.2% of the FA > 0.7 voxels have a peak ratio

larger than 0.01, which is a significant amount and thus lends further support to the proposed method.

5.4.5 Potential impact for group studies

Selecting the correct RF is an important factor for SD approaches (Parker et al. 2013). Typically, most SD methods perform the RF calibration in a data driven way and for each data set separately. An alternative approach is to model the RF (e.g., defined as the DW signal profile that corresponds to an axially symmetric diffusion tensor with a specific FA and MD value) and keep it constant for each data set. With the advent of SD based quantitative measures, such as ‘apparent fiber density’ (Raffelt et al., 2012) and ‘hindrance modulated orientational anisotropy’ (Dell’Acqua et al., 2013), there are several factors related to the RF calibration step that could impact the validity of these new measures in group studies.

It is, for instance, not clear how pathology or age-related changes would affect the RF calibration and, in turn, how that would drive the final estimates of the diffusion metrics. This concern raises the question whether the RF should be computed for each data set separately, at a group level, or even across groups. And if the latter, there are many ways to achieve an ‘average RF’, adding more unknowns to the equation. Averaging, however, may not be optimal either, since the group RF could differ from the true RF for each subject. Alternatively, differences in diffusion properties between groups could even be reflected in the calibrated RF itself and, hence, comparing the RF between groups may provide us with useful information. Which RF calibration approach would be optimal for group studies remains an open question and would be an interesting topic for future investigations. Note that the above concerns are not specific to our proposed recursive framework for RF calibration, but more general to SD methods.

5.4.6 Limitations and possible improvements

Although our method has proven to overcome some of the limitations of FA method, there is room for improvement. The method still results in a single RF for the whole brain, which is likely an oversimplification. Theoretically, the recursive method could be applied on distinct brain regions to calibrate an individual RF that best ‘fits’ that particular region. In this way, we would not stick to one RF for the whole brain, but instead tune the RF for distinct brain regions. This would, however, result in a trade-off between spatial specificity and RF stability due to lower number of included WM voxels. More importantly, one may wonder whether the RF will vary only spatially, or whether it is more likely to vary from one to another fiber population in the same region.

Although CSD is the method of choice for SD in this work, there are also other implementations for SD approaches (Alexander, 2005; Anderson, 2005; Dell’Acqua et al., 2007; Dell’Acqua et al., 2010; Descoteaux et al., 2009). To demonstrate that our approach is more generally applicable, we performed an extra analysis in which the sharpening deconvolution transform (SDT) method of Descoteaux et al. (2009) was used. In contrast to the CSD method, where the RF represents the diffusion-weighted signal of a single fiber orientation, the SDT approach uses the dODF of a one-fiber profile as the RF. As such, our

algorithm was initialized with a dODF of a fat tensor ($FA = 0.05$, see Fig. 5.11a) and the PR threshold was set to 0.01. Spherical deconvolution of the Q-ball dODFs with the final calibrated RF then produces the fODFs (Fig. 5.11b). The framework can also be extended to non-SH-based methods. The (damped) Richardson-Lucy (RL) SD (Dell'Acqua et al., 2007; Dell'Acqua et al., 2010), for instance, does not use a spherical harmonic basis. Choosing the RF for RL-based SD may serve as a trade-off between angular resolution and noise stability with sharper RF profiles providing higher angular resolutions at the cost of lower stability (Parker et al., 2013). Hence, the recursive RF calibration procedure proposed in this work could also offer a valuable fine-tuning for the RL-based SD method.

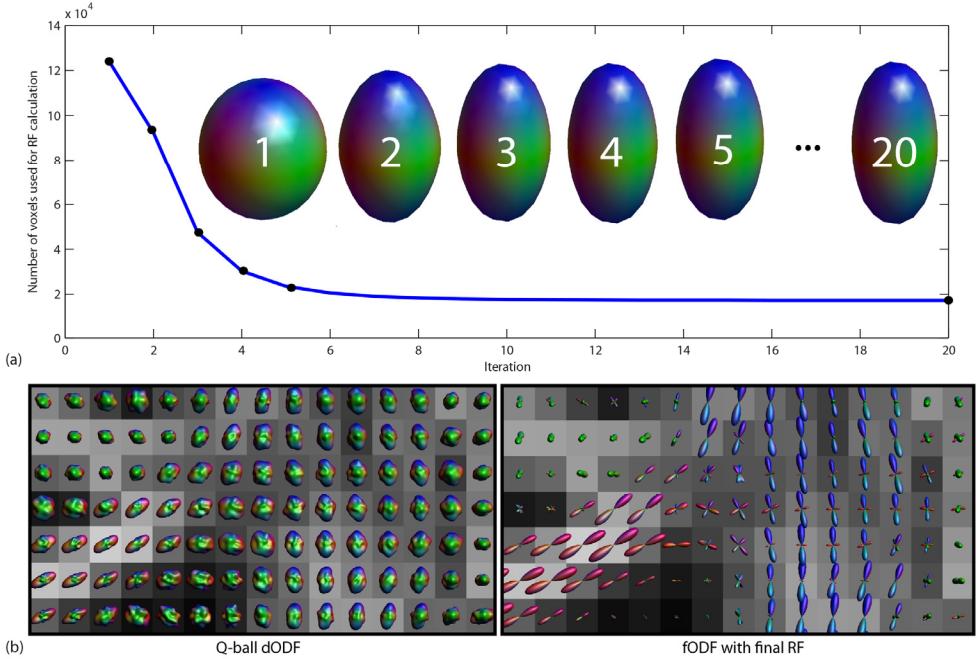


Fig. 5.11: Q-ball dODF sharpening, where the RF is calibrated with our recursive approach. (a) Number of voxels used for RF calibration as a function of iteration number, showing convergence after roughly 8 iterations. The RF represents in this case the dODF of a one-fiber profile and is shown as function of iteration. (b) Left: Q-ball dODFs; right: sharpened fODFs after SDT.

5.5 CONCLUSION

In this paper, a novel RF calibration method for SD approaches was presented. This method was compared to the commonly used method of calculating a mean RF of selected voxels with a range of FA values. Our recursive method has proven to overcome some of the limitations of the previously mentioned FA method and is completely independent of the diffusion tensor model. Voxels are selected for RF calibration when their fODFs have a sufficiently small or absent second peak as defined by the PR threshold, which allows us to make an informed and balanced trade-off between sensitivity and specificity of the fODF peaks. The method is robust and for SNR levels around 20, we recommend a PR threshold of 0.01. We have shown that

Chapter 5: Recursive Calibration of the Fiber Response Function for Spherical Deconvolution of Diffusion MRI Data

CSD with the recursive RF calibration approach yields plausible results in a variety of imaging data sets, both *in vivo* and *ex vivo*.

ACKNOWLEDGEMENTS

We would like to thank Laurena Holleran and Karla L. Miller for providing the *ex vivo* human brain data, and Tim B. Dyrby for the use of monkey spinal cord data. Furthermore, the authors thank Eleftherios Garyfallidis for the useful discussion on the sharpening deconvolution transform method.

6 CHARACTERIZING DIFFUSION MRI PROFILES OF SINGLE-FIBER POPULATIONS: AN INTER-SUBJECT AND INTER-TRACT ANALYSIS

Based on:

C.M.W. Tax, D.S. Novikov, S. St-Jean, E. Garyfallidis, M.A. Viergever, M. Descoteaux, A. Leemans, “Characterizing diffusion MRI profiles of single-fiber populations: An inter-subject and inter-tract analysis”, (to be submitted soon)

ABSTRACT

While a lot of efforts in diffusion MRI have been directed towards modelling of so-called ‘crossing-fibers’ in the brain, many fundamental questions on the microstructure in ‘single-fiber’ population (SFP) configurations still remain unanswered. For instance, studying SFPs can infer information on its fiber organization to distinguish different tissue compartments, on our ability to compute microstructural parameters like axon diameter, on the choice to define an appropriate response function for spherical deconvolution, and on realistic settings to design simulation experiments. Currently, however, most diffusion MRI methods define general properties of the SFP for the whole brain - typically derived from a region in the corpus callosum (CC) - and assume homogeneity in terms of diffusion profile characteristics across different fiber bundles. In this chapter, we challenge this assumption and investigate the spatial distribution of SFPs and their specific characteristics with diffusion MRI data from the HCP and MASSIVE databases. Our findings demonstrate that microstructural properties of the single-fiber diffusion profile vary significantly between tracts, where the inter-tract variability exceeds the inter-subject variability.

6.1 INTRODUCTION

With diffusion MRI (dMRI), the diffusion process of water particles can be measured to infer information on tissue architecture and microstructure. A popular feature of dMRI is the ability to extract fiber orientations for tractography to investigate brain connectivity. Diffusion tensor imaging (DTI), for example, can give an estimate of the underlying fiber direction by means of the principal direction of diffusion (Basser et al., 1994). However, DTI is unable to reveal multiple distinctly differently oriented fiber populations in a voxel, and a myriad of studies have therefore focused on optimal ways to solve the so-called ‘crossing-fibers problem’ over the past fifteen years (e.g. Dell’Acqua et al., 2007; Descoteaux et al., 2007; Descoteaux et al., 2009; Frank, 2002; Hess et al., 2006; Jeurissen et al., 2013; Ramirez-Manzanares et al., 2007; Tournier et al., 2004; Tournier et al., 2007; Tuch et al., 2002; Wedeen et al., 2005).

While a lot of attention has been given to crossing fibers in the brain, many fundamental properties of the microstructure in single-fiber population (SFP) configurations are still unknown. Because common DTI features such as fractional anisotropy (FA) and mean diffusivity (MD) are not specific to SFP tissue microstructure, other dMRI approaches are being developed to infer more specific information. For example, dMRI has been used to infer information on the organization of SFPs into distinct fiber compartments (Ferizi et al., 2015; Fieremans et al., 2011; Panagiotaki et al., 2012), to investigate the amount of neurite dispersion (Ferizi et al., 2013; Tariq et al., 2016; Zhang et al., 2012), and to test the ability of measuring microstructural parameters like axon diameter (Assaf et al., 2008) and amount of myelin. While some approaches attempt to recover crossing fibers and fiber-specific microstructural information at the same time (e.g. Ankele and Schultz, 2015; Assaf and Basser, 2005; Zucchelli et al., 2016), there appears to be a limit to the amount of parameters that can be estimated reliably from dMRI data. Even in the case of relatively simple models of tissue microstructure, parameter estimation suffers from poor precision and multiple physically plausible solutions can exist (Jelescu et al., 2016a). The remedy to this problem is usually to separate the orientation and microstructure estimation parts or to fix parameters to a predefined value that is then assumed to be the same throughout the brain. This spatial homogeneity of parameter values, however, is often used to simplify computations, but may not necessarily reflect the underlying microstructural reality.

In addition to studying tissue microstructure, there are other reasons why characterizing SFPs is useful. Information about the signal in SFPs can aid the design of an appropriate response function for spherical deconvolution (SD) approaches, a popular strategy to tackle crossing fibers (Dell’Acqua et al., 2010; Roine et al., 2015; Tax et al., 2014c; Tournier et al., 2004; Tournier et al., 2007). Furthermore, the development of dMRI methodology relies heavily on the appropriate design of simulation experiments. Incorporating reliable estimates of SFP signal characteristics as prior knowledge in such simulations is then crucial to predict and evaluate experimental findings in real tissue in an unambiguous way.

In order to characterize SFPs, it is important to localize voxels in which the condition is fulfilled that fiber bundles are relatively parallel and straight at the length scale of the voxel dimensions. Furthermore, localizing SFPs could also be useful to spatially constrain modelling approaches that assume the presence of a single dominant fiber orientation in a voxel. Many studies that investigate characteristics of SFPs and take advantage of such constraints use the

corpus callosum (CC) in the midsagittal plane where the fibers run relatively parallel (Ferizi et al., 2014; Panagiotaki et al., 2012). However, there is increasing evidence that even CC pathways are not strictly parallel and exhibit a degree of curvature and dispersion that cannot be ignored (Kleinnijenhuis et al., 2015; Mikula et al., 2012; Ronen et al., 2014). In addition, focusing only on the CC only ignores the fact that SFP properties may vary throughout the brain (Assaf et al., 2008; Fieremans et al., 2010; Fieremans et al., 2011; Reisert et al., 2014). Several alternative approaches that have looked beyond the CC to identify SFPs have used DTI constraints, such as high fractional anisotropy (FA) (Tournier et al., 2007) or a high linearity coefficient (Fieremans et al., 2011; Vos et al., 2012; Westin et al., 2002). While identifying SFP voxels is a necessary step in these works, it remains unclear to what extent these thresholds on DTI indices can affect the results. For example, a high FA may still occur in voxels with crossing fibers - e.g. when the volume fraction of one fiber population is relatively low. As these principles cannot be extended easily to high b -values due to assumption violations of the DT model and the introduction of estimation biases, it is desirable to avoid the diffusion tensor model for identifications of SFPs altogether (Tax et al., 2014c).

In this chapter, we study the characteristics of SFPs and their spatial variability across different fiber tracts. We adopt an alternative definition of an SFP that is not based on the DT: a voxel contains an SFP if the fiber orientation density function (fODF) only exhibits a single significant peak. More specifically, we use a previously proposed framework to recursively exclude crossing fiber voxels with constrained spherical convolution (CSD) (Tax et al., 2014c). This allows us to create an SFP map in which the fODF exhibits a single peak across b -values. We further localize the SFPs by intersecting the SFP map with tract masks of different bundles derived from tractography. By deriving features from a range of dMRI models, we investigate characteristics of SFPs between tracts and across subjects. The focus of this manuscript was not to compare different models or to discuss the validity of assumptions underlying these models, but to explore the features and their variability that can be derived from these models. Preliminary results of this work were presented at the ISMRM (Tax et al., 2015b).

6.2 METHODS

6.2.1 Data and processing

In this work we used the MASSIVE dataset (Froeling et al., 2016) and ten Human Connectome Project datasets (Glasser et al., 2013; Sotiropoulos et al., 2013; Van Essen et al., 2013). The MASSIVE dataset has a very high angular resolution with [250, 500, 500, 500] directions on the spheres and includes b -values of 500, 1000, 2000, and 3000 s/mm^2 , along with 170 $b = 0\ s/mm^2$ images. The spatial resolution is 2.5 mm isotropic. Images were corrected for signal drift, subject motion, eddy current distortions, and susceptibility induced deformations using *ExploreDTI* (Froeling et al., 2016; Leemans et al., 2009; Leemans and Jones, 2009; Vos et al., 2016b). The preprocessed HCP datasets have a higher spatial resolution of 1.25 mm isotropic, and 90 directions on the hemispheres with b -values 1000, 2000, and 3000 s/mm^2 , along with 18 $b = 0\ s/mm^2$ images. A white matter (WM) mask was derived for each dataset from the T1-weighted image using FSL-Fast (Zhang et al., 2001).

6.2.2 SFP localization by recursive CSD on multiple shells

The localization of SFPs in this work is done by recursively applying constrained SD (CSD). CSD is one of the preferred SD approaches to reconstruct the fODF from the data and a predefined response function (RF) (Tournier et al., 2007). Here, the RF was obtained with the recursive calibration approach described in Tax et al. (2014). In short, this recursive algorithm is initiated with a coarse RF (in this case a diffusion tensor with $MD = 0.7 \cdot 10^{-3} \text{ mm}^2/\text{s}$ and $FA = 0.2$), and CSD with $l_{max} = 8$ is performed on voxels within a WM mask. In each iteration, an ‘improved’ RF is computed by taking the mean rotated signal of voxels that contain only one genuine fiber orientation distribution (fODF) peak. Specifically, voxels with *peak ratio* second fODF peak / first fODF peak smaller than a preset peak ratio threshold t (here, $t = 0.05$) are considered to be SFP voxels. This process is repeated until the convergence criterion are met. An *SFP map* can be created by including voxels that are used to compute the final RF and still exhibit an SFP upon CSD with the final RF (Tax et al., 2014). To take into account the b -value dependency, we used the recursive framework to create an SFP map for every shell. A final SFP map was created by taking the intersection of the SFP maps over all shells.

To be able to extract more tract-specific information, SFP voxels were considered to belong to a particular tract if they intersected with the tract mask. The tract mask included every voxel through which a streamline of the corresponding tract passes. Deterministic tractography on the fODF resulting from CSD on the $b = 3000 \text{ s/mm}^2$ shell was performed with the following parameters: fODF threshold 0.1, angle threshold 35° , step size 1 mm (Jeurissen et al., 2011). We defined regions of interest on the MASSIVE and one HCP dataset to reconstruct the cingulum (CG), corticospinal tract (CST), inferior fronto-occipital fasciculus (IFOF), superior longitudinal fasciculus (SLF), uncinate fasciculus (UNC), occipital projection of the corpus callosum (CCO), and frontal projection of the corpus callosum (CCF) using the guidelines in Wakana et al. (2007). These tracts were reconstructed for the remaining HCP subjects using atlas based tractography as described in Lebel et al. (2008), with the first HCP subject as template. Spurious pathways were manually removed.

6.2.3 SFP characterization: raw signal

The signals of all SFP voxels were rotated so that the fODF peak resulting from CSD on the $b = 3000 \text{ s/mm}^2$ shell (Tournier et al., 2013b) coincides with the z-axis. Signals were interpolated on each shell using spherical harmonics ($l_{max} = 8$) to determine the signal decay in the direction parallel and the average signal perpendicular to the peak direction.

6.2.4 SFP characterization: model features

A wide range of dMRI models and representations is available to further characterize SFPs, and obtain more specific information on the diffusion process and tissue microstructure. To explore the characteristics of SFPs throughout the brain, we fit a subset of several popular models: 1) diffusion kurtosis imaging (DKI) model on the shells up to $b = 2000 \text{ s/mm}^2$ using REKINDLE with $\kappa = 6$ (Jensen et al., 2005; Tax et al., 2015c); 2) CSD on the $b = 3000 \text{ s/mm}^2$ shell with maximum spherical harmonics order $l_{max} = 8$ (Tax et al., 2014c;

Tournier et al., 2007); and 3) 3D simple harmonic oscillator based reconstruction and estimation (3D-SHORE) basis on all shells with maximum radial order $N_{max} = 6$ and Laplacian regularization with regularization parameter 0.2 (Fick et al., 2016; Garyfallidis et al., 2014; Merlet and Deriche, 2013; Özarslan et al., 2009).

We derive the following features and scalar maps: 1) diffusion tensor measures from the DKI fit, i.e. FA, mean diffusivity (MD), axial diffusivity (AD), and radial diffusivity (RD); 2) kurtosis tensor measures from the DKI fit, i.e. kurtosis anisotropy (KA), mean kurtosis (MK), axial kurtosis (AK), radial kurtosis (RK) (Poot et al., 2010); 3) white matter tract integrity (WMTI) measures from the DKI fit and a white matter model, i.e. axonal water fraction (AWF), axial extra-axonal diffusivity (AxEAD), radial extra-axonal diffusivity (radEAD), and tortuosity (Fieremans et al., 2011); 4) an fODF measure, i.e. peak apparent fiber density (AFD) (Dell'Acqua et al., 2013; Raffelt et al., 2012)) from CSD; and 5) ensemble average propagator (EAP) measures, i.e. mean-squared displacement (MSD), q -space inverse variance (QIV), return-to-origin probability (RTOP), return-to-axis probability (RTAP), and return-to-plane probability (RTPP) (Fick et al., 2016; Özarslan et al., 2013; Wu et al., 2008b) from the 3D-SHORE fit.

6.3 RESULTS

6.3.1 SFP localization by recursive CSD calibration on multiple shells

Fig. 6.1 shows the SFP map for every shell of the MASSIVE dataset overlaid on an FA map (left 4 columns). The number of voxels identified as SFP voxels slightly decreases with b -value (also in the HCP data, results not shown), but overall the same areas could be found. Fig. 6.1 right column shows the final SFP map. SFPs were located in the CC, optic radiation, cerebellar peduncle, fornix, and internal capsule, amongst others. Fig. 6.2 shows the final SFP maps for the ten HCP datasets. Overall, the SFP locations are consistent between HCP datasets and between the HCP and MASSIVE data. For the HCP data, the number of SFP voxels in the frontal part of the CC varies across subjects.

6.3.2 SFP characterization: raw signal

Fig. 6.3a shows the 3D mean signal decay over all SFP voxels for the MASSIVE dataset. Figs. 6.3b and 6.3c show respectively the signal decays of each individual voxel in the direction parallel and perpendicular to the fiber. The thick line represents the mean signal decay for each b -value. In both axial and radial directions, the signal decay is clearly non-monoexponential. Some axial signal decay curves of the HCP data show ‘bumps’ which indicates that the interpolated signal is not monotonically decreasing.

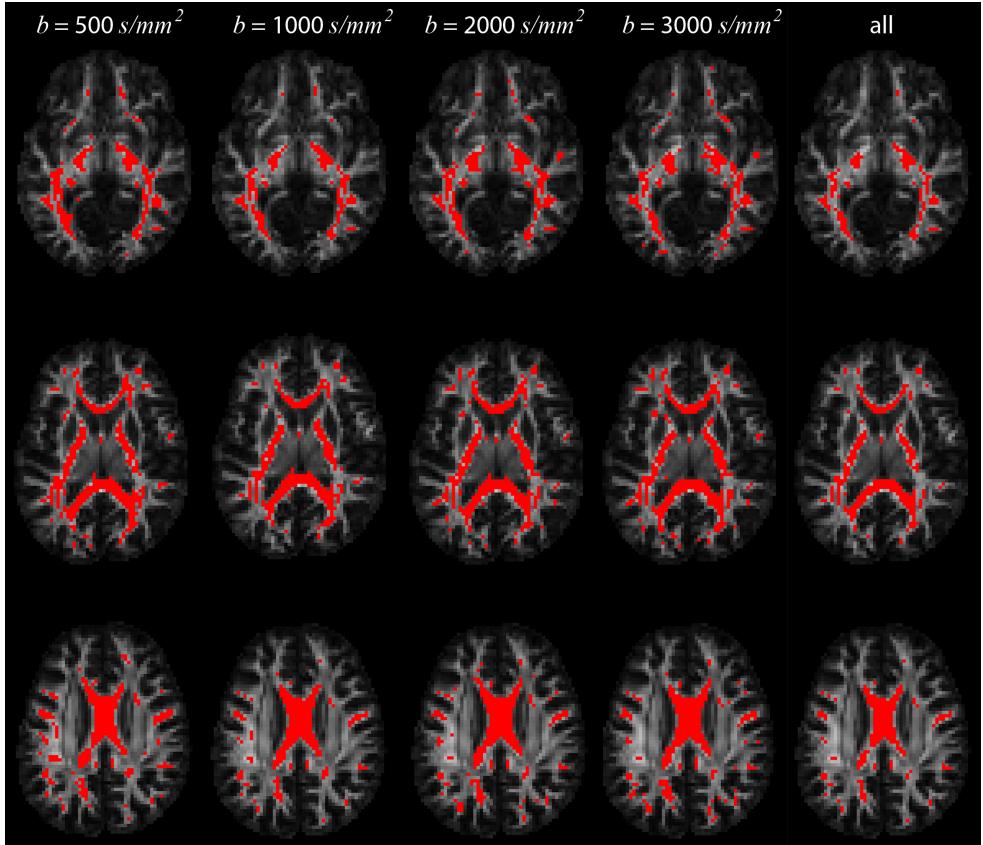


Fig. 6.1: Localization of SFPs in the MASSIVE data across b -values in three slices.

Fig. 6.4 shows results for SFPs located on different tracts. The reconstructed tracts are shown in Fig. 6.4a for the MASSIVE dataset and an HCP dataset. Fig. 6.4b displays signal decay curves for SFPs in the corresponding tracts for the MASSIVE dataset. Visual inspection already reveals a difference in the characteristics of decay curves such as spread, amount, and shape for different tracts. Fig. 6.4c shows the decay curves of SFPs intersecting with the CST tract mask for the MASSIVE dataset (left) and the ten HCP subjects (right). Curve characteristics are visually similar between HCP subjects, but slightly different from the MASSIVE dataset.

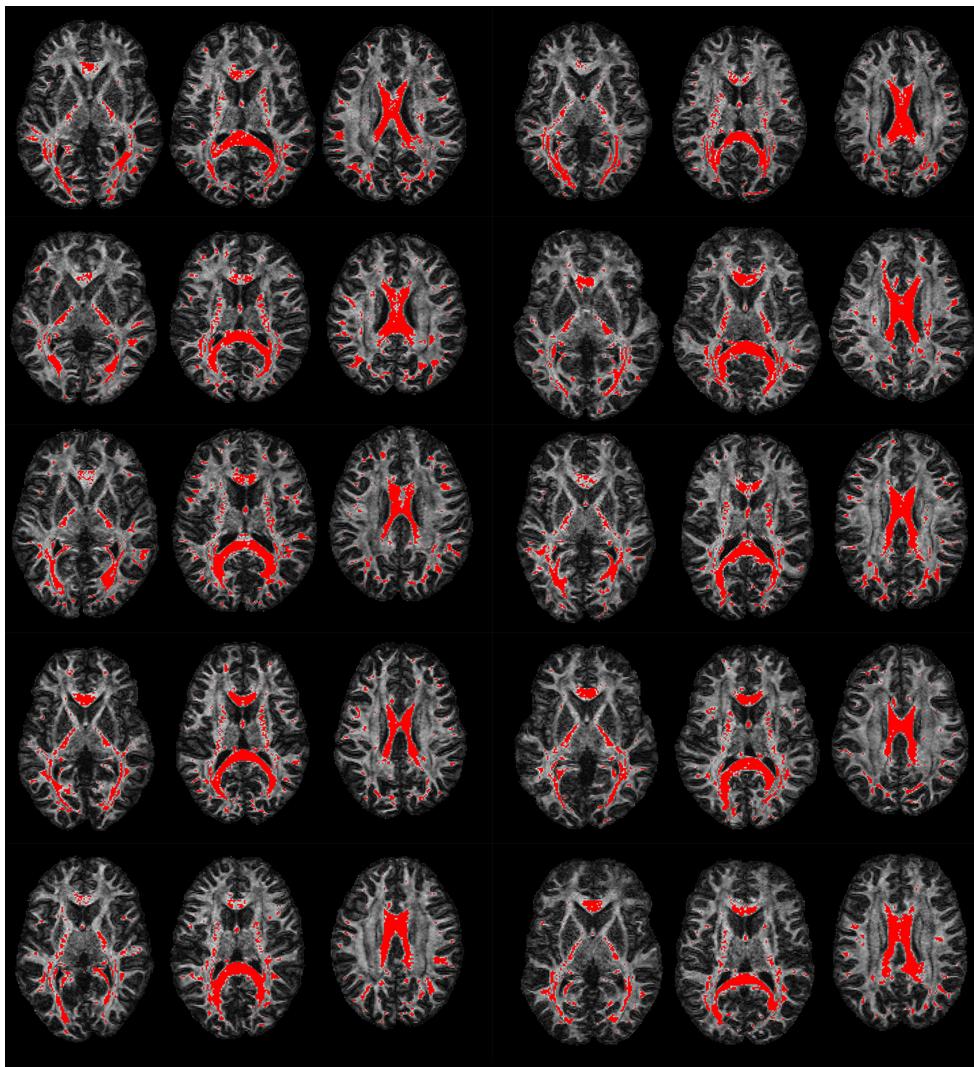


Fig. 6.2: SFP maps for ten HCP subjects in three slices.

6.3.3 SFP characterization: model features

Fig. 6.5 shows DTI (a) and DKI (b) features for different tracts (colors correspond to Fig. 6.4a). The left column shows median values only across the SFP voxels on the tract. As a comparison, the right column shows the median value across all the voxels in the tract mask. Mean or median tract values are often used in group studies to investigate pathology, but are less ‘tract-specific’ since also voxels with crossing fibers are taken into account. Each point is the median value for a subject, where the asterisk represents the MASSIVE data and each circle represents an HCP subject. The plots in Fig. 6.5 show variation of the features between tracts, both for the SFPs and whole tract mask. The mean value and variability of the features across subjects is

different for different tracts. Overall, the variability between tracts already qualitatively appears larger than the variability between subjects. For example, the median AD for all subjects is higher in SFPs of the occipital part of the corpus callosum (CCO) than of the SLF, and the median RK for all subjects is higher in the CST than the uncinate fasciculus when considering the whole tract mask. The variability between subjects is roughly similar when considering only SFPs on a tract compared to considering the whole tract, even though the former involved fewer voxels. As an example, we tested whether there was an overall significant difference between tracts for the AD with a repeated measures analysis of variance. In both the SFP and whole tract scenarios there was a significant effect of tract ($F(6,54) = 44.5, p = 3.6 \cdot 10^{-19}$, effect size (partial η^2) = 0.83 and $F(6,54) = 108.2, p = 2.8 \cdot 10^{-28}$, effect size = 0.92, respectively). The fingerprint of typical values for the tracts is, however, markedly different between the two scenarios: the FA, AD, MK, KA, and RK are generally lower for the whole tract mask whereas the RD was generally higher. MD showed the least difference between the two scenarios. The features are the most variable across subjects in the frontal part of the corpus callosum (CCF). Overall, the MASSIVE and HCP data features follow the same trend across tracts, but the AK is systematically lower for the MASSIVE data.

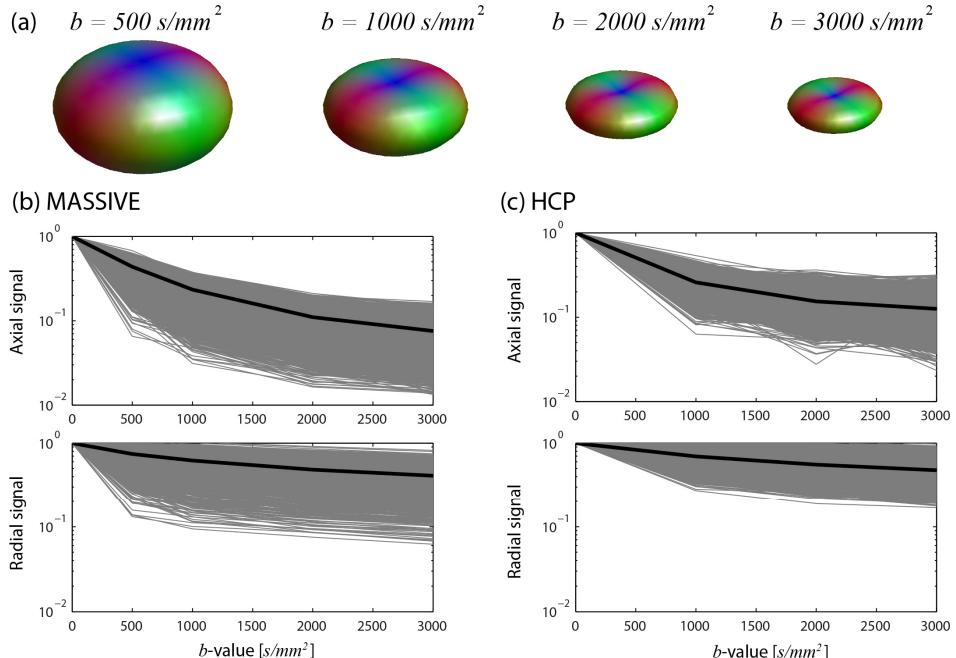


Fig. 6.3: (a) 3D mean signal decay over all SFP voxels in the MASSIVE data as a function of b -value. (b) Axial and radial signal decay for each individual SFP voxel (grey), and the average signal decay (black) in the MASSIVE data. (c) Signal decay for a HCP dataset.

Chapter 6: Characterizing Diffusion MRI Profiles of Single-Fiber Populations: An Inter-Subject and Inter-Tract Analysis

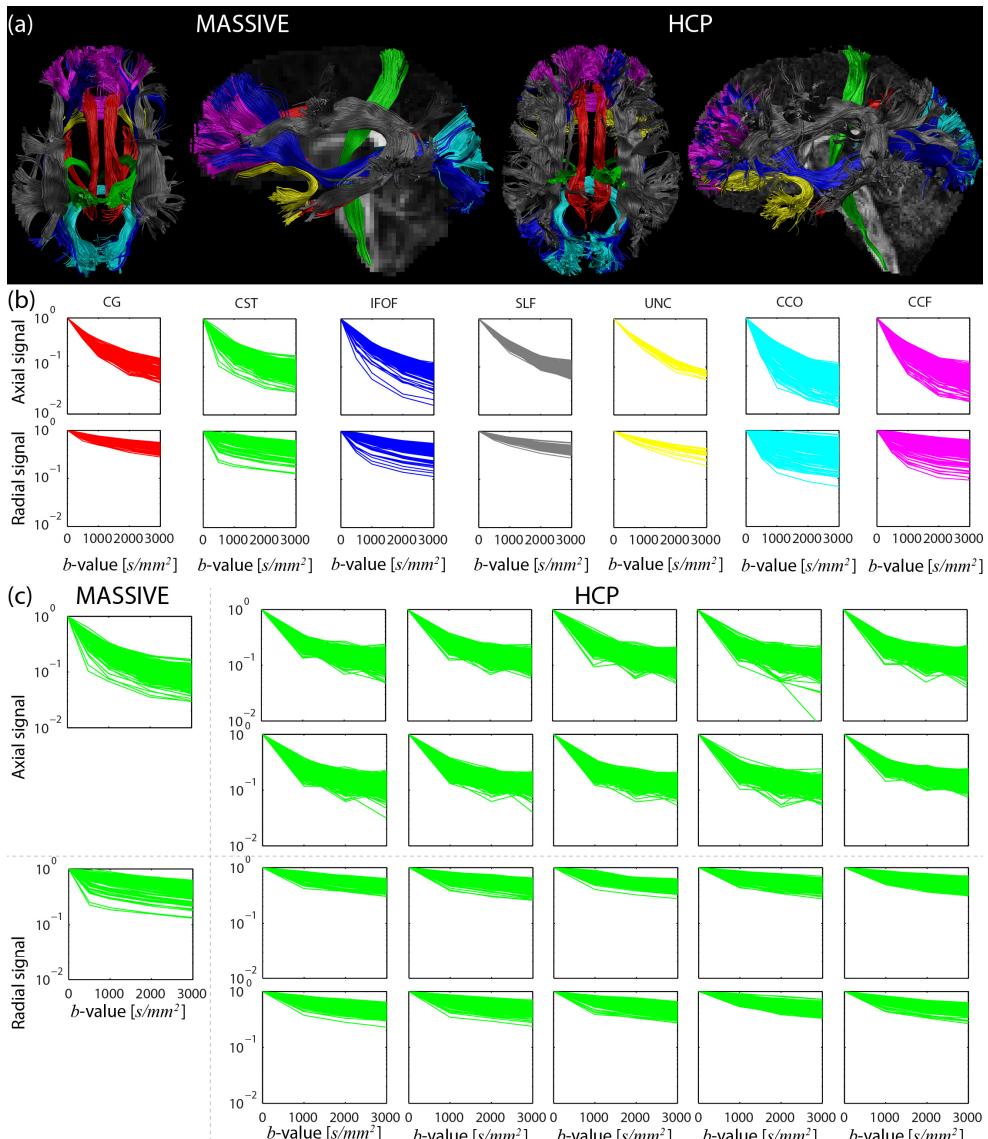


Fig. 6.4: (a) Reconstructions of tracts for the MASSIVE (left) and one HCP dataset (right): cingulum (CG, red), corticospinal tract (CST, green), inferior fronto-occipital fasciculus (IFOF, dark blue), superior longitudinal fasciculus (SLF, grey), uncinate fasciculus (UNC, yellow), occipital projection of the corpus callosum (CCO, light blue), and frontal projection of the corpus callosum (CCF, purple). (b) Signal decay curves for the MASSIVE dataset of SFPs located on different tracts, colors correspond to (a). (c) Signal decay curves of SFPs intersecting with the CST mask, for the MASSIVE dataset (left) and the ten HCP datasets (right).

Fig. 6.6 shows WMTI features (a), an fODF feature (b), and EAP features (c) for SFP voxels on different tracts in HCP subjects. These results confirm the finding that the difference of median features between tracts is in some cases higher than the inter-subject variability. As an example, we statistically confirmed that there was a significant effect of tract for the AFD with a repeated

Less Confusion in Diffusion MRI

measures analysis of variance ($F(6,54) = 78.2, p = 7.7 \cdot 10^{-25}$, effect size = 0.90). Similar trends are observed for the other metrics.

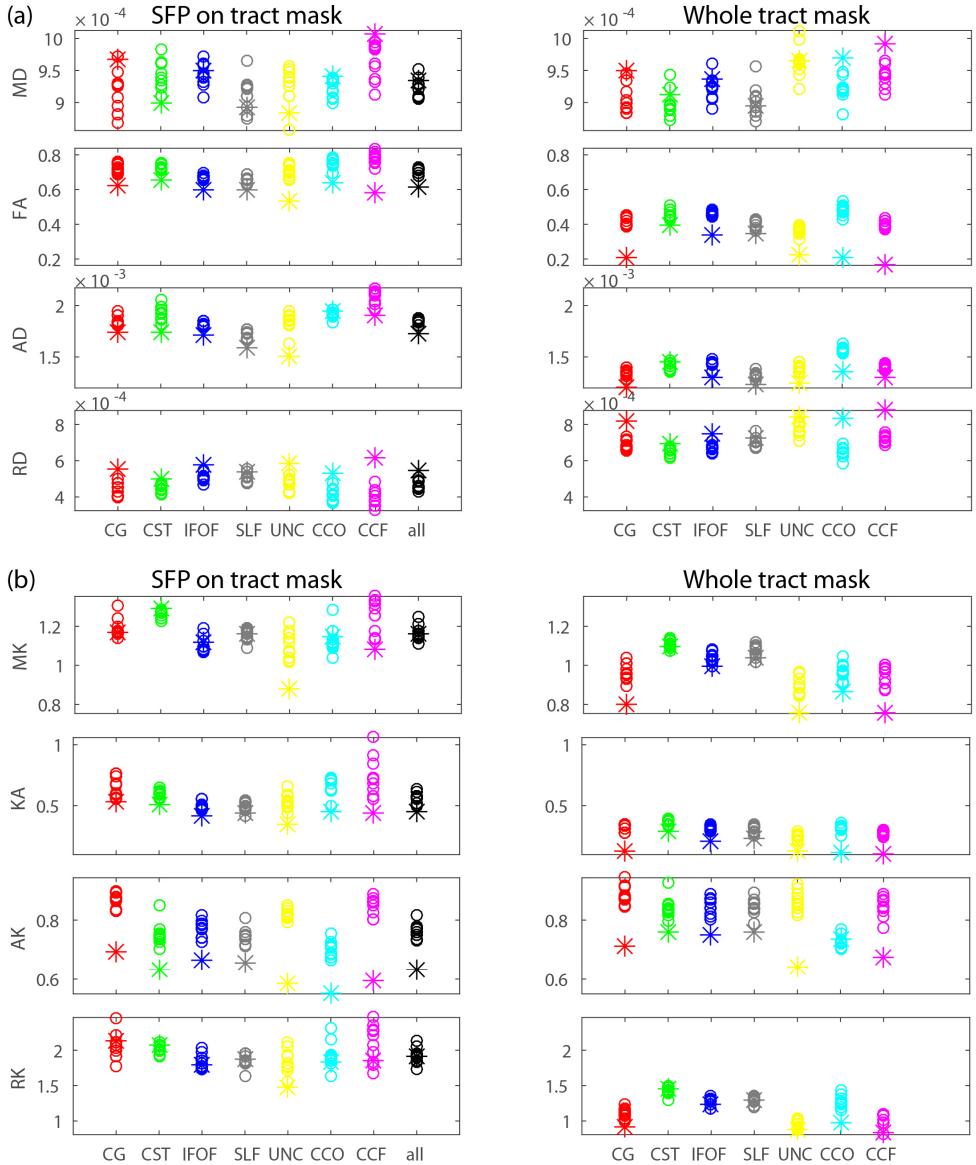


Fig. 6.5: Features derived from the DKI fit for different tracts, represented by different colors. Each point is the median value of the corresponding feature for a single subject, where the asterisk represents the MASSIVE data and the circles represent the HCP datasets. Left column: Median of SFP voxels on a given tract. Right column: Median of all voxels in the tract mask. (a) Features derived from the diffusion tensor. (b) Features derived from the kurtosis tensor. MD = mean diffusivity in mm^2/s , FA = fractional anisotropy, AD = axial diffusivity in mm^2/s , RD = radial diffusivity in mm^2/s , MK = mean kurtosis, KA = kurtosis anisotropy, AK = axial kurtosis, RK = radial kurtosis. CG = cingulum, CST = corticospinal tract, IFOF = inferior fronto-occipital fasciculus, SLF = superior longitudinal fasciculus, UNC = uncinate fasciculus, CCO = occipital part of the corpus callosum, CCF = frontal part of the corpus callosum, all = median over all SFPs.

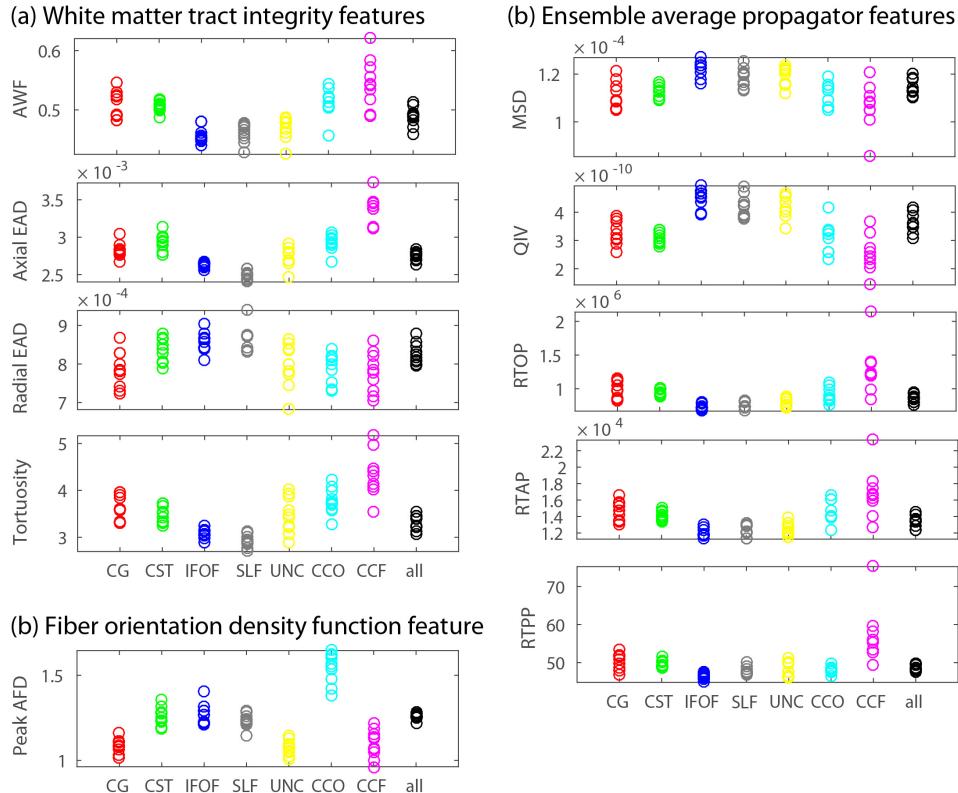


Fig. 6.6: Features for different tracts, represented by different colors. Each point is the median value of the corresponding feature for a single subject. (a) Features derived from white matter model of Fieremans et al. (2011). (b) Feature derived from the FODF. (c) Features derived from the EAP. AWF = axonal water fraction, EAD = extra-axonal diffusivity in mm^2/s , AFD = apparent fiber density, MSD = mean-squared displacement in mm^2 , QIV = q -space inverse variance in mm^5 , RTOP = return-to-origin probability in mm^{-3} , RTAP = return-to-axis probability in mm^{-2} , RTPP = return-to-plane probability in mm^{-1} . CG = cingulum, CST = corticospinal tract, IFOF = inferior fronto-occipital fasciculus, SLF = superior longitudinal fasciculus, UNC = uncinate fasciculus, CCO = occipital part of the corpus callosum, CCF = frontal part of the corpus callosum, all = median over all SFPs.

6.4 DISCUSSION

In this work we have investigated the spatial distribution of SFPs and their specific characteristics with diffusion MRI data from the HCP and MASSIVE databases. We localized SFP voxels by excluding crossing fiber voxels with CSD on multiple b -value shells. The signal in these SFP voxels can be further characterized by fitting any dMRI model or representation. Here, we have analyzed the variability of features from a subset of dMRI models throughout the brain.

6.4.1 Definition of a single fiber population

We have explored the definition of an SFP as a voxel in which the fODF only exhibits one significant peak across b -values. This definition is not based on manual localization of structures such as the CC, or on the DTI model. The idea is that the peak ratio threshold provides a more intuitive means to select SFP voxels than DTI indices, that it can be extended to higher b -values, and that it can be tuned to take into account spurious peaks (Tax et al., 2014c). However, the definition of SFPs that we used is open to discussion. The fODF in a voxel will have a single significant peak in a number of configurations other than straight parallel fibers. For example, in the case of bending and fanning fibers the fODF is expected to still exhibit a single peak which is broader than in the case of straight parallel fibers. In addition, fibers that cross or kiss will result in a single fODF peak when their angle is smaller than the resolving power of the dMRI technique used to obtain the fODF. For CSD as used here this is around 45° (Tournier et al., 2007). On the one hand, the definition of an SFP used in this work is thus not the same as the strict definition of an SFP as straight parallel fibers. On the other hand, it is questionable whether such strict SFP configurations actually exist in brain tissue at the scale investigated, since even the CC exhibits curvature and dispersion (Kleinnijenhuis et al., 2015; Mikula et al., 2012; Ronen et al., 2014). Therefore, relaxing the definition of an SFP to ‘relatively straight and parallel fibers’ in the context of brain dMRI at this spatial scale might be necessary, and the fODF then provides a more general way to capture such different configurations.

Assuming that the fODF with a single significant peak can represent a range of fiber configurations, there are a number of ways in which these configurations could potentially be further disentangled. For example, in addition to a peak ratio threshold, an additional criterion could be imposed that only selects voxels with low fiber dispersion (e.g. based on the integral of the fODF and ratio of the peak amplitude). However, if one is interested in investigating fiber dispersion in SFPs, it is desirable to leave these voxels in the SFP map. Alternatively, neighborhood information can be used to infer information on the asymmetry of ODFs, potentially revealing configurations such as bending or fanning fibers (Barmoutis et al., 2008; Bastiani et al., 2016; Cetin et al., 2015). Overall, the problem of not being able to distinguish different complex fiber configurations from each other is a problem general to the field, and further developments in this direction can also be beneficial for the framework presented here.

6.4.2 Localization of single fiber populations

In this work we have used CSD with the recursive calibration method to localize voxels with a single significant fODF peak across all b -value shells (Tax et al., 2014c). In this algorithm, different choices for parameter settings and selection criteria can be made. The peak ratio threshold was proposed as an intuitive parameter to select SFPs based on the magnitude of the largest and second largest peak. If the value is set too low the algorithm might exclude true SFPs because of the presence of spurious peaks, and if the value is set too high the algorithm might include voxels with crossing fibers. In practice, the ‘optimal’ parameter setting might also depend on SNR and acquisition settings. We ran the algorithm with different peak ratio thresholds on the MASSIVE and HCP datasets, and visually inspected SFP maps across datasets and b -values. For simplicity we here chose one parameter setting (0.05) that selected

the majority of the CC and other areas where SFPs are likely to be present (parts of the internal capsule and the optic radiation) for all datasets and b -values. Alternative criteria for the selection of fODFs with a single peak are possible, for example by looking at absolute peak values or comparing the AFD of the largest peak to the sum of the AFDs of all other positive lobes (Smith, 2016). We found that the simpler peak ratio threshold combined with a WM mask performed good in our case, and other criteria might be further investigated in future work. Another setting involved the choice of initial response function. In analogy to previous work we used a near-isotropic response function corresponding to a diffusion tensor with low FA. Here we chose an initial FA of 0.2 to prevent CSD from producing degenerate fODFs at the largest b -value shell of the HCP dataset. Based on visual inspection on lower b -value shells of the HCP data and all shells of the MASSIVE dataset, the final SFP maps did not change significantly when this threshold was lowered. Alternative choices for the initial response function are possible, for example to better deal with the intensity scaling. However, as we used the ratio between peaks to select voxels for the next iteration, a scaling of the first-iteration fODFs resulting from the scaling of the initial response function theoretically does not influence the result.

A final SFP map was created by taking the intersection of the SFP maps per shell to exploit the information present in multiple b -shells. The number of SFP voxels found decreased with b -value, likely because of the increased ability to resolve crossing fibers at higher b -value, but also due to the lower SNR at higher b -values which may cause spurious peaks. For the HCP data this effect was more visible (results not shown) than for the MASSIVE dataset (Fig. 6.1): the increase of amount of directions for increasing b -value in the MASSIVE dataset might have compensated for this effect. Alternative ways to obtain a final SFP map are possible. We compared our recursive exclusion of crossing fiber voxels, which does not reconsider a voxel once it is excluded, with a peak ratio threshold on all the voxels after final CSD and with an absolute peak threshold of 0.1 on all the voxels after final CSD. The results in Supplementary Fig. S6.1 show that the first two scenarios (recursive vs peak ratio threshold on all the voxels) result in the same SFP map both for the MASSIVE and HCP dataset. A relative versus absolute threshold results in a similar SFP map with the same structures. Other adaptations to the creation of the SFP map include a relaxation of the constraint that the fODF has to exhibit a single peak for all shells to a subset of the shells. Alternatively, multi-shell spherical deconvolution (Cheng et al., 2014; Jeurissen et al., 2014) can be performed to do recursive calibration on all shells simultaneously. These options can be explored in future work.

SFP locations were visually consistent between subjects (Fig. 6.2). SFP locations found in the HCP data were mostly also found in the MASSIVE data. Due to the larger voxel size and corresponding partial volume effects some SFP locations were not found as extensively in the HCP data as in the MASSIVE data. An example is the frontal part of the CC, where the amount of SFP voxels also varied between HCP subjects. This could for example be caused by Gibbs ringing (Perrone et al., 2015; Veraart et al., 2016), CSF partial volume effects, a high curvature, or acquisition or processing imperfections. In addition, the MASSIVE and HCP datasets were acquired with a different phase encoding direction (i.e. anterior-posterior vs left-right), which might have attributed to the differences (Kennis et al., 2016).

To further localize an SFP voxel on a particular tract, we intersected the tract mask with the SPF mask. Using the tract mask to compute summary statistics of the corresponding tract (e.g. mean FA) is a common analysis strategy. Instead of considering all the voxels in the tract mask, we have here focused on features only of SFPs in different tracts. By computing mean indices only from SFPs in a tract instead of averaging over all voxels, the results are to a lesser extent biased by crossing fiber voxels and other partial volume effects. This can reveal potential between-tract or between-subject differences that might otherwise remain undetected. An SFP voxel is, however, not always tract specific: when pathways of multiple fiber tracts run more or less parallel through a voxel, one SFP voxel can be attributed to multiple tracts.

6.4.3 Characterization of single fiber populations

Figs. 6.3 and 6.4 show axial and radial signal decay curves for SFPs located on different tracts, and SFPs located on the same tract in different datasets. Here, we interpolated the signal decays in the respective directions from spherical harmonics fits on each shell. The signal decay both in the axial and radial direction appears to be non-monoexponential. To test whether this was not an effect solely of the noise floor, we determine the signal decay in the same directions after an attempt to adjust for the non-Gaussianity of the original signal. The non-central chi distributed signal is mapped to its equivalent value of a Gaussian distribution as described in Koay et al. (2009a); St-Jean et al. (2016). Although a slight change in the signals of individual SFP voxels could be observed (Supplementary Fig. S6.2), the curves still showed a non-monoexponential behavior both in the radial and axial direction.

Many attempts have been made to characterize decay curves and their non-monoexponential behavior (Mulkern et al., 2009). The fitting of biexponential curves has been much debated since it is highly challenging, and even though it fits the data well the assignment of the diffusivities to histological compartments seems impossible (Kiselev and Il'yasov, 2007; Mulkern et al., 2009; Niendorf et al., 1996). Several works, however, have found some experimental and theoretical justification for this model in WM: the slow compartment would correspond to the highly restricted water diffusion inside the axis and the fast compartment to the less hindered diffusion the extra-axonal space (Assaf and Basser, 2005; Fieremans et al., 2010). Fig. S6.3 shows results of a biexponential fit on axial and radial decay curves of SFP voxels intersecting with the tract mask of the CST. Here we used a non-linear least squares fit (Levenberg-Marquardt) on the data of all SFPs on a tract simultaneously where the weights were determined by the reciprocal of the variance for each b -value. We have tested other approaches to biexponential fitting, for example a fit for every individual voxel before taking the mean which was highly unstable and therefore not used, and a fit on the mean over all SFP voxels under consideration with different optimization options as in Fieremans et al. (2011); Kiselev and Il'yasov (2007). The results for these options were more or less the same as shown in Fig. S6.3. In this work we used a maximum b -value close to Kiselev and Il'yasov (2007) but with fewer different b -values and a smaller maximum b -value and fewer different b -values than in Fieremans et al. (2011).

The numbers in each plot indicate parameter estimates for the slow and fast diffusing compartments (D_s and D_f respectively), with the volume fraction f corresponding to the slow diffusing compartment. The diffusion coefficients differ between compartments in a given direction. The parameter estimates were similar between HCP subjects and relatively similar

between the MASSIVE and HCP datasets. The largest discrepancy occurred in D_s for the axial signal decay and D_f for the radial signal decay. The estimates of the diffusion coefficients in the radial direction were in the same order of magnitude, but larger than the values reported in Fieremans et al. (2011). This could be attributed to a lower maximum b -value in this work, a different tract under investigation, or other factors. The parameter estimates for the slow and fast diffusing compartments suggest that diffusion coefficients can differ between compartments in a given direction, contrary to some assumptions in literature (Panagiotaki et al., 2012; Zhang et al., 2012) . At this point, however, it is problematic to assign the slow and fast diffusing compartments to intra- and extra-axonal space since the estimated fractions of the decay curves are not consistent between both directions.

We fitted different other dMRI models to the data in SFP voxels to get a feeling for the characteristics of SFPs and their variability throughout the brain. Figs. 6.5 and 6.6 show that the variability of most features is larger between various tracts than between subjects. When comparing median DTI and DKI features for SFP voxels located on a tract or for the whole tract mask, we can see that these values are significantly biased by crossing fibers (Fig. 6.5). As expected, FA, AD, and KA of SFPs only are much higher compared with the whole tract average, whereas RD is lower in this comparison. While most features are comparable between the MASSIVE and HCP datasets, the AK is not. This may be due to sensitivity of AK to noise as it is evaluated in a single direction (a similar effect as the noisier axial decay curves in Fig. 6.4 compared to the radial decay curves), and the fact that this effect is different between different acquisitions. AD, RD, AK, and RK are directional features related to the direction of the first eigenvector of the tensor. Assuming that the first eigenvector better represents the direction of the underlying fiber configuration in SFP voxels compared to crossing fibers, the terms radial and axial have a less ambiguous interpretation in this context.

The values in Fig. 6.6 roughly correspond to those presented in previous works (e.g. Fieremans et al. (2011) for WMTI values and Fick et al. (2016) for EAP values). The WMTI parameters in Fig. 6.6 are derived from the DKI fit and aim to give a more specific interpretation of DTI and DKI features in terms of tissue microstructure. The underlying white matter tissue model assumes that the intra-axonal space (IAS) and extra-axonal space (EAS) are two non-exchanging compartments, and that the diffusion in the IAS and EAS can be modeled by compartment specific diffusion tensors (Fieremans et al., 2011). The assumptions underlying this model are valid when the crossing angle of two axonal bundles is smaller than 30°, which is slightly lower than the resolving power of CSD. The derived features are variable between tracts, in agreement with the maps in Fieremans et al. (2011). The axial EAD seems to be the most discriminative between tracts, taking into account the inter-subject variability.

The AFD was derived in the peak direction of the fODF. The AFD measure is based on the assumption that intra-axonal radial diffusion is restricted and that the extra-axonal radial signal is strongly attenuated at b -values typically used for CSD (Raffelt et al., 2012). Under these assumptions, the magnitude of the fODF provides a measure of apparent fiber density in that direction. Also here, the peak AFD was highly variable between tracts, while the inter-subject variability was relatively low.

The EAP features derived from the SHORE basis also showed variability between tracts, which exceeds the inter-subject variability for most tracts. Of these features, the RTAP and RTTP

depend on the direction of the underlying first eigenvector of the diffusion tensor. Under the assumptions of parallel axons, infinitesimally small gradient pulse length, long enough diffusion times, and that the signal only originates from the intra-axonal compartment, the RTAP is related to the mean diameter of the axons (Özarslan et al., 2013). However, these conditions are hard to fulfill in real tissue, where dispersion almost always occurs and the diffusion signal also originates from the extra-axonal compartment.

Fig. S6.4 shows a plot where the median value of each tract of each HCP subject is visualized as a point in 3D space according to its axial EAD, AK, and AFD. We have chosen these features here based on a visual inspection of variability between tracts in Figs. 6.5 and 6.6. Different clouds of points can be recognized, visually confirming that tracts have a ‘fingerprint’ in terms of the features of their SFPs that is relatively consistent between subjects.

The computed indices can be used as reference for simulation studies. In future work a wider range of models and descriptive features can be estimated from the diffusion profiles of these SFPs. For example, an interesting feature not investigated in this work is the degree of neurite dispersion (Zhang et al., 2012). The MASSIVE decay curves and signals in SFPs and SFPs per tract will be made available to the community at <http://massive-data.org/>.

6.5 CONCLUSION

In this work we have studied characteristics of SFPs and their variability throughout the brain. We did this by recursively excluding crossing fiber voxels with CSD on multiple b -value shells. The signal in these SFP voxels was subsequently further characterized by fitting various dMRI models. Our results show that features derived from different models show already qualitatively a consistent variability between tracts, which exceeds the inter-subject variability. This was confirmed statistically for several measures. The proposed framework and identified SFP signals will be made publically available to facilitate further characterization of SFPs. In addition to providing useful microstructural information, these findings could for example be used in the design of simulation experiments.

6.6 SUPPLEMENTARY MATERIAL

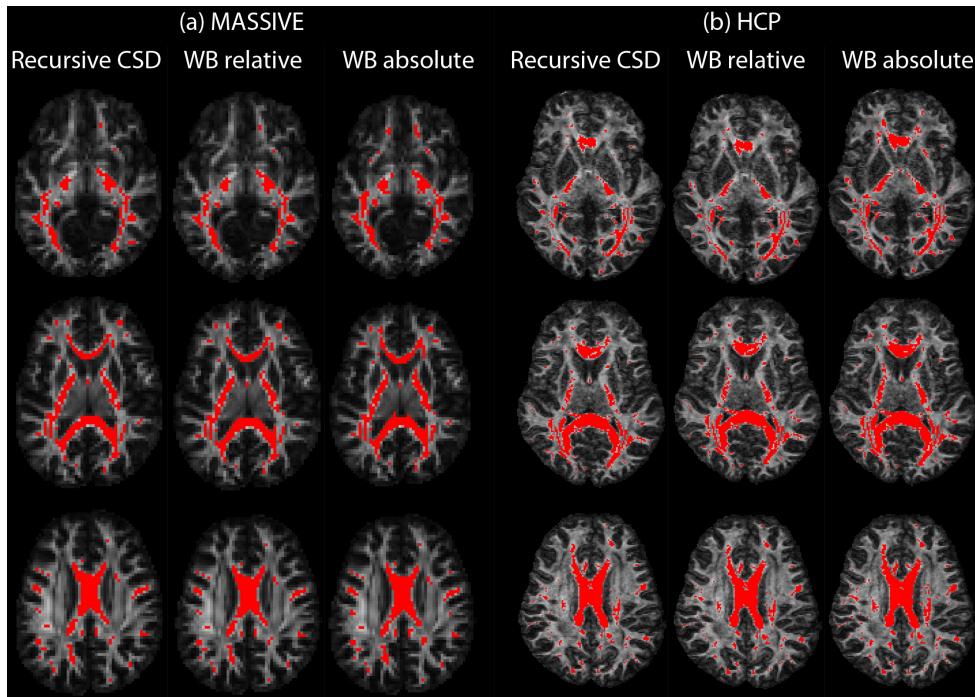


Fig. S6.1: Comparison of different options to create an SFP map for the MASSIVE (a) and one HCP (b) dataset. The recursive CSD is the approach used throughout the manuscript. whole brain (WB) relative refers to a peak ratio threshold on all the voxels after final CSD. WB absolute refers to an absolute threshold of 0.1 on all the voxels after final CSD.

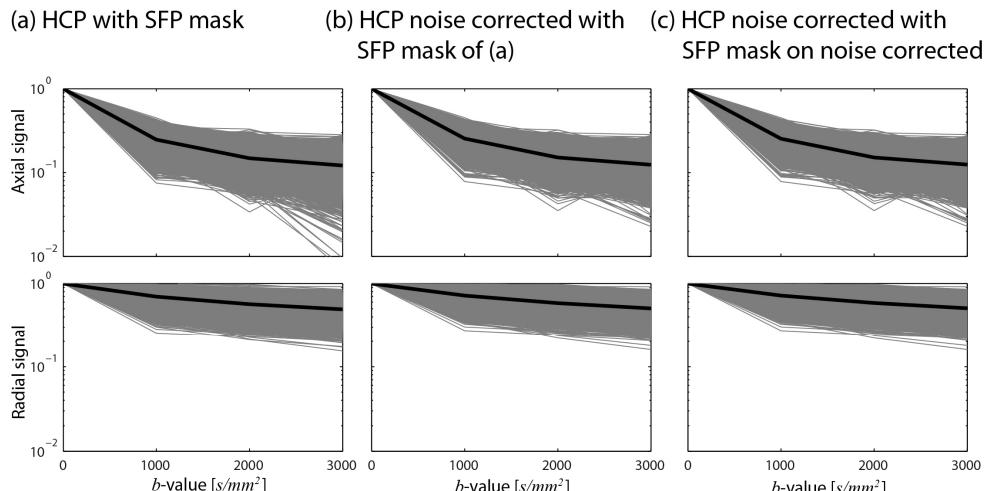


Fig. S6.2: Decay curves prior (a) and after (b,c) noise bias correction. In (b) the same SFP map is used as in (a), whereas in (c) we repeated the creation of the SPF mask on the actual noise corrected data.

Less Confusion in Diffusion MRI

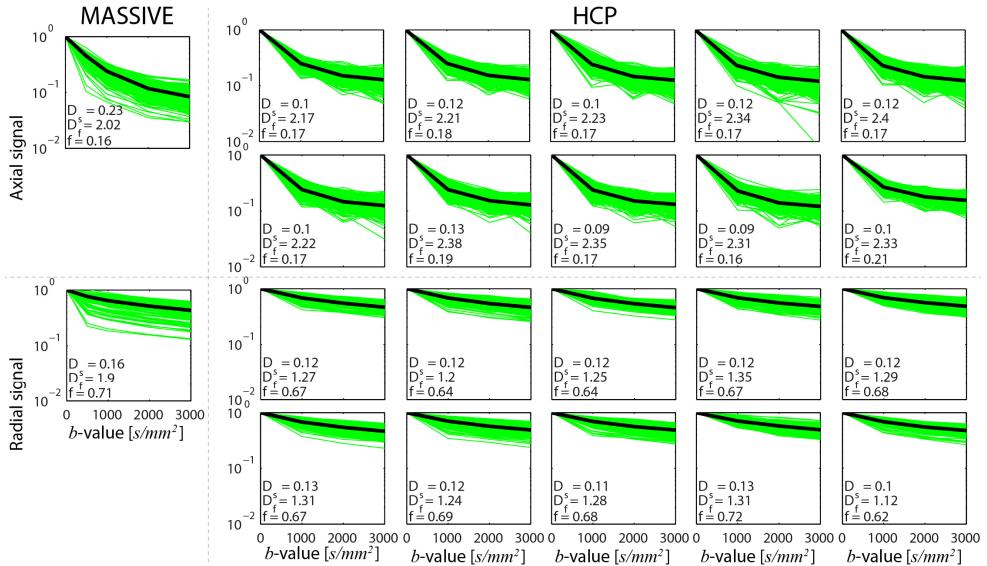


Fig. S6.3: Biexponential fit on the axial (top) and radial (bottom) decay curves for the MASSIVE (left) and HCP (right) datasets.

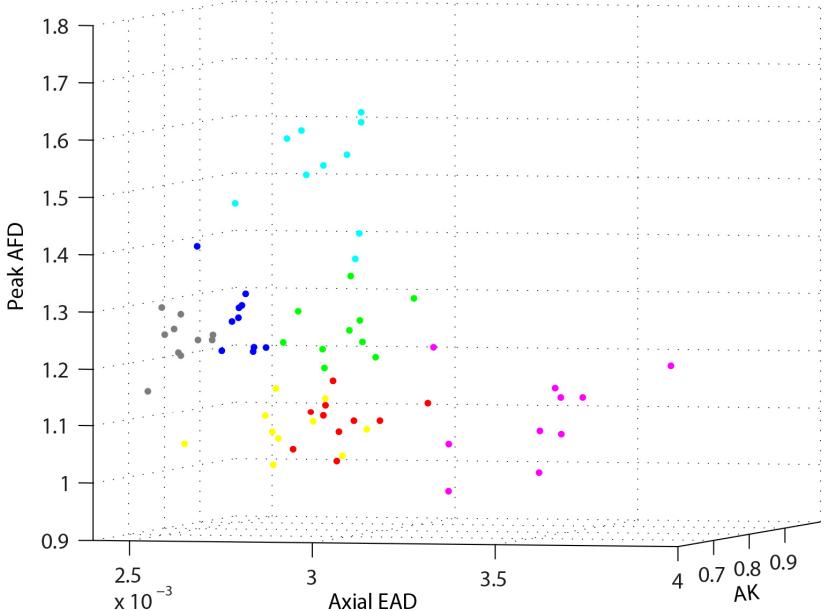


Fig. S6.4: Each point represents a tract (colored according to Fig. 4a) of one HCP subject as a function of its peak AFD, axial EAD, and AK.

7 SHEET PROBABILITY INDEX (SPI): CHARACTERIZING THE GEOMETRICAL ORGANIZATION OF THE WHITE MATTER WITH DIFFUSION MRI

Based on:

C.M.W. Tax, T.C.J. Dela Haije, A. Fuster, C.F. Westin, M.A. Viergever, L.M.J. Florack, A. Leemans, "Sheet Probability Index (SPI): Characterizing the geometrical organization of the white matter with diffusion MRI", NeuroImage (in press)

ABSTRACT

The question whether our brain pathways adhere to a geometric grid structure has been a popular topic of debate in the diffusion imaging and neuroscience society. Wedeen et al. (2012ab) proposed that the brain's white matter is organized like parallel sheets of interwoven pathways. Catani et al. (2012) concluded that this grid pattern is most likely an artifact, resulting from methodological biases that cause the tractography pathways to cross in orthogonal angles. To date, ambiguities in the mathematical conditions for a sheet structure to exist (e.g. its relation to orthogonal angles) combined with the lack of extensive quantitative evidence have prevented wide acceptance of the hypothesis. In this chapter, we formalize the relevant terminology and recapitulate the condition for a sheet structure to exist. Note that this condition is not related to the presence or absence of orthogonal crossing fibers, and that sheet structure is defined formally as a surface formed by two sets of interwoven pathways intersecting at arbitrary angles within the surface. To quantify the existence of sheet structure, we present a novel framework to compute the sheet probability index (SPI), which reflects the presence of sheet structure in discrete orientation data (e.g. fiber peaks derived from diffusion MRI). With simulation experiments we investigate the effect of spatial resolution, curvature of the fiber pathways, and measurement noise on the ability to detect sheet structure. In real diffusion MRI data experiments we can identify various regions where the data supports sheet structure (high SPI values), but also areas where the data does not support sheet structure (low SPI values) or where no reliable conclusion can be drawn. Several areas with high SPI values were found to be consistent across subjects, across multiple data sets obtained with different scanners, resolutions, and degrees of diffusion weighting, and across various modelling techniques. Under the strong assumption that the diffusion MRI peaks reflect true axons, our results would therefore indicate that pathways do not form sheet structures at every crossing fiber region but instead at well-defined locations in the brain. With this framework, sheet structure location, extent, and orientation could potentially serve as new structural features of brain tissue. The proposed method can be extended to quantify sheet structure in directional data obtained with techniques other than diffusion MRI, which is essential for further validation.

7.1 INTRODUCTION

A three-dimensional Manhattan street grid or the intricate streets of Victorian London, which configuration reflects our brain's organization best? This debate added three Science publications to the list of diffusion MRI (dMRI) literature (Catani et al., 2012; Wedeen et al., 2012a; Wedeen et al., 2012b). Wedeen et al. (2012b) analyzed adjacency and crossings between cerebral fiber pathways of the brain using dMRI and found that the pathways form a 3D grid structure. More specifically, the authors used diffusion spectrum imaging (DSI (Callaghan et al., 1990; Wedeen et al., 2005)), which infers information on the diffusion probability density function (PDF) by extensively sampling q -space in a Cartesian fashion, to reconstruct a so-called path neighborhood with tractography. This path neighborhood can be computed by tracking pathways from a small seed region, and subsequently computing the paths incident on these paths. It was found that the pathways in such a neighborhood cross nearly orthogonally in 2D grid- or sheet-like structures (similar to the "warp and weft of a fabric") that are layered in 3D space "like pages of a book" (Wedeen et al., 2012a; Wedeen et al., 2012b). This sheet structure was consistently recognized across species and scales, and throughout the white matter.

Catani et al. (2012) suggested that the observed grid pattern is most likely an artifact, attributed to the limitations of DSI used in Wedeen et al. (2012b). The authors showed that diffusion orientation distribution functions (dODFs) as derived from the DSI-PDFs have inherently low angular resolution, and therefore have a limited ability to resolve crossing fibers with small angles. They concluded that this bias towards orthogonal angles negatively impacts the tractography results in Wedeen et al. (2012b) and inadvertently makes "the grid structure of interwoven sheets a very likely configuration" (Catani et al., 2012). By using another dMRI technique called spherical deconvolution (SD) (Dell'Acqua et al., 2007; Dell'Acqua et al., 2010; Tournier et al., 2007), which has a higher angular resolution through the direct reconstruction of the fiber ODF (fODF), they were able to show that non-orthogonal crossings represent a large percentage of the total crossings in white matter ($> 88\%$ in a group study of 10 subjects). In addition, Catani et al. noted that the results in Wedeen et al. (2012b) are mainly *qualitative* and that dMRI-based pathways cannot be equated to true axons.

Wedeen et al. (2012a) rebutted the technical concerns regarding DSI and claimed to find further support for the sheet-structure theory in classic degeneration studies. In addition, they agreed that non-orthogonal angle crossings do exist, and stated that Catani et al. did not address the main finding of their study: the existence of sheet structure. This structure "does not depend on fiber orthogonality or the absence thereof" and the authors stated that "there are no mechanisms known whereby technical limitations will create it as an artifact" (Wedeen et al., 2012a).

This debate has gained a lot of attention from the diffusion and neuroscience communities. While still considered controversial by many, the existence of sheet structure could have significant impact on models of structural and functional brain connectivity, embryogenesis, and development. It could for example play an important role in axonal path-finding during embryogenesis by guiding growing fibers, thereby greatly reducing the complexity of such processes (Wedeen et al., 2012b). The prevalence and geometry of sheet structures in the brain can potentially also be a novel feature to characterize brain structure, complementing the wide range of existing microstructural and geometrical measures (e.g. Assaf et al., 2008; Astola et

al., 2011; Dell’Acqua et al., 2013; Fieremans et al., 2011; Leemans et al., 2006; Raffelt et al., 2012; Savadjiev et al., 2012; Tax et al., 2012; Zhang et al., 2012).

To date, however, there is no general consensus on the degree to which sheet structure is present in the brain. The lack of a clear exposition of the relevant mathematical concepts may have contributed to this ongoing debate. In particular, there still exist ambiguities regarding the exact definition of a sheet structure, the conditions for it to exist, and its relation to orthogonal angles. In addition, the evidence for the existence of sheet structure was mainly *qualitative*, and no extensive *quantitative* analysis was performed so far. In this work, we first formalize the terminology to clarify the definition of sheet structure as proposed by Wedeen et al. (2012b) and the condition for its existence; this is done in Section 7.2. Subsequently, we propose a robust method to compute a sheet probability index (SPI), which indicates to what extent the data supports a sheet structure, in Section 7.3.1 and 7.3.2. An intuitive way to visualize the SPI throughout the brain is described in Section 7.3.3. We evaluate this method with simulations and real dMRI data sets as described in Sections 7.3.4 and 7.4. Finally, we use the proposed method to investigate and discuss some of the claims made in Wedeen et al. (2012ab) and Catani et al. (2012) in Sections 7.4 and 7.5. Note that even though we use dMRI data here to investigate the existence of sheet structure, our approach can be extended to compute the SPI in other types of directional data such as polarized-light imaging data (Aixer et al., 2001). Preliminary results of this work have been presented at the ISMRM (Tax et al., 2014a; Tax et al., 2016a) and the BASP workshop (Tax et al., 2015a).

7.2 THEORY

In this section, we present the theoretical background that is required to understand some of the key mathematical concepts in relation to the investigation of the brain’s sheet structure. In Section 7.2.1, we present a definition of sheet structure in terms of integral curves of vector fields, and in Section 7.2.2, we discuss the relevant measure used to assess the presence of this sheet structure, i.e. the Lie bracket. In Section 7.2.3, we explain the Frobenius theorem, which formalizes the necessary and sufficient condition for the sheet structure to exist. Note that we adopt a rather informal mathematical language and omit technical definitions and proofs to make the main concepts accessible to a broader audience. We refer to Appendix A7 for a list of symbols and Appendix B7 for a more formal summary.

7.2.1 Definition of sheet structure

Several definitions and interpretations of ‘sheets’ in the brain exist in the literature (Kindlmann et al., 2007; Schultz et al., 2010; Vilanova et al., 2004; Yushkevich et al., 2008; Zhang et al., 2003). Some major white matter tracts resemble a thin sheet-like structure by itself: well-known examples are the corona radiata and the corpus callosum. Such *single-fiber-direction* sheets have typically been represented by a sheet-like skeleton (Smith et al., 2006) or a surface (Yushkevich et al., 2008). In contrast, *crossing* or *intertwining-fiber sheets* as proposed by Wedeen et al. (2012b) (hereafter shortened to sheets) are not segregated structures, but are composed of two tracts that cross each other on the same surface in certain regions along their trajectories. As a result, tubular-shaped tracts can in theory still form sheets at locations where

they intersect with other structures. Here, we focus on the latter definition of sheet structure and on how such sheets can be detected in directional data where two structures cross.

Consider a set of vectors at each position of the brain $M \subset \mathbb{R}^3$, which can for example be obtained from the dODF or fODF using any dMRI reconstruction technique. With streamline tractography one integrates a smooth three dimensional *vector field* V defined on a subset $N_V \subset M$, generating streamlines (or integral curves) $\Phi^V(s, p)$ ¹ (with s denoting arc length and $p = \Phi^V(0, p)$ the initial spatial position), such that for all $p \in N_V$ the following integrability holds (Fig. 7.1a):

$$\frac{d}{ds} \Phi^V(s, p) \Big|_{s=0} = V_p. \quad [7.1]$$

Here, $V_p \in \mathbb{R}^3$ denotes the vector at position p . This idea of integrability can potentially be extended to two vector fields V and W : an integral surface $S \subset N_V \cap N_W$ called the *sheet structure* is defined as the surface S whose tangent plane at p is parallel to the plane spanned by V_p and W_p for all $p \in S$, see Fig. 7.1b. The requirements for such a surface to exist are formalized in the Frobenius theorem, which states that the vector fields V and W should interact in a ‘nice’ way according to their Lie bracket.

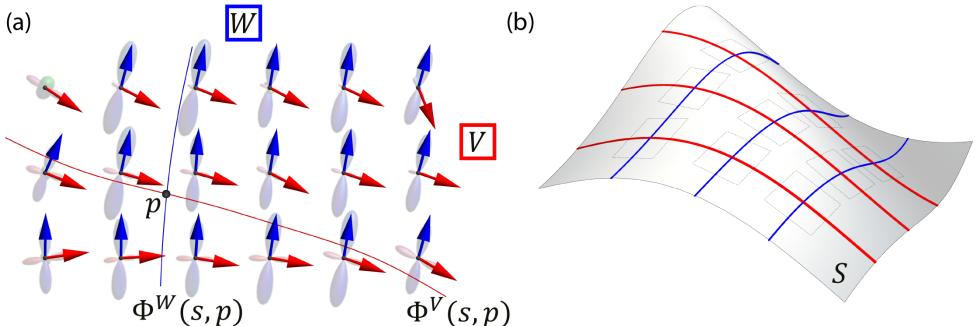


Fig. 7.1: (a) An example set of fiber ODFs (semi-transparent) along with their peak directions (arrows), which form the vector fields V (red) and W (blue). By integrating these vector fields, one can reconstruct at each initial position p the integral curves $\Phi^V(s, p)$ (red curve) and $\Phi^W(s, p)$ (blue curve) where s denotes arc length parameterization. (b) The tangent plane of an integral surface S at any point $p \in S$ is parallel to the plane spanned by V_p and W_p (indicated by the dashed squares annotated on S).

7.2.2 The Lie bracket

The *Lie bracket*, a mathematical concept from differential topology, is a bilinear operator on two vector fields V and W and defines a third vector field denoted by $[V, W]$. Intuitively, the Lie bracket at a given location p can be understood as the deviation from p when trying to

¹ Here we adopt the flow operator notation Φ^V to represent the streamlines, which has advantages over traditional notation in the remaining Theory sections.

move around in a small loop along the integral curves of V and W with p as the starting position (Fig. 7.2). Such a loop consists of four ‘legs’ (i.e., a quadrilateral) and starts at point p by following the integral curve $\Phi_s^V(p)$ ¹ along V for some distance s . After arriving at point $\Phi_s^V(p)$ (the end point of the first leg of the loop), one continues along the vector field W for the same distance s . From the resulting end point of this second leg $\Phi_s^W \circ \Phi_s^V(p)$, where \circ is the symbol used to indicate the concatenation operation, one moves ‘backwards’, along the integral curves of $-V$ and $-W$ in order, again along each leg with the same distance s . The end point of this loop is then written as $(\Phi_s^{-W} \circ \Phi_s^{-V} \circ \Phi_s^W \circ \Phi_s^V)(p)$ or with the following shorthand notation

$$\alpha_p(s) := (\Phi_s^{-W} \circ \Phi_s^{-V} \circ \Phi_s^W \circ \Phi_s^V)(p). \quad [7.2]$$

In other words, α_p is the curve that comprises the end points of all loops starting in the point p (see Fig. 7.2a for a schematic overview). For a specific value of s , the difference vector R_p (also called the *closure*) is then defined as

$$R_p(s) := \alpha_p(s) - p. \quad [7.3]$$

The formal definition of the Lie bracket (Misner et al., 1973; Spivak, 1979) in p , $[V, W]_p$, follows from the Taylor expansion $\alpha_p(s) = \alpha_p(0) + \alpha'_p(0)s + 1/2 \alpha''_p(0)s^2 + O(s^3)$ (where $\alpha_p(0) = p$ and $\alpha'_p(0) = 0$, and where prime denotes differentiation with respect to s) and taking the limit

$$\lim_{s \rightarrow 0} \frac{R_p(s)}{s^2} = \frac{1}{2} \alpha''_p(0) =: [V, W]_p. \quad [7.4]$$

See Appendix B7 for details. At this point we remark that by choosing s sufficiently small, there is a linear relationship between the Lie bracket and the closure (we will use this relation in Section 7.3.1.3).

7.2.3 The Frobenius theorem

The requirement for a sheet structure to exist can be understood intuitively by following the line of thinking in the previous section, and is depicted in Fig. 7.2c and d. The Lie bracket was introduced as the local deviation from p after a small flow over V and W , and if these vector fields are tangent to a sheet in a neighborhood of p then the end point $\alpha_p(s)$ of this flow must lie on the sheet structure as well (Fig. 7.2c). In the limit of smaller and smaller loop sizes this leads to the Frobenius theorem (Lang, 1995; Spivak, 1979), which tells us that two vector fields

¹ For notational convenience, we use the flow operator $\Phi_s^V(p) \equiv \Phi^V(s, p)$ if we are considering streamlines of a fixed arc length s .

V and W form a sheet structure in a neighborhood of point p if and only if the Lie bracket $[V, W]_p$ lies in the plane spanned by V_p and W_p , cf. Fig. 7.2d. As a comparison, Fig. 7.2b shows the case (corresponding to Fig. 7.2a) in which the vector fields V and W do not form a sheet structure.

In practice we can check this condition by examining the component of $[V, W]_p$ normal to V_p and W_p , the *normal component of the Lie bracket* (Wedeen et al., 2012b). We do this by computing the orthogonal projection of the vector $[V, W]_p$ onto the outer product $V_p \times W_p$ (normalized):

$$[V, W]_p^\perp := [V, W]_p \cdot \frac{V_p \times W_p}{\|V_p \times W_p\|}. \quad [7.5]$$

The sheet structure S is then given by the set of all points $p \in N_V \cap N_W$ where $[V, W]_p^\perp = 0$. Note that this condition does not involve the presence of orthogonal angles.

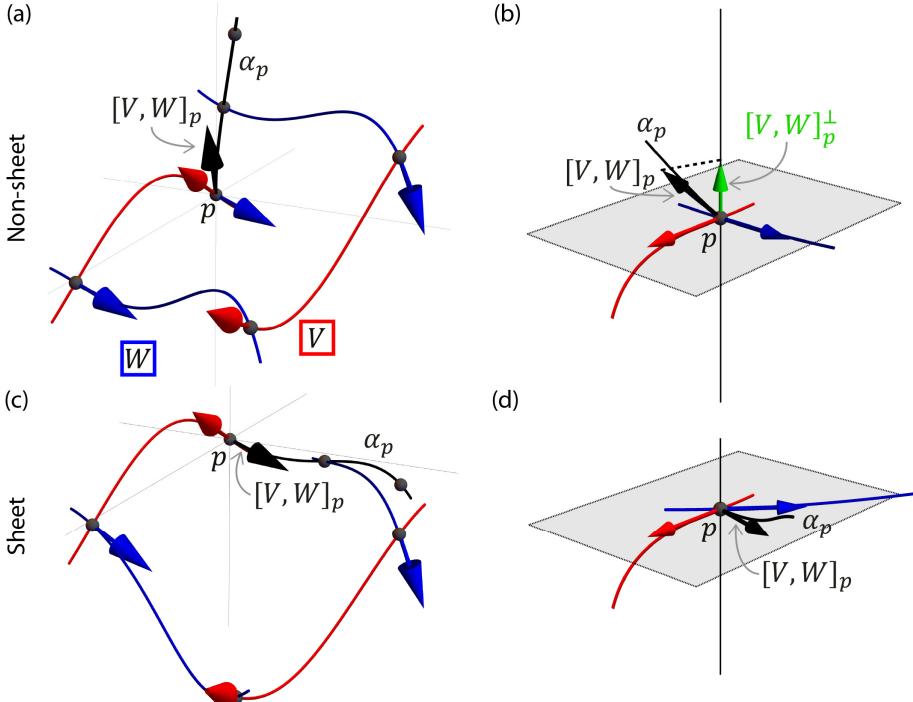


Fig. 7.2: (a) A loop composed of integral curves of V (red arrows) and W (blue arrows) with p as the starting position. α_p is the curve formed by the end points of all loops (by varying the distance s) starting in the point p . $[V, W]_p$ is the Lie bracket at point p , and has a relation with the difference vector $R_p(s) := \alpha_p(s) - p$ according to Eq. [7.4]. In this scenario, the vector fields V and W cannot be integrated to form a sheet structure. (b) The Lie bracket from (a) does not lie in the plane spanned by V_p and W_p (grey). Hence the normal component of the Lie bracket $[V, W]_p^\perp$ defined in Eq. [7.5] is non-zero (green arrow), and the vector fields cannot be integrated to form a sheet structure. The vertical black line is perpendicular to the plane spanned by V_p and W_p . (c) A loop in a scenario where the vector fields V (red arrows) and W (blue arrows) do form a sheet structure. In this case α_p is (locally) a curve on the sheet structure. (d) The Lie bracket $[V, W]_p$ from (c) lies in the plane spanned by the vectors at p , so that the normal component is zero and the vector fields V and W can be integrated to form a two-dimensional sheet.

7.3 MATERIALS AND METHODS

We can examine the existence of sheet structure by 1) estimating the integral curves and loops in Eq. [7.2], 2) estimating the Lie bracket based on Eq. [7.4], and 3) extracting the normal component of the estimated Lie bracket (Eq. [7.5]) as an indicator of sheet structure. Note that we can perform these computations for every pair of vector fields in a neighborhood. However, the definitions in Section 7.2 assume separate, continuous, and smooth unit vector fields, while our input data is a (possibly incomplete) set of unsorted vectors per discrete position, perturbed by noise and generally without consistent sign attributes (e.g. when derived from the dODF or fODF). These issues greatly challenge the actual computation of a Lie bracket.

In Section 7.3.1 we describe the approach to compute the *discrete* Lie bracket, inspired by the qualitative reconstruction of path neighborhoods in Wedeen et al. (2012b). Our method can deal with noisy vector fields and addresses the problem of sorting vectors in a neighborhood of a point p . In Section 7.3.2 we further address the issue of noise by deriving a *sheet probability index* (SPI) from multiple computations of the discrete Lie bracket of a pair of vector fields per point. In Section 7.3.3 we define the *sheet tensor*, which allows us to visualize the Lie bracket for every pair of vector fields. Finally, in Section 7.3.4 we describe the simulated and acquired MRI data used for the experiments.

7.3.1 The discrete Lie bracket

We propose here to calculate a discrete Lie bracket by approximating the integral curves in Eq. [7.2] with tractography (Tax et al., 2015a; Tax et al., 2014a; Wedeen et al., 2014), and by computing a large number of loops with configurations as in Fig. 7.2. The tractography process and averaging of multiple loop configurations implicitly and partially deals with noise in the Lie bracket computation.

Similar to conventional tractography, we have to make the assertion that each vector is an element of a smooth vector field. Whereas conventional tractography looks for the vector that aligns most with the incoming direction to propagate a tract, here we have to keep track of the whole *frame* of vectors (defined as all vectors at a point) during tracking to be able to switch between different vector fields in a loop. In Section 7.3.1.1 we describe the clustering of frames (the process of assigning vectors to specific vector fields), which takes care of both the sorting and of possible sign inconsistencies in the vector data. In Section 7.3.1.2 we outline the *frame tractography* that performs clustering during tract propagation. Finally, the estimation of the Lie bracket is discussed in Section 7.3.1.3.

7.3.1.1 Clustering of frames

The purpose of frame clustering is to assign each vector of a frame at a given position to a vector field (Fig. 7.3). To this end, the frame is compared to an already clustered frame (an ordered set of vectors) at a nearby position. The vectors of the non-clustered frame are matched (permuted) such that the total angle between pairs of vectors is minimized. The frame clustering algorithm can be found in Appendix C7.1.

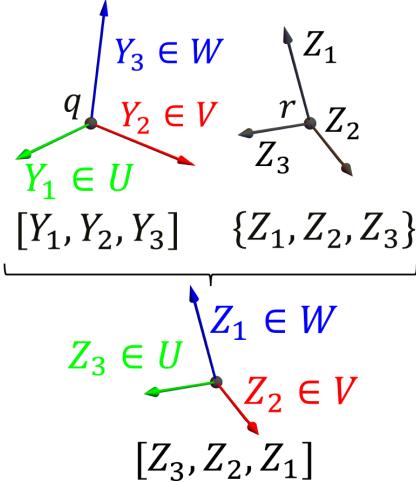


Fig. 7.3: Example of the clustering of frames. We have an ordered set of vectors $[Y_1, Y_2, Y_3]$ at a position q , and 3 vectors $\{Z_1, Z_2, Z_3\}$ in some point r near q . We assume that they are assigned to the vector fields U , V , and W as follows: $U_q = Y_1, V_q = Y_2$, and $W_q = Y_3$. Frame clustering yields the ordered set $[Z_3, Z_2, Z_1]$.

7.3.1.2 Frame tractography

Given a step size Δh and a discrete distance h (assuming arc length parameterization), we can approximate the flow along a vector field $X \in \{U, V, \dots\}$ with streamline tractography (Mori and van Zijl, 2002). The algorithm explained in Appendix C7.2 is similar to other deterministic tractography algorithms, but keeps track of the vector fields defined in the neighborhood. Note that we use nearest neighbor interpolation unless stated otherwise. From here on, approximations are marked by a circumflex, i.e., $\widehat{\Phi}_h^X$ denotes the approximate flow along X for a distance h , corresponding to the true flow Φ_h^X .

7.3.1.3 The closure and the Lie bracket

To calculate the discrete Lie bracket we reconstruct approximate flow loops, which can be used to obtain estimates \widehat{R}_p of the difference vectors R_p . In practice, we will compute difference vectors for a large number of loops with several configurations (i.e., variations on Eq. [7.2], see Appendix B7). More concretely, given a point p we will consider the set of difference vector estimates $\{\widehat{R}_1, \widehat{R}_2, \widehat{R}_3\}$ resulting from the following loop configurations (Fig. 7.4):

$$\begin{aligned}\widehat{R}_1 &:= (\widehat{\Phi}_{h_2}^{-W} \circ \widehat{\Phi}_{h_1}^{-V} \circ \widehat{\Phi}_{h_2}^W \circ \widehat{\Phi}_{h_1}^V)(p) - p, \\ \widehat{R}_2 &:= p - (\widehat{\Phi}_{h_1}^{-V} \circ \widehat{\Phi}_{h_2}^{-W} \circ \widehat{\Phi}_{h_1}^V \circ \widehat{\Phi}_{h_2}^W)(p), \\ \widehat{R}_3 &:= (\widehat{\Phi}_{h_2}^W \circ \widehat{\Phi}_{h_1}^V)(p) - (\widehat{\Phi}_{h_1}^V \circ \widehat{\Phi}_{h_2}^W)(p).\end{aligned}\quad [7.6]$$

Here, h_1 and h_2 are the flow distances along the integral curves of V and W , respectively. We choose $h_1, h_2 \in \{-h_{max}, -h_{max} + \Delta h, -h_{max} + 2\Delta h, \dots, -\Delta h, \Delta h, \dots, h_{max}\}$, where h_{max} is the maximum distance (a parameter in our algorithm). Note that we thus sample all four ‘quadrants’ surrounding point p . The reconstruction of a loop for computation of the difference vector estimate \widehat{R}_1 in Eq. [7.2] (the other difference vectors are computed in a similar manner) is described by the algorithm outlined in Appendix C7.3.

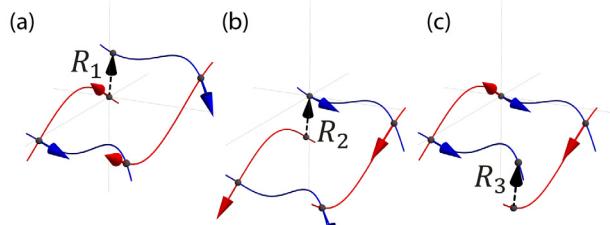


Fig. 7.4: Loops that lead to $R_1 := (\hat{\Phi}_{h_2}^{-W} \circ \hat{\Phi}_{h_1}^{-V} \circ \hat{\Phi}_{h_2}^W \circ \hat{\Phi}_{h_1}^V)(p) - p$, $R_2 := p - (\hat{\Phi}_{h_1}^{-V} \circ \hat{\Phi}_{h_2}^{-W} \circ \hat{\Phi}_{h_1}^V \circ \hat{\Phi}_{h_2}^W)(p)$ and $R_3 := (\hat{\Phi}_{h_2}^W \circ \hat{\Phi}_{h_1}^V)(p) - (\hat{\Phi}_{h_1}^V \circ \hat{\Phi}_{h_2}^W)(p)$.

When all difference vectors are estimated, we can compute an estimate of the Lie bracket $[\widehat{V}, \widehat{W}]_p$ with a simple linear least squares fit corresponding to Eq. [7.4]

$$\mathbf{R} = \mathbf{H}\boldsymbol{\beta} + \boldsymbol{\epsilon}. \quad [7.7]$$

Here, \mathbf{R} is a $K \times 3$ matrix with the difference vectors $\widehat{\mathbf{R}}$, \mathbf{H} is a $K \times 1$ matrix with the products of the used h_1 and h_2 for the corresponding difference vector, $\boldsymbol{\beta}$ is a 1×3 matrix with the Lie bracket components $[\widehat{V}, \widehat{W}]_p^i$, and $\boldsymbol{\epsilon}$ is a $K \times 3$ matrix with errors, where K is the total number of difference vectors. The errors $\boldsymbol{\epsilon}$ are assumed to be normally distributed. The normal component of the estimated Lie bracket follows from Eq. [7.5]. In practice, we only compute $[\widehat{V}, \widehat{W}]_p$ when the number of successfully estimated loops and corresponding difference vectors exceeds a minimum threshold.

7.3.2 Sheet probability index

The algorithm described in Section 7.3.1 allows us to compute estimates of the Lie bracket normal component $[\cdot, \cdot]^\perp$ for every combination of vectors at every position in the brain, and according to the Frobenius theorem (Section 7.2.3) a combination of vector fields supports the sheet conjecture if $[\cdot, \cdot]^\perp = 0$. Due to the occurrence of noise, however, the sheet-constraint is rarely exactly fulfilled, and a single estimate does not provide information on its variability. This makes it difficult to quantify to what degree the local structure effectively resembles a sheet.

Ideally, repeated MRI acquisitions could be used to approximate the variance of the estimated Lie bracket normal component. By assuming a normal distribution (in practice verified using a Shapiro-Wilks test) with data-derived mean μ and standard deviation σ , we can calculate the integral probability P_λ inside the region $[-\lambda, \lambda]$ (where we can tune the parameter λ) for the estimated distribution $N(\mu, \sigma^2)$. P_λ produces a value that lies between 0 and 1 which we coin the *sheet probability index* (SPI) of the local sheet structure. Choosing a higher value for λ means the SPI is less sensitive to small deviations from zero. In practice it is often difficult to acquire a large number of repeated dMRI sets, so we consider residual bootstraps as an alternative (see Section 7.3.4.2.3).

The introduction of the SPI does not only address the issue of noise, but it also makes the interpretation of the Lie bracket normal component much more intuitive. A high value for the SPI corresponds to a high likelihood of sheet structure ($[\cdot, \cdot]_p^\perp$ is likely close to 0), while a low value indicates that there are significant deviations from sheet structure ($[\cdot, \cdot]_p^\perp$ likely differs significantly from 0).

7.3.3 Sheet tensors

Investigation of consistent sheet structures in a spatial neighborhood asks for an appropriate way to visualize the SPI throughout the brain. The SPI can be computed for every pair of vector fields, i.e., n vector fields generate $\binom{n}{2}$ SPI's. In this work we propose to visualize the local sheet structure throughout the brain by means of a *sheet tensor*. Given a pair of vector fields V and W , the sheet tensor at location p is defined as

$$T_p = \frac{P_\lambda}{\beta_1} (V_p \otimes V_p + W_p \otimes W_p) = \frac{P_\lambda}{\beta_1} (V_p V_p^T + W_p W_p^T). \quad [7.8]$$

Here β_1 denotes the largest eigenvalue of the tensor $(V_p V_p^T + W_p W_p^T)$ and \otimes denotes tensor product. The sheet tensor can then be represented by an ellipsoid whose third eigenvector is normal to the span of V and W , and which defines the color of the ellipsoid in the well-known RGB scheme (normal in left-right direction gives a red tensor, normal in inferior-superior direction gives a blue tensor, and normal in anterior-posterior direction gives a green tensor) (Zhang et al., 2006b). Furthermore, the size of the ellipsoid is determined by the SPI, where a larger SPI gives larger ellipsoids, and the division by β_1 fixes the largest semi-axis of the ellipsoid for a given SPI. The shape represents the angle between V and W . Fig. 7.5 shows sheet tensors for different angles and different SPI. The sheet tensor allows us to visualize the SPI for every pair of vector fields at a given location, and can thus also reveal *crossing sheets*.

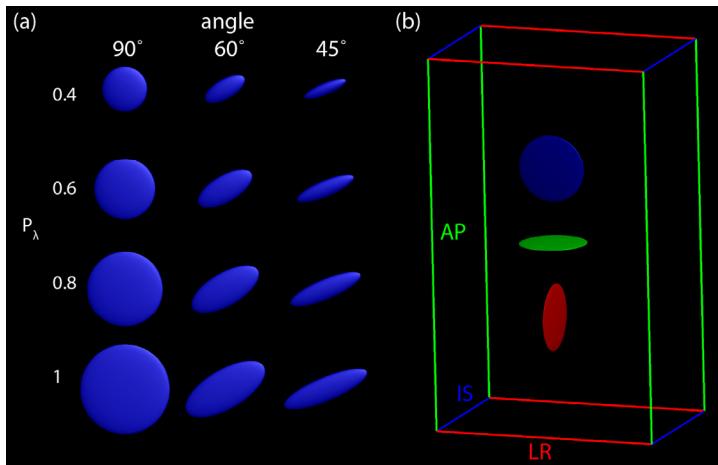


Fig. 7.5: (a) Sheet tensors with a normal direction pointing towards the reader, for different angles between U_p and V_p and different sheet probabilities. Here, V is always oriented in left-right direction. (b) Sheet tensors with different orientations are colored according to their third eigenvector. AP is the anterior-posterior direction, IS is inferior-superior, and LR is left-right.

7.3.4 Data

We will evaluate our framework with different types of data: analytical vector fields, dMRI simulations, and real dMRI data. These test data sets increase in degree of complexity, allowing us to investigate different aspects of the implemented methodology.

7.3.4.1 Analytical vector field simulations

We define three vector fields that are tangent to a sphere with radius ρ (U and V are tangent to the upper hemisphere, W is tangent to the lower hemisphere, see Fig. 7.6a):

$$\begin{aligned} U &= (-\sin \phi_1, \cos \phi_1 \cos \theta_2, \cos \phi_1 \sin \theta_2), \\ V &= (\cos \phi_2 \cos \theta_1, -\sin \phi_2, \cos \phi_2 \sin \theta_1), \\ W &= (\cos \phi_2 \cos \theta_2, -\sin \phi_2, -\cos \phi_2 \sin \theta_2). \end{aligned} \quad [7.9]$$

Here $\theta_i = \tan^{-1} x^i / \sqrt{\rho^2 - (x^1)^2 - (x^2)^2}$ and $\phi_i = \cos^{-1} (x^i) / \rho$ (with $i = 1, 2$), and $x = (x^1, x^2, x^3)$ denotes Cartesian coordinates. The integral curves of these vector fields have constant curvature $\kappa = 1/\rho$.

Vector fields U and V form a sheet, so that $[U, V]_p^\perp = 0 \forall p \in \{x \in \mathbb{R}^3 | (x^1)^2 + (x^2)^2 < \rho^2, x^3 = z\}$. U and W generally do not form a sheet, and the normal component of the Lie bracket $[U, W]_p$ at $p \in \{x \in \mathbb{R}^3 | (x^1)^2 + (x^2)^2 < \rho^2, x^1 \neq 0, x^2 \neq 0\}$ is given by

$$[U, W]_p^\perp = \frac{6x^1x^2(-\rho^2+(x^1)^2+(x^2)^2)}{\sqrt{(\rho^2-(x^1)^2)(\rho^2-(x^2)^2)(\rho^6-8\rho^2(x^1)^2(x^2)^2+4(x^1)^2(x^2)^2((x^1)^2+(x^2)^2))}}. \quad [7.10]$$

A plot of $[U, W]_p^\perp$ as a function of x^1 and ρ is shown in Fig. 7.6b, where we take $x^2 = -x^1$ so that $[U, W]_p^\perp$ is generally larger than zero. By evaluating Lie bracket estimates along these lines $x^2 = -x^1$ for fixed curvature κ , we can evaluate the performance of the algorithm as a function of the magnitude of the Lie bracket normal component. By varying ρ we can similarly evaluate our Lie bracket estimates as a function of the curvature of the integral curves. Note that these combinations of vector fields generally cross in non-orthogonal angles.

The vector fields are discretized by sampling them on a Cartesian grid with period δ (corresponding to the voxel size). We add noise to the discrete vector fields (N_n noise iterations) by drawing random samples of a Watson distribution (Chen et al., 2015b; Mardia and Jupp, 2009) with probability density function

$$f_W(\pm \tilde{V}_q; V|_q, k) = M\left(\frac{1}{2}, \frac{3}{2}, k\right)^{-1} e^{k(\tilde{V}_q \cdot V_q)^2}. \quad [7.11]$$

Here, $M(1/2, 3/2, \cdot)^{-1}$ is the Kummer function (Mardia and Jupp, 2009), \tilde{V}_q denotes the perturbed vector at location q , and $k > 0$ is a concentration parameter (here referred to as the ‘SNR level’, higher k results in lower perturbation).

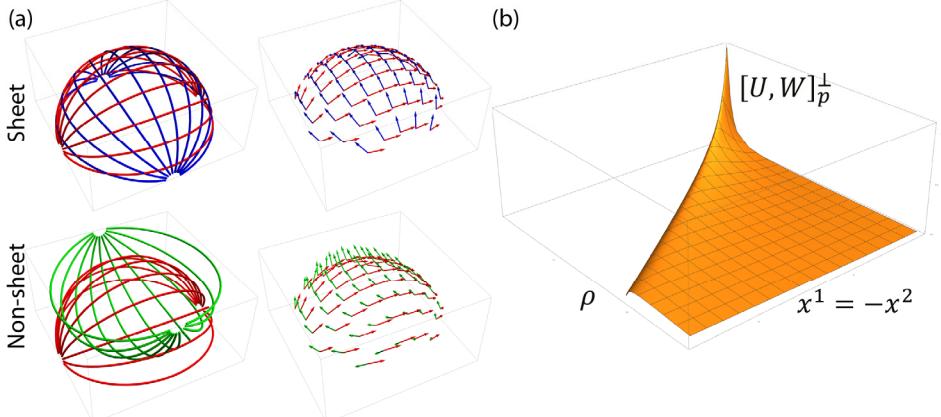


Fig. 7.6: (a) Vector fields U (red), V (blue), and W (green), where U and V form a sheet and U and W do not. The left column shows a subset of integral curves, and vectors sampled on the upper hemisphere are shown on the right. This pattern of vector fields is repeated in the vertical direction. (b) Plot of $[U, W]_p^{\perp}$ as a function of ρ and x^1 , with $x^2 = -x^1$.

7.3.4.2 Diffusion MRI data

7.3.4.2.1 Simulations

dMRI signals were simulated using a ZeppelinStickDot model (Ferizi et al., 2014) with the fiber direction defined by the noise free vector fields described in the previous section. N_n noise iterations were generated using the Rician distribution. We simulate two types of datasets: single shell datasets with 90 directions and $b = 3000 \text{ s/mm}^2$ suitable for SD, and Cartesian sampled datasets with 514 directions (maximum b -value of 10000 s/mm^2) and one $b = 0 \text{ s/mm}^2$ point suitable for DSI (protocol corresponds to the MGH HCP DSI data, see section 7.3.4.2.2).

7.3.4.2.2 Real data

We use different dMRI data sets with varying spatial and angular resolutions, diffusion weightings, and sampling schemes to investigate our framework: 1) the $b = 3000 \text{ s/mm}^2$ shell with 90 diffusion directions of three subjects of the WU-Minn Human Connectome Project (HCP) with an isotropic voxel size of 1.25 mm (Glasser et al., 2013; Sotiroopoulos et al., 2013; Van Essen et al., 2013); 2) the $b = 3000 \text{ s/mm}^2$ shell with 500 diffusion directions of the MASSIVE database with an isotropic voxel size of 2.5 mm (Froeling et al., 2016); 3) the separate shells ($b = \{1000, 3000, 5000, 10000\} \text{ s/mm}^2$ with $\{64, 64, 128, 256\}$ directions) of one subject of the MGH HCP with an isotropic voxel size of 1.5 mm (Setsompop et al., 2013); and 4) a Cartesian sampled data set (514 directions) with b -values up to 10000 s/mm^2 of the MGH-USC HCP with an isotropic voxel size of 2 mm (<http://www.humanconnectomeproject.org/data/inventory/>).

7.3.4.2.3 Processing

Data sampled on a single shell was processed using constrained spherical deconvolution (CSD, $l_{max} = 8$) (Tournier et al., 2007) in *ExploreDTI* (Leemans et al., 2009). The response function for the simulated data was generated from the ZeppelinStickDot model, and the response

function for real data was computed using recursive calibration (Tax et al., 2014c). Peaks were extracted using a Newton optimization algorithm (Jeurissen et al., 2011) with an FODF peak threshold of 0.1, and a maximum number of 3 peaks. To compute the SPI, we used the N_n noise iterations for simulated data, and generated N_b residual bootstrap realizations for simulated and real dMRI data from a single set of noisy measurements (Jeurissen et al., 2011). The peaks extracted from the different bootstrap realizations were clustered using the method described in Section 7.3.1.1, taking the peaks extracted from the original data as reference frames.

Cartesian sampled data was analyzed using the DSI model (Wedgeon et al., 2005), which was reconstructed with Diffusion ToolKit (DTK) using default settings (Wang et al., 2007). The algorithm readily provides a set of peaks at each position, obtained from the local maxima of a roughly uniform sampling (181 points) of a hemisphere, from which we take at most 3 vectors per position based on the dODF magnitude. No bootstrapping could be performed, so in this case only one Lie bracket was computed for every pair of vector fields.

7.4 RESULTS

The results for analytical vector field simulations are presented in Section 7.4.1, for dMRI simulations in Section 7.4.2, and for real dMRI data in Section 7.4.3.

7.4.1 Analytical vector field simulations

With the analytical vector fields we will systematically investigate different aspects of the Lie bracket implementation: the influence of discretization (the finite voxel size δ), the noise (different settings of the concentration parameter k), the curvature κ (by varying ρ), and the Lie bracket magnitude. We use nearest-neighbor interpolation of the vector fields and $N_n = 50$ noise iterations here.

7.4.1.1 The influence of spatial resolution and noise

Fig. 7.7 shows results for different voxel sizes $\delta = \{0.5, 1, 2\} \text{ mm}$ (a-c), different settings for h_{max} (rows) and different SNR levels k (the concentration parameter in Eq. [7.11], higher k indicates lower perturbation). Here we consider a relatively simple case: since we know which vector belongs to which vector field, we skip the clustering step and show results that are not affected by clustering errors. We set the curvature $\kappa = 1/\rho = 1/26 \text{ mm}^{-1}$ and estimate the normal component of the Lie bracket at $p = (10, -10, 0)$ to have a Lie bracket magnitude significantly deviating from zero for the given radius ρ .

Each plot shows the mean and range of the estimates $\widehat{[\cdot, \cdot]}_p^\perp$ in the case of sheet (grey, $[U, V]_p^\perp = 0 \text{ mm}^{-1}$) and non-sheet (black, $[U, W]_p^\perp = 0.031 \text{ mm}^{-1}$). The range becomes smaller with higher k (the noiseless case $k = \infty$ is also plotted) in all cases. The precision of the estimates increases with increasing h_{max} (smaller error bars), and the accuracy increases for $h_{max} = 3 \text{ voxels}$ compared to $h_{max} = 1 \text{ voxel}$, but remains similar when further increasing to $h_{max} = 5 \text{ voxels}$. We can see that $h_{max} = 1 \text{ voxel}$ is generally too low to obtain a reasonable accuracy and precision, and to distinguish sheet from non-sheet. The precision is similar for

approximately the same h_{max} in mm (see for example the approximately equal error bars in the cases $h_{max} = 3$ voxels, $\delta = 1$ mm and $h_{max} = 5$ voxels, $\delta = 0.5$ mm).

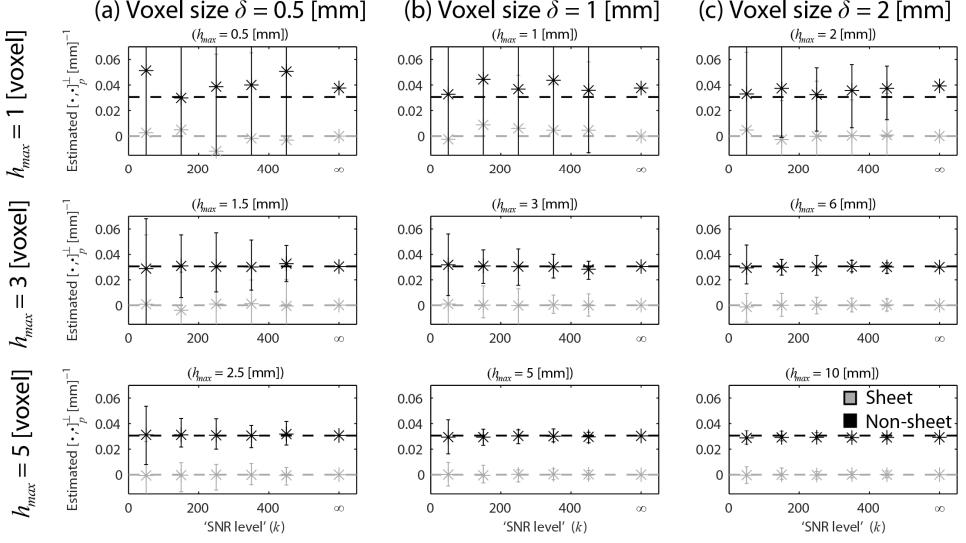


Fig. 7.7: $\widehat{[\cdot, \cdot]}_p^\perp$ for different voxel sizes $\delta = \{0.5, 1, 2\}$ mm (a-c, columns) and different settings for $h_{max} = \{1, 3, 5\}$ voxels (the corresponding h_{max} in mm is noted above each plot). Each plot shows the mean and range of the estimates in the case of sheet (grey, $[U, V]_p^\perp = 0$ indicated by the dashed line) and non-sheet (black, $[U, W]_p^\perp = 0.031$) for different SNR levels (i.e. the concentration parameter k , higher k means a lower perturbation of the vectors). We used 50 noise iterations, $\kappa = 1/p = 1/26 \text{ mm}^{-1}$, and $p = (10, -10, 0)$.

7.4.1.2 The influence of the Lie bracket normal component magnitude

Fig. 7.8a shows the mean and range of the estimates $\widehat{[\cdot, \cdot]}_p^\perp$ for different points $p = (x^1, -x^1, 0)$, where the Lie bracket normal component magnitude $[U, W]_p^\perp$ (non-sheet) varies while the curvature remains constant at $\kappa = 1/26 \text{ mm}^{-1}$. $[U, V]_p^\perp$ (sheet) is evaluated at the same points for reference. We set $k = 350$, $\delta = 1$ mm, and $h_{max} = 5$ voxels. Here and in further analyses, we apply clustering of the vector fields as described in Section 7.3.1.1 (using the known vector fields as prior information gave similar results, not shown here).

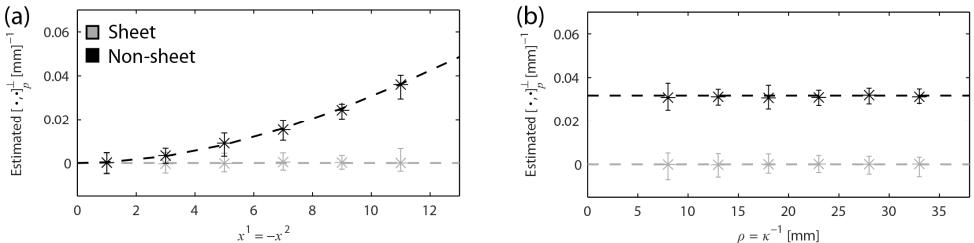


Fig. 7.8: (a) Mean and range of $\widehat{[\cdot, \cdot]}_p^\perp$ for different points $p = (x^1, -x^1, 0)$ to vary the Lie bracket normal component magnitude $[U, W]_p^\perp$ ($\kappa = 1/26 \text{ mm}^{-1}$). (b) Mean and range of $\widehat{[\cdot, \cdot]}_p^\perp$ is shown for different curvatures $\kappa = 1/\{8, 13, 18, 23, 28, 33\}$. Different curvatures were achieved by changing ρ in Eq. [7.9], and $[U, W]_p^\perp$ was kept constant using Eq. [7.10] by adapting the point of evaluation $x^2 = -x^1$. In both experiments, $[U, V]_p^\perp$ (dashed lines) is evaluated at the same points for reference and $k = 350$, $\delta = 1$ mm, $h_{max} = 5$ voxels.

The estimates $\widehat{[\cdot, \cdot]}_p^\perp$ correspond very well to the true $[\cdot, \cdot]_p^\perp$ for all p in both the sheet and non-sheet case. The range of the estimates remains relatively constant for all cases. The sheet-case can be distinguished from the non-sheet-case for $[U, W]_p^\perp \gtrsim 0.015 \text{ mm}^{-1}$.

7.4.1.3 The influence of curvature

Fig. 7.8b shows the mean and range of the estimates $\widehat{[\cdot, \cdot]}_p^\perp$ for different curvatures $\kappa^{-1} = [8, 13, 18, 23, 28, 33]$, where we keep the Lie bracket normal component magnitude $[U, W]_p^\perp$ (non-sheet) constant by evaluating at different points $p = (x^1, -x^1, 0)$ (obtained by solving Eq. [7.10]). $[U, V]_p^\perp$ (sheet) is evaluated at the same points for reference. We set $k = 350$, $\delta = 1 \text{ mm}$, and $h_{\max} = 5 \text{ voxels}$.

The accuracy and precision of the estimates do not seem to depend heavily on the curvature at the considered SNR level and scale, since both the mean and range of the estimates remain relatively constant. We evaluated radii as small as 8 mm , which starts to approximate cortical folding radii. We note here that to detect even smaller radii, a smaller voxel size is required in order to have enough neighborhood information to probe the structure.

7.4.2 Diffusion MRI simulations

With the diffusion MRI simulations we will investigate more realistic noise scenarios (i.e. we can simulate realistic noise on the actual dMRI images instead of perturbing vectors), the influence of the interpolation technique (nearest neighbor vs fODF interpolation), and the influence of dMRI reconstruction technique (CSD vs DSI). In addition, we will explore the effect of using bootstraps instead of real noise iterations for the calculation of the sheet probability index. We use $N_n = N_b = 50$ noise iterations/bootstrap realizations here.

7.4.2.1 The influence of noise and interpolation

Here, we extract fODFs and peak directions using CSD from the single shell simulated data. Fig. 7.9a shows the mean and range of the normal component of the Lie bracket for different settings of h_{\max} (rows) and different SNR. We set the curvature $\kappa = 1/\rho = 1/26 \text{ mm}^{-1}$, voxel size $\delta = 1 \text{ mm}$, and evaluate $\widehat{[\cdot, \cdot]}_p^\perp$ at $p = (10, -10, 0)$. We use nearest-neighbor interpolation here, supplemental Fig. S7.1 shows the same results with linear interpolation on the fODF spherical harmonic coefficients.

Similar to the vector field simulations in Fig. 7.7, the range of the estimates becomes smaller with higher SNR (the noiseless case is also plotted) and the precision increases with increasing h_{\max} . Nearest neighbor interpolation and fODF interpolation give similar results in terms of both the accuracy and precision of the estimates for higher SNR (≥ 20). We hypothesize that a large number of loops reduces the influence of error propagation along a tract, and ‘smooth out’ some interpolation errors. For an SNR of 10, however, we found that 1 or 2 outliers cause the large range in Fig. 7.9a at $h_{\max} = 3 \text{ voxels}$. This does not occur when using fODF interpolation (Fig. S7.1), but the mean of the estimates still corresponds very well to the true value in both cases. For the sake of computational time and cost we use nearest neighbor interpolation in the following.

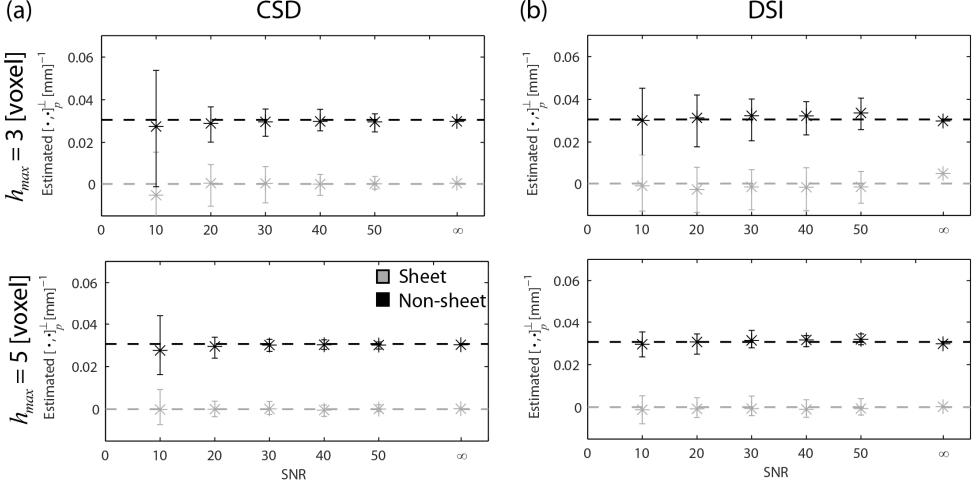


Fig. 7.9: Mean and range of $\widehat{[\cdot, \cdot]}_p^\perp$ for different settings of h_{max} (rows) and different SNR. We set the curvature $\kappa = 1/\rho = 1/26 \text{ mm}^{-1}$, voxel size $\delta = 1 \text{ mm}$ (giving $h_{max} = \{3, 5\} \text{ voxels} = \{3, 5\} \text{ mm}$), and evaluate $\widehat{[\cdot, \cdot]}_p^\perp$ at $p = (10, -10, 0)$. Dashed lines indicate the true Lie bracket normal component. (a) Peaks extracted from single shell data using CSD. (b) Peaks extracted from Cartesian sampled data using DSI.

7.4.2.2 The influence of dMRI technique: CSD vs DSI

In Fig. 7.9b, we extracted diffusion ODFs and peak directions using DSI from the Cartesian sampled simulated data. It shows the mean and range of $\widehat{[\cdot, \cdot]}_p^\perp$ for different settings of h_{max} (rows) and different SNR. We set the curvature $\kappa = 1/\rho = 1/26 \text{ mm}^{-1}$, voxel size $\delta = 1 \text{ mm}$, and evaluate $\widehat{[\cdot, \cdot]}_p^\perp$ at $p = (10, -10, 0)$.

The mean of the estimates with DSI is in good agreement with the true values. For $\text{SNR} > 10$, the Lie bracket estimates from CSD are more precise than the estimates resulting from DSI (smaller error bars, most obvious at $h_{max} = 3 \text{ voxels}$), even though the simulated CSD datasets have over five times fewer measurements (90 vs 514) and a lower diffusion weighting (maximum $b = 3000 \text{ s/mm}^2$ compared to $b = 10000 \text{ s/mm}^2$).

7.4.2.3 Sheet probability index

Here, we used the peak directions resulting from CSD on the single shell simulated data. Fig. 7.10 shows the mean and range of the estimates $\widehat{[\cdot, \cdot]}_p^\perp$ and the SPI P_λ for different points $p = (x^1, -x^1, 0)$ to vary the Lie bracket normal component magnitude $[U, W]_p^\perp$. $[U, V]_p^\perp$ is evaluated at the same points for reference. We set $\text{SNR} = 20$, $\delta = 1 \text{ mm}$, $\kappa = 1/26 \text{ mm}^{-1}$ and $h_{max} = 5 \text{ voxels}$.

The estimates $\widehat{[\cdot, \cdot]}_p^\perp$ reflect the true $[\cdot, \cdot]_p^\perp$ well for all x^1 in both the sheet- and the non-sheet case for the noise iterations. At $x^1 = 11$ the paths start to come in the vicinity of the vector field edge where the vector fields of the non-sheet pair make angles much smaller than the resolving power of CSD, this has a stronger effect on the bootstraps than on the noise iterations (hence the deviation). Overall the bootstraps prove good alternatives to real noise iterations. P_λ decreases in the non-sheet case when the true $[\cdot, \cdot]_p^\perp$ deviates more from zero.

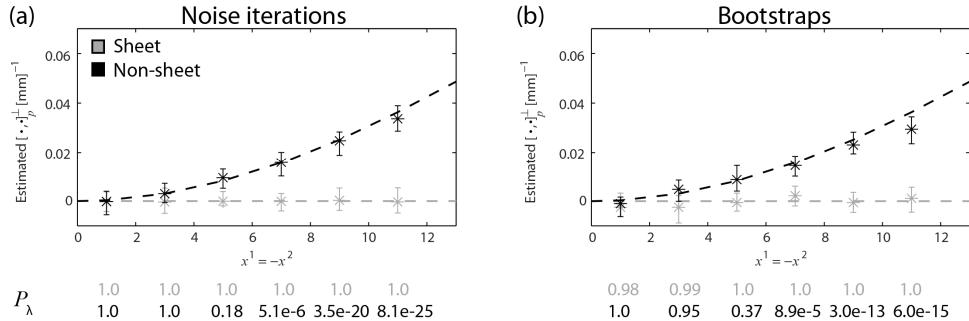


Fig. 7.10: Mean and range of $\widehat{[\cdot, \cdot]}_p^\perp$ and the SPI P_λ for different points $p = (x^1, -x^1, 0)$ to vary the Lie bracket normal component magnitude $[U, W]_p^\perp$, $[U, V]_p^\perp$ is evaluated at the same points for reference. We set SNR = 20, $\delta = 1$ mm, $h_{max} = 5$ voxels, and $\kappa = 1/26$ mm⁻¹.

7.4.3 Diffusion MRI real data

In this section we present SPI and sheet tensor maps for real diffusion MRI data (we use $N_b = 20$ bootstrap realizations if not mentioned otherwise). In Section 7.4.3.1 we explain the interpretation of these maps and their relation to the Lie bracket by means of an example data set, and in Section 7.4.3.2 we investigate inter-subject variability. In the remaining Sections 7.4.3.3, 7.4.3.4, and 7.4.3.5, we consider the influence spatial resolution, diffusion weighting, and dMRI technique, respectively.

7.4.3.1 Sheet probability index

Fig. 7.11a (left) shows Lie bracket normal component estimates of the two largest fODF peaks in a single slice of a WU-Minn HCP dataset. An FA color map of the same slice is shown for reference in Fig. 7.11b. Dark blue areas indicate $\widehat{[\cdot, \cdot]}_p^\perp \ll 0$, red areas $\widehat{[\cdot, \cdot]}_p^\perp \gg 0$, and light blue/green/yellow areas $\widehat{[\cdot, \cdot]}_p^\perp$ close to zero. The order of magnitude of $\widehat{[\cdot, \cdot]}_p^\perp$ is in agreement with our simulations. Three areas are highlighted with arrows. The areas indicated by the green and blue arrow look spatially continuous, whereas the area indicated by the grey arrow looks noisy. The two largest fODF peaks (used to create this image) in neighboring voxels do not necessarily belong to the same vector fields, we therefore have to consider the Lie bracket normal component for every pair of vector fields in each voxel. This further clarifies our motivation to use sheet tensors for visualization in the following since multiple sheet tensors can be visualized in each voxel. Histograms of the normal component for the bootstraps at these locations are shown in Fig. 7.11a on the right: high SPI (green arrow, $\widehat{[\cdot, \cdot]}_p^\perp$ concentrated around zero), ‘medium’ SPI (grey arrow, $\widehat{[\cdot, \cdot]}_p^\perp$ spread), and low SPI (blue arrow, $\widehat{[\cdot, \cdot]}_p^\perp$ concentrated but not around zero). The histograms illustrate that the normality assumption used to calculate the SPI is reasonable. Figs. 7.11c and d show SPI maps for the largest fODF peaks.

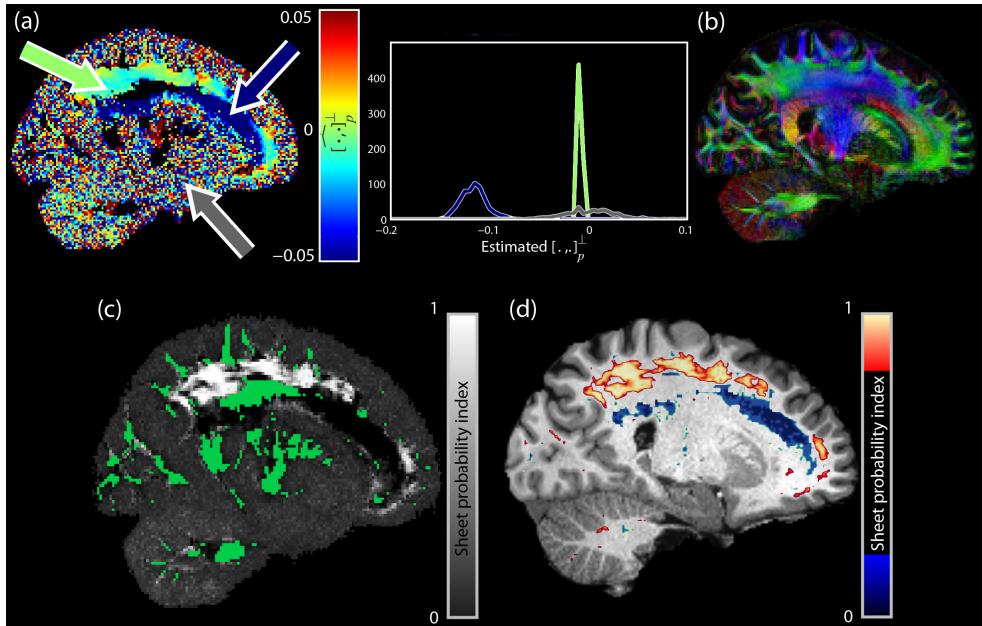


Fig. 7.11: (a) A single bootstrap of the Lie bracket normal component (two largest fODF peaks) in a single slice, together with histograms of the normal component at the indicated locations in high- (green arrow), medium- (grey arrow), and low-sheet probability (blue arrow) area. (b) A direction color-encoded FA map of the slice shown in (a) provided for reference. (c) The corresponding SPI map (maximum per voxel) with $\lambda = 0.008$. The green voxels only contain one peak and thus no Lie bracket can be computed. (d) The high- ($P_\lambda > 0.5$) and low-sheet probability areas ($P_\lambda < 0.1$) shown as an overlay on an anatomical scan.

Fig. 7.12a shows for one HCP subject sheet tensors on different coronal (top), sagittal (middle), and axial (bottom) slices. Here, $h_{max} = 5$ voxels is used, and sheet tensors with $P_\lambda < 0.2$ are not shown. When navigating through the brain slice-by-slice, these high-sheet probability areas seem to form continuous structures throughout the brain (see Supplementary Video S7.2). Fig. S7.3 shows similar results for $h_{max} = 3$, where we can recognize the same sheet areas (sometimes slightly less pronounced). Two high-SPI areas (green rectangle on coronal slice and red rectangle on sagittal slice) are detailed in Figs. 7.12b and c. The streamlines shown are a subset of the paths reconstructed to compute the Lie bracket in a voxel in the center of the high-SPI area (the paths $(\hat{\Phi}_{h_2}^W \circ \hat{\Phi}_{h_1}^V)(p)$ and $(\hat{\Phi}_{h_1}^V \circ \hat{\Phi}_{h_2}^W)(p)$, to be specific). Fig. 7.12b shows a sheet formed by the corpus callosum (CC) and the corticospinal tract (CST) in the left hemisphere (see also Supplementary Video S7.4). In addition, the white arrow highlights an area in which crossing sheets are found. Fig. 7.12c shows a more medial and sagittally oriented sheet structure, formed by parts of the CC/CST and anterior-posterior oriented association fibers. Details of a low SPI area (cyan rectangle on axial slice in Fig. 7.12a) are shown in Fig. 7.12d. This case highlights an important potential pitfall when using only visual and qualitative analysis to investigate sheet structures: Even though this structure much looks like a sheet from a superior point of view, it is clearly not a sheet from a lateral and posterior point of view (the fibers ‘diverge’ from each other and are not located on a surface, as can be seen in the views (1) and (2)). Our quantitative method indeed finds a low SPI in this area. Several high SPI areas in the brainstem could also be recognized, e.g. on the fourth coronal slice from the left in Fig. 7.12a.

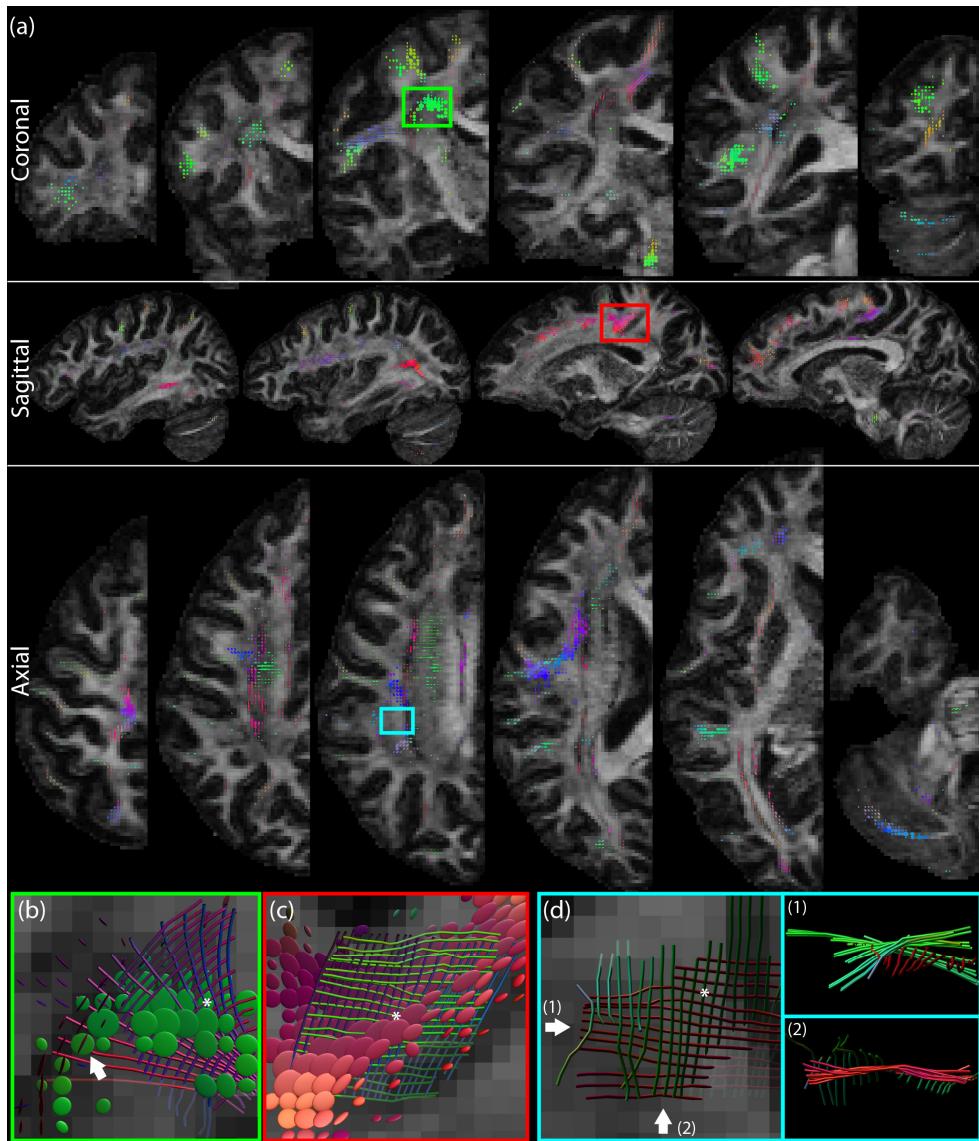


Fig. 7.12: (a) Sheet tensors ($\lambda = 0.008$) on different coronal (top), sagittal (middle), and axial (bottom) slices. Ellipsoids with $P_\lambda < 0.2$ are not shown for clarity, and the sheet tensors are colored according to Fig. 7.5b.. (b) High-SPI area with streamlines (paths $(\hat{\Phi}_{h_2}^W \circ \hat{\Phi}_{h_1}^V)(p)$ and $(\hat{\Phi}_{h_1}^V \circ \hat{\Phi}_{h_2}^W)(p)$ with $h_{max} = 5$ voxels used to compute the Lie bracket in a voxel marked by a white asterisk) of the CC and the CST in the left hemisphere. Non-orthogonal angles can be recognized, and the white arrow indicates crossing sheets. (c) A medial and sagittally oriented sheet structure, formed by parts of the CC/CST and anterior-posterior oriented association fibers. (d) A low SPI area in which the fibers look like a sheet from a superior view, but clearly diverge when inspecting other viewpoints (1) and (2).

7.4.3.2 Inter-subject variability

Fig. 7.13 shows results for 3 HCP subjects (first three rows, $h_{max} = 5 \text{ voxels} = 6.25 \text{ mm}$) and the MASSIVE dataset (last row, $h_{max} = 2.5 \text{ voxels} = 6.25 \text{ mm}$). For each subject, corresponding coronal, sagittal, and axial slices are shown in the different columns (two different slices per viewpoint). The arrows highlight examples of high-SPI areas that visually appear consistent across subjects.

7.4.3.3 The effect of spatial resolution

Fig. 7.13 compares the MASSIVE dataset (2.5 mm isotropic voxels) with the HCP datasets (1.25 mm isotropic voxels), where we kept the maximum distance h_{max} constant at 6.25 mm . In the MASSIVE dataset the same high-SPI areas can be recognized as in the HCP data.

Supplementary Fig. S7.5 shows results for the same subject as in Fig. 7.12 (voxel size $1.25 \text{ mm}, h_{max} = 5 \text{ voxels} = 6.25 \text{ mm}$) and Fig. S7.3 (voxel size $1.25 \text{ mm}, h_{max} = 3 \text{ voxels} = 3.75 \text{ mm}$), but now we downsampled the data spatially (voxel size $2.5 \text{ mm}, h_{max} = 2.5 \text{ voxels} = 6.25 \text{ mm}$). The same sheet structures can still be recognized, but some finer scale structures get lost.

7.4.3.4 The effect of diffusion weighting

Fig. 7.14 shows maps of the SPI for different diffusion weightings ($b = \{1000, 3000, 5000, 10000\} \text{ s/mm}^2$) of the MGH HCP data set, where we use CSD to extract the fODF peaks for every shell separately ($h_{max} = 5 \text{ voxels}$). At $b = 1000 \text{ s/mm}^2$ the SPI was significantly lower, which is the direct consequence of the decreased ability to resolve crossing fibers. $b = 10000 \text{ s/mm}^2$ results in the most extensive high-SPI areas, although most of these regions could already be recognized at a b -value of 3000 s/mm^2 .

7.4.3.5 The effect of dMRI technique: CSD vs DSI

Fig. 7.15a shows results of a single Lie bracket computation in tissue on the MGH DSI dataset ($h_{max} = 3 \text{ voxels} = 6 \text{ mm}$). Here, we visualize a tensor if $|\widehat{[\cdot, \cdot]}_p^\perp| \leq 0.008$ (all tensors have the same size), and we color the voxel red if the minimum $|\widehat{[\cdot, \cdot]}_p^\perp|$ in that voxel is larger than 0.025 (which would indicate that there is likely no sheet structure locally, see e.g. Fig. 7.10). The arrows indicate high-SPI areas that could also be identified in the previous experiments. Many high $|\widehat{[\cdot, \cdot]}_p^\perp|$ estimates are found in the grey matter, but the rectangles indicate example areas in the white matter where most likely no sheet exists. The paths in these areas are visualized in Figs. 7.15b and c.

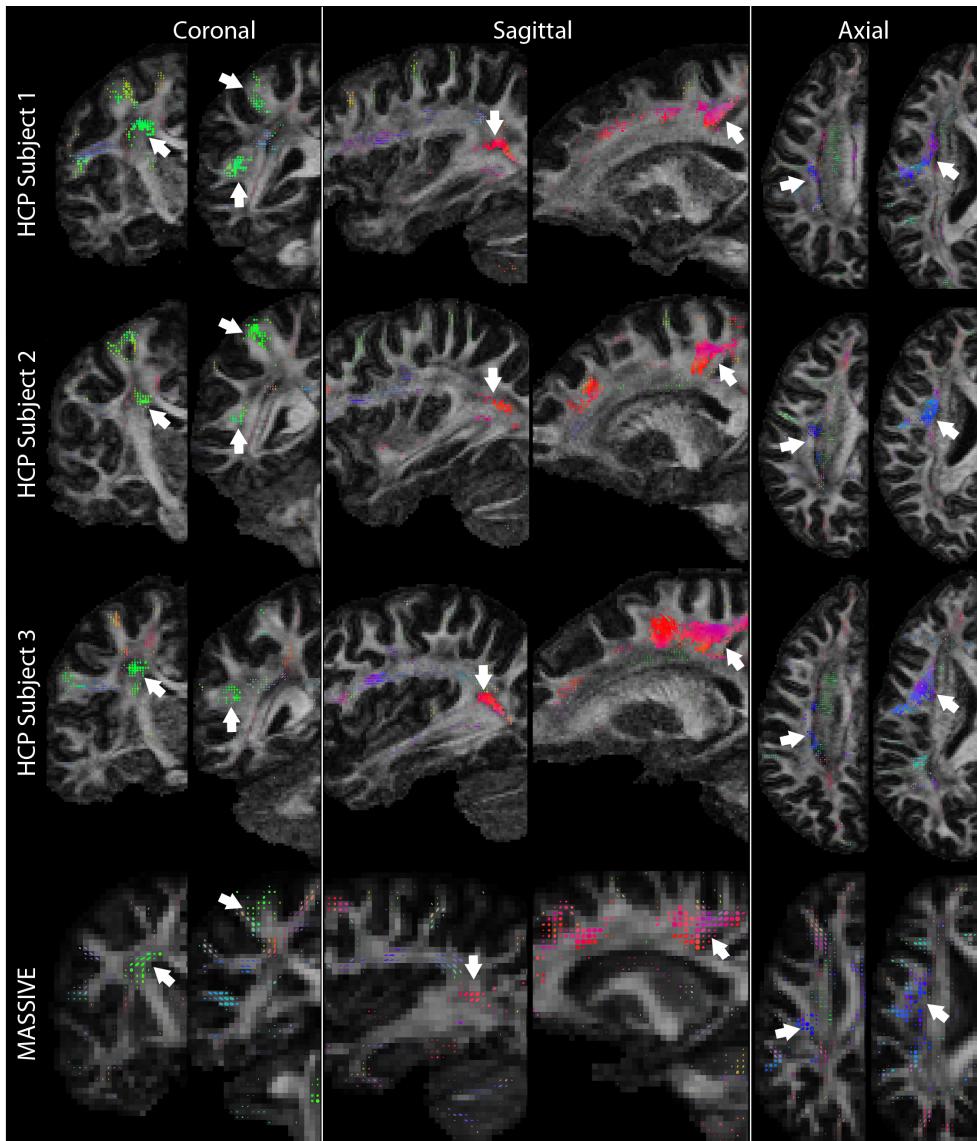


Fig. 7.13: Visual comparison of sheet structures between subjects and spatial scales (tensors with $P_\lambda < 0.2$ are not shown for clarity, colors according to Fig. 7.5b, and we set $\lambda = 0.008$). Examples of visually similar sheet structures are indicated by the arrows.

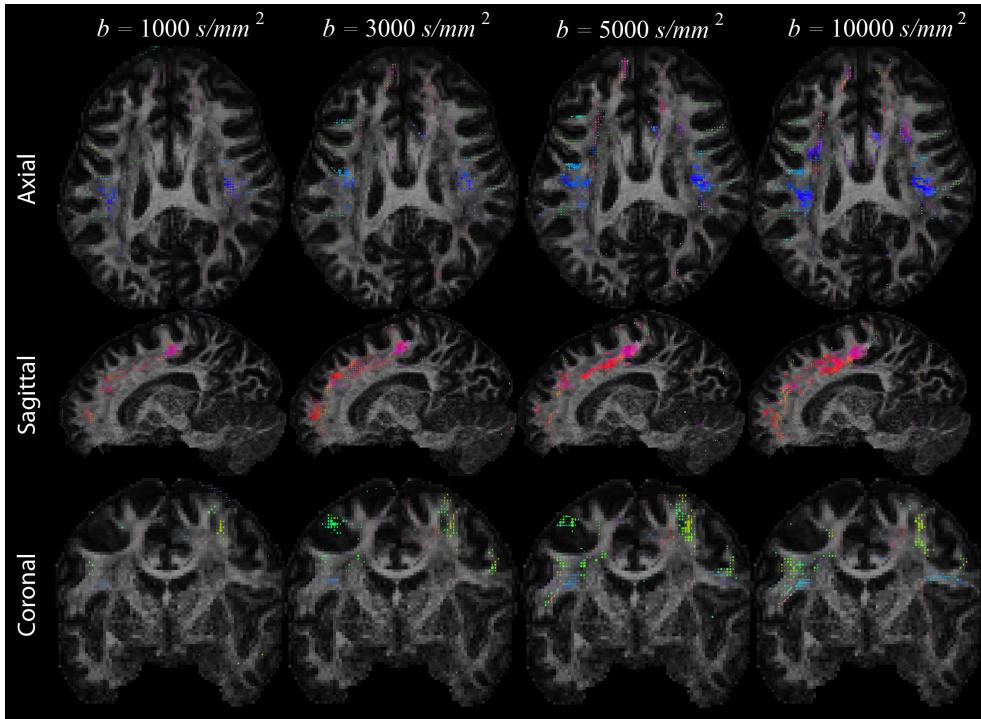


Fig. 7.14: The detection of sheet structure with different diffusion weightings in the MGH HCP data set, overlaid on the FA of the $b = 1000 \text{ s/mm}^2$ shell (tensors with $P_\lambda < 0.2$ are not shown for clarity, colors according to Fig. 7.5b, and we set $\lambda = 0.008$).

7.5 DISCUSSION

The hypothesis that brain pathways cross nearly orthogonally forming two-dimensional sheet-like structures is an active topic of debate (Catani et al., 2012; Wedeen et al., 2012a; Wedeen et al., 2012b). To date, there is no consensus on the large-scale existence of sheet structure, partly because the conditions for sheet structure are unclear (e.g., whether or not it depends on orthogonal angles), and, more importantly, because extensive quantitative proof is still lacking. In this work, we have focused on the definition of sheet structure defined as a surface formed by interwoven pathways, which does not depend on the angle of crossing fibers (see also Fig. S7.6 for a plot of the angle against the SPI). We have recapitulated the Frobenius theorem and investigated the discrete Lie bracket as a quantitative indicator of sheet structure. We performed extensive validation of the resulting algorithm by quantifying the effects of different settings and parameters. Finally we presented an investigation into the extent of sheet structure presence in the human brain for different spatial resolutions, dMRI models, and other acquisition parameters.

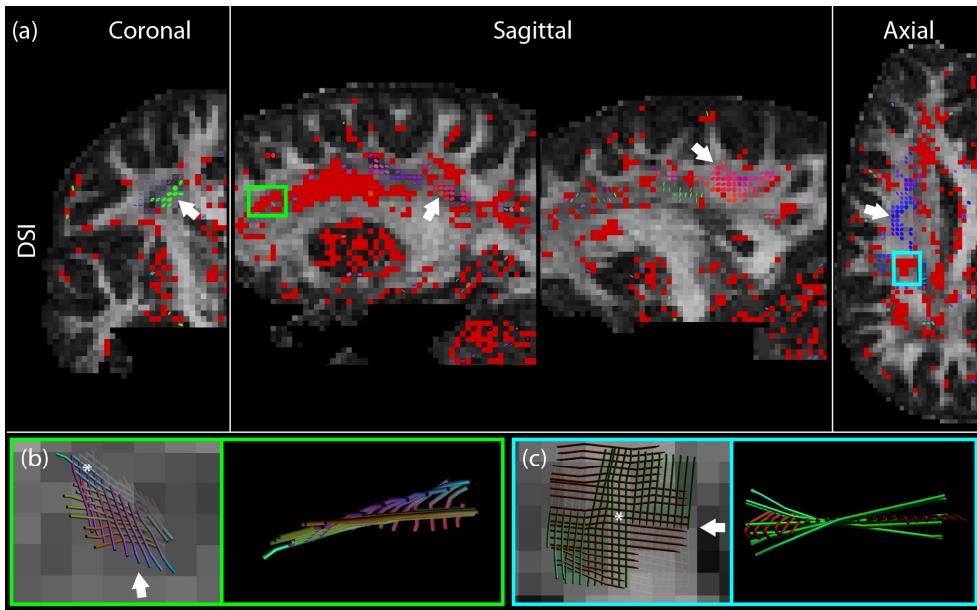


Fig. 7.15: (a) A map of $|\langle \cdot, \cdot \rangle_p^\perp|$ for the MGH DSI dataset ($h_{max} = 3$). A tensor is visualized if $|\langle \cdot, \cdot \rangle_p^\perp| \leq 0.008$ with colors according to Fig. 7.5b (tensors have the same size, arrows indicate example areas), and we color the voxel red if the minimum $|\langle \cdot, \cdot \rangle_p^\perp|$ in that voxel is larger than 0.025. Rectangles show spatially continuous example areas of high minimal $|\langle \cdot, \cdot \rangle_p^\perp|$, the corresponding paths for the voxel marked by a white asterisk are visualized in (b) and (c). The pathways clearly do not form a sheet but instead ‘diverge’ from the plane when inspecting them from a rotated view (orientation indicated by arrows).

7.5.1 Sheet or no sheet? Consistency with previous work

The first question that comes to mind is whether we can now prove or disprove the ubiquitous existence of sheet structure in the brain. Our simple simulations show that the discrete Lie bracket can distinguish between *vector fields* that do and do not form a sheet structure (Figs. 7.7 and 7.8). Also in the case of *vector fields derived from diffusion MRI simulations*, the method is able to correctly identify data representing a sheet structure (Figs. 7.9 and 7.10). We found that the performance was dependent on factors such as noise, voxel size, curvature, true Lie bracket normal component magnitude, and the chosen dMRI technique, which we discuss in the following sections. To be able to extend the findings from our simulations to the brain, however, we need to make the strong assumption that the vectors (or ODF peak directions) represent true underlying fiber directions, and that tractography correctly reconstructs true underlying bundles. Although these assumptions are often considered to be more or less valid in many connectivity studies, it is well-known that tractography is subject to many limitations and challenges (Jones, 2010a), also a concern raised by Catani et al. (2012).

Based on our results from real dMRI data (Figs. 7.10 – 7.15), we can only state that the *data supports the existence of sheet structures at several locations in the brain at the investigated scales*, with the SPI indicating the likelihood. Tractography pathways at locations with high SPI values were visually confirmed to form a sheet by reconstructing the path neighborhood as in

Wedgeen et al. (2012b) (Fig. 7.12bc). Relying only on visual inspection of (layers of) pathways, however, holds an important pitfall: where paths *seem to form a sheet* from a particular point of view (since a grid pattern is easy to recognize by the human eye), they may not be an actual sheet, which becomes more apparent when the view is rotated (Fig. 7.12d and 7.15bc). This discrepancy clearly shows the added value of quantitative analysis in the investigation of such structures in the brain. Whereas Wedeen et al. (2012ab) state that “no brain pathways were observed without sheet structure”, our results indicate that this is not the case: crossing fiber regions with very low SPI could be identified at this scale (Figs. 7.10-15). At some locations no reliable conclusion can be drawn on the existence of sheet structure for various reasons (e.g. only one fiber population could be reconstructed, the normality condition of the different bootstraps was not fulfilled, or the SPI was not clearly ‘high’ or ‘low’ (Fig. 7.11c) e.g. due to a high standard deviation in the Lie bracket normal component estimates).

Certain areas with high SPI values were found to be qualitatively consistent across subjects (Fig. 7.13), indicating that our framework provides reproducible results. In agreement with Wedeen et al. (2012ab), we found high SPI values in crossing regions of the corpus callosum with the cingulum (see supplementary video S7.2) and SLF 1-3 (e.g. Fig. 7.12c), and observed the continuous (grid) character of these major longitudinal pathways (as opposed to them being clearly distinct). We also found high SPI values in crossing regions of the corpus callosum and the corticospinal tract (Fig. 7.12b). The existence of this sheet structure has been much debated, and here we find that the data supports sheet structure at the location where these pathways cross (Wedgeen et al., 2012a) with non-orthogonal angles (Catani et al., 2012). However, we did not find a high SPI at crossings between callosal pathways and the fornix due to a high standard deviation of the Lie bracket normal component estimates. We could therefore not draw a reliable conclusion whether or not the data supports sheet structure at this location. Further extensive localization of sheet structures and investigation of the involved pathways is subject to future work.

Is the sheet structure something that can be trivially found in the brain, or is it a ‘special’ configuration? Wedeen et al. (2015a) remarked that the sheet structure “is mathematically specific and highly atypical, having prior probability ≈ 0 ”, and that “there are no mechanisms known whereby technical limitations will create it as an artifact”. Indeed, most configurations of vector fields do not form a sheet structure, and in this sense the sheet structure is thus special. On the other hand, one configuration in which two thicker bundles trivially form sheets is when they are both straight. Even though this may approximately occur in some regions (e.g. in the cingulum/corpus callosum), tracts exhibit a significant curvature at the scale we have investigated, and the results we have obtained also show high SPI in regions with high curvature in the streamlines of both vector fields (e.g. Fig. 7.12b). At this point, we cannot ascribe this phenomenon to a more straightforward alternative geometry.

We did not further investigate the issue of pathways making sharp turns (Wedgeen et al., 2012a), which we consider a separate topic; it cannot easily be addressed by current tractography algorithms or by the frame tractography used here because of necessary smoothness constraints.

7.5.2 The issue of scale

The discrete Lie bracket, and thus the derived SPI and sheet tensor maps, are *locally* defined in terms of the surrounding structure. The term ‘local’ here implies that spatial scale is an important factor in the method. The flow distance h_{max} (that determines the extent to which the neighborhood is taken into account), the voxel size δ , and the curvature of the streamlines affect the performance of the algorithm in different but related ways.

In Fig. 7.7 we show that for a fixed voxel size and fixed curvature, an increase in h_{max} improves the accuracy and precision of the method. This is likely a result of the corresponding increase of the number of data points in the least squares fit (Eq. [7.7], where K is determined by h_{max} and the fixed step size $\Delta h = \delta/2$). This also motivates the incorporation of multiple loop configurations and the exploration of all four ‘quadrants’ surrounding point p (Section 7.3.1.3). Though increasing h_{max} generally has a positive effect, its value is naturally limited by the domain of definition of the vector fields. If a significant number of the loops extend beyond this domain, the accuracy and precision can be expected to drop. This means we should not choose the value of h_{max} to be much greater than the expected size of the sheet structure, which leads to the interpretation of h_{max} as a *sheet structure scale parameter*: h_{max} serves as an approximate lower bound to the size of the sheet structures that can be detected with the algorithm. Note that taking $h_{max} \leq 1$ voxel does not lead to reliable estimates of the Lie bracket (Fig. 7.7), so the voxel size δ is, not unexpectedly, a hard limit on the size of detectable sheet structures. Related to the note that “grid structure was maintained at all scales, from the single voxel, to the lobe, to the hemisphere” (Wedge et al., 2012b), we can thus conclude that it is only possible to reliably detect sheet structures larger than the voxel scale.

Fig. 7.7 furthermore shows that the accuracy and precision varies with the voxel size δ and with h_{max} in voxels, but remains relatively constant when h_{max} is defined in millimeters (the product of the former two) at the scales considered. This gives the definition of h_{max} as a sheet structure scale parameter a more intuitive physical interpretation. In our real data experiments, we kept h_{max} constant at 6 – 7 mm. The optimal detection of a given sheet structure then involves tuning of these parameters: for example smaller or highly curved sheet structures require smaller voxel sizes (to be able to set $h_{max} > 1$ voxel) at the cost of a lower precision when keeping h_{max} in voxels constant, or at the cost of higher computational demands when calculating more paths for a higher h_{max} in millimeters. This lower precision at smaller voxel sizes (Fig. 7.7) occurs because the tractography error *in terms of voxels* remains more or less constant for a given SNR and h_{max} in voxels. Deviations in the Lie bracket, having units mm^{-1} , will thus be larger for a smaller voxel size in terms of millimeters. This implies that the normal component of the Lie bracket should be larger in order to still be able to distinguish sheet from non-sheet for a given SNR and h_{max} in voxels at a smaller voxel size. The physical limits of detecting sheets at particular scales have to be examined further in future work. At the scale investigated in our work, curvature does not have a significant effect (Fig. 7.8b).

The angle threshold is another parameter that can be varied in our algorithm. In this work, we have set a constant angle threshold of 35° for the whole brain. In tractography, however, a single threshold might not be optimal for all brain pathways and should be adapted to the curvature of the tract relative to the voxel size and the SNR of the data, amongst others

(Chamberland et al., 2014). This reasoning can be extended to the frame tractography used in our algorithm: the angle threshold being set too low might result in suboptimal Lie bracket normal component estimates and a failure to detect curved sheets (Fig. S7.7). Conversely, if the angle threshold is set too high, this might result in tracts taking a wrong turn, outliers, and a lower accuracy and precision.

7.5.3 The impact of diffusion weighting and dMRI method

The impact of the dMRI method, using DSI versus e.g. SD, has also been a big part of the debate (Catani et al., 2012; Wedeen et al., 2012a). This discussion centers on the ability of these models to accurately resolve the orientations of fiber populations. Although the initial concern was that DSI “does not allow separation of fibers that cross at non-orthogonal angles, thus making a grid structure of interwoven sheets a very likely configuration” (Catani et al., 2012), we find instead that the ability to robustly detect fiber populations mainly influences the precision and accuracy of the Lie bracket estimates (and thus the SPI), but it does not necessarily *promote* sheet structure.

The first factor of importance is the ability to detect *crossing fibers*, since 1) the Lie bracket cannot be computed in voxels with a single fiber population, and 2) paths end prematurely if peaks of a vector field are missing, reducing the number of difference vectors K and thus potentially reducing the accuracy and precision (a similar effect as shown in Fig. 7.7, where a lower h_{max} in voxels equals a lower amount of reconstructed paths). We visually confirmed that in the DSI experiment (Fig. 7.15) a lower amount of crossing fibers was detected than in a CSD experiment with similar or lower spatial resolution (e.g. MASSIVE data in Fig. 7.13), resulting in a lower amount of voxels where the Lie bracket could be computed. This is in agreement with Catani et al. (2012), where it was stated that DSI likely has a lower angular resolution. A second criterion that is of importance here is robustness to noise, or the accuracy and precision of the peak estimates. We investigated this effect using simulated dMRI data (based on the best scoring model of experiments in Ferizi et al. (2014), which also included high b -values). The results shown in Fig. 7.9 suggest that the performance of DSI and CSD is comparable for a broad range of SNR with CSD having a higher precision, contradicting the statement in Wedeen et al. (2012a) that “DSI should present the lower risk of bias”.

A comparison between the CSD results of the MASSIVE data with voxel size $\delta = 2.5 \text{ mm}$ (Fig. 7.13, bottom row) and the MGH DSI data with $\delta = 2 \text{ mm}$ (Fig. 7.15) reveals similar large scale high-SPI areas. These sheet areas could also be observed in the Wu-Minn HCP subjects (Fig. 7.13, first three rows). Neither CSD nor DSI results in the detection of sheet structure at *every* crossing fiber location in the brain. Based on our findings, we conclude that DSI has no bias towards detecting sheet structure, and also reveals non-sheet areas (Fig. 7.15).

A higher diffusion weighting generally causes an increase in the number of voxels with a significantly high SPI value (Fig. 7.14). Using CSD we find very little high-SPI areas in the $b = 1000 \text{ s/mm}^2$ shell of the MGH HCP data set, while for higher b -values the percentage of voxels that are likely to form sheet structures increases. This is consistent with the fact that the angular resolution increases with higher b -values, confirming the importance of a reasonable “diffusion resolution” (Wedeen et al., 2012a). Increased diffusion weighting, however, also comes at the cost of a lower SNR. Generally, a b -value of 3000 s/mm^2 is used for CSD, and

the vast majority of the sheets detected at this b are retained when moving to higher diffusion weighting.

With the ever increasing amount of proposed diffusion models, the reliable extraction of fiber directions is still an active area of research. To assess the presence of sheet structure in dMRI data, one requires a robust and reliable means to determine these peak directions. Our method is general and not limited to a particular dMRI technique or acquisition scheme, and we therefore believe that its performance can be improved with ongoing technical developments in the field.

7.5.4 Further methodological considerations and potential improvements

The error term ϵ in Eq. [7.7] includes errors due to the approximation of Eq. [7.4], and errors in the streamline tractography. The approximation errors depend on h_1 and h_2 and on the underlying vector fields. These errors are small and turn out to be negligible compared to other sources of errors. The tractography errors result from measurement noise, interpolation, curvature of the tracts, and step size Δh (which is linked to the voxel size in our case), among others. Our experiments indicate that noise has the largest effect; in the case of infinite SNR the estimate is accurate. Regarding interpolation, we opted for simple nearest neighbor interpolation of the vector fields for reasons of speed and computational efficiency (both for DSI and for CSD). In the case of CSD we performed additional experiments using fODF interpolation (strictly speaking the actual diffusion measurements would have to be interpolated), which is more precise (Fig. 7.9a and Fig. S7.1) but also more computationally intensive because peak extraction has to be performed at every step. In future work, more advanced tractography algorithms could be used to reconstruct the loops and estimate the Lie bracket (e.g. using more complex integration schemes or combining model fitting and tractography (e.g. Daducci et al., 2015b; Reisert et al., 2014)). The question of whether the added benefit outweighs the additional computation time (the streamlines computed in the algorithm are fairly short) remains to be answered.

To estimate the Lie bracket and its normal component at point p we reconstruct multiple loop configurations, in all quadrants surrounding point p , and with a range of walking distances. Currently we do a simple linear least squares fit (Eq. [7.7]) on the difference vectors of these loops. We do not take into account the variance of the measurements or potential outliers that arise, for example, when the tractography takes a wrong path during the reconstruction of a loop. This can potentially be improved by doing a (robust) iteratively weighted least squares fit, with the weight dependent on the variance of the values for a given h_1h_2 from corresponding loops. In addition, the variance of the difference vectors from single loops potentially holds information on the sheet probability; however, it does not give information on the uncertainty of the underlying peaks.

To compute the SPI maps in the case of CSD we use residual bootstrapping, since the acquisition of repeated dMRI data sets is mostly not feasible. In Fig. 7.10 we show comparable results between the SPI obtained with residual bootstrapping and the SPI obtained with true noise iterations (in the case of simulations). To the best of our knowledge no bootstrapping method exists for DSI, so the reported DSI results are based on a single noise iteration. This makes the quantification and investigation of sheet structure more difficult since the choice of

threshold has a hard effect on the visualization of the sheet tensors, sometimes resulting in regions with noisy (isolated) sheet tensors and high Lie bracket normal components (Fig. 7.15). This further illustrates the necessity of more extensive evaluation of measurement variance (i.e., the computation of the SPI as opposed to considering only a single measurement) for reliable quantification of sheet structure.

When computation of the SPI was possible (in the case of CSD), we set $\lambda = 0.008$ and we excluded the small percentage of voxels that did not have normally distributed Lie bracket normal components over the bootstrap iterations. In future work, the distribution of normal components per voxel and the optimal way to extract an SPI from this (e.g. detection of outliers, fitting, threshold settings) could be investigated more thoroughly. In this work, the value for λ was chosen based on the variability in simulation experiments, and the used setting resulted in regions of smoothly varying SPI with sheet tensors of a similar orientation (i.e. normal) in a certain neighborhood. This suggests that continuous sheet structures could perhaps be better visualized as actual surfaces; tractography could be extended to sheetography by means of a surface propagation process. There are examples of surface reconstruction approaches for DTI data, e.g. Vilanova et al. (2004); Zhang et al. (2003), which compute streamsurfaces at points where the DTI tensor has a high planarity coefficient C_p (Westin et al., 2002). We found points in the data where the SPI and the C_p were both high, and where the reconstructed surfaces corresponded well with the information represented by the sheet tensor. This is however no strict prerequisite for the presence of sheet structure. There are places with a high C_p without sheet structure (i.e., low SPI due to spatial incoherence of the data), and places with a low C_p that do show evidence of sheet structure (i.e., high SPI, for example in the case of several sheet structures crossing in a voxel). A comparison between the SPI and different DTI shape measures is included in the supplementary materials (Fig. S7.8).

In this manuscript we view the Lie bracket in terms of the integral curves of vector fields (Eq. [7.4]), which is equivalent to a combination of differential operators on the vector fields known as the commutator (see Eq. [A7.2] in Appendix B7 for details). This raises the question of whether there might be an alternative way to calculate the Lie bracket that does not require the reconstruction of many loops, which is computationally expensive. A direct discretization of Eq. [A7.2] would involve the computation of vector field derivatives, which is a complicated operation. A finite difference implementation does not give stable results in the case of noise. A very recently proposed method to estimate the Lie bracket for diffeomorphic registration purposes computes central difference Jacobians of discrete vector fields by considering them as bandlimited signals in the Fourier domain (i.e. truncating the high frequency components) (Zhang and Fletcher, 2015). However, in addition to discretization and noise, our application has other major challenges: 1) all peaks in a certain neighborhood would have to be clustered into distinct vector fields to be able to compute the Jacobian; 2) there are potential sign inconsistencies between neighboring peaks; and 3) there is no guarantee that all peaks of a vector field exist in a certain neighborhood. The first two challenges are addressed in our algorithm by clustering the peaks into vector fields ‘on the fly’ during the proposed frame tractography. The third challenge still affects our method: In the case of missing peaks the path is terminated, resulting in fewer difference vectors to compute the Lie bracket. This has an important influence on the precision of the estimates (see e.g. the experiments for different h_{max} in Fig. 7.7). Future work will be directed towards investigating whether the alternative

definition of the Lie bracket in terms of the Jacobians can be used to obtain an estimate of the SPI, omitting the computationally expensive reconstruction of many paths per voxel required for the current method.

7.5.5 Future perspectives

Our results indicate that areas with high SPI values are relatively consistent between healthy subjects, and we therefore hypothesize that they could be used as new structural features of the brain. The extensiveness, orientation, and spatial distribution of sheet structures could be altered in the case of pathology. For example, if these sheets truly occur in the brain like “the warp and weft of fabric” (Wedeen et al., 2012a), it might be the case that space occupying lesions could dislocate the whole sheet structure as opposed to individual pathways. Another interesting feature could be the angle between pathways that form sheets through the shape of sheet tensors; it was shown that there was a significant crossing-angle difference in the frontal connections between a schizophrenia and healthy control group (Pasternak et al., 2012). The hypothesis that sheet structures have a close connection to development, axonal path finding, and the chemotactic gradients of early embryogenesis (Wedeen et al., 2012b) can now be investigated in a quantitative fashion. Whereas we only visually confirmed the consistency of areas with high SPI values between subjects in this work, a quantitative evaluation should be performed in future work. Such an inter-subject or inter-group evaluation would require a way to register these structures towards each other. Preliminary results of a quantitative comparison in which we registered the FA images of HCP subjects and calculated the overlap of thresholded maximum SPI maps indicated only a moderate overlap (Fig. S9). However, in analogy to recent insights in tract-based analyses, it is likely not optimal to work in voxel coordinate space and register scalar volumes such as FA, since such methods can for example not distinguish between nearby but differently oriented tracts (O'Donnell et al., 2009). Instead, point correspondences should be found ‘in the space of sheet structures’. Ongoing developments in registration of tensor fields (e.g. Zhang et al. (2006a), which could potentially be applied to sheet tensors), ODFs and multi-fascicle models (e.g. Raffelt et al., 2011; Taquet et al., 2014), and tractography data (e.g. Garyfallidis et al., 2015; O'Donnell et al., 2012) could contribute towards this end.

Although we find that the dMRI data investigated in this work supports the existence of sheet-like structures at certain locations in the brain, it should be noted that the dMRI data reflects just a few aspects of the true underlying structure and its derived tracts do not correspond to true axons. Ideally, the existence of sheet structure should also be validated with a ‘gold standard’, such as histology, and quantified with other techniques that can map brain structure orientations. Exciting new technologies such as CLARITY (Chung et al., 2013) and polarized light imaging (Aixer et al., 2001) could provide more insight into the existence of sheet structures and the scale on which they exist. Our method can be used in combination with such techniques: the Lie bracket computation is based on vector fields and could therefore be extended to directional data derived from these other techniques.

7.6 CONCLUSION

The extensive presence of sheet structures in the brain has been debated since its proposal, mainly due to its unclear relation to orthogonal angles and a lack of quantitative characterization. In this work we have explored the necessary and sufficient condition for a sheet structure to exist, which involves the computation of the Lie bracket of vector fields. We have proposed a method to compute the Lie bracket throughout the brain taking into account challenges such as discretization, noise, and clustering of vector fields. We have proposed a novel metric based on the Lie bracket, the sheet probability index (SPI), which indicates the extent to which the data supports sheet structure. In simple vector field and diffusion MRI simulations the method is able to quantitatively distinguish sheet from non-sheet structure, with spatial resolution and SNR being important factors that influence the accuracy and precision. Real diffusion MRI data experiments reveal a high SPI at various locations in the brain at the investigated scale, but also low SPI areas were found. Several high SPI areas could consistently be recognized across subjects, scanners, diffusion MRI techniques (i.e. CSD vs DSI), and spatial resolutions. Neither CSD nor DSI finds a high SPI at every location in the brain, and we find no bias towards sheet structure for DSI. Since tractography pathways do not represent true axons, validation of sheet structure with other technologies is necessary, and the proposed method can be extended to quantify sheet structure in other directional data. We hypothesize that sheet structure location, extent, and orientation could serve as new and important structural features of the brain.

ACKNOWLEDGEMENTS

The authors thank many colleagues in the field for useful discussions. Maxime Chamberland is acknowledged for help with visualization using the FiberNavigator.

7.7 APPENDICES

7.7.1 Appendix A7: List of symbols

U, V, W, X, Y, Z, \dots	Vector fields, where a field X is defined on $N_X \subset \mathbb{R}^3$	S_{ij}	The cosine similarity between the i -th and j -th vector in an ordered set
$p, q \in M \subset \mathbb{R}^3$	Positions in the brain M	h	Discrete flow distance along an integral curve
X_p	Vector X at point p	Δh	Step size in a discrete flow
$S \subset \mathbb{R}^3$	Sheet structure	h_{max}	The maximum distance in a discrete flow
$\Phi^X(s, p), \Phi_s^X(p)$	The flow operator, along X for a distance s and with initial position p	R, H, β, ϵ	Matrices
R_p	The vector given by the deviation from p after specific flows, called the closure	P_λ	The sheet probability index with parameter λ
$[V, W]_p$	Lie bracket of V and W at point p	$N(\mu, \sigma^2)$	Normal distribution with mean μ and standard deviation σ
$[V, W]_p^\perp$	Normal component of $[V, W]_p$	x^i	Cartesian coordinates
$\widehat{\Phi}_s^X$	Discrete approximation to the flow Φ_s^X	T_p	Sheet tensor at p
$\widehat{[V, W]}_p$	Discrete Lie bracket of V and W at point p	$X_p \otimes X_p, X_p X_p^T$	Tensor product between two identical vectors X_p
\widehat{R}	Estimated closure	κ, ρ	The curvature κ , and its reciprocal ρ
$[a, b, \dots]$	Ordered set of elements	δ	Voxel size
P	A permutation of an ordered set	N_n, N_b	The number of noise or bootstrap iterations
E	Similarity energy	b	The b -value

7.7.2 Appendix B7: Mathematical background

There are three concepts that are explained or mentioned in the main text, which we feel might benefit from a more detailed exposition. These are the formal definition of the Lie bracket and its relation to the definition in terms of flows along vector fields (Section 7.3.1.2), the relation between the Lie bracket and the sheet structure (Sections 7.3.1.1 and 7.3.1.3), and the definition of three-dimensional grid structure (e.g. Section 7.1). In this appendix we discuss these more in-depth.

7.7.2.1 Appendix B.1 The Lie bracket and its relation to flow

In a *differentiable manifold* M , a *tangent vector* V_p at $p \in M$ is identified with the directional derivative (denoted by ∇_V) of any suitably differentiable function $f: M \rightarrow \mathbb{R}$. Hence we can write $V_p f = \nabla_V f(p) = \sum_{i=1}^3 V^i \frac{\partial f}{\partial x^i}(p)$, given the components $\{V^1, V^2, V^3\}$ of V relative to the coordinates $\{x^1, x^2, x^3\}$ on M . Let $T_p M$ denote the space of all tangent vectors at $p \in M$, and let the *tangent bundle* TM represent the union $\{T_p M \mid p \in M\}$, i.e., the set of all tangent vectors in M . A vector field V is now defined as the smooth mapping $V: M \rightarrow TM: p \mapsto V_p$. The Lie bracket of two (non-zero) vector fields V and W on M is defined by its action on smooth functions f :

$$[V, W]_p(f) = V_p W(f) - W_p V(f) = \nabla_V \nabla_W f(p) - \nabla_W \nabla_V f(p). \quad [\text{A7.1}]$$

The Lie bracket satisfies the defining properties of a derivative (linearity and Leibniz' product rule), and so $[V, W]_p$ can be regarded as a directional derivative, i.e. an element of the vector space $T_p M$. If one is interested specifically in the Lie bracket $[V, W]$ in terms of the vector fields V and W , Eq. [A7.1] can be used to obtain

$$[V, W] = \nabla_V W - \nabla_W V, \quad [\text{A7.2}]$$

which is commonly known as the commutator of the vector fields, and where $\nabla_V W$ denotes the derivative of W along V .

To relate the Lie bracket to flows $\Phi_t^V: M \rightarrow M: p \mapsto \Phi_t^V(p)$ we write, using Taylor's theorem,

$$\Phi_t^V(p) = \sum_{k=0}^{\infty} \frac{t^k}{k!} \left(\frac{d^k}{dt^k} \Big|_{t=0} \Phi_t^V(p) \right) = p + tV_p + \frac{1}{2} t^2 V_p V + O(t^3). \quad [\text{A7.3}]$$

Note that because $\frac{d}{dt} \Big|_{t=0} f(\Phi_t^V(p)) = V_p(f)$ for any suitably differentiable f (compare this to the definition $\frac{d}{dt} \Big|_{t=0} \Phi_t^V(p) = V_p$ used in the main text) it follows that $\frac{d}{dt} \Big|_{t=0} W_{\Phi_t^V(p)} = V_p W$. Taking $s(t) = t$ this gives

$$\begin{aligned} (\Phi_{s(t)}^W \circ \Phi_t^V)(p) &= \Phi_t^V(p) + s(t)W_{\Phi_t^V(p)} + \frac{1}{2}s(t)^2 W_{\Phi_t^V(p)} W + O(s(t)^3) \\ &= p + tV_p + \frac{1}{2} t^2 V_p V + O(t^3) + s(t)W_p + s(t)tV_p W + O(s(t)t^2) \\ &\quad + \frac{1}{2}s(t)^2 W_p W + O(s(t)^2 t) + O(s(t)^3), \end{aligned} \quad [\text{A7.4}]$$

where we repeatedly apply Eq. [A7.3], such that

$$(\Phi_t^W \circ \Phi_t^V)(p) = p + t(V_p + W_p) + t^2 V_p W + \frac{1}{2} t^2 (V_p V + W_p W) + O(t^3). \quad [\text{A7.5}]$$

Continuing this process to compute $(\Phi_{s(t)}^{-V} \circ \Phi_t^W \circ \Phi_t^V)(p)$ and $(\Phi_{s(t)}^{-W} \circ \Phi_t^{-V} \circ \Phi_t^W \circ \Phi_t^V)(p)$, we then find

$$\begin{aligned} (\Phi_t^{-W} \circ \Phi_t^{-V} \circ \Phi_t^W \circ \Phi_t^V)(p) &= p + t^2 (V_p W - W_p V) + O(t^3) \\ &= p + t^2 [V, W] + O(t^3), \end{aligned} \quad [\text{A7.6}]$$

after which Eq. [7.4] follows readily. Validity of the alternative definitions in Eq. [7.6] can be derived analogously. If V and W are defined on subsets N_V and N_W of M as in the main text, then the Lie bracket is only defined on the intersection $N_V \cap N_W$. In the remaining sections of this appendix we will continue with V and W defined on the whole of M .

7.7.2.2 Appendix B.2 Sheet structure and the Lie bracket

The flow $\Phi_t^V(p)$ defines a curve $t \mapsto \Phi_t^V(p)$. The curves obtained when varying p form the *leaves* of a one-dimensional *foliation* of M ; they are disjoint immersed submanifolds that partition M , which are tangent to V for all p and t . The existence of such a one-dimensional foliation given V is ensured by smoothness of V , but for higher-order foliations that are tangent to multiple vector fields there are additional constraints. These are summarily provided by the Lie bracket, which guarantees the existence of a foliation if the set of all vector fields that are locally tangent to the sheet, is closed under the Lie bracket.

For the existence of a two-dimensional foliation given two locally non-collinear vector fields V and W , it is sufficient if for all $p \in M$ we have $[V, W]_p \in \text{span}(V_p, W_p) \subset T_p M$, which is equivalent to the requirement $[V, W]^\perp = 0$ used in the main text. The fact that $[V, W]^\perp = 0$ is a necessary and sufficient condition for the existence of a foliation (with leaves that are locally perpendicular to both V and W) is known as the Frobenius theorem (Lang, 1995; Spivak, 1979). If V and W form a foliation, then all points on a leaf can be connected with a flow along these vector fields. The two-dimensional foliation of (a subset) of \mathbb{R}^3 is called the *sheet structure*.

7.7.2.3 Appendix B.3 Grid structure

Grid structure as introduced by Wedeen et al. (2012b) is finally defined as three independent two-dimensional foliations, which form a three-dimensional foliation. Given three independent vector fields U , V , and W on M , a sufficient condition for the existence of this structure is that the Lie bracket normal component of each pair of vector fields is zero. As was the case with sheet structure, grid structure ensures that any two points on the leaves of this foliation can be connected with a flow along the vector fields. In the considered case of \mathbb{R}^3 , this means that any two points in the entire space can be connected by means of such a flow: the flows along U , V , and W are coordinate lines.

7.7.3 Appendix C7: Algorithms

7.7.3.1 Appendix C7.1 Clustering of frames

Consider an *ordered set* of n vectors $[Y_1, \dots, Y_n]$ at a position q , and m vectors $\{Z_1, \dots, Z_m\}$ in some point r near q ¹. The ordered set $[Y_1, \dots, Y_n]$ serves as a *reference frame*, i.e., we assume that n vector fields U, V, \dots are present in the local neighborhood of q that satisfy $U_q = Y_1, V_q = Y_2, \dots$. The aim of the clustering algorithm is to find a permutation of the frame (an ordered set $[Z_{P_1}, \dots, Z_{P_n}, 0, \dots]$) that corresponds to the reference frame, so that we can take $U_r = Z_{P_1}, V_r = Z_{P_2}, \dots$ for some *permutation* P of $[1, \dots, n]$ (see Fig. 3 for a schematic example). Here, P_i denotes the index given by the i^{th} element of P , and $Z_{P_i} = 0$ implies that no matching vector was found.

Clustering is done by maximizing (over all permutations P) a *similarity measure* E that represents the total element-wise similarity between the frames $[Y_1, \dots, Y_n]$ and $[Z_{P_1}, \dots, Z_{P_n}]$. In the algorithm below we will use the ‘total cosine similarity’ for E , which is defined as the sum of the cosines of the angles between corresponding vectors. The steps of our algorithm are as follows:

Algorithm for the clustering of frame $\{Z_1, \dots, Z_m\}$ given an ordered frame $[Y_1, \dots, Y_n]$

- 1) Compute the cosine similarity S_{ij} between Y_i and Z_j for all $i \in \{1, \dots, n\}, j \in \{1, \dots, m\}$ (recall that Y_i and Z_j are unit or zero vectors):

$$S_{ij} = |Y_i \cdot Z_j|. \quad [\text{A7.7}]$$

- 2) For every n -permutation P of $[1, \dots, m]$ (e.g. for $n = 2$ and $m = 3$ these are the permutations $[1,2], [2,1], [1,3], [3,1], [2,3]$, and $[3,2]$), compute the similarity energy of the permutation by:

$$E_P = \sum_{i=1}^n S_{iP_i} \quad [\text{A7.8}]$$

(e.g. for $P = [2,3]$ in the example above, which associates Z_2 to Y_1 and Z_3 to Y_2 , we have $E_P = S_{12} + S_{23}$). Note that taking the sum (as opposed to the mean) favors the assignment of more vectors.

- 3) Determine the permutation P for which E_P is maximal, and define the reordered set of vectors $[Z_{P_1}, \dots, Z_{P_n}]$.
- 4) If $Y_i \cdot Z_{P_i}$ has a negative sign, set Z_{P_i} to $-Z_{P_i}$.

¹ If $m < n$, we append $n - m$ zero-vectors to the list $\{Z_1, \dots, Z_m\}$, so in the following we can take $m \geq n$.

- 5) Apply an angle threshold on Z_{P_i} . If $|Y_i \cdot Z_{P_i}| < t$ (here set to $\text{acos } 35^\circ$) for some threshold $t \in \mathbb{R}$, set Z_{P_i} to 0.
- 6) Return $[Z_{P_1}, \dots, Z_{P_n}]$.

7.7.3.2 Appendix C7.2 Frame tractography

Algorithm for the approximate flow $\hat{\Phi}_h^X$ along vector field X from point q_0

- 1) While the number of taken steps l is less than or equal to a preset number of steps $L = \lfloor h/\Delta h \rfloor$ (where $\lfloor \cdot \rfloor$ denotes the floor function), do the following:
 - a) Move from the current position q_l in the direction X_{q_l} with step size Δh . The new position q_{l+1} is given by

$$q_{l+1} = q_l + \Delta h \cdot X_{q_l}. \quad [\text{A7.9}]$$
 - b) Identify the frame $\{Z_1, \dots, Z_m\}$ at location q_{l+1} .
 - c) Determine $[U_{q_{l+1}}, V_{q_{l+1}}, \dots]$ from $\{Z_1, \dots, Z_m\}$ with the algorithm described in Section 7.3.1.1, using $[U_{q_l}, V_{q_l}, \dots]$ as a reference frame. If $X_{q_l} = 0$, i.e., if X is not defined at the new location, the propagation is terminated.
- 2) Return $q_L = \hat{\Phi}_h^X(q_0)$ and the corresponding frame $[U_{q_L}, V_{q_L}, \dots]$.

7.7.3.3 Appendix C7.3 The closure and the Lie bracket

Algorithm for the reconstruction of difference vector \hat{R}_1 at point p

- 1) Identify the frame $[Y_1, \dots, Y_n]$ at p , and set $U_p = Y_1, V_p = Y_2, \dots$. In practice, determining $\{Y_1, \dots, Y_n\}$ involves some kind of interpolation (e.g., nearest neighbor, trilinear) and potentially the computation of peak directions from dMRI data if one does not start from a set of pre-extracted peaks (see Section 7.3.4.2.3 for details). Here, we stick to a general formulation and assume that the frame $[Y_1, \dots, Y_n]$ is given or can be calculated.
- 2) Select two vectors Y_i and Y_j from $\{U, V, \dots\}$ for the computation of the Lie bracket, which could be, for example, peaks along two predefined tracts. Here, we take the example $Y_i \in V$ and $Y_j \in W$.
- 3) Starting from $q_0 = p$, compute the end point of the first leg $\hat{\Phi}_{h_1}^V(p)$ and the corresponding frame using the algorithm described in Section 7.3.1.2. If the tracking algorithm fails, the reconstruction algorithm is terminated.
- 4) Repeat step 3 with $q_0 = \hat{\Phi}_{h_1}^V(p)$ to compute the end point of the second leg $\hat{\Phi}_{h_2}^W \circ \hat{\Phi}_{h_1}^V(p)$.
- 5) Repeat step 3 with $q_0 = (\hat{\Phi}_{h_2}^W \circ \hat{\Phi}_{h_1}^V)(p)$ to compute the end point of the third leg $(\hat{\Phi}_{h_1}^{-V} \circ \hat{\Phi}_{h_2}^W \circ \hat{\Phi}_{h_1}^V)(p)$.
- 6) Repeat step 3 with $q_0 = (\hat{\Phi}_{h_1}^{-V} \circ \hat{\Phi}_{h_2}^W \circ \hat{\Phi}_{h_1}^V)(p)$ to compute the end point of the final leg $(\hat{\Phi}_{h_2}^{-W} \circ \hat{\Phi}_{h_1}^{-V} \circ \hat{\Phi}_{h_2}^W \circ \hat{\Phi}_{h_1}^V)(p)$.
- 7) Return $\hat{R}_1 = (\hat{\Phi}_{-h_2}^W \circ \hat{\Phi}_{-h_1}^V \circ \hat{\Phi}_{h_2}^W \circ \hat{\Phi}_{h_1}^V)(p) - p$.

When computing \hat{R} for all h_1 and h_2 , (parts of) paths are taken multiple times. The positions q and the corresponding frames for each path are therefore stored in a lookup table.

7.8 SUPPLEMENTARY MATERIAL

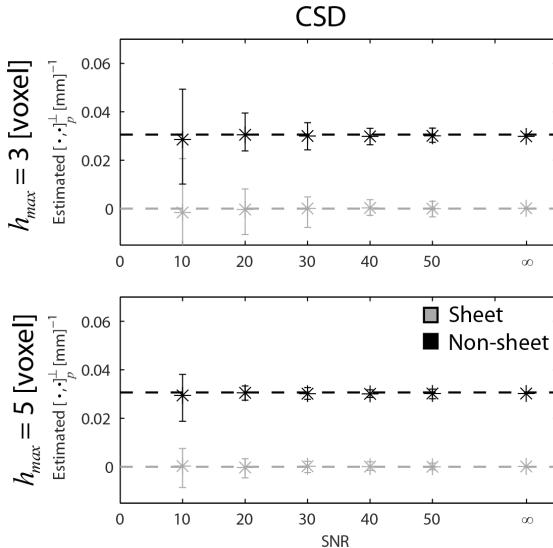


Fig. S7.1: Mean and range of $\widehat{[\cdot, \cdot]}_p^\perp$ for different settings of h_{\max} (rows) and different SNR. We set the curvature $\kappa = 1/\rho = 1/26 \text{ mm}^{-1}$, voxel size $\delta = 1 \text{ mm}$, and evaluate $\widehat{[\cdot, \cdot]}_p^\perp$ at $p = (10, -10, 0)$. Here we used trilinear interpolation of the fODF spherical harmonics coefficients and extracted the peaks during tractography. This is in contrast to Fig. 7.9, where peaks were pre-extracted and nearest neighbor interpolation was used. Dashed lines indicate the true Lie bracket normal component.

Vid. S7.2: Sheet tensor visualization on all (sagittal) slices of a data set of the HCP, showing large regions with consistent sheet structures. $h_{\max} = 5 \text{ voxels}$, and sheet tensors with $P_\lambda < 0.2$ ($\lambda = 0.008$) are not shown. Link to the video: https://www.dropbox.com/s/jwal6lstb0oy8pq/S7_2.wmv?dl=0.

Chapter 7: Sheet Probability Index (SPI): Characterizing the Geometrical Organization of the White Matter with Diffusion MRI

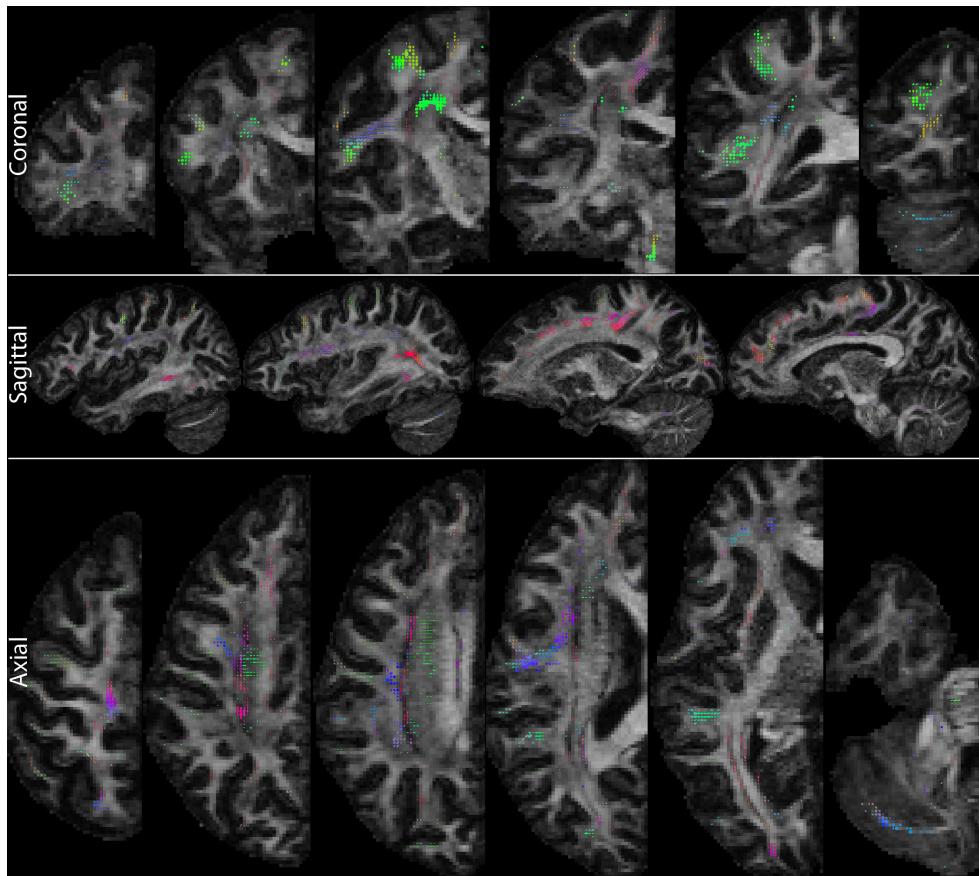


Fig. S7.3: Sheet tensors (satisfying $P_\lambda > 0.2$, $\lambda = 0.008$) on different coronal (top), sagittal (middle), and axial (bottom) slices, computed with $h_{max} = 3$ voxels.

Vid. S7.4: The path neighborhood of a sheet formed by the CC and the CST in the left hemisphere, in a data set of the HCP. Link to the video: https://www.dropbox.com/s/upro8fkowgxm4xx/S7_4.mp4?dl=0.

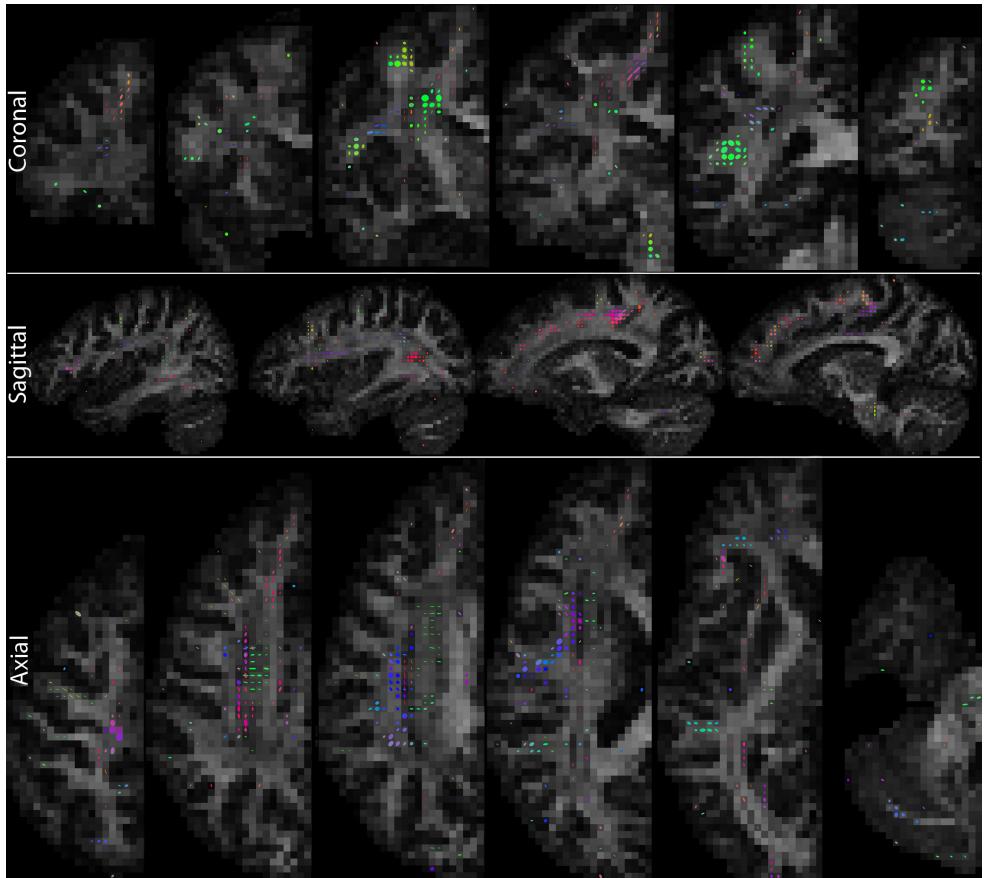


Fig. S7.5: Sheet tensors (satisfying $P_\lambda > 0.2$, $\lambda = 0.008$) on different coronal (top), sagittal (middle), and axial (bottom) slices, computed with $h_{max} = 2.5$ voxels on a sub-sampled data set.

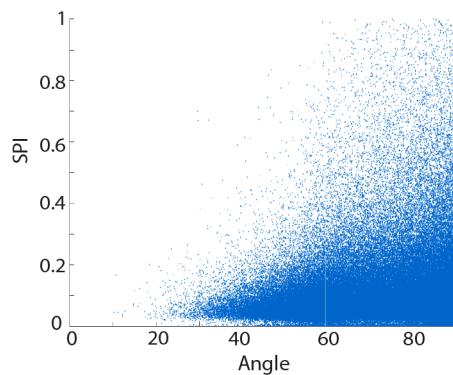


Fig. S7.6: SPI plotted against angle for an HCP dataset. High SPI can be encountered for angles ranging from 40-90°, which is in the range of the angular resolution of the used reconstruction technique. There are only slightly more voxels with high SPI for angles closer to 90°.

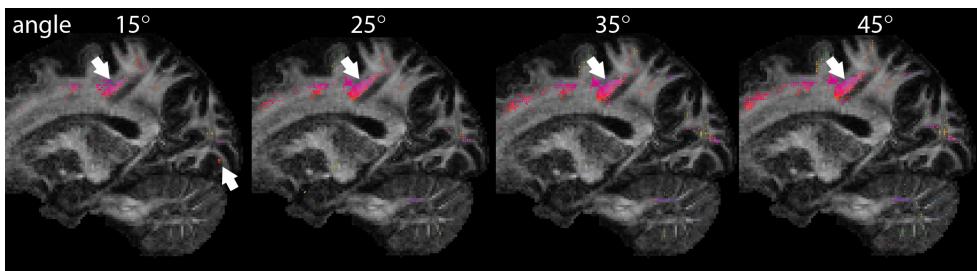


Fig. S7.7: SPI for an HCP dataset for different angle thresholds. A too low angle threshold (15°) might result in failure to detect curved sheets or spurious high SPI values because of early termination of paths (white arrows).

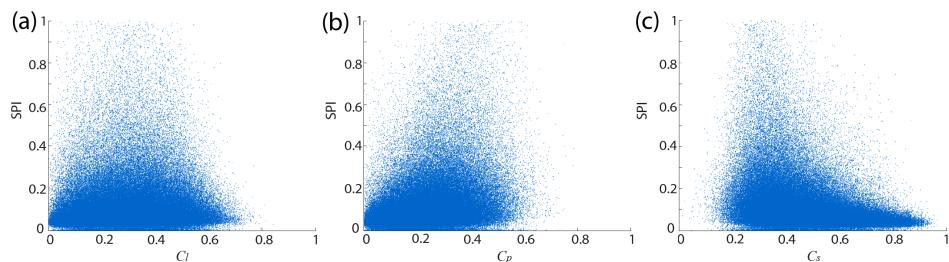


Fig. S7.8: Scatter plots of the SPI and three DTI measures linear coefficient C_l (a), planar coefficient C_p (b), and spherical coefficient C_s (c) for an HCP dataset.

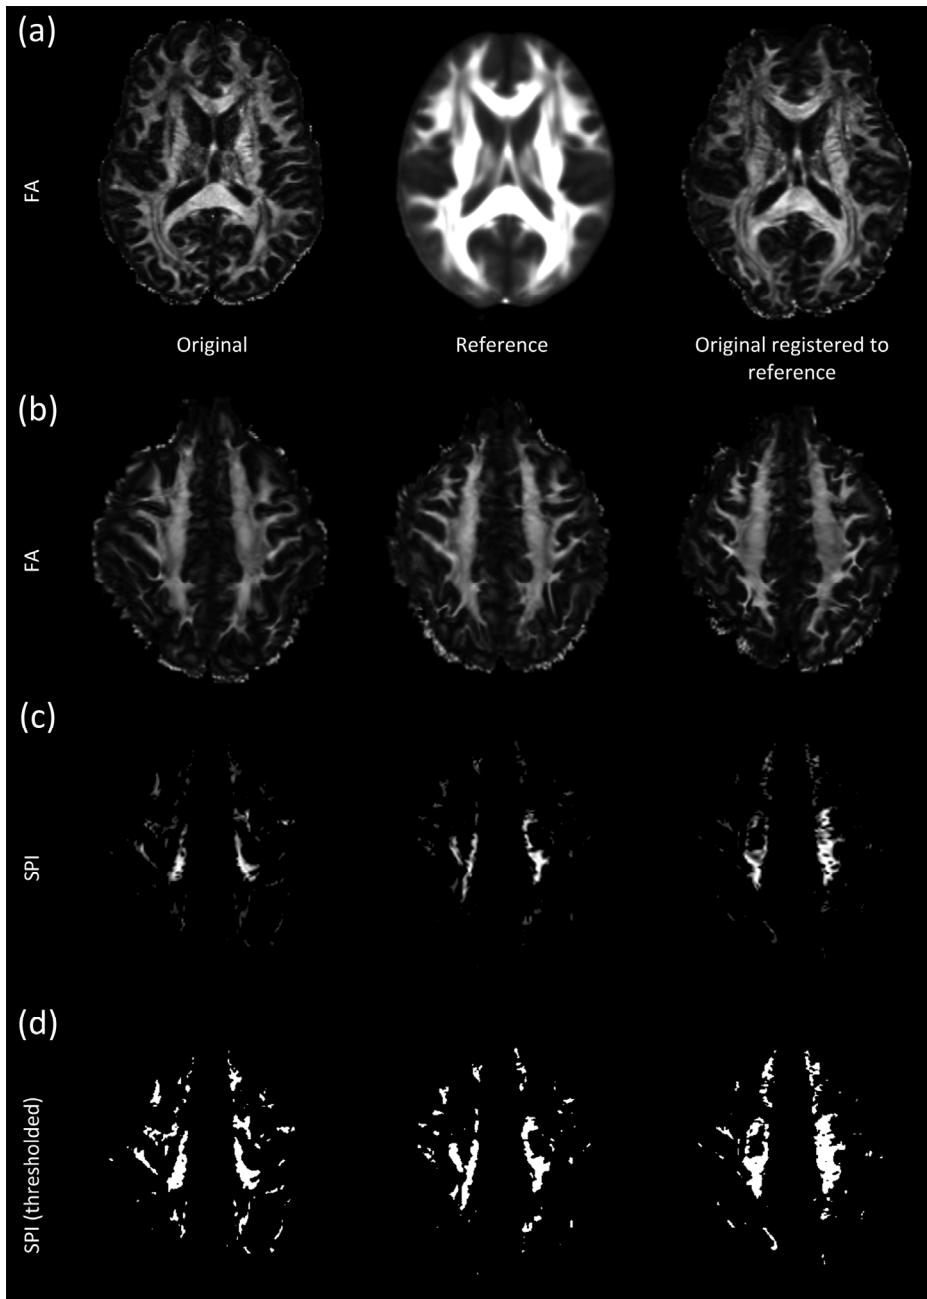


Fig. S9: Preliminary results of a quantitative comparison between three HCP subjects. (a) An FA map for one subject before (left) and after (right) registration to the FMRIB58 template (middle) using Elastix (Klein et al., 2010) with the Oxford-optimized settings described in De Groot et al. (2013). (b) Registered FA data sets used in the manuscript. (c) The voxel-wise maximum-SPI maps corresponding to (a). (d) The binarized maps ($P_\lambda > 0.2$) corresponding to (b) and (c) used in similarity computations. The mean Dice similarity between the three data sets is approximately 0.28.

8 QUANTIFYING THE BRAIN'S SHEET STRUCTURE WITH NORMALIZED CONVOLUTION

Based on:

C.M.W. Tax, C.F. Westin, T.C.J. Dela Haije, A. Fuster, M.A. Viergever, E. Calabrese, L.M.J. Florack, A. Leemans, "Quantifying the brain's sheet structure with normalized convolution" (submitted)

ABSTRACT

The hypothesis that brain pathways form 2D sheet-like structures layered in 3D as “pages of a book” has been a topic of debate in the recent literature. This hypothesis was mainly supported by a qualitative evaluation of “path neighborhoods” reconstructed with diffusion MRI (dMRI) tractography. Notwithstanding the potentially important implications of the sheet structure hypothesis for our understanding of brain structure and development, it is still considered controversial by many for lack of quantitative analysis. A means to quantify sheet structure is therefore necessary to reliably investigate its occurrence in the brain. Previous work has proposed the Lie bracket as a quantitative indicator of sheet structure, which could be computed by reconstructing path neighborhoods from the peak orientations of dMRI orientation density functions. Robust estimation of the Lie bracket, however, is challenging due to high noise levels and missing peak orientations. We propose a novel method to estimate the Lie bracket that does not involve the reconstruction of path neighborhoods with tractography. This method requires the computation of derivatives of the fiber peak orientations, for which we adopt an approach called normalized convolution. With simulations and experimental data we show that the new approach is more robust with respect to missing peaks and noise. We also demonstrate that the method is able to quantify to what extent sheet structure is supported for dMRI data of different species, acquired with different scanners, diffusion weightings, dMRI sampling schemes, and spatial resolutions. The proposed method can also be used with directional data derived from other techniques than dMRI, which will facilitate further validation of the existence of sheet structure.

8.1 INTRODUCTION

A recent debate on the existence of ‘sheet structures’ in the brain has gained much attention from the neuroscience and diffusion MRI (dMRI) communities (Catani et al., 2012; Wedeen et al., 2012a; Wedeen et al., 2012b). While the term sheets had already been suggested in several contexts before (Kindlmann et al., 2007; Schultz et al., 2010; Smith et al., 2006; Vilanova et al., 2004; Yushkevich et al., 2008; Zhang et al., 2003), Wedeen et al. (2012b) proposed a different and specific definition of brain sheet structure: a composition of two sets of tracts that locally cross each other on the same surface. Such crossing-fiber sheets should not be confused with white matter tracts that approximate thin single-fiber sheets by themselves (e.g. parts of the corpus callosum and the corona radiata). According to Wedeen et al. this geometric structure was, although “mathematically specific and highly atypical”, deemed to be “characteristic of the brain pathways”, “not unfathomable but unlimited”, and “consistent with embryogenesis” (Wedeen et al., 2012a; Wedeen et al., 2012b).

To support the sheet structure hypothesis, Wedeen et al. (2012b) reconstructed so-called path neighborhoods with dMRI tractography: pathways were tracked from a seed point, and paths incident on these pathways were subsequently computed. With this method, sheet structures were observed pervasively in white matter and found consistently across species and scales. In a response to this finding, Catani et al. (2012) suggested that the observed “grid pattern is most likely an artifact attributable to the limitations of their method”. These limitations included the lower angular resolution of the dMRI method used in Wedeen et al. (2012b), which has a negative impact on the tractography reconstructions of path neighborhoods. In addition, Catani et al. (2012) emphasized that dMRI reconstructions of pathways cannot be equated to true axons. Finally, they stated that comprehensive quantitative analysis is necessary to reliably assess the occurrence of sheet structure in the brain and in different individuals. This would allow the investigation of sheet structure extent, location, and orientation as novel quantitative features of brain structure.

Quantitative analysis of sheet structure requires a measure that is specific to the existence of this geometry. Previous work has shown that the Lie bracket, a mathematical concept from differential geometry, captures information on the existence of sheet structure (Tax et al., 2015a; Tax et al., 2014a; Wedeen et al., 2012b; Wedeen et al., 2014). Specifically, the Lie bracket can be computed from two vector fields, and is a vector field itself. If the two constituent vector fields are tangent to the orientations of two associated tracts, then the Lie bracket is an indicator of sheet structure between these tracts. In dMRI, such vector fields can be derived, for example, from the peak directions of diffusion or fiber orientation density functions (dODFs or fODFs).

Estimating a Lie bracket from dMRI peak directions is greatly challenged by the sensitivity of dMRI to noise and the inherent spatial discretization. In addition, dMRI peaks are not clustered into distinct vector fields *a priori*. In previous work we have proposed to compute the Lie bracket by reconstructing an extensive amount of path neighborhoods (Tax et al., 2015a; Tax et al., 2014a), inspired by the qualitative reconstructions in Wedeen et al. (2012b). More specifically, by taking small ‘loops’ along the tractography pathways of two vector fields, we could obtain an estimate of the Lie bracket. In this so-called flows-and-limits approach, dMRI peaks were clustered into distinct vector fields during tractography. By computing the Lie

bracket for repeated dMRI measurements or for different bootstraps of a single measurement, we were able to derive a sheet probability index (SPI) (Tax et al., 2016a). SPI values close to 1 indicate that the underlying vector fields are highly supportive of sheet structure.

The previously proposed flows-and-limits approach for estimating the Lie bracket suffers from several limitations. First, it relies on tractography and storage of a large amount of loops, and is thus computationally very expensive. Second, computing the Lie bracket with this method suffers from the same intrinsic limitations as tractography, notably its sensitivity to parameter settings (e.g., step size). Finally, because the algorithm has to keep track of the assignment of peaks to vector fields during tractography, it is unclear how to proceed when peaks are missing (e.g. as a result of noise or of modelling imperfections). The unwanted termination of loops in such cases may severely impact the accuracy and precision of the Lie bracket estimates.

In this chapter we propose to examine the existence of sheet structure based on a conceptually different but equivalent approach to defining the Lie bracket. Instead of using the flows-and-limits approach, we here use the fact that the vector fields are defined on subsets of the 3D Euclidean space. The Lie bracket can thus be expressed in terms of derivatives of the vector fields with respect to the three spatial variables (which we will henceforth denote by ‘the coordinate approach’). Taking derivatives of vector fields is, however, complicated: a straightforward finite difference implementation is far from stable for our purpose. Therefore, we adopt a method called normalized convolution (Knutsson and Westin, 1993) to calculate vector field derivatives and subsequently the Lie bracket. This approach was designed to cope with noise and uncertainty in data, and is thus more robust with respect to missing peaks.

This chapter is organized as follows. In Section 8.2 we will first briefly recapitulate the mathematical condition for the existence of sheet structure, involving the Lie bracket of vector fields (Tax et al., 2015a; Tax et al., 2014a; Wedeen et al., 2012b; Wedeen et al., 2014). In addition, we will present the flows-and-limits and coordinate definitions in more detail, and discuss their implementations. For the coordinate approach, this involves the implementation of normalized convolution and the clustering of vector fields in a spatial neighborhood. Finally, we will present the experiments on vector field simulations, dMRI simulations, and real data. In Sections 8.3 and 8.4, we will describe and discuss the corresponding results. Preliminary results of this work have been presented at the ISMRM (Tax et al., 2016b).

8.2 METHODS

Section 8.2.1 explains which condition two vector fields have to fulfill in order to form sheet structures. Section 8.2.2 describes the two definitions of the Lie bracket and their corresponding algorithms (i.e., the flows-and-limits and coordinate approaches). The computation of vector field derivatives with normalized convolution is outlined in Section 8.2.2.1 and the approach for clustering of vector fields in a spatial neighborhood is presented in Section 8.2.2.2.

8.2.1 The Lie Bracket as an indicator of sheet structure

Consider the set of directions (or vectors) at each position of the brain $M \subset \mathbb{R}^3$, which we assume to be representative of the underlying local fiber architecture. Tractography processes

then assume that in a specific part of the brain $N_V \subset M$ there exists a smooth 3D vector field V that can be integrated into integral curves (or streamlines) $\Phi^V(s, p)$. Here, p denotes the starting position, s the arc length, and $V(p) = \frac{d}{ds} \Phi^V(s, p) \Big|_{s=0}$ the vector at location p . Similarly, two smooth vector fields V and W in a brain region $N_V \cap N_W$ can potentially be integrated into integral surfaces or sheet structures.

The Lie bracket of two vector fields V and W defines a third vector field $[V, W]$ that holds information on the presence of sheet structure. More specifically, to investigate whether two vector fields locally form a sheet structure we compute the normal component of the Lie bracket at a point p which is then perpendicular to both $V(p)$ and $W(p)$ (Fig 8.1a):

$$[V, W]^\perp(p) := [V, W](p) \cdot \frac{V(p) \times W(p)}{\|V(p) \times W(p)\|}. \quad [8.1]$$

According to the Frobenius theorem (Lang, 1995; Spivak, 1979), the condition for sheet structure is then fulfilled at all points $p \in N_V \cap N_W$ where $[V, W]^\perp(p) = 0$ (Tax et al., 2015a; Tax et al., 2014a; Wedeen et al., 2012b). In other words, the Lie bracket $[V, W](p)$ must lie in the plane spanned by $V(p)$ and $W(p)$. Note that we will use shorthand notations $[\cdot, \cdot]_p = [\cdot, \cdot](p)$ and $[\cdot, \cdot]_p^\perp = [\cdot, \cdot]^\perp(p)$.

8.2.2 Computation of the Lie bracket

In this Section, we explain the two definitions of the Lie bracket and their corresponding implementations.

8.2.2.1 Flows-and-limits definition

In previous work, we have used the definition of the Lie bracket as the deviation from p , or closure $R(p) = (r^1, r^2, r^3)(p)$ (with r^i the components) when trying to move around in an infinitesimal loop along the integral curves of V and W (Fig. 8.1b) (Tax et al., 2015a; Tax et al., 2014a; Wedeen et al., 2012b; Wedeen et al., 2014). With streamline tractography, loop configurations can be reconstructed to estimate their corresponding closures:

- $\hat{R}_1(p) := (\hat{\Phi}_{h_2}^{-W} \circ \hat{\Phi}_{h_1}^{-V} \circ \hat{\Phi}_{h_2}^W \circ \hat{\Phi}_{h_1}^V)(p) - p$
- $\hat{R}_2(p) := p - (\hat{\Phi}_{h_1}^{-V} \circ \hat{\Phi}_{h_2}^{-W} \circ \hat{\Phi}_{h_1}^V \circ \hat{\Phi}_{h_2}^W)(p)$
- $\hat{R}_3(p) := (\hat{\Phi}_{h_2}^W \circ \hat{\Phi}_{h_1}^V)(p) - (\hat{\Phi}_{h_1}^V \circ \hat{\Phi}_{h_2}^W)(p).$

[8.2]

Here, Φ_h^V denotes the flow along V for some fixed time or distance h (with arc length parameterization), $\hat{\Phi}_h^V$ the estimate of the flow as approximated by tractography, \circ indicates function composition, and h_1 and h_2 are the walking distances along the integral curves of V and W , respectively.

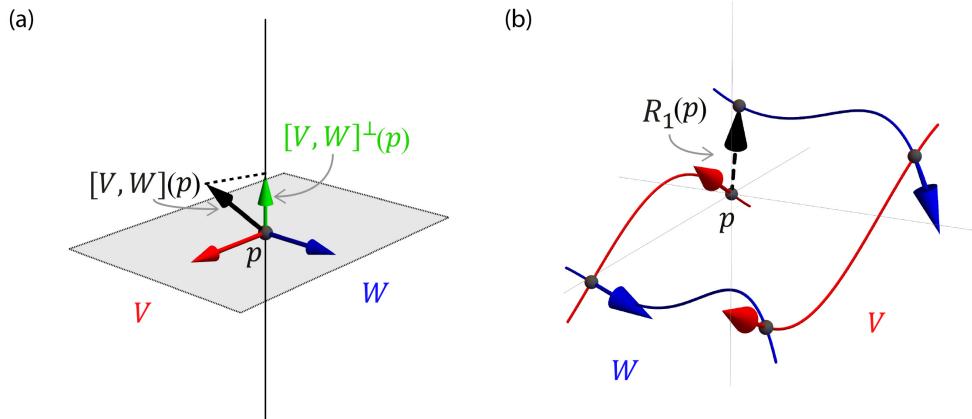


Fig. 8.1: (a) To investigate the existence of sheet structure, we consider the relation of the Lie bracket $[V, W](p)$ (black arrow) to the plane spanned by vectors $V(p)$ (red arrow) and $W(p)$ (blue arrow). More specifically, we evaluate the normal component of the Lie bracket $[V, W]^\perp(p)$ (green arrow). When it is zero, sheet structure exists. (b) Definition of the Lie bracket as the closure $R(p)$ when trying to move around in an infinitesimal loop along the integral curves of V and W . Here, the loop corresponding to the closure R_1 is displayed (Eq. [8.2]).

We choose $h_1, h_2 \in \{-h_{max}, -h_{max} + \Delta h, -h_{max} + 2\Delta h, \dots, -\Delta h, \Delta h, \dots, h_{max}\}$ with h_{max} the maximum walking distance and Δh a step size, resulting in K difference vectors (1200 for $h_{max} = 5$ voxel size and $\Delta h = 0.5$ voxel size, where we hereafter shorten voxel size to voxels).

If h_1 and h_2 are sufficiently small, there exists an approximately linear relationship between the product $h_1 h_2$ and a difference vector component r^i ($i \in [1, 2, 3]$), with the corresponding component of the Lie bracket $[V, W]^i$ being the steepness of the line:

$$r^i \approx h_1 h_2 [V, W]^i. \quad [8.3]$$

We can compute an estimate of the Lie bracket $\widehat{[V, W]}_p$ with a simple linear least squares fit

$$\mathbf{R} = \mathbf{H}\boldsymbol{\beta} + \boldsymbol{\epsilon}. \quad [8.4]$$

Here, \mathbf{R} is a $K \times 3$ matrix with the estimated closures $\widehat{R}(p)$, \mathbf{H} is a $K \times 1$ matrix with the products of the used h_1 and h_2 for the corresponding difference vector, $\boldsymbol{\beta}$ is a 1×3 matrix with the Lie bracket components $[V, W]^i_p$, and $\boldsymbol{\epsilon}$ is a $K \times 3$ matrix with normally distributed errors.

8.2.2.2 Coordinate definition

The Lie bracket can also be computed with respect to a coordinate system. In our case, V and W are vector fields on a subset of the 3D Euclidean space (i.e., $M \subset \mathbb{R}^3$) and can therefore each be represented by a vector-valued function in standard Cartesian coordinates (x^1, x^2, x^3) . The Lie bracket is then given by

$$[V, W] = J_W \cdot V - J_V \cdot W, \quad [8.5]$$

where J_V is the Jacobian matrix of the vector-valued function $V = (v^1, v^2, v^3)^T$:

$$J_V = \begin{pmatrix} \frac{\partial v^1}{\partial x^1} & \frac{\partial v^1}{\partial x^2} & \frac{\partial v^1}{\partial x^3} \\ \frac{\partial v^2}{\partial x^1} & \frac{\partial v^2}{\partial x^2} & \frac{\partial v^2}{\partial x^3} \\ \frac{\partial v^3}{\partial x^1} & \frac{\partial v^3}{\partial x^2} & \frac{\partial v^3}{\partial x^3} \end{pmatrix}. \quad [8.6]$$

$J_W \cdot V$ can be interpreted as the directional derivative or ‘rate of change’ of W in the direction of V . The Lie bracket indicates the difference of the rate of change of W in the direction of V and the rate of change of V in the direction of W , hence its definition as the commutator of vector fields.

To compute the Lie bracket we thus need an approximation of the Jacobians J_V and J_W of the vector fields. Differentiation of discretized and noisy vector fields is, however, a notoriously unreliable operation. A straightforward finite differences implementation amplifies high frequency noise, and does not provide a stable solution in our case. More importantly, for our application it is not guaranteed that a vector field has a vector at every location within a neighborhood of the point of operation, because of noise or modelling imperfections. Interpreting such a missing vector as a vector with zero magnitude will introduce severe errors. We therefore adopt the normalized convolution approach (Knutsson and Westin, 1993) which is able to cope with noisy, discretized, and missing data and automatically takes care of normalization during convolution. In the next subsections, we will explain this approach in detail and outline the implementation for our application.

8.2.2.2.1 Normalized convolution

The problem of ill-posed differentiation can be addressed by convolution (Florack, 2013; Koenderink and Van Doorn, 1990). General convolution can be defined as

$$\{B \tilde{\otimes} V\}(p) := \sum_{\xi} B(\xi) \otimes V(p - \xi), \quad [8.7]$$

where B denotes a real linear filter basis, \otimes is the tensor product, a tilde \sim above the operator (e.g. $\tilde{\otimes}$) denotes convolution with that operator, and $\xi = x - p$ is the local spatial coordinate.

8.2.2.2.1.1 Normalization and connection to linear least squares

One problem with directly applying Eq. [8.7] is that the result is scaled when the basis is not orthonormal. For example, applying a Gaussian derivative basis function, the value resulting from the operation will depend on the standard deviation of the Gaussian unless the basis function is normalized appropriately. In our application, the actual value of the normal component of the Lie bracket is important, and we thus have to take into account a normalization factor. A normalized version of Eq. [8.7] is

$$\{B \tilde{\otimes} V\}_N = \{B \tilde{\otimes} B\}^{-1} \{B \tilde{\otimes} V\}, \quad [8.8]$$

where $^{-1}$ denotes the inverse.

Intuitively, convolution obtains a new description of the original representation of a neighborhood in terms of a new set of basis functions. Here, the original representation is a vector for each voxel accompanied by a basis of impulse functions at each voxel location. Ideally, the new set of basis functions is chosen in such a way that the neighborhood can be better understood. In fact, Eq. [8.8] obtains such a representation that is optimal in the least square sense: it produces the linear least squares (LLS) estimates for the coefficients corresponding to the new basis B . To clarify this, we will explicitly write out Eq. [8.8] for our application.

To estimate the first order derivatives in the Jacobian of a vector field V we choose

$$B(\xi) = \begin{pmatrix} 1 \\ \xi^1 \\ \xi^2 \\ \xi^3 \end{pmatrix}, \text{ which gives } B \otimes B = \begin{pmatrix} 1 & \xi^1 & \xi^2 & \xi^3 \\ \xi^1 & \xi^1\xi^1 & \xi^1\xi^2 & \xi^1\xi^3 \\ \xi^2 & \xi^1\xi^2 & \xi^2\xi^2 & \xi^2\xi^3 \\ \xi^3 & \xi^1\xi^3 & \xi^2\xi^3 & \xi^3\xi^3 \end{pmatrix}. \quad [8.9]$$

This holds a close connection to calculating derivatives of a ‘smoothed’ signal by fitting low-degree polynomials to every component v^i of V in a neighborhood of p (Savitzky and Golay, 1964). In this case, each polynomial includes only terms up to first order, i.e. $Y^i(\xi) = \alpha_1^i + \alpha_2^i\xi^1 + \alpha_3^i\xi^2 + \alpha_4^i\xi^3$. We now write this problem as a regression equation

$$V = B\alpha + \epsilon, \quad [8.10]$$

where, V is an $N^3 \times 3$ matrix with the vectors V in a neighborhood of p with size $N \times N \times N$, B is a $N^3 \times 4$ matrix with the basis B evaluated in each point of the neighborhood of p , α is a 4×3 matrix with the coefficients for every polynomial Y^i , and ϵ is an $N^3 \times 3$ matrix with normally distributed errors. The least squares solution for the coefficients α can then be obtained by solving the normal equations, yielding

$$\alpha = (B^T B)^{-1} B^T V, \quad [8.11]$$

which corresponds to Eq. [8.8]. We can now see that the result of convolution with the basis in Eq. [8.9] is a regularized vector and an estimate of the Jacobian at the point p ($\xi = 0$), calculated from the fitted polynomials ($\hat{v}^i(p) = Y^i(0)$, $\frac{\partial \hat{v}^i}{\partial x^j}\Big|_p = \frac{\partial Y^i}{\partial \xi^j}\Big|_{\xi=0}$):

$$\{B \tilde{\otimes} V\}_N^T = \alpha^T = \begin{pmatrix} \alpha_1^1 & \alpha_2^1 & \alpha_3^1 & \alpha_4^1 \\ \alpha_1^2 & \alpha_2^2 & \alpha_3^2 & \alpha_4^2 \\ \alpha_1^3 & \alpha_2^3 & \alpha_3^3 & \alpha_4^3 \end{pmatrix} = \begin{pmatrix} \hat{v}^1 & \frac{\partial \hat{v}^1}{\partial x^1} & \frac{\partial \hat{v}^1}{\partial x^2} & \frac{\partial \hat{v}^1}{\partial x^3} \\ \hat{v}^2 & \frac{\partial \hat{v}^2}{\partial x^1} & \frac{\partial \hat{v}^2}{\partial x^2} & \frac{\partial \hat{v}^2}{\partial x^3} \\ \hat{v}^3 & \frac{\partial \hat{v}^3}{\partial x^1} & \frac{\partial \hat{v}^3}{\partial x^2} & \frac{\partial \hat{v}^3}{\partial x^3} \end{pmatrix} = (\hat{V} \quad \hat{J}_V). \quad [8.12]$$

8.2.2.2.1.2 Normalized convolution and connection to weighted linear least squares

Knutsson and Westin (1993) considered a more general case of Eq. [8.8], and proposed to assign a scalar component to both the data and the operator that represents their certainty. More specifically, the vector field V is accompanied by a non-negative scalar function c representing its certainty, and the operator filter basis B is accompanied by a non-negative scalar function a representing its applicability (the operator equivalent to certainty). If we view a and c as additional independent variables, a generalized version of normalized convolution can be expressed (by abuse of notation) as

$$\{aB \tilde{\otimes} cV\}_N = \{aB \otimes B \cdot c\}^{-1} \{aB \tilde{\otimes} cV\}, \quad [8.13]$$

Here \cdot denotes the scalar product and $\tilde{\cdot}$ convolution with the scalar product (Knutsson and Westin, 1993). Explicitly writing out the terms for our application gives:

$$\{aB \tilde{\otimes} cV\}_N$$

$$= \begin{pmatrix} \{a \cdot c\} & \{a\xi^1 \cdot c\} & \{a\xi^2 \cdot c\} & \{a\xi^3 \cdot c\} \\ \{a\xi^1 \cdot c\} & \{a\xi^1 \xi^1 \cdot c\} & \{a\xi^1 \xi^2 \cdot c\} & \{a\xi^1 \xi^3 \cdot c\} \\ \{a\xi^2 \cdot c\} & \{a\xi^1 \xi^2 \cdot c\} & \{a\xi^2 \xi^2 \cdot c\} & \{a\xi^2 \xi^3 \cdot c\} \\ \{a\xi^3 \cdot c\} & \{a\xi^1 \xi^3 \cdot c\} & \{a\xi^2 \xi^3 \cdot c\} & \{a\xi^3 \xi^3 \cdot c\} \end{pmatrix}^{-1} \begin{pmatrix} \{a \cdot cv^1\} & \{a \cdot cv^2\} & \{a \cdot cv^3\} \\ \{a\xi^1 \cdot cv^1\} & \{a\xi^1 \cdot cv^2\} & \{a\xi^1 \cdot cv^3\} \\ \{a\xi^2 \cdot cv^1\} & \{a\xi^2 \cdot cv^2\} & \{a\xi^2 \cdot cv^3\} \\ \{a\xi^3 \cdot cv^1\} & \{a\xi^3 \cdot cv^2\} & \{a\xi^3 \cdot cv^3\} \end{pmatrix}.$$

For the certainty and applicability functions we choose

$$c(q) = \begin{cases} 1 & \text{if } V(q) \neq (0,0,0)^T \\ 0 & \text{if } V(q) = (0,0,0)^T \end{cases}, \quad [8.14]$$

$$a(r) = \begin{cases} \cos^\beta \left(\frac{\pi r}{2r_{max}} \right) & \text{if } r < r_{max} \\ 0 & \text{otherwise} \end{cases}. \quad [8.15]$$

Here, $q \in \mathbb{R}^3$, and the certainty c is thus set to zero when a vector is missing. Furthermore, r is the distance to the neighborhood center, and $a(r)$ is shown in Fig. 8.2 for different β and r_{max} , which determine the ‘weight’ of data points in the neighborhood. Here, β determines the power or the ‘sharpness’ of the function. In practice, we set $r_{max} = 0.5N\delta$, with N and δ the kernel and voxel sizes, respectively.

Normalized convolution obtains a new description of the original representation, but now the original basis is accompanied by a strength or certainty c . A missing vector can then be interpreted as a missing basis impulse function (i.e. by setting the certainty to zero), rather than a vector with zero magnitude. Normalized convolution obtains a description in the new basis such that the local weighted mean square error is minimal, corresponding to a weighted linear least squares approach (WLLS):

$$\boldsymbol{\alpha} = (\mathbf{B}^T \mathbf{W} \mathbf{B})^{-1} \mathbf{B}^T \mathbf{W} \mathbf{V}. \quad [8.16]$$

The weights in the diagonal matrix \mathbf{W} are then given by

$$W_{ii} = a(\xi_i)c(p - \xi_i), \quad [8.17]$$

where the subscript i refers to the i^{th} local spatial coordinate in the $N \times N \times N$ neighborhood. To summarize, the result of normalized convolution (Eq. [8.12]) in an $N \times N \times N$ spatial neighborhood of a point p is a WLLS estimate of the vector and Jacobian at point p . These estimates can be directly substituted into Eq. [8.5] to obtain an estimate of the Lie bracket.

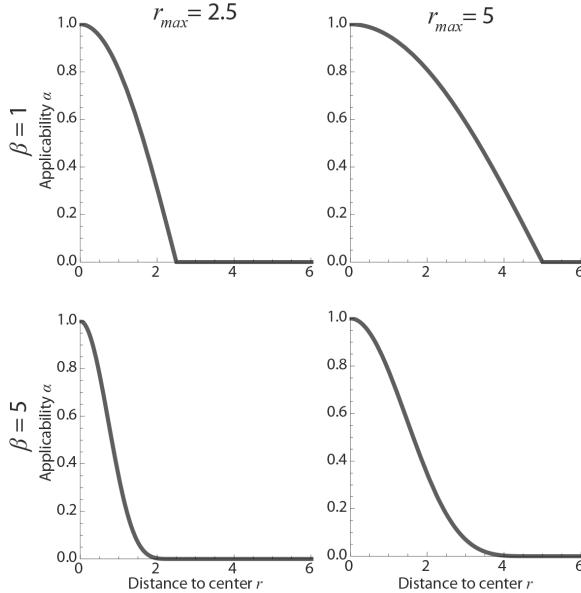


Fig. 8.2: Applicability function $a(r)$ for different r_{\max} and β (Eq. [8.15]).

8.2.2.2 Clustering of peaks in a spatial neighborhood

In the computation of the Lie bracket we have thus far assumed that two vector fields exist in a neighborhood of the point of operation that can be derived from dMRI data. However, dMRI data typically allow the extraction of a variable number of peaks per voxel without a notion to which local vector field (or fiber population) they belong. We will call the set of peaks at a spatial location a *frame*, and the frames in a spatial neighborhood thus have to be clustered. We match each vector of a frame at a point in the spatial neighborhood to the frame of the center voxel (the *reference frame*). In this work, we cluster the frames in an $N \times N \times N$ spatial neighborhood using a front-propagation approach that starts from the center of the neighborhood (Jbabdi et al., 2010). As the similarity measure we use the total cosine similarity, i.e. we find the configuration for which the total angle difference with a previously clustered frame at a nearby location is minimal. We set an angle threshold to constrain the maximum angle between corresponding vectors of frames. The clustering algorithm is detailed in Appendix A8.

8.2.3 Evaluation of sheet structure

Once the Lie bracket of a pair of vector fields has been computed with the flows-and-limits method or the coordinate method, we can examine the existence of sheet structure by extracting the normal component of the estimated Lie bracket (Eq. [8.1]) and determine whether it is ‘sufficiently close’ to zero. This condition is hard to evaluate, however, since a single noisy estimate of the Lie bracket normal component does not provide information on its variability. If we assume repeated MRI measurements that yield multiple Lie bracket estimates, a *sheet probability index* (SPI) can be calculated (Tax et al., 2016a). This index was introduced to address the issue of noise and to provide an intuitive sheet indicator. In short, the estimates are checked for normality using a Shapiro-Wilks test and the mean μ and standard deviation σ are estimated. Subsequently, the integral probability P_λ inside the region $[-\lambda, \lambda]$ can be calculated for the estimated distribution $N(\mu, \sigma^2)$. We can tune the parameter λ where a higher value for λ will allow larger deviations from zero. An SPI value close to 1 corresponds to a high likelihood of sheet structure, whereas an SPI close to 0 indicates that there are significant deviations from sheet structure. In practice, if only one dataset is acquired, we generate residual bootstraps to calculate the SPI (see Section 8.2.4.2.3).

The SPI can be visualized for pairs of vector fields throughout the brain by means of a *sheet tensor*. Given a pair of vector fields V and W , the sheet tensor at location p is defined as

$$T(p) = \frac{P_\lambda}{\beta_1} (V(p) \otimes V(p) + W(p) \otimes W(p)). \quad [8.18]$$

Here β_1 denotes the largest eigenvalue of the tensor $(V(p) \otimes V(p) + W(p) \otimes W(p))$. The sheet tensor can then be represented by an ellipsoid, of which the third eigenvector is normal to the span of V and W . The ellipsoid is colored according to the third eigenvector by using the well-known diffusion tensor imaging (DTI) coloring scheme of Zhang et al. (2006b). The SPI divided by β_1 determines the size of the ellipsoid (a larger SPI results in a tensor with a larger semi-axis) and the angle between V and W determines its shape (a lower angle results in a more anisotropic sheet tensor).

8.2.4 Data

We evaluated the two proposed methods on different types of data: analytical vector fields, dMRI simulations, and real dMRI data.

8.2.4.1 Analytical vector field simulations

We defined three vector fields that are tangent to a sphere with radius ρ (U and V are tangent to the upper hemisphere, W is tangent to the lower hemisphere, see Appendix B8). Vector fields U and V form sheets (i.e. $[U, V]_p^\perp = 0$), whereas U and W do not form a sheet except at $x^1 = 0, x^2 = 0$. The normal component of the Lie bracket $[U, W]_p^\perp$ is dependent on the radius ρ and the spatial coordinates. These simulations allowed us to evaluate the methods for Lie bracket calculation as a function of curvature $\kappa = 1/\rho$ and non-commutativity of the pair U and W . The vector fields were discretized on a Cartesian grid with voxel size δ and noise was added by drawing random samples of a Watson distribution (Chen et al., 2015b; Mardia and Jupp, 2009)

with concentration parameter $k > 0$ ('SNR level', higher k results in smaller perturbations). We generated 50 noise iterations in the experiments.

8.2.4.2 Diffusion MRI data

8.2.4.2.1 Simulations

dMRI signals were simulated using a ZeppelinStickDot model (Ferizi et al., 2014) with the fiber direction defined by the noise-free vector fields described in the previous section. We simulate single-shell datasets with 90 directions and $b = 3000 \text{ s/mm}^2$. Noise was added according to a Rician distribution.

8.2.4.2.2 Real data

We used the following dMRI data sets: 1) the $b = 3000 \text{ s/mm}^2$ shell with 90 diffusion directions of a subject of the WU-Minn Human Connectome Project (HCP) with an isotropic voxel size of 1.25 mm (Sotiroopoulos et al., 2013; Van Essen et al., 2013); and 2) two ex-vivo monkey brain datasets of the same brain with $b = 4000 \text{ s/mm}^2$, one dataset had an isotropic voxel size of 0.3 mm and 60 diffusion directions, while the other dataset had an isotropic voxel size of 0.4 mm and 120 diffusion directions (Calabrese et al., 2015b). The SNR of both datasets was measured in the cortex on the $b = 0$ image, and was 131 for the 0.4 mm and 96 for the 0.3 mm datasets, respectively.

8.2.4.2.3 Processing

The monkey brain datasets were corrected for motion and eddy current distortions using *ExploreDTI* (Leemans et al., 2009; Leemans and Jones, 2009). All datasets were processed using constrained spherical deconvolution (CSD, $l_{max} = 8$) (Tournier et al., 2007). The response function for the simulated data was generated from the ZeppelinStickDot model, and the response function for real data was computed using recursive calibration (Tax et al., 2014c). Peaks were extracted using a Newton optimization algorithm (Jeurissen et al., 2011) with a maximum number of 3 peaks and an fODF peak threshold of 0.1. To compute the SPI, we generated 50 noise iterations for simulated data, and 20 residual bootstrap realizations for real dMRI data from a single set of noisy measurements (Jeurissen et al., 2011). The peaks extracted from the different bootstrap realizations were clustered using the method described in Section 8.2.2.2.2, with the peaks extracted from the original data representing the reference frames. The SPI was calculated within a white matter mask dilated by one voxel. For the HCP data the mask was derived from the T1-weighted image using FSL-Fast (Zhang et al., 2001), and for the monkey brain datasets it was derived from the FA and MD images since no T1-weighted image was available ($FA > 0.4$ or $FA > 0.15$ with $MD < 2.5 \cdot 10^{-4} \text{ mm}^2/\text{s}$ adapted for ex-vivo data (Sherbondy et al., 2008)).

8.3 RESULTS

Section 8.3.1 presents the results for analytical vector field simulations, Section 8.3.2 for dMRI simulations, and Section 8.3.3 for real data.

8.3.1 Analytical vector field simulations

With the analytical vector fields we will first investigate the influence of the parameters in the coordinate implementation (kernel size N and β) on the accuracy and precision of the Lie bracket normal component estimates. Subsequently, we will compare the coordinate and flows-and-limits implementations for varying voxel sizes, SNR levels, and peak dropout fractions. These experiments all use the predefined clustering of vector fields, i.e. the vector fields as they were simulated. We will also investigate the performance of the clustering method, which is described in Section 8.2.2.2 and Appendix A8.

8.3.1.1 Coordinate implementation: Influence of spatial resolution and parameter settings

Fig. 8.3 shows the mean absolute error (upper row) and the range (maximum value minus minimum value, bottom row) of the Lie bracket normal component estimated values for the non-sheet vector field pair U and W (see Section 8.2.4.1 and Appendix B8). The columns a-c show the results for different voxel sizes $\delta = \{0.5, 1, 2\} \text{ mm}$. Each color plot shows the results for different settings of kernel size N and β . The following settings were used: curvature $\kappa = 1/\rho = 1/26 \text{ mm}^{-1}$, ‘SNR level’ $k = 250$, and $p = (10, -10, 0)$ yielding a non-zero Lie bracket normal component $[U, W]_p^\perp = 0.031 \text{ mm}^{-1}$.

The mean absolute error and range of the estimates generally decrease for 1) increasing voxel size at a given parameter setting, 2) increasing kernel size N , and 3) decreasing β . This means that incorporating neighborhood information over a larger scale in mm is generally beneficial for the accuracy and precision of the Lie bracket normal component estimate. There is a limit to this, however: at $\delta = 2, N = 11, \beta = 1$ the mean absolute error of the estimates starts to increase again. In the remainder of the simulation experiments, where we simulate $\delta = 1$, we will use $\beta = 1$.

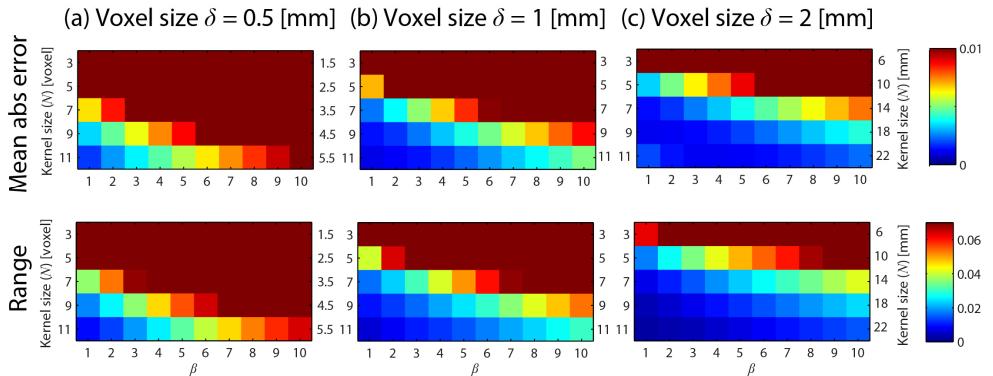


Fig. 8.3: Mean absolute error (first row) and range (second row) of estimates $[\widehat{U}, \widehat{W}]_p^\perp$ for different voxel sizes $\delta = \{0.5, 1, 2\} \text{ mm}$ (a-c, columns). Each color plot shows the results for different choices for the kernel size N (indicated in *voxel* on the left side of each graph and in *mm* on the right side of each graph) and different settings for β . We used $\kappa = 1/\rho = 1/26 \text{ mm}^{-1}$, $k = 250$, and $p = (10, -10, 0)$.

8.3.1.2 Coordinate vs. flows-and-limits implementation: Influence of spatial resolution and noise

Fig. 8.4 shows results for different voxel sizes $\delta = \{0.5, 1, 2\} \text{ mm}$ (a-c), for the coordinate implementation (top row, $N = 11$ voxels and $\beta = 1$) and the flows-and-limits implementation (bottom row, $h_{max} = 5 \text{ voxel}$ and $\Delta h = 0.5 \text{ voxel}$). Each plot shows the range and mean for different SNR levels k for the sheet pair (U and V , grey) and the non-sheet pair (U and W , black) defined in Section 8.2.4.1 and Appendix B. The ground truth value is indicated by the dashed line. We set the curvature $\kappa = 1/\rho = 1/26 \text{ mm}^{-1}$ and evaluate at point $p = (10, -10, 0)$.

The range of the estimates becomes smaller with 1) increasing SNR level k and 2) increasing voxel size for both implementations. In all cases, the coordinate implementation has a higher precision (lower variability) than the flows-and-limits implementation. This indicates that it potentially has a higher power to distinguish sheet configurations from non-sheet configurations. At a voxel size of $\delta = 2 \text{ mm}$, the estimates start to be slightly biased (i.e. start deviating from their true value) for both implementations and the used parameter settings.

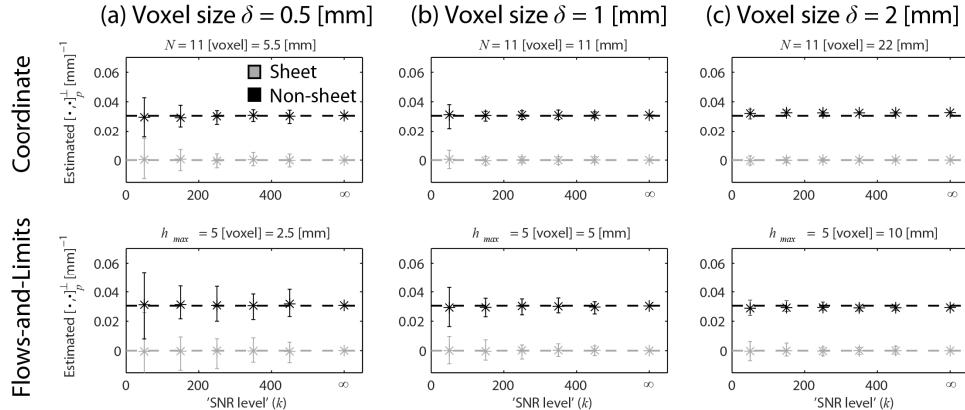


Fig. 8.4: Mean (asterisks) and range (error bars) of $\widehat{[\cdot, \cdot]}_p^\perp$ for different voxel sizes $\delta = \{0.5, 1, 2\} \text{ mm}$ (a-c, columns) and the two different implementations (rows). We used $N = 11 \text{ voxel}$ and $\beta = 1$, and $h_{max} = 5 \text{ voxel}$ and $\Delta h = 0.5 \text{ voxel}$ (the corresponding values in mm are given above each plot). Each plot shows the estimates in the case of sheet (grey, $[U, V]_p^\perp = 0$ indicated by the dashed line) and non-sheet (black, $[U, W]_p^\perp = 0.031$) for different SNR levels k . We set $\kappa = 1/\rho = 1/26 \text{ mm}^{-1}$, and $p = (10, -10, 0)$.

8.3.1.3 Coordinate vs. flows-and-limits implementation: Influence of dropouts and parameter settings

Fig. 8.5 shows the mean and range of the estimates $\widehat{[\cdot, \cdot]}_p^\perp$ as a function of the dropout fraction: For every noise iteration we randomly remove a certain fraction of peaks of each vector field. The rows show the results for the coordinate implementation (top, $\beta = 1$) and the flows-and-limits implementation (bottom) for different parameter settings of each method (columns). We used $\kappa = 1/\rho = 1/26 \text{ mm}^{-1}$, $p = (10, -10, 0)$, $\delta = 1 \text{ mm}$, and $k = 250$.

It becomes immediately apparent that the flows-and-limits implementation is much more sensitive to dropouts than the coordinate implementation: the range of the estimates is much larger for the former method. While for a limited use of spatial neighborhood information

($N = 3$ or $h_{max} = 1$ voxel) neither method is able to distinguish sheets from non-sheets, the coordinate implementation performs markedly better at higher N even for high dropout fractions. In fact, in this case the dropout fraction does not affect the precision and mean of the estimates (up to the tested fraction). For the flows-and-limits implementation the precision of the estimates decreases with dropout fraction, while the mean remains close to the true value.

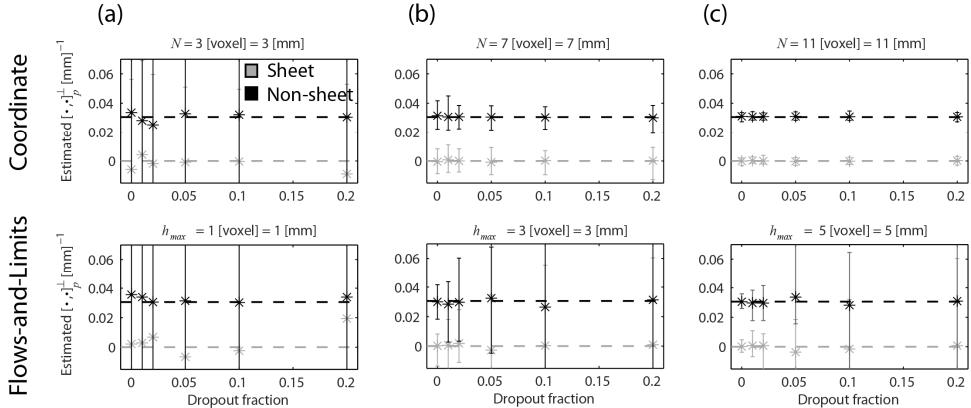


Fig. 8.5: Mean and range of $||\cdot||_p^1$ for different ‘spatial neighborhood settings’ (a-c, columns): $N = \{3,7,11\}$ voxel with $\beta = 1$ for the coordinate implementation (top row) and $h_{max} = \{1,3,5\}$ voxel with $\Delta h = 0.5$ voxel for the flows-and-limits implementation (the corresponding values in mm are noted above each plot). Each plot shows the estimates in the case of sheet (grey) and non-sheet (black) for different dropout fractions. We set $\kappa = 1/\rho = 1/26 \text{ mm}^{-1}$, $k = 250$, and $p = (10, -10, 0)$.

8.3.1.4 Coordinate implementation: Evaluation of clustering of frames in a spatial neighborhood

The results presented so far use the predefined clustering of the vector fields and are thus not affected by clustering errors. Here, we evaluate the clustering algorithm presented in Section 8.2.2.2 and Appendix A8 on non-sheet vector field pairs U and W with different radii ($\rho = \{5, 15, 25\}$ mm) and a third constant vector field in the z-direction. For each radius we generated 50 noise instances with $k = \{50, 150, 250, 350, 450, \infty\}$ and different dropout fractions $\{0, 0.01, 0.02, 0.05, 0.1, 0.2\}$. We count the number of peaks that were not clustered correctly, where we distinguish between peaks that did not get a label at all and peaks that got an incorrect label.

Supplementary Fig. S8.1 shows the peak clustering error as a fraction of the total number of peaks simulated and for different angle threshold settings. In all erroneous cases the algorithm failed to assign the peaks to a vector field; a wrong assignment to a vector field did not occur in the evaluated cases. The number of non-clustered peaks generally increased with dropout fraction. In fact, this led to a larger fraction of dropouts than initially simulated, but did not have a significant effect on the Lie bracket estimates in analogy with Fig. 8.5 upper row (results not shown).

8.3.2 Diffusion MRI simulations

In this Section we present a comparison of the two methods for dMRI simulated data. Sections 8.3.2.1 to 8.3.2.3 outline the influence of SNR, curvature, and commutativity of the non-sheet vector field on the Lie bracket normal component estimates.

8.3.2.1 Coordinate vs. flows-and-limits implementation: Influence of SNR

Fig. 8.6a shows the results for different SNR levels for the coordinate (top, $\beta = 1$ and $N = 11$ voxels) and flows-and-limits (bottom, $h_{max} = 5$ voxels) implementations. We used $\kappa = 1/\rho = 1/26 \text{ mm}^{-1}$, $p = (10, -10, 0)$, and voxel size $\delta = 1 \text{ mm}$.

Higher SNR increases the precision of the estimates for both methods. For an SNR of 10-20 the estimates were more precise for the coordinate implementation, while for higher SNR this effect is less pronounced. The mean of the estimates was close to the true value in all cases.

8.3.2.2 Coordinate vs. flows-and-limits implementation: The influence of curvature

Fig. 8.6b presents the mean and precision of the Lie bracket normal component estimates $\widehat{[\cdot, \cdot]}_p^\perp$ as a function of curvature κ of the integral curves ($\rho = \kappa^{-1} = \{8, 13, 18, 23, 28, 33\} \text{ mm}$). Results are shown for the coordinate (top, $\beta = 1$ and $N = 11$ voxels) and flows-and-limits (bottom, $h_{max} = 5$ voxels) implementations. We simulated an SNR of 20 and a voxel size $\delta = 1 \text{ mm}$. We evaluated the Lie bracket normal component at different points p for the different curvatures to keep the magnitude of the Lie bracket constant (i.e. 0 and 0.031 for the sheet and non-sheet vector field pairs, respectively). Note that for each curvature the point of evaluation was the same for the sheet and non-sheet vector field pairs.

The precision of the estimates was higher for the coordinate implementation in all cases. For a radius larger than 23 mm , the mean of the estimates is in line with the true value. For a radius smaller than 23 mm in the non-sheet case, the mean of the estimates is only slightly off for the coordinate implementation, while it significantly deviates from the true value for the flows-and-limits implementation.

8.3.2.3 Coordinate vs. flows-and-limits implementation: The influence of commutativity

Fig. 8.6c shows the results for different points $p = (x^1, -x^1, 0)$, where the Lie bracket normal component magnitude of the non-sheet vector field pair (i.e. its ‘non-commutativity’) varies while the curvature remains constant at $\kappa = 1/26 \text{ mm}^{-1}$. Here, $x^1 = \{1, 3, 5, 7, 9, 11\} \text{ mm}$. The Lie bracket was evaluated at the same points for the sheet vector field pair as a reference. Results are shown for the coordinate (top, $\beta = 1$ and $N = 11$ voxels) and flows-and-limits (bottom, $h_{max} = 5$ voxels) implementations. We used an SNR of 20 and voxel size $\delta = 1 \text{ mm}$.

In all cases, the estimates have a higher precision for the coordinate implementation than for the flows-and-limits implementation. The mean of the estimates corresponds well to the true value in all cases.

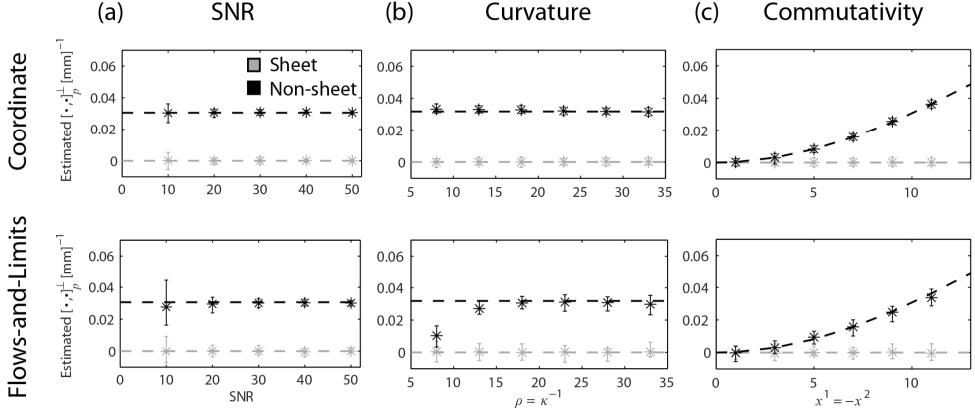


Fig. 8.6: Mean and range of $\widehat{[\cdot, \cdot]}_p^\perp$ for the coordinate implementation (top row $N = 11$ voxels, $\beta = 1$) and the flows-and-limits implementation (bottom row $h_{max} = 5$ voxels, $\Delta h = 0.5$ voxel) in the case of sheet (grey) and non-sheet (black). In every experiment, we vary one factor while keeping others constant: (a) varying SNR, constant $\kappa = 1/\rho = 1/26$ mm $^{-1}$, $p = (10, -10, 0)$, $\delta = 1$ mm; (b) varying curvature and p , constant $[\cdot, \cdot]_p^\perp$, SNR = 20, $\delta = 1$ mm; (c) varying p and $[U, W]_p^\perp$ (non-sheet case), constant SNR = 20, $\kappa = 1/\rho = 1/26$ mm $^{-1}$, $\delta = 1$ mm.

8.3.3 Diffusion MRI real data

8.3.3.1 In vivo HCP data

We first evaluated the performance of the clustering algorithm with real dMRI data. Supplementary Fig. S8.2 shows the peak clustering of an $11 \times 11 \times 11$ neighborhood at three different anatomical locations, where we used an angle threshold of 35° . Overall the clustering looks plausible, and the integral curves corresponding to these vector fields would exhibit significant curvature at this scale (integral curves are not shown, with some abuse of terminology we will say that the vector fields exhibit curvature). However, at some locations the clustering appeared to be challenging: near the edges of the neighborhood, discontinuities occurred in a few occasions. Since the operator applicability reduces towards the edges, the effect of such wrongly clustered peaks generally stays limited.

Fig. 8.7 displays results of the Lie bracket normal component and SPI estimation for the HCP data, where we compare the flows-and-limits ($h_{max} = 5$ voxels) and coordinate ($\beta = 1, N = 11$ voxels) implementations. Fig. 8.7a shows maps of the standard deviation of $\widehat{[\cdot, \cdot]}_p^\perp$ over the 20 bootstraps for both methods (two largest peaks in the WM). Overall, the flows-and-limits approach maps look brighter, which indicates that the standard deviation of the estimates is higher. We used these standard deviation maps to find an appropriate value for λ in the SPI calculation (see Section 8.2.3). To this end, we identified spatially continuous regions with a Lie bracket normal component close to zero and a low standard deviation (an example area is indicated with yellow arrows in Fig. 8.7a). We chose λ to be approximately two times the standard deviation in such areas, i.e. $\lambda = 0.008$. We observed that in these areas, the standard deviation is relatively similar between the flows-and-limits and coordinate approaches. Note that areas in which only a single fiber population can be identified, have a standard deviation of zero since no Lie bracket calculation was possible.

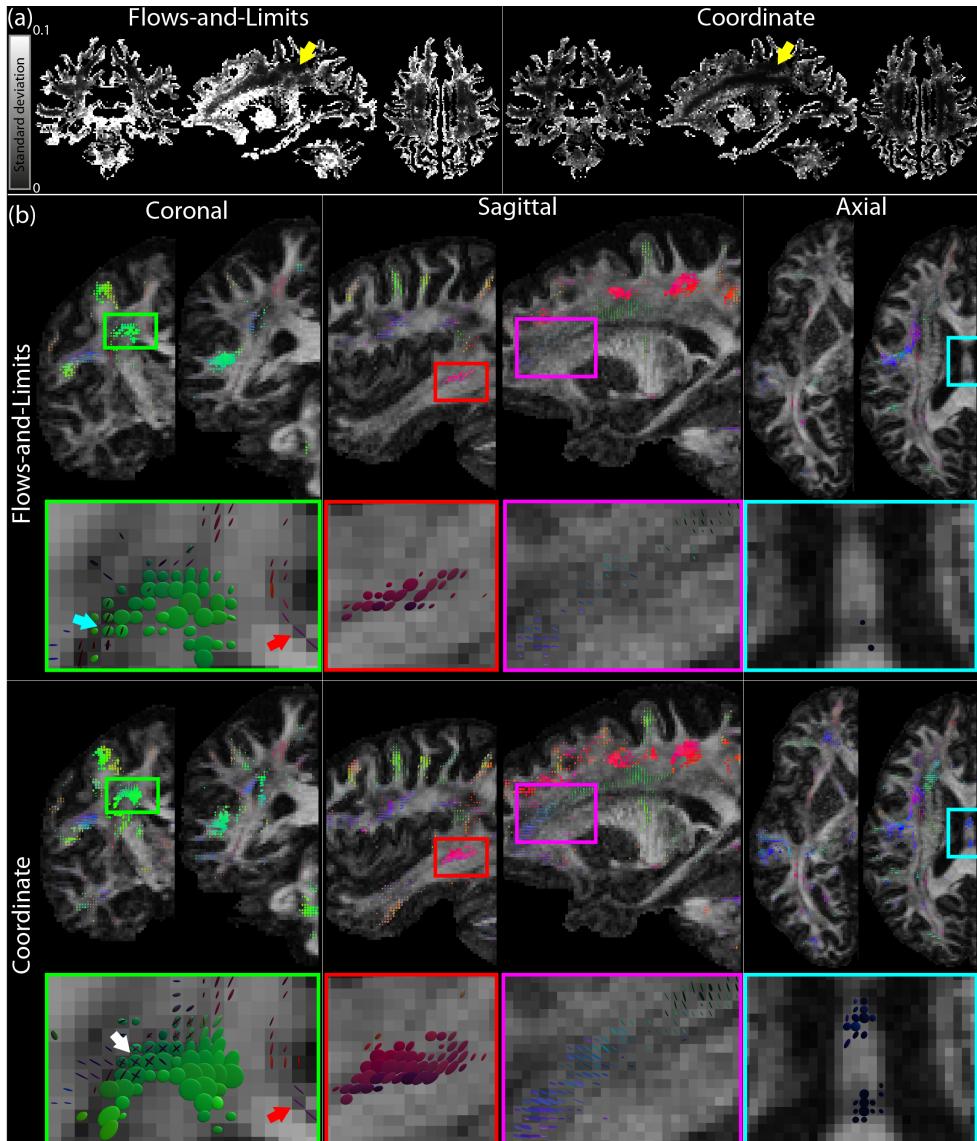


Fig. 8.7: Lie bracket and SPI computation on real *in vivo* HCP data: a comparison between the flows-and-limits and coordinate methods. (a) Standard deviation of the Lie bracket normal component estimates. Yellow arrows indicate areas where the mean is close to zero and the standard deviation is spatially uniform and low. The values in these areas were used to find an appropriate value for λ . (b) Sheet tensors on coronal, sagittal, and axial slices, for the flows-and-limits and coordinate implementations ($\lambda = 0.008$, sheet tensors with SPI < 0.2 are not shown for clarity). White arrow: crossing sheets found with the coordinate implementation, but not with the flows-and-limits method; blue arrow: crossing sheets found with the flows-and-limits method; red arrow: high SPI between the cingulum and the corpus callosum.

Fig. 8.7b shows sheet tensors on different coronal, sagittal, and axial slices. We compare the flows-and-limits method (two top rows) with the coordinate implementation (two bottom rows). In most cases, high SPI areas that could be found with the flows-and-limits method were also identified with the coordinate method. An example is the high SPI area highlighted with a green

rectangle on a coronal slice and the enlarged image below. The green sheet tensors indicate a high SPI between pathways of the corpus callosum and the corticospinal tract. On a coarse scale these regions look similar, but the enlarged image shows some subtle differences. The white arrow indicates a location where all three crossing fiber populations have a high SPI with the coordinate method: they form crossing sheets. These regions did not have a high SPI for each pair of fiber populations when using the flows-and-limits implementation. In turn, the crossing sheets that were identified with the flows-and-limits method (more inferior, indicated with the blue arrow) were not identified with the coordinate method: the SPI was significantly lower or the Lie bracket normal component estimates for the corresponding fiber populations were not always normally distributed and therefore no SPI was calculated. Finally, the red arrow indicates a high SPI in a region where the cingulum and corpus callosum cross. This high SPI region is consistently found with both methods.

Several high SPI areas that are found with both methods are more ‘extensive’ when using the coordinate implementation. Examples are the areas indicated with the red and purple rectangles in the sagittal slices of Fig. 8.7b. The area highlighted in the red box shows a high SPI which likely occurs between pathways of the arcuate fasciculus and association pathways. The purple box in the coordinate case highlights a high SPI area in which the sheet tensors seem to be layered. While this layering appears fairly continuous for the coordinate implementation, it is not obviously present for the flows-and-limits implementation.

Some high SPI areas are exclusively found with the coordinate method. Examples are the region in the brainstem on the coronal slice and the area highlighted by the blue box in the axial slice (see Fig. 8.7b). A zoom of this area is also depicted below the axial slices, and shows a region where parts of the corpus callosum and the fornix pathways cross. The coordinate implementation finds a high SPI coherently over a number of voxels, whereas the flows-and-limits implementation hardly finds any high SPI voxels.

8.3.4 Ex vivo animal data

Fig. 8.8 shows SPI maps for the ex vivo monkey datasets: the 0.4 mm voxel size data in Fig. 8.8a and the 0.3 mm voxel size data in Fig. 8.8b. Here, we used the coordinate approach to estimate the Lie bracket. For each dataset, sheet tensors on coronal (top), sagittal (middle) and axial (bottom) slices are displayed. Here, we set r_{max} in mm the same for both datasets, i.e. $r_{max} = 0.5 \cdot 11 \cdot 0.3\text{ mm}$ based on $N = 11$ for the 0.3 mm dataset. The standard deviation of the estimated Lie bracket normal components was generally larger for the 0.3 mm than for the 0.4 mm dataset (results not shown). An appropriate value for λ was obtained in a similar manner as described in Section 8.3.3.1 ($\lambda = 0.015$ for the 0.3 mm dataset and $\lambda = 0.020$ for the 0.4 mm dataset). Many high-SPI areas can consistently be recognized in both datasets. The 0.3 mm dataset shows more discontinuous high-SPI areas, likely because the normality condition is not fulfilled locally.

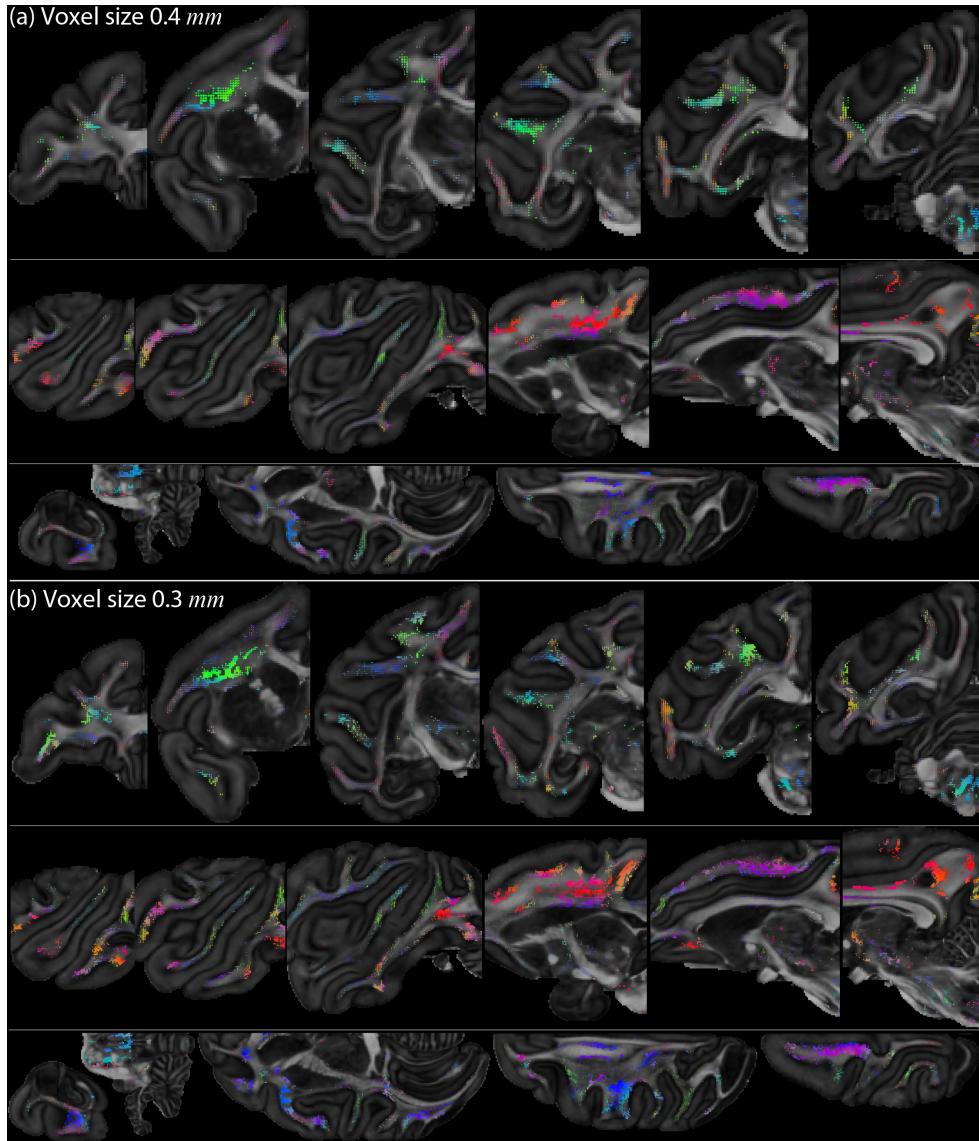


Fig. 8.8: SPI maps for the ex vivo monkey brain. (a) Dataset with a voxel size of 0.4 mm isotropic and 120 directions ($\lambda = 0.015$). (b) Dataset with a voxel size of 0.3 mm isotropic and 60 directions ($\lambda = 0.020$). Sheet tensors with $P_\lambda < 0.2$ are not shown for clarity.

8.4 DISCUSSION

The existence of sheet structure at multiple scales throughout the brain is an ongoing topic of debate, mainly because obtaining proof with reliable quantification is challenging. Previous work has proposed the SPI as a novel quantitative measure to investigate the extent at which

directional data support sheet structure (Tax et al., 2015a; Tax et al., 2014a). The SPI is based on the calculation of the Lie bracket of two vector fields. In this work, we proposed the coordinate method to calculate the Lie bracket, which can be applied to directional data such as peaks resulting from dMRI. We compared this method to the previously proposed flows-and-limits approach on the basis of vector field simulations, dMRI simulations, and real dMRI data. We have shown results of SPI calculations on various datasets of different species, scanners, spatial resolutions, and acquisition settings. In the following sections, we will discuss the different choices we made for the clustering of peaks and the calculation of derivatives with normalized convolution. We will also compare the flows-and-limits and coordinate approaches in terms of performance and implementation/computation considerations. Finally, we will discuss several remaining challenges and outline potential avenues for future work.

8.4.1 Clustering of vector fields

Calculation of the Lie bracket requires the frames in voxels to be clustered into distinct vector fields, which poses a challenging sorting problem in dMRI. The clustering of peaks is also an issue in most tractography algorithms: a tract is propagated along a direction that has the smallest angle with the incoming direction. However, tractography algorithms generally do not take into account the entire frame at the current and previous locations. In the flows-and-limits implementation this is taken into account, that is, the vectors of a frame are clustered into vector fields during the tractography process. The angle threshold on the peaks that are followed can be interpreted directly as a constraint on the curvature of the tract. Sorting peaks along a tract seems a more natural way of clustering than directly comparing peaks in nearby voxels, which mostly do not lie along the same pathway. This comparison of nearby voxels is done in the front-propagation approach for the coordinate implementation. We have observed that the approach of clustering peaks along tracts has fewer problems to recognize the edge (i.e., the domain of definition) of a vector field than the front-propagation approach. For example, at the white-grey matter boundary, we found that the flows-and-limits reconstruction of loops terminated more often naturally because, for instance, peaks were missing or the angle between the peaks was too large. The peaks in grey matter are less reliable because we used a single white matter response function for CSD throughout the brain. A solution to this would be the use of a tissue-specific response function (Jeurissen et al., 2014; Roine et al., 2015). Instead, we used a mask here to take into account only the peaks within the white matter. On the other hand, the front-propagation approach has the advantage that it suffers less from missing peaks resulting from modelling imperfections or noise. Where the loop reconstruction in the flows-and-limits approach terminates in this case, the front-propagation approach can use the information from more than one six-connected voxel to cluster the peaks in a given voxel.

We opted for a front-propagation clustering approach for the coordinate implementation because a direct comparison of all voxels in the $N \times N \times N$ neighborhood with the center voxel could give errors in the case of strong curvature of the vector fields. An angle threshold of 35° seemed to be a good trade-off between including incorrect peaks into a vector field and excluding correct peaks from the vector field in the configurations investigated (e.g. see Fig. S8.1 for excluding peaks). However, we still observed some discontinuities of clustered vectors towards the edges of the spatial neighborhood in real data (Fig. S8.2). In future work this angle threshold could be made dependent on the voxel size and the expected curvature of the vector

fields, and it could be investigated in a wider range of configurations. The influence of incorrectly clustered peaks towards the edges is minimized by the decreasing applicability of the filter basis as the distance from the center voxel increases. Overall, the sorting of peaks is an interesting and challenging problem by itself. Future work could be directed to improve this process, for example by exploiting the use of tractography clustering in this context (e.g. O’Donnell and Westin, 2007), or by taking into account microstructural information for each population and assuming that this is continuous along trajectories (e.g. Girard et al., 2015).

8.4.2 Coordinate implementation: Parameter choices

The normalized convolution approach requires a choice of applicability function (including parameter settings for r_{max} or N and β), certainty function, and filter basis. For the applicability function we chose a power of the cosine function with a period of $4r_{max}$. We chose the cosine function because it is zero at r_{max} and thus has no truncation jumps, but other choices could be explored in future work. Fig. 8.3 shows that a lower value for β decreases the mean absolute error and improves the precision of the estimates up to a certain point. We generally scaled r_{max} according to the kernel size N and the voxel size δ (i.e. $r_{max} = 0.5N\delta$), because adding extra zeros at the edges would only increase the computation time. Increasing N for a given voxel size leads to a higher precision (Figs. 8.3 and 8.5), but comes at the cost of a lower accuracy beyond a given point (Fig. 8.3). Also, increasing N beyond the domain of definition of one of the vector fields might increase the risk of clustering errors (see also Section 8.4.1). In the monkey brain datasets, we chose the same r_{max} in mm for the two spatial resolutions. This means that the effective kernel size N in voxels used for the 0.4 mm dataset is smaller, which is expected to come with a decrease in precision (Fig. 8.3) of the estimates for a given SNR. However, the SNR of the 0.4 mm dataset is larger than that of the 0.3 mm dataset (131 vs 96 on the $b = 0$ images), and the number of acquired directions was higher (120 vs 60), which in fact increased the precision of the estimates for the 0.4 mm dataset (see Section 8.3.3.2). This motivates the data-driven way to defining the λ parameter for SPI calculation.

We currently set the certainty of a vector either to zero or to one depending on whether it is absent or present. However, theoretically it can take any value between zero and one, and could, for example, be based on the angle with its six-connected neighbors. So instead of setting a hard threshold on the angle that two vectors in neighboring voxels are allowed to make in order to belong to the same vector field, the certainty could be defined as a function of this angle. Other factors such as peak magnitude and SNR could also be used to tune the certainty of vectors.

In practice, the filter basis in Eq. [8.9] is sufficient to calculate the first order derivatives of the vector fields and can be extended to calculate higher-order derivatives. Future work could explore the effect of fitting higher-order polynomials on the accuracy and precision of the estimated Jacobians and vectors at the center of the neighborhood.

8.4.3 Flows-and-limits vs. coordinate implementation: Performance

We have investigated the effect of different factors on the performance of the flows-and-limits and coordinate methods, viz. spatial neighborhood, spatial resolution, missing peaks, SNR, curvature, and commutativity.

Our results show that the precision of the Lie bracket normal component estimates (Figs. 8.3 to 8.7) is consistently higher for the coordinate implementation than for the flows-and-limits implementation, which can be attributed mainly to the higher degree of neighborhood information used in the former. In our experiments we aim to allow both methods to probe neighborhood information at a similar distance from the center voxel: an h_{max} of 5 voxels would result in a total walking distance of 11 voxels along one vector field if it was perfectly parallel to the discretization grid (hence, we compare to a kernel size of $N = 11$ voxels). However, the vector fields are generally curved and the flows-and-limits approach only probes vectors that are on the same pathway, whereas the coordinate approach can utilize information of the whole $N \times N \times N$ neighborhood. The results with a given h_{max} and $N = 2h_{max} + 1$ for the respective methods can thus not strictly be compared one-to-one. Further increasing h_{max} might lead to a higher precision, but it comes at the cost of increased computational demands (since more pathways have to be stored), a lower accuracy beyond a given point (similar to e.g. Fig. 8.3 for the coordinate implementation), and the increased risk of clustering errors when h_{max} extends beyond the domain of definition of one of the vector fields (see Section 8.4.1). Another explanation for the better performance of the coordinate implementation is that it gives regularized versions of the vectors at location p where the Lie bracket is to be computed. This is expected to reduce the influence of noise on the computation of the Lie bracket normal component, as seen e.g. in Fig. 8.4. In contrast, the flows-and-limits implementation uses the original (noisy) vectors at point p to compute the Lie bracket normal component. We note that for the real data in Fig. 8.7 the estimated standard deviation was approximately the same between the two methods in the spatially continuous region indicated with yellow arrows. Hence, we used the same λ for the SPI computation.

Figs. 8.3 and 8.4 indicate that the precision of the Lie bracket normal component estimates decreases for smaller voxel size while keeping SNR and spatial neighborhood in voxels constant. Because the extent of the spatial neighborhood in millimeters then decreases, the perturbation of the vectors becomes more significant relative to the quantity that we aim to measure, the rate of change of one vector field along the other vector field. This perturbation is more or less independent of the voxel size. This effect is more prominent for the flows-and-limits approach than for the coordinate implementation.

Curvature and commutativity did not affect the precision of the estimates (Fig. 8.6). A higher curvature, however, did influence the accuracy of the flows-and-limits non-sheet estimates. We observed that this was a combined effect of 1) the design of the vector field, 2) the ability of the used dMRI technique (CSD in this case) to resolve crossing fibers at small angles, and 3) the sensitivity of the flows-and-limits implementation to what we call edge effects. At a radius as small as 8 mm, the edge of the spatial neighborhood for the Lie bracket estimation starts to come in the vicinity of the edge of the simulated vector fields (see Appendix B8). At this edge, the vectors of the non-sheet pair have a crossing angle beyond the angular resolution of CSD. As such, only one peak can be found and many paths are terminated. This effect is larger in one side of the neighborhood, to wit the side near the edge of the vector field, where $[\cdot, \cdot]_p^\perp$ becomes larger. As each path is weighted equally in the LLS fit in Eq. [8.10], the result is more determined by the paths in the direction away from the edge (where $[\cdot, \cdot]_p^\perp$ becomes smaller). This effect does not occur in the same simulations with analytical vector fields (results not shown), and is thus not caused by the flows-and-limits method alone: the method is just more

sensitive to these missing peaks at the edge. In Fig. 8.6c, for $x^1 > 11$ edge effects will also start to play a role in the flows-and-limits implementation (results not shown). While it remains the question whether the edges of our simulated vector fields are representative of a scenario that truly occurs in the brain, it shows that the coordinate implementation is more robust to these effects. Missing peaks also affect the accuracy and precision of the flows-and-limits estimates significantly. We note that missing peaks do not only have an effect on the quantification of sheet structure with the flows-and-limits method, but they also make visual examination challenging because pathways will terminate prematurely.

8.4.4 Flows-and-limits vs. coordinate implementation: Computational issues

The flows-and-limits and coordinate methods are two conceptually different but equivalent approaches to defining the Lie bracket. They do differ greatly in implementation: whereas the flows-and-limits approach is based on tractography, the coordinate implementation works directly on the vector fields. The flows-and-limits approach requires the reconstruction and storage of points and clustered directions of many loops, and it is therefore demanding in terms of memory if many voxels are processed simultaneously. The coordinate approach based on normalized convolution is easier to implement, but has to be performed voxel-by-voxel because the vector fields have to be clustered in each neighborhood. In this work, the computation speed of the coordinate approach was increased tremendously by making use of multi-core processing capabilities.

8.4.5 Further improvements and future work

The computation and visualization of the SPI involves multiple steps, and this manuscript specifically focused on the estimation of the Lie bracket. We already discussed potential improvements for the clustering of vector fields, but other steps in the pipeline could also be improved. Future work could focus on finding alternatives for the normality condition, which might be relaxed, and on automatic tuning of the value of λ . In addition, within a voxel, all peaks from the different bootstraps are currently clustered towards the peaks of the original (non-bootstrapped) data. This clustering procedure could be potentially improved by an iterative scheme that tries to update the mean of the clusters, such as with a k-means clustering principle with k the number of peaks within a voxel. Furthermore, visualization of the sheet tensors is currently based on the original (noisy) vectors, whereas instead the regularized vectors obtained with the coordinate implementation could be used. This could potentially circumvent the presence of unexpected shapes of sheet tensors with high SPI such as those originating from the crossing between the fornix and corpus callosum pathways (see Fig. 8.7). Whereas nearly orthogonal crossings in this region were expected to result in circular flat sheet tensors, some sheet tensors had a more ellipsoidal profile. This was likely caused by spurious peaks resulting from partial volume effects with CSF. Finally, future work could be directed towards trying to avoid the peak extraction altogether, and calculate a sheetness indicator more directly from the ODFs.

8.5 CONCLUSION

Reliable quantification of the brain's sheet structure is needed to assess its presence and configurational properties such as extent and location across individuals (Catani et al., 2012). In this chapter, we have proposed the coordinate implementation to estimate the normal component of the Lie bracket, which is an indicator of sheet structure between two vector fields. We have compared this method to the previously proposed flows-and-limits approach by performing simulations and experiments with real data. We show that the coordinate approach can achieve a higher precision of the Lie bracket normal component estimates. It is more robust against missing peaks resulting from noise or modelling imperfections, increasing the accuracy. Our results demonstrate that the SPI can be quantified with the proposed approach for datasets from different species, scanners, voxel sizes, and diffusion gradient sampling schemes, and in both *in vivo* and *ex vivo* datasets.

8.6 APPENDICES

8.6.1 Appendix A8: Clustering of frames in an $N \times N \times N$ spatial neighborhood

We will first describe the algorithm to cluster two frames based on their total cosine similarity, and subsequently the front-propagation approach to cluster frames in a spatial neighborhood. We consider an ordered set of n vectors $[Y_1, \dots, Y_n]$ at a position q , and m vectors¹ $\{Z_1, \dots, Z_m\}$ in some point r close to q . The ordered set $[Y_1, \dots, Y_n]$ defines the reference frame, i.e., it is assumed that n vector fields U, V, \dots are present in the spatial neighborhood of q satisfying $U(q) = Y_1, V(q) = Y_2, \dots$. The clustering algorithm aims to find a permutation of the frame (an ordered set $[Z_{P_1}, \dots, Z_{P_n}]$) that corresponds to the reference frame, so that we can set $U(r) = Z_{P_1}, V(r) = Z_{P_2}, \dots$ for some permutation P of $[1, \dots, n]$. Here, P_i denotes the index given by the i^{th} element of P , and $Z_{P_i} = 0$ implies that no matching vector was found.

Algorithm for the clustering of frames $\{Z_1, \dots, Z_m\}$ given an ordered frame $[Y_1, \dots, Y_n]$

- 1) Compute the cosine similarity S_{ij} between Y_i and Z_j for all $i \in \{1, \dots, n\}, j \in \{1, \dots, m\}$ (recall that Y_i and Z_j are unit or zero vectors):

$$S_{ij} = |Y_i \cdot Z_j|. \quad [\text{A8.1}]$$

- 2) For every n -permutation P of $[1, \dots, m]$, compute the similarity energy of the permutation by:

¹ If $m < n$, we append $n - m$ zero-vectors to the list $\{Z_1, \dots, Z_m\}$, so in the following we can take $m \geq n$.

$$E_P = \sum_{i=1}^n S_{iP_i}. \quad [\text{A8.2}]$$

- 3) Determine the permutation P for which E_P is maximum, and define the reordered set of vectors $[Z_{P_1}, \dots, Z_{P_n}]$.
- 4) If $Y_i \cdot Z_{P_i}$ has a negative sign, set Z_{P_i} to $-Z_{P_i}$.
- 5) Apply an angle threshold on Z_{P_i} . If $|Y_i \cdot Z_{P_i}| < t$ for some threshold $t \in \mathbb{R}$, set Z_{P_i} to 0.
- 6) Return $[Z_{P_1}, \dots, Z_{P_n}]$.

Algorithm for the clustering of frames in an $N \times N \times N$ spatial neighborhood of point p

- 1) Start from point p with ordered frame $[Y_1, \dots, Y_n]$ satisfying $U(p) = Y_1, V(p) = Y_2, \dots$
- 2) While there are still non-clustered frames in the neighborhood, do the following for each of the K voxels clustered in the previous iteration (denoted by its location q_k , $k \in \{1, \dots, K\}$, for the first iteration $q_1 = p$):
 - a) Identify the 6-connected voxels denoted by their locations $\{r_1, \dots, r_6\}$.
 - b) Keep the L voxels that have not been visited before and that fall within the $N \times N \times N$ spatial neighborhood of point p , i.e. voxels $\{r_1, \dots, r_L\}$.
 - c) For every voxel r_l , $l \in \{1, \dots, L\}$, calculate the cosine similarities for each pair of vectors of the frames $[U(q_k), V(q_k), \dots]$ and $\{Z_1(r_l), \dots, Z_m(r_l)\}$ with Eq. [A8.2].
 - d) If voxel r_l is a 6-connected neighbor of more than one previously clustered voxel (we take as an example here two voxels q_{k_1} and q_{k_2} with $k_1 \neq k_2$), take the mean of the cosine similarities with the frames $[U(q_{k_1}), V(q_{k_1}), \dots]$ and $[U(q_{k_2}), V(q_{k_2}), \dots]$. This corresponds to taking the cosine similarity with the ‘mean frame’ $([U(q_{k_1}), V(q_{k_1}), \dots] + [U(q_{k_2}), V(q_{k_2}), \dots])/2$.
 - e) Cluster the frame $\{Z_1(r_l), \dots, Z_m(r_l)\}$ to this ‘mean frame’ using the algorithm for clustering of two frames described previously.
- 3) Return U, V, \dots

8.6.2 Appendix B8: Analytical vector field simulations

We define three vector fields U, V , and W , where U and V are tangent to the upper hemisphere with radius ρ and W is tangent to the lower hemisphere with radius ρ (see Fig. A8.1 (Tax et al., 2016a; Tax et al., 2016b)):

$$\begin{aligned} U &= (-\sin \phi_1, \cos \phi_1 \cos \theta_2, \cos \phi_1 \sin \theta_2), \\ V &= (\cos \phi_2 \cos \theta_1, -\sin \phi_2, \cos \phi_2 \sin \theta_1), \\ W &= (\cos \phi_2 \cos \theta_2, -\sin \phi_2, -\cos \phi_2 \sin \theta_2). \end{aligned} \quad [\text{A8.3}]$$

Here $\theta_i = \tan^{-1} \frac{x^i}{\sqrt{\rho^2 - (x^1)^2 - (x^2)^2}}$ and $\phi_i = \cos^{-1} \frac{(x^i)}{\rho}$ (with $i = 1, 2$). The integral curves of these vector fields have constant curvature $\kappa = 1/\rho$.

Vector fields U and V form a sheet, so that $[U, V]_p^\perp = 0 \forall p \in \{x \in \mathbb{R}^3 | (x^1)^2 + (x^2)^2 < \rho^2, x^3 = z\}$. U and W generally do not form a sheet, and the normal component of the Lie bracket $[U, W]_p$ at $p \in \{x \in \mathbb{R}^3 | (x^1)^2 + (x^2)^2 < \rho^2, x^1 \neq 0, x^2 \neq 0\}$ is given by

$$[U, W]_p^\perp = \frac{6x^1x^2(-\rho^2+(x^1)^2+(x^2)^2)}{\sqrt{(\rho^2-(x^1)^2)(\rho^2-(x^2)^2)(\rho^6-8\rho^2(x^1)^2(x^2)^2+4(x^1)^2(x^2)^2((x^1)^2+(x^2)^2))}}. \quad [\text{A8.4}]$$

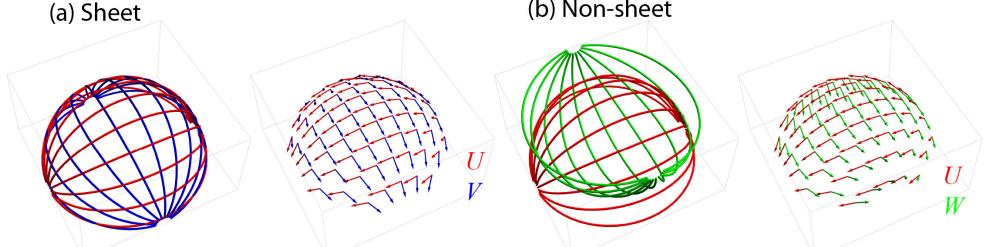


Fig. A8.1: Integral curves and sampled vectors of pair U and V that form a sheet (a), and pair U and W that do not form a sheet (b).

8.7 SUPPLEMENTARY MATERIAL

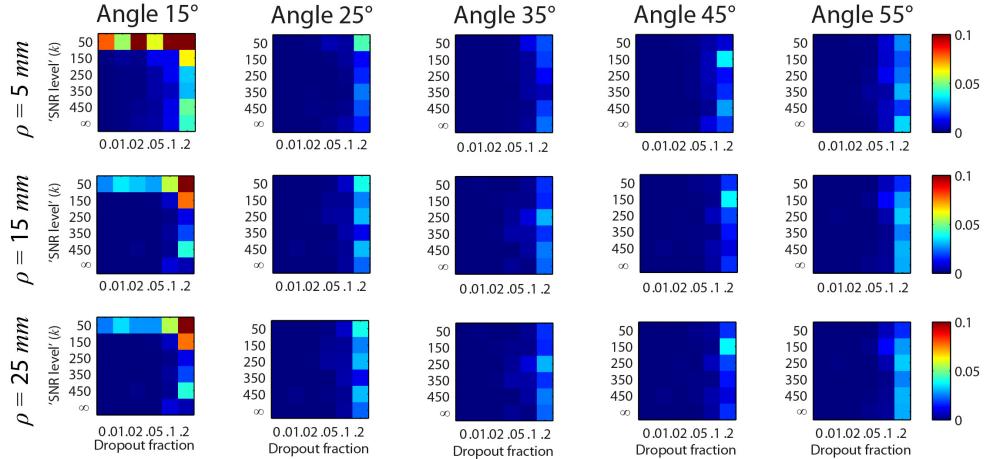


Fig. S8.1: Percentage of non-correctly clustered peaks for different angle thresholds (columns) and radii of vector fields U and W (rows).

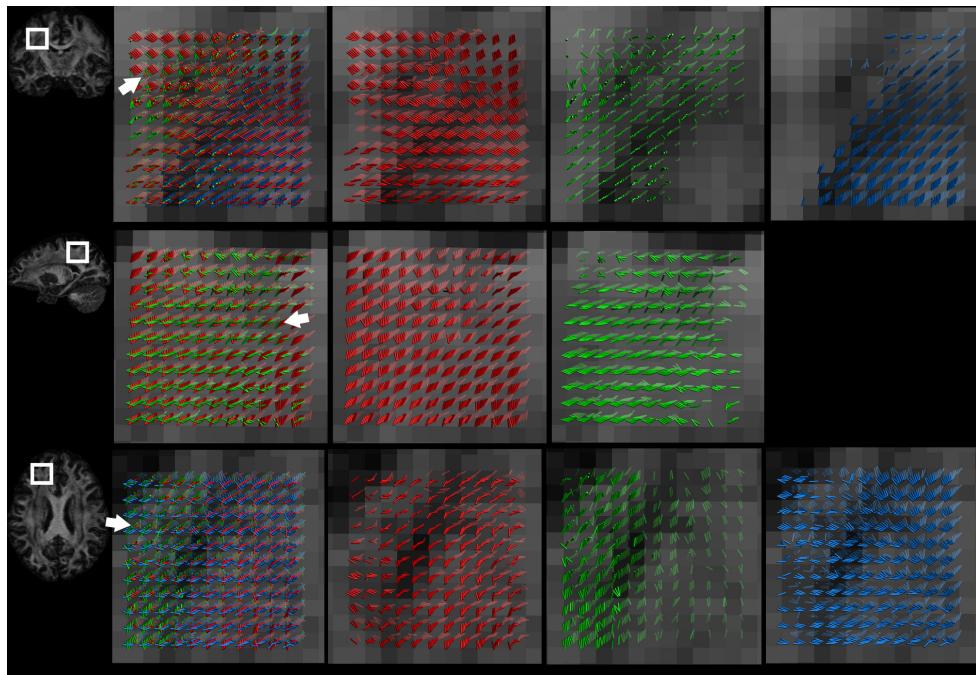


Fig. S8.2: Clustering of vector fields in an $11 \times 11 \times 11$ spatial neighborhood at three different locations in the HCP data (rows). The first (red), second (green), and third (blue) peaks are visualized together (second column) and separately (third to fifth column). Arrows indicate example areas where clustering appeared challenging, mostly at the edges.

9 SEEING MORE BY SHOWING LESS: ORIENTATION-DEPENDENT TRANSPARENCY RENDERING FOR FIBER TRACTOGRAPHY VISUALIZATION

Based on:

C.M.W. Tax, M. Chamberland, M. van Stralen, M.A. Viergever, K. Whittingstall, D. Fortin, M. Descoteaux, A. Leemans, "Seeing more by showing less: Orientation-dependent transparency rendering for fiber tractography visualization", PLoS One, 2015, nr. 10, vol. 10, p. e0139434

ABSTRACT

Fiber tractography plays an important role in exploring the architectural organization of fiber trajectories, both in fundamental neuroscience and in clinical applications. With the advent of diffusion MRI (dMRI) approaches that can also model ‘crossing fibers’, the complexity of the fiber network as reconstructed with tractography has increased tremendously. Many pathways interdigitate and overlap, which hampers an unequivocal 3D visualization of the network and impedes an efficient study of its organization. We propose a novel fiber tractography visualization approach that interactively and selectively adapts the transparency rendering of fiber trajectories as a function of their orientation to enhance the visibility of the spatial context. More specifically, pathways that are oriented (locally or globally) along a user-specified opacity axis can be made more transparent or opaque. This substantially improves the 3D visualization of the fiber network and the exploration of tissue configurations that would otherwise be largely covered by other pathways. We present examples of fiber bundle extraction and neurosurgical planning cases where the added benefit of our new visualization scheme is demonstrated over conventional fiber visualization approaches.

9.1 INTRODUCTION

Diffusion magnetic resonance imaging (dMRI) is a unique technique that can infer information about the architectural organization of tissue *in vivo* (Johansen-Berg and Behrens, 2009; Jones, 2010b). It works by sensitizing the MRI sequence to the random motion of water molecules in a particular direction. Each variation in measurement direction, degree of sensitivity to diffusion, and diffusion time yields an individual image that contains unique information. These images can be combined to obtain a more complete picture of the diffusion properties, generally resulting in a complex high-dimensional dataset.

The high dimensionality of dMRI data challenges not only image interpretation, but also its visualization (Leemans, 2010; Margulies et al., 2013). Therefore, dMRI reconstruction methods generally attempt to reduce the amount of data to meaningful features that can subsequently be visualized. One example of such a feature is the fiber orientation distribution function (fODF), which gives the probability of a fiber population in each direction. fODFs can be visualized at every location as spherical ‘glyph’ representations that have a magnitude proportional to the probability in every direction (Fig. 9.1a, right) (Vaillancourt et al., 2015). Whereas these glyph representations only contain local information, fiber tractography provides a way to visualize large-scale structures by virtually reconstructing trajectories, thereby further simplifying the data on a more global scale. Tractography visualizations are increasingly used in clinical applications, for example in neurosurgical procedures (Girard et al., 2012; Tax et al., 2014b) and in brain connectivity studies. However, in contrast to conventional 2-dimensional (2D) slice-visualizations that are often used in clinical practice, tractography visualizations provide 3D information that might be hard to interpret due to streamlines that are running in the direction of the viewing axis or ‘out of the plane’ (e.g. the geometry of the corticospinal tract (CST), which runs in inferior-superior orientation, is hard to interpret from a 2D axial slice-visualization). Therefore, several methods have been presented that focus on the 3D visualization of reconstructed fiber pathways; these are based on geometric principles like streamlines (Conturo et al., 1999), streamtubes (Zhang et al., 2003), hyperstreamlines (Jones et al., 2005; Reina et al., 2006; Vos et al., 2013), tuboids (Petrovic et al., 2007), triangle strips (Merhof et al., 2006), and streamribbons (Atkinson et al., 2008). In addition to these different geometric principles, the visualization methods also apply shadowing (Eichelbaum et al., 2013a), lighting (Peeters et al., 2006), and coloring.

Notwithstanding all the efforts to improve tractography visualization, displaying large tractography datasets remains highly challenging (Margulies et al., 2013; Presseau et al., 2015). Such visualizations are often cluttered due to overlapping pathways in 3D views and the ability to resolve ‘crossing fibers’ at a voxel level (Fig. 9.1a and Fig. 9.1b, left) (Descoteaux et al., 2009; Jeurissen et al., 2013). Therefore, it is desirable to even further focus on the relevant information available in tractography datasets. One approach is to strategically place seed regions to only extract bundles of interest (Chamberland et al., 2014). While these visualizations provide greater detail, they often lose important contextual information. Other approaches have been developed that try to address the challenge of visualizing both detail and context of streamlines (Margulies et al., 2013). Schurade et al. (2010) proposed a cutting surface to select pathways of interest based on their interaction with this surface, somewhat similar to real anatomical dissections. Some methods use light exchange between streamlines to

show more detail in a contextual visualization (Eichelbaum et al., 2013a; Everts et al., 2009). Calamante et al. (2012) (Calamante et al., 2010) proposed a technique to generate 2D ‘super resolved’ track density maps based on the intersection between dense tractography streamlines and a user-defined sub-millimeter grid.

In this chapter, we propose a new visualization approach that interactively and selectively visualizes pathways based on their orientation by applying orientation-dependent transparency rendering (Tax et al., 2014d). This approach renders pathways more opaque or transparent if they run parallel to a predefined opacity axis (Fig. 9.1b, right). In this way, a more surveyable picture of the 3D architectural organization of these trajectories can be obtained. The proposed approach allows us to better explore the underlying tissue configurations that would be covered when using conventional visualization approaches. We explore different aspects of the method, such as local vs. global transparency rendering and choice of opacity function, and show example applications. A real-time software tool is made available that allows interactive orientation-dependent transparency rendering in large tractography datasets.

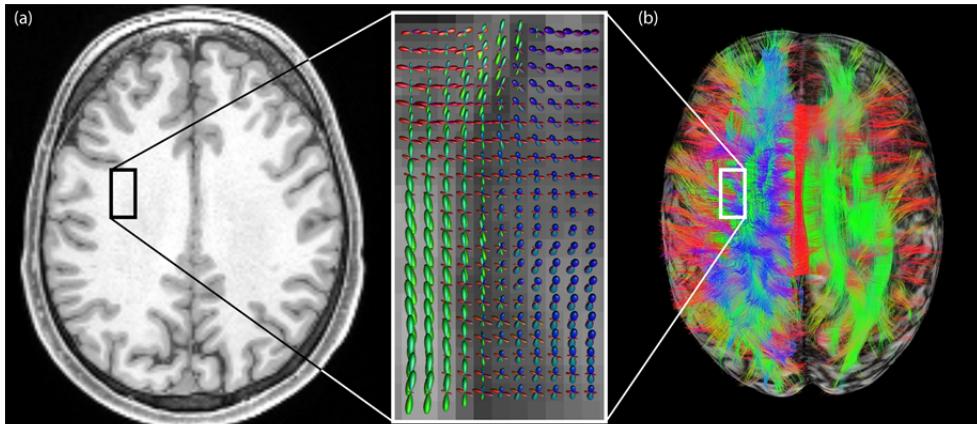


Fig. 9.1: Cluttered fiber tractography visualizations due to the ability to resolve ‘crossing fibers’ at a voxel level and the great overlap of pathways. (a) At the location of the square in the anatomical T1-weighted image (left), the dODFs (visualized as spherical ‘glyph’ representations overlaid on a fractional anisotropy map in the magnified box on the right) reveal fiber crossings at the voxel level. (b) Left hemisphere: pathways greatly overlap, resulting in a cluttered view (superior) in which underlying pathway configurations are hidden. The square indicates the same location as in (a). Right hemisphere: orientation-dependent transparency rendering in which all streamline segments that run in the direction of the viewing axis are rendered transparent. In this way, the underlying structures are revealed and can be explored.

9.2 MATERIALS AND METHODS

9.2.1 Orientation-dependent transparency rendering

In this work, each streamline was rendered with an opacity related to its orientation (either on a local or global scale) with respect to a predefined *opacity axis*. We will first describe our definitions of local and global streamline orientation, and subsequently propose useful opacity functions for transparency rendering in different scenarios.

9.2.1.1 Local vs. global streamline orientation

A streamline is a curve $\gamma: \mathbb{R} \rightarrow \mathbb{R}^3$, $t \mapsto (x(t), y(t), z(t))$ with $x(t)$, $y(t)$, and $z(t)$ its components at position \mathbf{r} , and denoted as $\gamma(t) = \mathbf{r}(t)$. Tractography methods yield a finite sequence of points $(\gamma_i)_{i \in \{1, 2, \dots, N\}}$ to represent a streamline (Conturo et al., 1999; Mori et al., 1999), where short notation $\gamma(t_i) = \gamma_i = (x_i, y_i, z_i)$ is used. From this sequence of points we can compute the local and global orientations of a streamline.

Locally, streamline orientation \mathbf{n}_i at a point γ_i was approximated by the normalized vector connecting its neighbors (Fig. 9.2a):

$$\mathbf{n}_i = \begin{cases} \frac{\gamma_{i+1} - \gamma_{i-1}}{\|\gamma_{i+1} - \gamma_{i-1}\|}, & i \in \{2, \dots, N-1\} \\ \frac{\gamma_{i+1} - \gamma_i}{\|\gamma_{i+1} - \gamma_i\|} & i = 1 \\ \frac{\gamma_i - \gamma_{i-1}}{\|\gamma_i - \gamma_{i-1}\|} & i = N. \end{cases} \quad [9.1]$$

Globally, streamline orientation was approximated in two different ways: 1) by calculating the normalized vector that connects the endpoints of a streamline (Fig. 9.2b, left):

$$\mathbf{n}_{\text{endpoints}} = \frac{\gamma_N - \gamma_1}{\|\gamma_N - \gamma_1\|}, \quad [9.2]$$

and 2) by calculating the scatter matrix \mathbf{S} , which is the average of the dyadics of all local orientations \mathbf{n}_i within a streamline:

$$\mathbf{S} = \frac{1}{N} \sum_{i=1}^N \mathbf{n}_i \mathbf{n}_i^T. \quad [9.3]$$

In the latter case, the first eigenvector of \mathbf{S} , i.e., \mathbf{s}_1 , approximates the global streamline orientation:

$$\mathbf{n}_{\text{scatter}} = \mathbf{s}_1. \quad [9.4]$$

This is similar to considering the vectors \mathbf{n}_i as samples of a 3-dimensional bipolar Watson distribution on the unit sphere S^2 (Fig. 9.2b right), with $\mathcal{S} = \{\mathbf{n} | \mathbf{n} \in \mathbb{R}^3, \|\mathbf{n}\| = 1\}$ (Appendix A) (Mardia and Jupp, 2009). The maximum likelihood estimate of the mean axis $\pm \hat{\mu}$ of the Watson distribution is the first eigenvector \mathbf{s}_1 of \mathbf{S} .

In addition, the eigenvalues $\beta_1, \beta_2, \beta_3$ of \mathbf{S} can provide useful information on the shape of the data (Mardia and Jupp, 2009). When β_1 is significantly larger than β_2 and β_3 , the data is well described by a bipolar/bimodal distribution, and we can consider \mathbf{s}_1 to be a reasonable estimate of the streamline orientation. To investigate whether this assumption holds, one could look at the linear coefficient

$$c_l = \frac{\beta_1 - \beta_2}{\beta_1 + \beta_2 + \beta_3}, \quad [9.5]$$

which ranges from 0 to 1, and is close to 1 in the case of $\beta_1 \gg \beta_2 \cong \beta_3$ (Westin et al., 2002). The c_l measure can be computed for each streamline, and can be used as an extra tuning parameter for global transparency rendering (see Section 9.2.1.3).

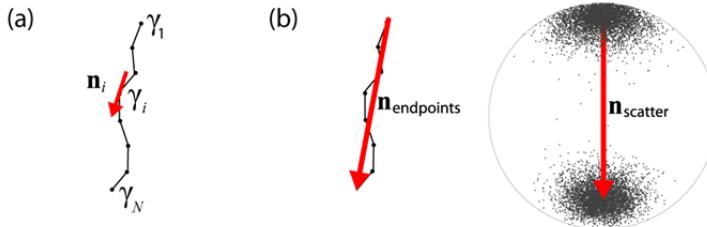


Fig. 9.2: Definitions of local (a) and global (b) streamline orientation. (a) Locally, streamline orientation \mathbf{n}_i at a point γ_i was approximated by the normalized vector connecting its neighbors. (b) Globally, streamline orientation was approximated by 1) calculating the normalized vector that connects the endpoints of a streamline ($\mathbf{n}_{\text{endpoints}}$, left), and 2) calculating the first eigenvector \mathbf{s}_1 of the scatter matrix \mathcal{S} , which is the maximum likelihood estimate of the mean axis of a 3-dimensional bipolar Watson distribution on the unit sphere ($\mathbf{n}_{\text{scatter}}$, right).

9.2.1.2 Opacity functions

The opacity (α) at a point γ_i lies in the interval $[0, 1]$, where an α value of 1 signifies fully opaque, and a value of 0 signifies fully transparent. In this study, the opacity is considered a function of the streamline orientation \mathbf{n} (with $\mathbf{n} = \mathbf{n}_i$ in the local and $\mathbf{n} = \mathbf{n}_{\text{endpoints}}$ or $\mathbf{n} = \mathbf{n}_{\text{scatter}}$ in the global case) and an *opacity axis* \mathbf{t} . We propose two kinds of opacity function (Fig. 9.3): 1) the opacity is decreasing when the streamline orientation is increasingly parallel to the opacity axis (i.e. when the inner product $|\mathbf{n} \cdot \mathbf{t}|$ increases, or the angle θ between \mathbf{n} and \mathbf{t} decreases, Fig. 9.3a);

$$\alpha_{\text{decreasing}}(\gamma_i) = (1 - |\mathbf{n} \cdot \mathbf{t}|)^c; \quad [9.6]$$

or 2) the opacity is increasing when the streamline orientation is increasingly parallel to the opacity axis (Fig. 9.3b):

$$\alpha_{\text{increasing}}(\gamma_i) = |\mathbf{n} \cdot \mathbf{t}|^c. \quad [9.7]$$

Here, c is a constant that tunes the ‘steepness’ of the curve and the ‘cut-off’ value of $|\mathbf{n} \cdot \mathbf{t}|$ where the opacity starts to approach 0.

To motivate our choices of opacity function, we distinguish two scenarios: 1) the opacity axis coincides with the *viewing axis* \mathbf{v} , which is the axis normal to the screen ($\mathbf{t} \parallel \mathbf{v}$), and 2) the opacity axis does not coincide with the viewing axis ($\mathbf{t} \nparallel \mathbf{v}$). In the first case, streamlines that run parallel to the viewing axis do not provide any additional information to the user but obscure underlying streamlines instead. Therefore, it is useful to render such streamlines less opaque (using $\alpha_{\text{decreasing}}$) since the 2D projection of such a line results in a single point. In the

second case, it is useful to visualize streamlines that run in a particular direction and render the other ones transparent (using either $\alpha_{\text{decreasing}}$ or $\alpha_{\text{increasing}}$). In this case, the opacity axis is fixed in a predefined direction and the rendered streamlines can be viewed from different angles.

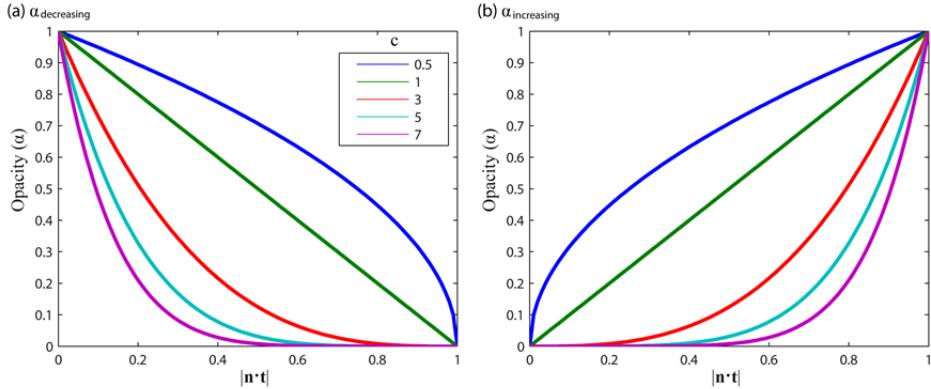


Fig. 9.3: Two opacity functions and the influence of tuning constant c . (a) $\alpha_{\text{decreasing}}$: the opacity is decreasing when the streamline orientation is increasingly parallel to the opacity axis (i.e. when the inner product $|n \cdot t|$ is increasing, or the angle θ between n and t is decreasing), and (b) $\alpha_{\text{increasing}}$: the opacity is increasing when the streamline orientation is increasingly parallel to the opacity axis.

9.2.1.3 Streamline dispersion in global transparency rendering

In the global case, the c_l measure indicates how ‘dispersed’ the pathway is (i.e., a low value of c_l indicates that the streamline does not run in one dominant direction). When a pathway travels in many different directions, rendering the streamline opaque might be desired to be able to see it in all views. Therefore, we investigated the influence of a ‘ c_l threshold’ T_{c_l} on the opacity function: If the c_l of a pathway is below the threshold, render the current pathway opaque (i.e., $\alpha = 1$); otherwise, apply the opacity function $\alpha_{\text{decreasing}}$ or $\alpha_{\text{increasing}}$.

9.2.1.4 Real-time implementation

An interactive implementation of the orientation-dependent transparency rendering method is made available in the Fibernavigator (www.github.com/chamberm/fibernavigator) (Chamberland et al., 2014). This software allows real-time visualization of whole brain fiber tractography datasets, as well as instantaneous segmentation of user-defined bundles. The computation is done in C++ while the rendering is done with calls to OpenGL and GLSL shaders.

For real-time orientation-dependent transparency rendering, the line segments of the streamlines are first ordered according to their mean distance from the point of observation to prevent rendering artifacts. As the viewing axis changes, this ordering is recalculated on the fly, which makes the efficiency of the method proportional to the number of line segments that needs be sorted. To accelerate this depth-sorting step, tractography datasets were compressed by removing redundant points in regions where a streamline is almost linear (tolerance error of 0.01 mm was used) (Presseau et al., 2015). Finally, an opacity value is attributed to each point according to Eq. [9.6] or Eq. [9.7].

9.2.2 Data acquisition and processing

To illustrate our method, two tractography datasets were used. The first dataset was a Human Connectome Project (HCP) dMRI dataset with an isotropic voxel size of 1.25 mm (Sotiroopoulos et al., 2013; Van Essen et al., 2013). Only the $b = 3000 \text{ s/mm}^2$ shell with 90 diffusion directions was used. Whole brain deterministic tractography was performed with *ExploreDTI* (Jeurissen et al., 2011; Leemans et al., 2009) using constrained spherical deconvolution (CSD, $l_{max} = 8$ (Tournier et al., 2007)) and recursive calibration of the response function (Tax et al., 2014c). The second dataset was from a tumor patient (42 year old, male) with an oligoastrocytoma anaplastic tumor (Grade III, WHO) located in the left prefrontal cortex. The study was approved by the Internal Review Board of the Centre Hospitalier Universitaire de Sherbrooke (CHUS) and performed according to their guidelines. Written informed consent was obtained from the patient for the use of anonymized data. Acquisition parameters were: voxel size $2 \times 2 \times 2 \text{ mm}^3$, $b = 1000 \text{ s/mm}^2$, 64 directions (further details have been previously described in Bernier et al. (2014)).

9.3 RESULTS

This section is organized according to the scenarios we have considered previously: either the opacity axis coincides with the viewing axis, or the opacity axis does not coincide with the viewing axis. We use these scenarios to show different aspects of orientation-dependent transparency rendering, including local vs global streamline orientation, choice of opacity function, and the role of streamline dispersion. Table 9.1 shows an overview of the different options, of which we will show some in this section (references to the corresponding figures can be found in Table 9.1). The possibility to interactively rotate the opacity and viewing axes and switch between different settings can be appreciated in the supplementary video available online (<https://www.youtube.com/watch?v=IzJ537KNpR0>).

Scale		Opacity function	Relation of viewing and opacity axis	
			$\mathbf{t} \parallel \mathbf{v}$	$\mathbf{t} \# \mathbf{v}$
Local	Local	$\alpha_{\text{decreasing}}$	Fig. 9.4b	
		$\alpha_{\text{increasing}}$		
	Global	$\alpha_{\text{decreasing}}$	Fig. 9.5	Fig. 9.6a
		$\alpha_{\text{increasing}}$		Fig. 9.6b

Table 9.1: Overview of the different options in orientation-dependent transparency rendering. We investigated different scales (either local or global transparency rendering), different opacity functions (either decreasing or increasing with increasing inner product of streamline orientation and opacity axis $|\mathbf{n} \cdot \mathbf{t}|$), and different relationships between opacity axis \mathbf{t} and viewing axis \mathbf{v} (either $\mathbf{t} \parallel \mathbf{v}$ or $\mathbf{t} \# \mathbf{v}$). We show some of these combinations in the referred figures.

9.3.1 Opacity axis coincides with viewing axis

9.3.1.1 Local streamline orientation

Fig. 9.4 compares conventional streamline rendering (a) with local orientation-dependent transparency rendering (b) where the opacity axis coincides with the viewing axis ($\alpha_{decreasing}$, Eq. [9.6]). Pathway segments that are more aligned with the viewing axis are rendered more transparent. This greatly improves the visualization of the underlying tissue configuration oriented perpendicular to the viewing axis. The middle and bottom rows show results for different tuning constants ($c = 3$ and $c = 7$). When increasing the tuning constant from $c = 3$ to $c = 7$, more streamline segments are rendered transparent due to the increased ‘steepness’ and lower $|n \cdot t|$ ‘cut-off’ value of the opacity function ($\alpha_{decreasing}$, see also Fig. 9.3a).

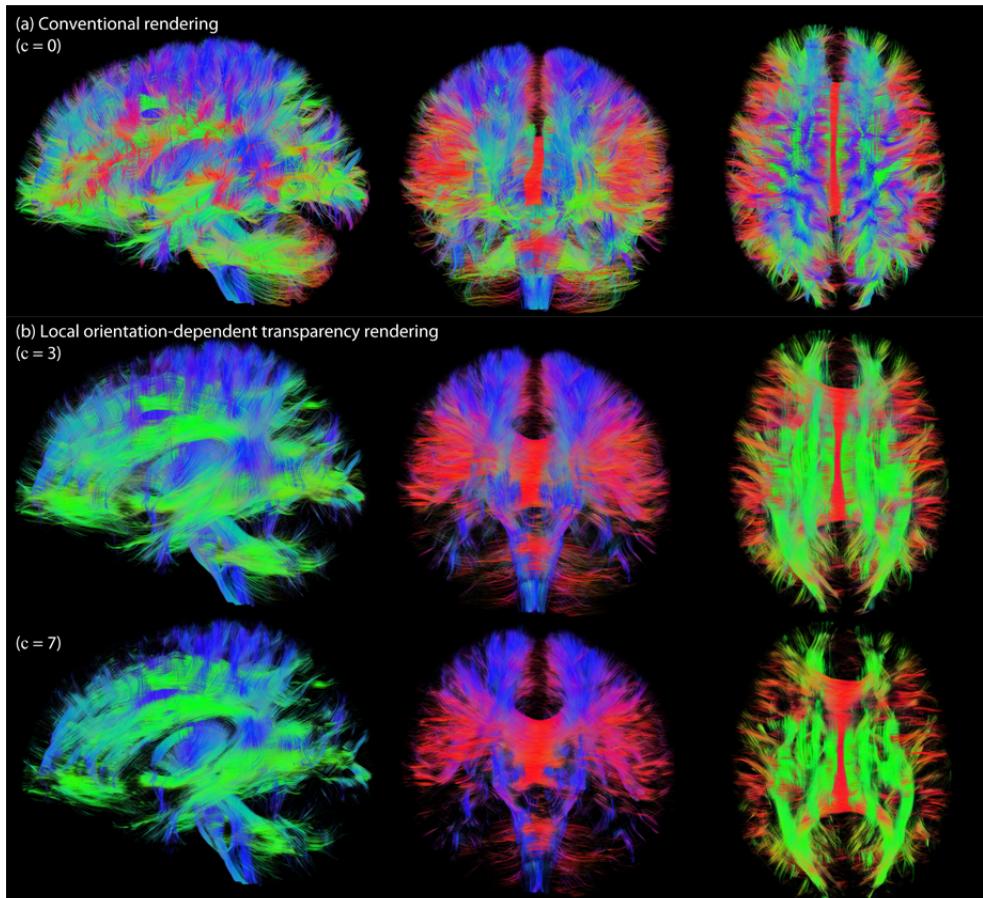


Fig. 9.4: Comparison of conventional streamline rendering (a) with local orientation-dependent transparency rendering (b). Here, the opacity axis coincides with the viewing axis and $\alpha_{decreasing}$ is applied. The middle and bottom rows show results for different tuning constants ($c = 3$ and $c = 7$).

9.3.1.2 Global streamline orientation

The results of global transparent streamline rendering are shown in Fig. 9.5. Whereas for local transparency rendering the opacity is different for every line segment (Fig. 9.4b), for global transparency rendering the opacity is the same along the whole streamline. The latter better preserves the continuous character of a streamline. For example, the whole cingulum (Cg) bundle is visible with global transparency rendering (Fig. 9.5b, yellow square in axial view), whereas parts that run in inferior-superior direction are rendered transparent with local transparency rendering (Fig. 9.4b, same location). Figs. 9.5a and b compare the two different methods to obtain the global streamline orientation, based on the endpoints ($\mathbf{n}_{\text{endpoints}}$, Eq. [9.2]) and based on the scatter matrix ($\mathbf{n}_{\text{scatter}}$, Eq. [9.4]), respectively. The main difference between the methods can be observed in the transparent rendering of U-shaped streamlines (e.g. in the corpus callosum (CC) and in the middle cerebellar peduncle, indicated with white squares in axial views and coronal views, respectively). In the CC, for example, the global orientation of the streamlines is left-right when only looking at the endpoints, whereas it is inferior-superior when looking at the first eigenvector of the scatter matrix (i.e., more line segments are oriented along inferior-superior direction). This explains why the U-shaped CC streamlines are still visible in Fig. 9.5a (blue streamline segments in the white squares), whereas they are rendered transparent in Fig. 9.5b (and thus do not ‘obscure’ other bundles). These scenarios already indicate that there might not be a well-defined global direction for such fibers, which we will investigate further in Section 9.3.2.2.

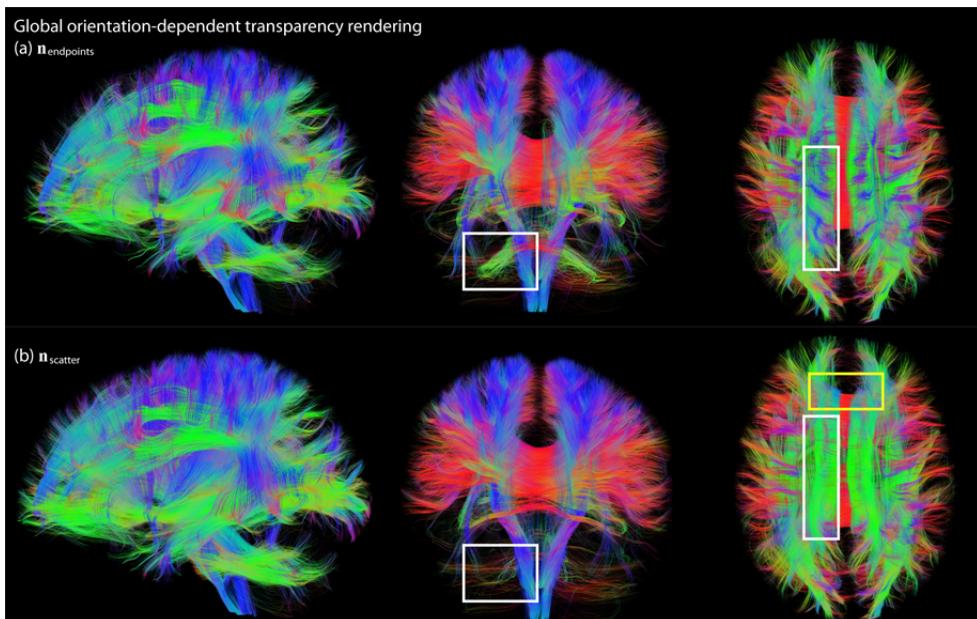


Fig. 9.5: Global transparent streamline rendering ($c = 3$, $\alpha_{\text{decreasing}}$), in which the opacity is the same along the whole streamline (yellow square shows that the whole Cg bundle is visible). In (a) the mean direction of each streamline is based on the endpoints ($\mathbf{n}_{\text{endpoints}}$, Eq. [9.2]), whereas in (b) the mean direction is calculated from the scatter matrix ($\mathbf{n}_{\text{scatter}}$, Eq. [9.4]). The white squares indicate differences between both methods.

9.3.2 Opacity axis does not coincide with viewing axis

9.3.2.1 Opacity functions

Fig. 9.6 shows an oblique view in which the opacity axis does not coincide with the viewing axis but is defined along the left-right (left), inferior-superior (middle), and antero-posterior (right) axes. Figs. 9.6a and 9.6b illustrate the difference between the two opacity functions $\alpha_{\text{increasing}}$ (Eq. [9.6]) and $\alpha_{\text{decreasing}}$ (Eq. [9.7]) on global transparency rendering, respectively. When applying $\alpha_{\text{increasing}}$, streamlines that increasingly coincide with the opacity axis are rendered more opaque. This can be used to only *show* pathways that run in a particular direction (Fig. 9.6a). On the other hand, when applying $\alpha_{\text{decreasing}}$, streamlines that increasingly coincide with the opacity axis are rendered more transparent. This can be used to only *eliminate* pathways that run in a particular direction (Fig. 9.6b).

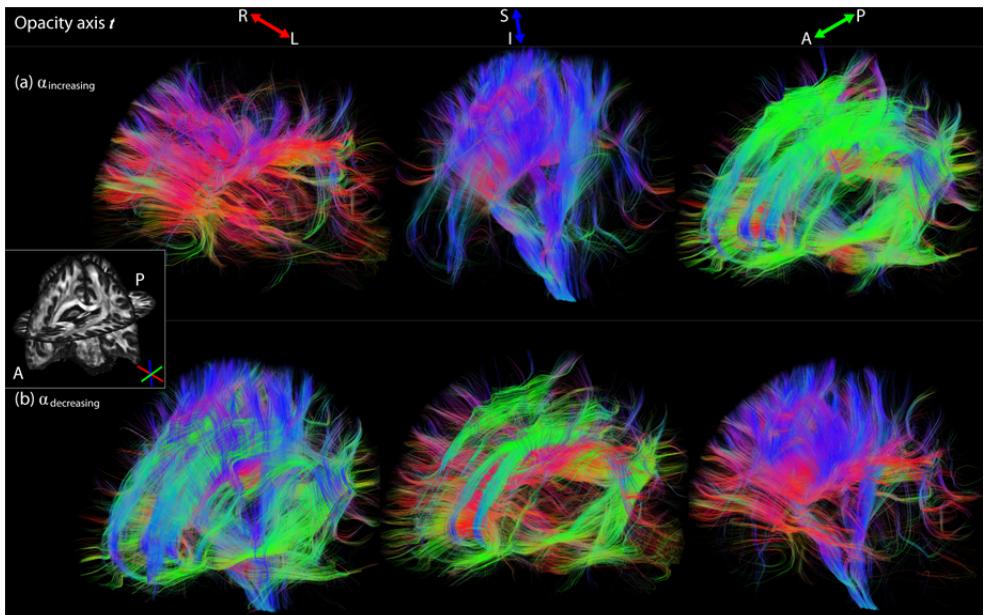


Fig. 9.6: Opacity axis is not aligned with the viewing axis, but fixed along the left-right (left), inferior-superior (middle), or antero-posterior (right) axis. Frames (a) and (b) illustrate the difference between the two opacity functions $\alpha_{\text{decreasing}}$ (Eq. [9.6]) and $\alpha_{\text{increasing}}$ (Eq. [9.7]) on global transparency rendering ($n_{\text{scatter}}, c = 10$), respectively. Streamlines can either be rendered opaque (a) or transparent (b) along a specific axis.

9.3.2.2 Streamline dispersion in global transparency rendering

Strongly dispersed streamlines do not have a single and well-defined global direction, and it might therefore be beneficial to render them opaque to be able to see them in all views (see also Fig. 9.5). Fig. 9.7 shows the results for taking into account streamline dispersion on a set of (both curved and straight) streamlines of the left inferior fronto-occipital fasciculus (iFOF), left corticospinal tract (CST), CC, and left fornix. Fig. 9.7a shows conventional rendering (no transparency) of these bundles in a sagittal, coronal, and axial view, respectively. Fig. 9.7b shows global orientation-dependent transparency rendering. Straight streamlines that have a global orientation along the viewing axis are rendered transparent (lateral projections of the CC,

IFOF, and CST in respectively the sagittal, coronal, and axial view), but also curved streamlines of the fornix are rendered (almost) transparent in both views, and the U-shaped streamlines of the CC are rendered transparent in the axial view. Fig. 9.7c shows global transparency rendering, but now we do take into account the dispersion of the streamlines. With this option, the curved fibers (fornix and U-fibers of the CC) remain visible in each view.

Fig. S9.1 illustrates the influence of T_{cl} on the visualization of streamlines in the CC and the Cg. Streamlines of the CC and the Cg are relatively curved, so the orientational distribution of their line segments will be more ‘dispersed’. The parameter T_{cl} can for example determine to what extent the fanning of the CC is visualized (Fig. S9.1a).

9.3.3 Applications

In the previous sections, we have shown that orientation-dependent transparency rendering greatly contributes to exploring underlying tissue configurations that would otherwise be covered by other pathways. In this section, we present two applications in which orientation-dependent transparency rendering is useful: the extraction of fiber bundles, and the combined visualization of streamlines with space occupying regions such as tumors.

9.3.3.1 Bundle localization and extraction

Fig. 9.8a shows a conventional whole brain fiber tractography rendering, in which the more superficial streamlines obscure the deeply located streamlines. In this view, it is for example difficult to accurately localize the inferior fronto-occipital fasciculus (iFOF). When using orientation-dependent transparency rendering, only streamlines running parallel to the opacity axis (in this case antero-posteriorly) will be maintained, and as a result, the iFOF is clearly identifiable (Fig. 9.8b). By interactively positioning regions of interest (ROIs) at the stem (Hau et al., 2014) of the iFOF and in V1 (Fig. 9.8c), the bundle of interest can be extracted.

9.3.3.2 Visualization of space occupying regions: Neurosurgical application

Fig. 9.9 shows a neurosurgical application of the orientation-dependent transparency rendering method. The left column shows the tumor location on an anatomical T1-weighted image, the middle column shows conventional streamline rendering, and the right column shows orientation-dependent transparency rendering. The superior (a), anterior (b) and lateral (c) views show that in conventional rendering, the tumor is largely covered by the massive amount of streamlines that are running parallel to the viewing axis. In contrast, orientation-dependent transparency rendering enables clearer visualization of the tumor mass by rendering streamlines that are running towards the viewing axis more transparent. Fig. 9.9d shows only the streamlines that touch the segmented tumor volume. By fixing the opacity axis along the antero-posterior direction, one can only render Cg streamlines opaque. The resulting view shows that the Cg pathway infiltrates the tumor area, which would otherwise have been covered by the superior projections of the CC.

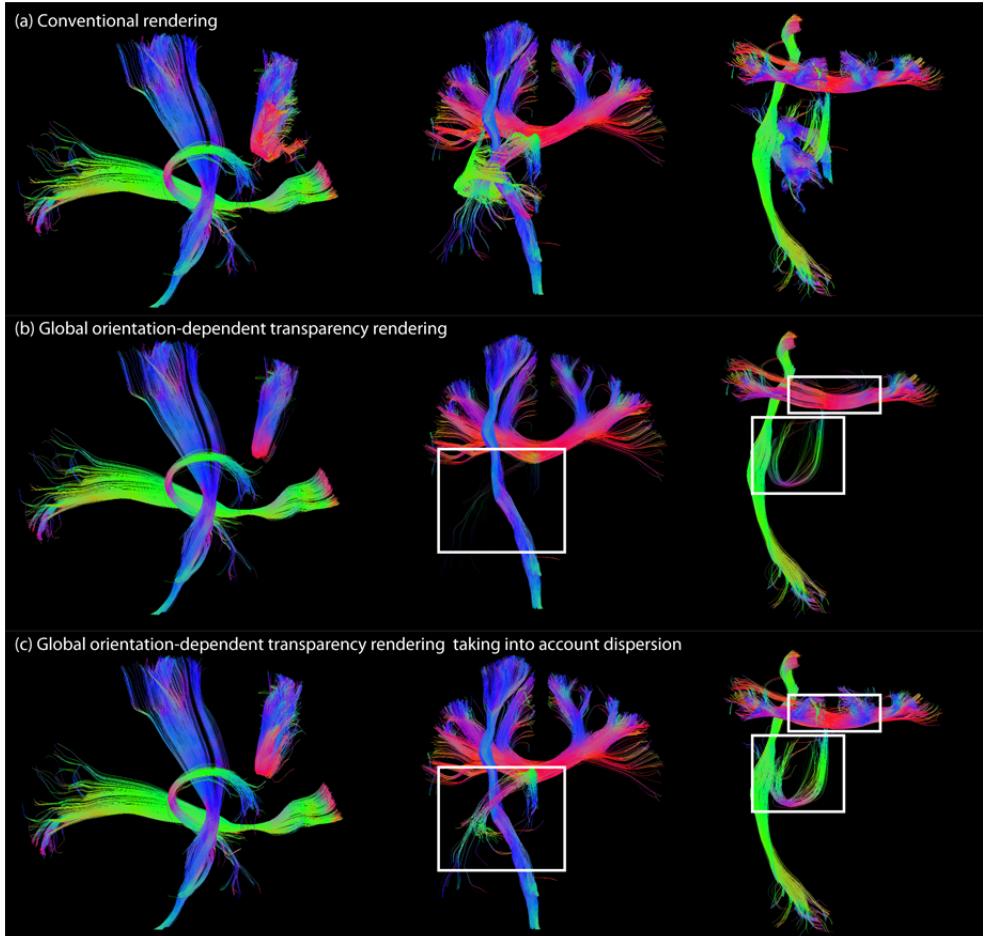


Fig. 9.7: The effect of taking into account streamline dispersion on a set of extracted bundles: left iFOF, left CST, CC, and left fornix. (a) Conventional rendering in a sagittal, coronal, and axial view, respectively. (b) Global orientation-dependent transparency rendering (opacity axis coincides with the viewing axis with $\alpha_{decreasing}$, $c = 3$, and $n_{scatter}$). (c) Global transparency rendering (same settings as in (b)) taking into account the dispersion of the streamlines ($T_{C_l} = 0.29$). White squares highlight streamlines that are rendered transparent when not taking into account streamline dispersion, but are visible when applying a c_l threshold.

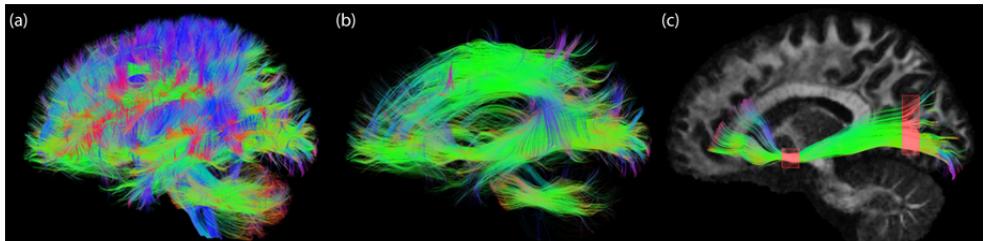


Fig. 9.8: Bundle extraction with orientation-dependent transparency rendering. (a) Conventional whole brain fiber tractography rendering. Using orientation-dependent transparency rendering, streamlines that are running parallel to the antero-posterior axis can easily be visualized (b, $\alpha_{increasing}$ with $c = 7$). From this visualization, one can easily extract the iFOF by positioning two ROIs (red boxes, stem and V1 area).

Chapter 9: Seeing More by Showing Less: Orientation-Dependent Transparency Rendering
for Fiber Tractography Visualization

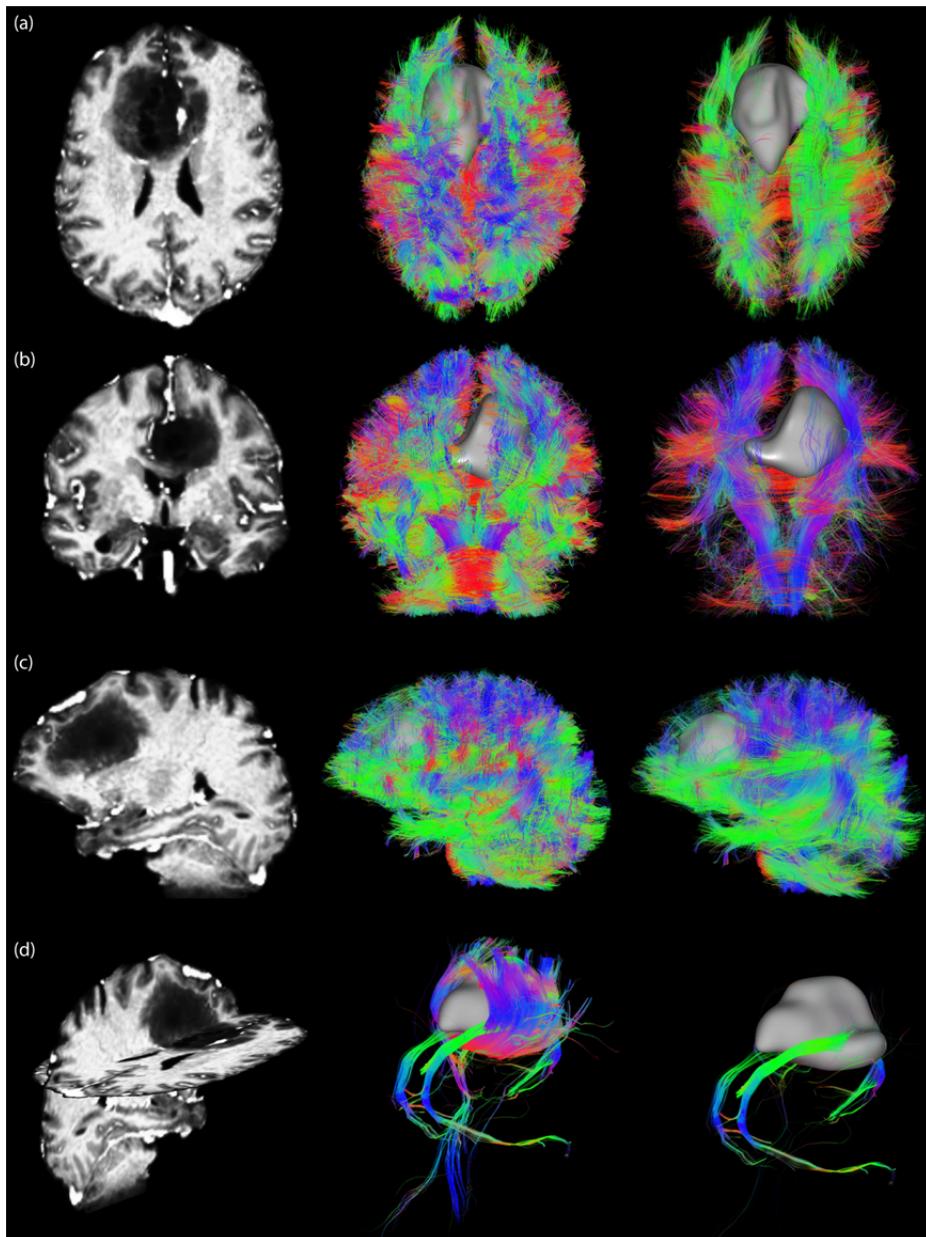


Fig. 9.9: Neurosurgical application of the orientation-dependent transparency rendering method. The left column shows a high resolution T1-weighted image for anatomical reference. The middle column shows conventional streamline rendering, where the tumor (grey) is largely covered by the massive amount of streamlines. The right column shows that orientation-dependent transparency rendering ($\alpha_{\text{decreasing}}$ with $c = 4$) allows for in-depth visualization of a tumor mass by rendering streamlines that are running towards the viewing axis transparent. Figs. (a), (b) and (c) show superior, anterior, and lateral views, respectively. In (d), the tumor volume is used as an ROI for the selection of relevant streamlines. By fixing the opacity axis along the antero-posterior direction ($\alpha_{\text{increasing}}$ with $c = 4$), one can render Cg streamlines opaque, showing that they penetrate the tumor area.

9.4 DISCUSSION

We propose a new visualization approach to interactively and selectively visualize pathways based on their orientation, resulting in a less cluttered and obscure picture of the 3D architectural organization of streamline trajectories. While other software packages have the possibility to attribute coloring and/or opacity levels to streamlines (e.g., Fedorov et al., 2012; Leemans et al., 2009; Tournier et al., 2012), we have explored different aspects of real-time orientation-dependent transparency rendering such as local vs global transparency, choice of opacity function, and the role of streamline-dispersion. In addition, we have shown two example applications. Below, we will discuss the limitations, potential improvements, and future perspectives of the method.

9.4.1 Methodology

9.4.1.1 Local vs. global streamline orientation

We have shown results for both local and global transparency rendering (Fig. 9.3b and Fig. 9.4, respectively). The advantage of global rendering is that the opacity is the same along the pathway, which better preserves its continuous character. The drawback is that globally opaque streamlines can still largely obscure underlying pathways. In addition, it can be challenging to define a robust and representative global orientation of a streamline, especially in the case of highly curved and dispersed fibers. We have investigated two approaches to define global pathway orientation (Fig. 9.5): either based on the endpoints, or based on the scatter matrix of the line segments of a pathway. The endpoints-approach gives useful information as to what parts of the brain are connected by a particular pathway, and is often used in dMRI (visualization) packages to define the global color of a pathway (e.g. Chamberland et al., 2014; Eichelbaum et al., 2013b). However, the location of the endpoint can be heavily influenced by accumulated errors in streamline propagation (e.g. due to noise), in which case it might be beneficial to use the scatter-approach. The endpoints- and scatter-approaches give different results in the case of curved U fibers (see Fig. 9.5). The scatter approach can give an indication of the dispersion/linearity of a pathway, which can be used as an extra parameter to tune opacity and to reduce the risk of hiding important neuroanatomical details (Fig. 9.7). Here, we have used the c_l measure, but there are other (dispersion) measures that can be used as an indicator of the shape of the distribution (e.g., Basser and Pajevic, 2000; Westin et al., 2002). More general, the opacity can be a function of the opacity axis \mathbf{t} and the eigenvectors and eigenvalues of the scatter matrix \mathbf{S} . In this way, not only the linearity of a pathway can be characterized, but for example also its planarity. By interactively switching between local and global transparency rendering, the different methods for global streamline orientation, and different values for T_{c_l} , one can adapt the visualization to the demands of a specific application.

9.4.1.2 Opacity functions

The opacity functions in this work are based on power functions, in which the exponent c (here we choose $c > 1$) can be used to tune the relationship between $|\mathbf{n} \cdot \mathbf{t}|$ and opacity α . While this opacity function is intuitive and elegant as it only has a single parameter and lies between 0 and 1 for $|\mathbf{n} \cdot \mathbf{t}| \in [0,1]$, it might be too simplistic and ad hoc. For example, increasing the exponent

c determines at the same time the ‘steepness’ of the curve and the ‘cut-off’ value of $|\mathbf{n} \cdot \mathbf{t}|$ where the opacity starts to approach 0. Other functions with more parameters can be used instead, such as sinusoidal functions or a linear function. A linear function ($\alpha(x) = ax + b$) might be more useful than sinusoidal functions since it has the freedom to easily tune the steepness and cut-off value while constraining the minimum and maximum α to be 0 and 1, respectively. Supplementary Fig. S9.2 shows an example of local transparency rendering with α a linear function of the angle $\theta = \text{acos}|\mathbf{n} \cdot \mathbf{t}|$, for different values of a and b . The cut-off value of $|\mathbf{n} \cdot \mathbf{t}|$ appears to have the largest influence on the visualization.

Instead of using predefined opacity functions and exploring the influence of their parameters on the visualization, one could also make a more informed decision about the preferred shape of the opacity function. To this end, one could inspect histograms of $|\mathbf{n} \cdot \mathbf{t}|$ as a function of the opacity axis to get an impression of the distribution of $|\mathbf{n} \cdot \mathbf{t}|$. Supplementary Fig. S9.3 shows such histograms for locally (a) and globally (b) defined orientations, for opacity axis in left-right (left), antero-posterior (middle), inferior-superior (right) direction. The opacity function could then be adapted interactively (e.g. by defining it as a piecewise linear function from user defined points) to emphasize or eliminate streamlines or streamline segments with particular orientations.

9.4.1.3 Real-time implementation

While the static images presented in this paper give a good impression of the usefulness of orientation-dependent transparency rendering, the method comes to its full justice when interactively rotating the view and exploring different settings. To this end, we have developed a real-time implementation which is made freely available in the Fibernavigator (Chamberland et al., 2014). Over $3 \cdot 10^6$ line segments (over $1 \cdot 10^6$ compressed) could be ordered and rendered real-time with an acceptable frame rate of over 20 frames-per-second. We found that transparency renderings were the most useful when the number of overlapping line segments was not too high (i.e., the tractography dataset was not too dense), hence we subsampled the first dataset from $\sim 8 \cdot 10^5$ streamlines to $\sim 4 \cdot 10^5$ streamlines. A possible improvement to the interactive method presented here is to accelerate the line segment depth-sorting step, enabling even smoother interaction with larger tractography datasets. Additionally, tube rendering (Zhang et al., 2003) can provide more perceptual depth, since lightning techniques can be more easily applied to tuboids than to thin lines. Future developments will include an upgraded version of the opacity-rendering method which will be implemented in MI-Brain (www.imeka.ca/mi-brain), a streamline-visualization and interaction tool based on the MITK platform (www.mitk.org).

9.4.2 Applications

9.4.2.1 Bundle localization and extraction

The extraction of streamline bundles is most commonly done by placing ROIs (Leemans et al., 2009; Tournier et al., 2012) on a reference image (e.g. T1, T2, direction encoded color FA, etc.) based on anatomical knowledge. This method has the disadvantage of not showing streamlines on the fly, and thus segmentation is often done in a ‘blind’ manner. In contrast, the ability to

interactively position 3D ROIs and real-time tractography allows for fast extraction of specific pathways and direct interaction with streamlines (Chamberland et al., 2014). Nevertheless, it can still be a complex task to place such ROIs when the streamline bundle of interest is obscured by other bundles. With orientation-dependent transparency rendering, one can generate a more efficient view of the current dataset, which can greatly help in ROI positioning. For example, the iFOF is a bundle that predominantly runs in the antero-posterior direction, and is often hidden by more superficial streamlines (Fig. 9.8). Using this a priori knowledge, it is possible to render streamlines that run in this particular direction more opaque. This greatly reduces the amount of streamlines displayed, while exposing large parts of the iFOF, and thus eases the ROI positioning process. In addition, orientation-dependent transparency rendering clearly reveals a recently reported new pathway in the orbitofrontal / prefrontal cortex (Heemskerk et al., 2015) (see e.g. Fig. 9.5, sagittal view). This pathway has previously been undocumented because of the complex ‘crossing fiber’ architecture in this area. New visualization techniques, such as transparency rendering, along with progress in acquisition and data processing aid in localization and identification of such previously undocumented structures.

9.4.2.2 Visualization of space occupying regions

In some applications, tractography results are visualized together with volume renderings. These volume renderings can become largely covered by the large amount of streamlines in tractography datasets. Combining volume renderings with orientation-dependent transparency rendering of streamlines gives a less cluttered view and allows for better exploration of the data, especially in the case of volume renderings that are located deeply in the brain. Here, we discuss two potential applications in which transparency rendering might be beneficial: a neurosurgical application (see also Fig. 9.9), and combined dMRI and functional MRI visualization.

Neurosurgical planning requires a generalized view of deviating and infiltrating streamlines in proximity of the tumor. We found that global transparency rendering was the most appropriate in this case, since the whole streamline has the same opacity (Fig. 9.9). Local transparency rendering might lead to misinterpretation in this case (e.g. local line segments interpreted as streamlines that terminate in the white matter). Furthermore, the ability to interact in real-time with the view and settings is necessary in pre-operative planning applications. More specifically, a neurosurgeon should be able to easily switch between conventional and opacity rendering at any time, so that no relevant pathways can be missed.

Traditional ways of coupling functional MRI (fMRI) and dMRI often come down to displaying fiber pathways that interconnect distant fMRI activation regions (Bernier et al., 2014; Chamberland et al., 2015; Tax et al., 2014b; Whittingstall et al., 2014) (i.e. brain networks). This provides a clear and simple view of the current network of interest. However, one may be interested in looking at a more extensive (whole brain) tractogram together with fMRI clusters, to see more contextual information. Therefore, it can become challenging to spot deeply-located fMRI regions (e.g. sub-cortical activations) that are surrounded by many streamlines. To overcome this problem, orientation-dependent transparency rendering can be applied. By rendering streamlines that are running parallel to the viewing axis more transparent, an

uncluttered view can be achieved for the 3D exploration of such deep fMRI regions while maintaining contextual information.

9.4.3 Recommendations and future perspectives

In a nutshell, orientation-dependent transparency rendering can be used to simplify the visualization of dense streamline datasets. We have proposed different opacity functions and opacity axes settings that can be combined to achieve a particular visualization. While we believe that the power of the method lies in the possibility to interact in real-time with these different views and settings, we found transparency rendering particularly useful in two scenarios. 1) The user has *a priori* knowledge on the general direction of a bundle of interest and wants to identify and extract this bundle. In this case, one can use $\alpha_{\text{increasing}}$ (Eq. [9.7]) with $\mathbf{t} \# \mathbf{v}$ to inspect the bundle from different viewpoints. Global transparency rendering gives the most intuitive visualization here, since the opacity is the same along the whole streamline. 2) The user wants to do a global or more local (e.g. only the centrum semiovale) exploration of the streamline network and wants to eliminate streamlines (or segments) that run in the line of view and obscure underlying configurations. In this case, one can use $\alpha_{\text{decreasing}}$ (Eq. [9.6]) with $\mathbf{t} \parallel \mathbf{v}$ and local or global opacity rendering to remove clutter. In the case of combined visualization of streamlines with volume renderings (such as tumors), we found both scenario 1) and 2) useful: the user can inspect whether a particular bundle intersects with the tumor volume, and the user can remove cluttering pathways in the line of view to obtain a more clear view of the streamlines in proximity of the tumor. While this work focuses on the presentation of the method and showcasing its different aspects, future work will further assess the (clinical) usefulness of this method in the form of a survey among highly qualified clinicians. In addition, the proposed rendering method will be integrated as part of neuronavigation software used for surgical interventions where neurosurgeons explore the wiring of the brain interactively with the tip of their neurotracking tool.

As fiber tractography yields complex 3D structural information (Fig. 9.10a), it is challenging to combine such information with conventional 2D slice-visualizations that are currently used in clinical practice. One way is to visualize 2D track density maps (Fig. 9.10b) (Calamante et al., 2010), which show the intersection of tractography streamlines with a particular slice. While this method results in striking high-resolution images, it abandons part of the 3D depth information, and the user still has to scroll through the slices to get a 3D overview. Orientation-dependent transparency rendering with $\alpha_{\text{decreasing}}$ (Eq. [9.6]) and $\mathbf{t} \parallel \mathbf{v}$ (Fig. 9.10c) can in a way be interpreted as a trade-off between conventional visualization and track density slice visualization: On the one hand it shows a less cluttered view by removing pathways (or segments) that run along the viewing axis, while on the other hand it maintains information on the 3D architecture of streamlines.

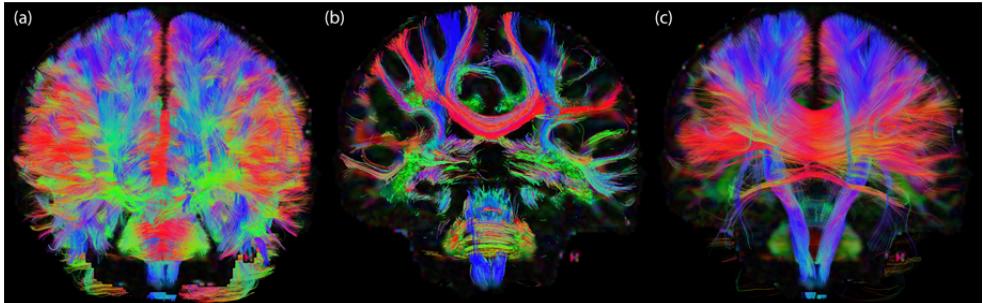


Fig. 9.10: Visual comparison between three streamline rendering techniques. (a) Conventional whole brain streamline rendering in which pathways greatly overlap and clutter the view. (b) Track density imaging (Calamante et al., 2010) in which only streamlines crossing a user-defined plane (in this case a coronal slice) are shown. (c) Orientation-dependent transparency rendering in which all pathways that run in the direction of the viewing axis are rendered transparent. Underlay: directionally-encoded fractional anisotropy map.

9.5 CONCLUSIONS

Orientation-dependent transparency rendering of streamlines as presented in this work avoids cluttered and obscure visualizations, and provides a way to interactively explore the 3D architectural organization of streamlines. We have explored global and local transparency rendering, different opacity functions, different settings for the opacity axis, and the role of streamline dispersion. Using orientation-dependent transparency rendering, streamlines that are oriented along the viewing axis can be rendered more transparent, thereby not obscuring underlying pathways. Alternatively, exclusively bundles that are oriented along particular axes can be visualized, virtually eliminating all other pathways. We have shown that the method is particularly useful in applications like bundle extraction and combined visualization with volume renderings such as tumors in neurosurgical planning.

ACKNOWLEDGEMENTS

The authors thank Hanne Kause for her contributions, Anneriet Heemskerk for her help with data processing and valuable discussions, and Jean-Christophe Houde and Marc-Alexandre Côté for their valuable comments.

9.6 APPENDICES

9.6.1 Appendix A9

The Watson distribution is a fundamental spherical distribution for axial data. The probability density function of a Watson distribution on S^2 is (Mardia and Jupp, 2009)

$$p(\pm \mathbf{n}_i; \boldsymbol{\mu}, \kappa) = M \left(\frac{1}{2}, \frac{3}{2}, \kappa \right)^{-1} e^{\kappa (\boldsymbol{\mu}^T \mathbf{n}_i)^2}. \quad [\text{A9.1}]$$

Here, $M(1/2, 3/2, \cdot)^{-1}$ is the Kummer function (Mardia and Jupp, 2009), μ ($\|\mu\| = 1$) is the population mean axis, and κ is a concentration parameter. For $\kappa > 0$, the distribution is bipolar and κ characterizes how strongly the unit vectors are concentrated around the mean orientation (larger κ means more concentrated around $\pm\mu$).

9.7 SUPPLEMENTARY INFORMATION

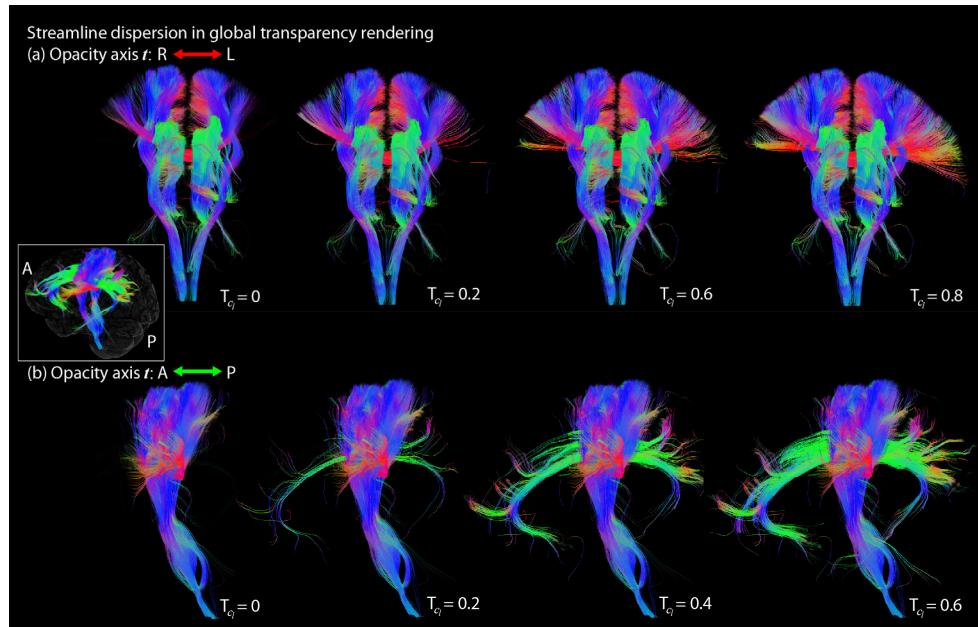


Fig. S9.1: Influence of T_{ci} on a set of three extracted bundles: the CC, the Cg, and the CST. In (a), the opacity axis is left-right oriented, which renders pathways with a global left-right direction (e.g. the lateral (fanning) projections of the CC) transparent. When applying a c_i threshold, streamlines that both have a left-right orientation and a high dispersion (i.e., low c_i) are rendered opaque. When increasing T_{ci} (towards the right), streamlines with an increasingly high c_i are rendered opaque. In (b), the opacity axis is antero-posterior oriented. Increasing T_{ci} results in the display of a larger amount of curved Cg streamlines. Opacity function $\alpha_{\text{decreasing}}$ with $c = 3$ was used in all figures.

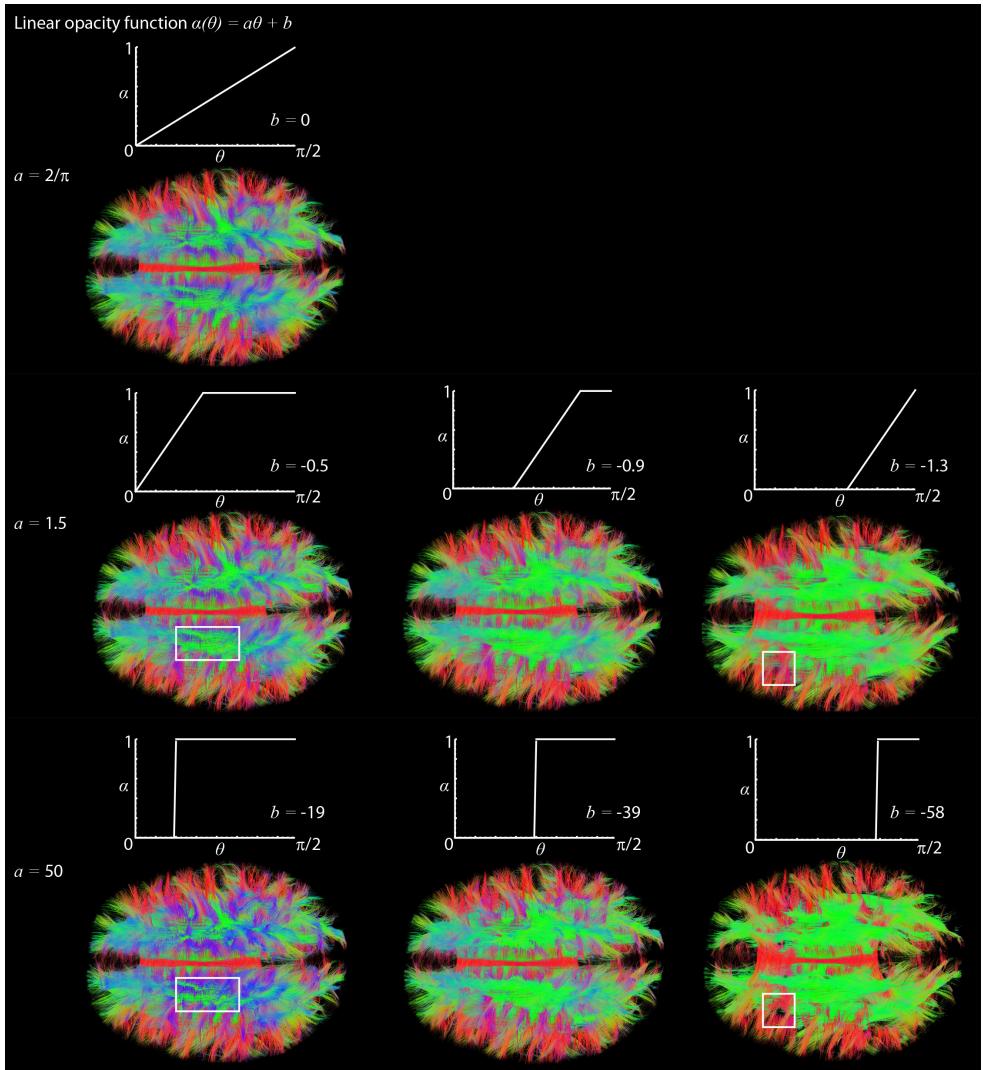
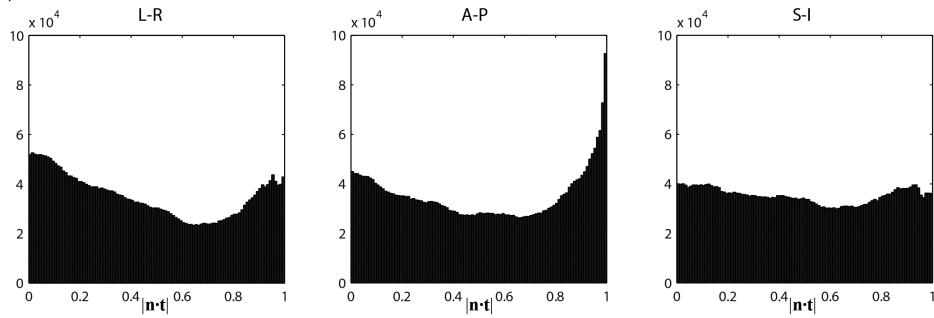


Fig. S9.2: Local transparency rendering with a piecewise linear function $\alpha(\theta) = a\theta + b$ if $\theta \geq -b/a \wedge \theta \leq (1-b)/a$, $\alpha(\theta) = 1$ if $\theta > (1-b)/a$, and $\alpha(\theta) = 0$ if $\theta < -b/a$, with $\theta = \text{acos}[n \cdot t] \in [0, \pi/2]$, the opacity increases when the angle increases. Graphs of the opacity function and the corresponding renderings are shown for different values of a and b . Parameter b appears to have the largest influence on the visualization, whereas a only smoothens the transitions between transparent and opaque streamline segments (see subtle differences highlighted by the white squares).

Chapter 9: Seeing More by Showing Less: Orientation-Dependent Transparency Rendering for Fiber Tractography Visualization

Histograms of $|n \cdot t|$

(a) Local



(b) Global

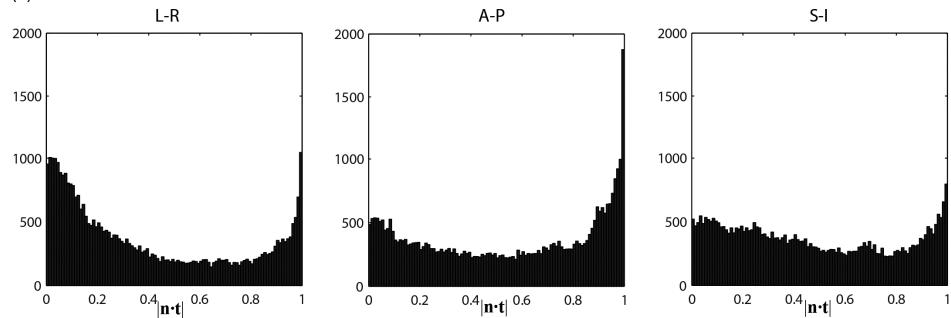


Fig. S9.3: Histograms of $|n \cdot t|$ as a function of the opacity axis (left-right (left), antero-posterior (middle), inferior-superior (right) direction). (a) Locally and (b) Globally defined orientations.

10 DISCUSSION

Based on:

C.M.W. Tax, S.B. Vos, A. Leemans, “Checking and Correcting DTI Data”, Editor(s): W. Van Hecke, L. Emsell, S. Sunaert, Springer New York, 2016, Diffusion Tensor Imaging: A Practical Handbook, pp. 127-150.

C.M.W. Tax, A. Leemans, “What’s new and what’s next in diffusion MRI processing” (to be submitted soon)

With dMRI data being high-dimensional and the processing pipeline being lengthy and complex, analyzing dMRI data for a particular application has become highly challenging for an end-user without extensive background in the field. In addition, numerous different processing strategies become available almost on a weekly basis without an extensive comparison to existing tools, making it difficult to choose the optimal approach for a given research question.

In this thesis, we have developed methodology and addressed fundamental questions with the aim to not further increase the confusion in diffusion MRI. We have given an overview of the dMRI pipeline (**Chapter 2**), and have presented a data evaluation framework to test new methodology (**Chapter 3**). In addition, we have developed methodology that can improve the performance of existing techniques such as DKI and SD that have already proven to be useful in various applications (**Chapters 4** and **5**), provided a way to quantitatively instead of qualitatively assess fundamental characteristics of the dMRI signal (**Chapter 6**) and brain geometry (**Chapters 7** and **8**), and developed a method to improve the visualization of dMRI tractography data (**Chapter 9**).

To improve the usability of our methods, we have designed the algorithms to make use of the underlying data; the outlier rejection with REKINDLE, recursive RF calibration, computation of the SPI, and orientation-transparency all rely on features of the data (i.e. residuals, FODF peak magnitudes, peak directions, and local or global tract directions, respectively). We aimed at avoiding ad-hoc algorithm parameters and instead introduced intuitive algorithm parameters that allow making a balanced trade-off between sensitivity and specificity (i.e. κ for REKINDLE, peak ratio threshold t for recursive calibration, scale parameters h_{max} or N for SPI computation, and opacity parameter c and dispersion threshold T_{c_l} for the orientation-dependent opacity rendering). In addition, we have extensively used simulation and real data evaluation frameworks such as the MASSIVE data to investigate the performance of our methods. In the next paragraphs, we discuss our work in the light of new developments, highlight possible improvements, and give potential directions to future developments in dMRI processing.

10.1 EVALUATION OF DIFFUSION MRI METHODOLOGY

A comprehensive framework for evaluating novel processing techniques and for an extensive comparison between such techniques is highly desired. Unfortunately, it is to date still challenging to validate dMRI and tractography. All evaluation and validation methodologies that are available have their own advantages and disadvantages. In **Chapter 3** we have presented the MASSIVE database as an in vivo real data evaluation framework that can be used to mimic acquisition scenarios in a clinical setting. With 8000 diffusion images accompanied by structural images, field maps, and noise maps, MASSIVE contains – to the best of our knowledge – the largest publicly available in vivo diffusion MRI dataset of a single subject. It complements other in vivo data repositories and databases, such as the Human Connectome Project (HCP) database which contains data of more subjects acquired with higher gradient performance but less densely sampled q -space (Van Essen et al., 2013), and the MyConnectome Project which contains repeated scans of a single subject over the course of 18 months (Poldrack et al., 2015). Real datasets and subsets thereof are paramount when

evaluating factors such as the precision, reproducibility, and stability of new data processing methodology. The drawback of real *in vivo* data is obviously that a ground-truth is lacking. Results on subsets of the data can at best be compared to the full dataset (a ‘silver standard’).

Alternatives exist to evaluate dMRI methodology. *Ex vivo* imaging of intact organs involves fewer restrictions regarding the scan time and can therefore potentially achieve a higher spatial and/or angular resolution than *in vivo* imaging. Many animal (e.g. Calabrese et al., 2014; Calabrese et al., 2015b; Lundell et al., 2011) and some human (e.g. Miller et al., 2011) specimens have been scanned with dMRI sequences and serve as useful references for *in vivo* studies. However, the diffusion properties in post mortem tissue are different than *in vivo* tissue, and a ground truth regarding the tissue architecture is still lacking.

Physical phantoms are suitable for validation of diffusion measures and fiber tractography results (Côté et al., 2013; Fillard et al., 2011; Pullens et al., 2010; Teipel et al., 2011; Walker et al., 2012). Biological phantoms such as vegetables (Boujraf et al., 2001) or animal nerve structures like rat spinal cord, garfish or lobster nerves (Beaulieu, 2002; Campbell et al., 2005) are sometimes used. However, in such organic matter it is more difficult to manipulate the natural geometry of the tissue in order to reflect more complex microstructural configurations (i.e. interdigitated crossing fibers), and the diffusion properties of such organic material may change over time. Hardware phantoms made of anisotropic media (e.g. capillaries or artificial fibers) can be tuned to resemble human white matter. Glass capillaries or PTFE (Teflon) capillaries (Lin et al., 2003) and are rigid, while for example hydrophobic fiber materials (with high FA, Lorenz et al., 2008) can be adapted to the desired geometry to create artificial fiber phantoms. Even though hardware phantoms are very valuable because of their known ground truth, they are often too simplistic.

Software phantoms are widely used because the ground truth is known. The generation of the signal is often based on the modelling of the diffusion process in a single fiber population. Crossing fibers can then be simulated by rotating the signals and taking a linear combination. Simulations can focus on generating the signal for a single voxel or for image volumes containing geometric structures that mimic fiber architecture (Leemans et al., 2005; Neher et al., 2014; Perrone et al., 2016). The knowledge to design such software phantoms comes from measurements in real data. Investigating the characteristics of single fiber populations and the geometry of the white matter is therefore important to make such simulations as realistic as possible. The findings in **Chapter 6** can be used to more realistically simulate single fiber populations. In addition, the geometry of future software phantoms could include sheet structures in line with our simulations in **Chapters 7 and 8**.

Other invasive methods are available to trace the white matter architecture and evaluate dMRI (see Axer (2010) for an overview). For example, Klingler dissections (Klingler, 1935) can be performed on fixed brains to follow single tracts of fibers, but it is difficult to map the 3D course of tracts and to investigate crossing fiber tracts. Tract tracing is based on cellular transport mechanisms between nerve cells. Manganese is a promising tracer since it can be detected with MRI (Pautler et al., 1998). Confocal laser scanning microscopy can provide high resolution 3D reconstructions of small specimens. Polarized light imaging (PLI) can selectively visualize anisotropic structures that have birefringent properties (e.g. the myelin sheaths) with polarized light (Axer et al., 2001). Slices of specimens can be scanned to recover fiber directions even down to single fiber dimensions. The resolution in the slice direction is,

however, much lower and in the order of $100 \mu\text{m}$. A further challenge is the registration between consecutive slices. CLARITY is a novel method to make brain tissue transparent and allows studying the structure of organs when combined with antibody or gene-labeling (Chung and Deisseroth, 2013).

Ideally, the previously discussed methods are combined to get a more complete picture of the performance of a new method. In this thesis, we used software phantoms to evaluate the outcome of our methods in relation to the simulated ground truth, and we often tested both on in vivo and ex vivo data to show their general applicability. In future work, the methods proposed in this thesis can be evaluated in a wider range of possible scenarios, which will be discussed into more detail for each method in the following sections.

10.2 MINIMIZING THE ADVERSE EFFECTS OF ARTIFACTS IN DIFFUSION MRI

It is of major importance to check data quality after acquisition and correct for artifacts during acquisition and image processing to ensure the reliability of all subsequent analyses. Artifacts can have serious implications for further analysis such as quantification in a multi-subject comparison and tractography. Quantitative measures can be compared between different groups with ROI-based analyses and voxel-based analyses (VBA) to reveal associations with clinical parameters, which has been the focus of a large amount of studies. Artifacts can affect the diffusion measures, as shown in Fig. 10.1a, where FA values (top two rows), and MD values (bottom row) are locally altered. Furthermore, artifacts can complicate proper analyses in particular areas. It is well-known that structures in the orbitofrontal cortex, for example, are prone to susceptibility artifacts, which makes these white matter areas less accessible to study. This might attribute to the fact that larger white matter tracts with densely packed neuronal fibers tend to be studied more than less prominent pathways. For example, the corpus callosum (CC) is a pathway that can readily be identified, which makes this pathway better suited for investigation in quantitative studies (White et al., 2008). One should be aware, however, that structures such as the CC can also be corrupted by artifacts, such as susceptibility distortions, Gibbs ringing, and interleave artifacts. Several artifacts cause difficulties in the registration of individual DWIs, which will eventually affect any subsequent analysis.

In addition to affecting quantitative measures, data artifacts and lack of proper correction tools can severely confound local fiber configurations and therefore fiber tractography results. Fig. 10.1b shows examples of the influence of bad quality data on tractography results using a deterministic algorithm, for different artifacts. The results clearly show the deviation in the reconstructed pathways when proper correction methods are not taken into account. Tracts can have a different geometry, and may even terminate in other brain areas.

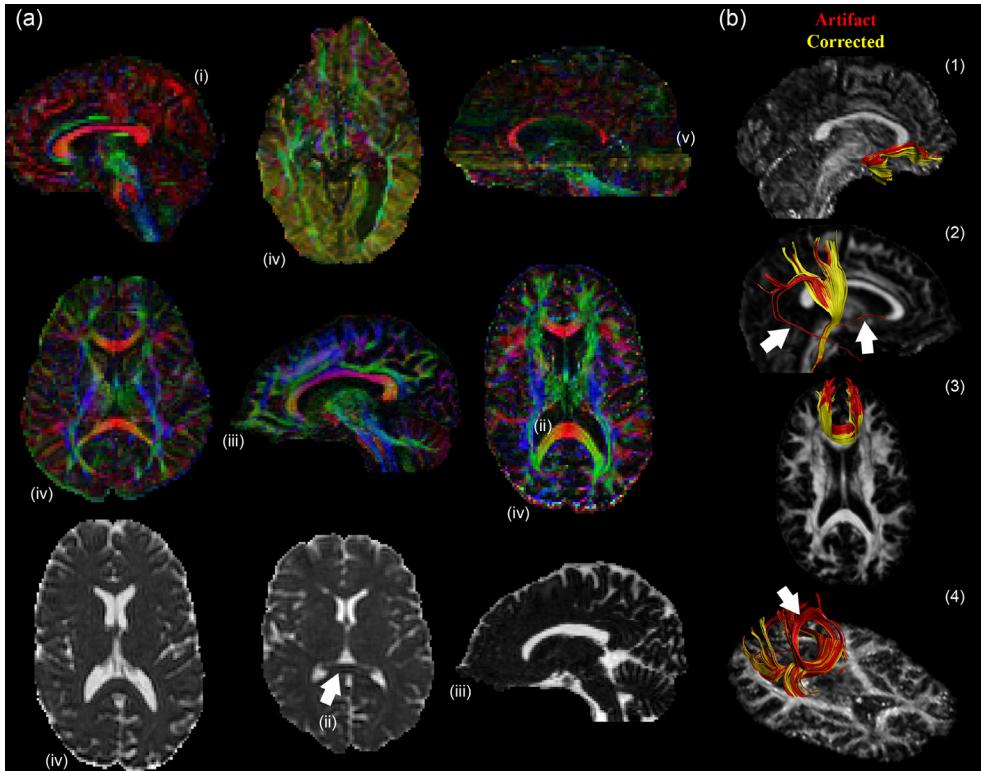


Fig. 10.1: (a) Effects of bad quality data on quantitative measures (FA top two rows, MD bottom row), showing (i) vibration artifact, (ii) Gibbs ringing, (iii) Susceptibility distortions anterior, (iv) Motion and eddy current distortions, (v) interleave artifacts. (b) Effects of bad quality data on tractography results. (1) Tractography of the uncinate fasciculus before (red) and after (yellow) motion and eddy current distortion. (2) Corticospinal tract in a set where we simulated an interleave artifact (red) and in the 'ground truth' set (yellow). The artifact results in clearly artificial tracts. (3) Frontal projections of the corpus callosum in a dataset with susceptibility distortions (red) and the corrected dataset (yellow). (4) Occipital projections of the corpus callosum. Artifactual parietal fibers connect the two hemispheres due to a vibration artifact (red), whereas this is not the case when the artifact is corrected (yellow).

Several of the artifacts discussed in **Chapter 2** can be corrected for at the processing stage. As such, they do not limit the analyses of dMRI data per se, but rather force the user to consciously consider the processing steps prior to the analyses. Most software packages available to date provide users with adequate options to do these correction steps. Robust estimation procedures can be used to further reduce the influence of defective images on model estimates. In **Chapter 4**, REKINDLE was proposed as a robust estimation framework to reduce the effect of outliers. REKINDLE has successfully been applied in a range of studies (e.g. Kennis et al., 2016; Odish et al., 2015a; Odish et al., 2015b; Roze et al., 2015). A recent comparison of REKINDLE with other DKI fitting using different methods showed that REKINDLE performed the best in terms of accuracy for a variety of metrics (MD, FA, MK, AK, RK) at high SNR, and for the majority of the metrics at lower SNR (Olson et al., 2015). In addition to the methods compared in their work (unconstrained singular value distribution (SVD), directional weighting and regularization (DWAR (Kuder et al., 2012)), and estimation with a sparsity constraint (Tabesh et al., 2011b), and the ones compared in our work, other robust estimation procedures have been recently proposed. Tobisch et al. (2016) combine the SHORE basis with iterative reweighting based on

the Geman-McClure M-estimator and l_1 regularization. They find in simulation studies that REKINDLE performs better in terms of sensitivity and specificity for an outlier percentage up to 10% in a crossing fiber configuration. They find an improved performance of their algorithm in case of a low number of acquired DW images, likely due to the sparsifying SHORE basis and L1-regularized fit. Collier et al. (2015) propose a linear iteratively reweighted linear least squares framework somewhat similar to REKINDLE. Differences include for example the reweighting for heteroscedasticity correction and reweighting based on the Geman-McClure function: this is not performed separately (i.e. the different loops in the REKINDLE algorithm), but in one step. Specifically, they compute the standardized residual in Eq. [4.13] by an estimate of the standard deviation based on the estimated signal and the MAD, rather than based on the MAD alone (Eq. [4.14]). While this may be faster, errors can be introduced if the estimated signal is influenced by outliers itself. To this end, we have introduced the two-step procedure to try to disentangle this effect. A close comparison of these methods is not performed thus far and is subject to future work.

Dealing with outliers is a highly challenging problem and it consists of different steps that can all be approached differently: 1) choice of model or representation, 2) performing a goodness of fit check to identify voxels which likely have outliers, 3) the reweighting process and reweighting function, 4) the outlier detection criterion, 5) rejection of outliers, 6) final fit. REKINDLE uses DKI as model for detection of outliers, and can be directly extended to use DTI. Other works use for example spherical harmonics (Pannek et al., 2012), or the SHORE basis (Tobisch et al., 2016). The requirement for the model is that it can adequately be fit by the data available (e.g. single shell vs multi shell). REKINDLE does not perform an a priori goodness of fit check, but this can be included if one has a reliable estimate signal variance a priori. The goodness of fit can be evaluated in multiple ways, e.g. checking whether the residuals are not larger than 3 times the signal variance (Chang et al., 2005). The reweighting is generally based on the residuals in the previous iteration, which can either be in linear- or nonlinear space for models that can be linearized. Most works use the Geman-McClure function to assign weights to different data points, but other weighting functions could be investigated in future work. The outlier detection criterion is mostly based on an estimate of the standard deviation of the error term. REKINDLE introduces the κ parameter here to balance between sensitivity and specificity of the outlier detection, but also an absolute threshold on the residuals can be used (Tobisch et al., 2016). An appropriate value for parameter κ can be determined by performing simulation experiments as in Chapter 4, or by visually inspecting outlier- or DKI feature maps. Future work can be directed towards investigating the influence of κ in a wider range of datasets, and how its optimal value is dependent on SNR and acquisition settings. When points are identified as outliers, they can either be rejected or replaced by predicted values before the final fit. The latter approach has shown to be beneficial in several works (e.g. Sharman et al., 2011; Tobisch et al., 2016). We have tested this approach in combination with REKINDLE and DTI, where we replaced the rejected points by an estimate of the robustly fitted DT, to be able to perform CSD afterwards. This approach scored well in the ISMRM Tractography challenge (results not shown here (Neher et al., 2016)). The final dMRI model or fitting approach can be chosen independently from the outlier detection procedure.

Outlier detection methods should be carefully evaluated in cases where the ground truth is known. For example, the REKINDLE framework was tested on simulated outliers involving an artificial signal increase or decrease of 50%. However, a signal perturbation can be defined in

different ways, and artifacts might cause various types of signal perturbations. These issues can be further addressed by using evaluation frameworks that simulate outliers arising from different types of artifacts (Graham et al., 2015; Perrone et al., 2016). In real data, robust methods have been proven useful to cope with a range of artifacts, including cardiac pulsation (Chang et al., 2005; Chang et al., 2012), artifacts in neonatal data (Collier et al., 2015), vibration artifacts (Scherrer and Warfield, 2012), subject motion, inter-slice instabilities, and other dropouts (Tax et al., 2015c). In the case of vibration artifacts, however, robust procedures do not seem to help in all cases (Berl et al., 2015). In the case of corruption of a large part of the image, e.g. due to interleaved artifacts, misregistration of the images might hide some of the outliers (Tax et al., 2016c). Even though voxels are not classified as outliers, they are potentially also unreliable because they provide diffusion information about different spatial locations. Misregistration of voxels within the WM might not provide a strong enough contrast difference to be classified as artifact. Even though not the whole volume is classified as an outlier, looking at the percentage outliers per DW image might still reveal this issue, and the volume can be removed for further analysis. In any case, the user should always remain critical when employing correction methods.

Even though there are various ways to reduce the influence of artifacts, it is of major importance to check quality requirements before acquiring data on clinical or research subjects, by using quality assurance (QA) tests. QA is concerned with the implementation of activities to fulfill quality requirements, such as comparison to a gold standard. Standard QA tests consist of gradient calibration (including linearity, uniformity, and amplitude), field mapping to minimize B0 inhomogeneities and eddy current compensation (De Santis et al., 2012a). This is most commonly done by imaging phantoms with different gradient directions and b -values, as they are suitable for validation of acquisition parameters. Hardware phantoms made of isotropic media (e.g. liquids of known diffusivity such as Dodecane) can for example be used to calibrate absolute gradient power.

10.3 MODELING AND PARAMETER ESTIMATION IN DIFFUSION MRI

While diffusion weighted imaging (DWI) is well established for some clinical applications, DTI has yet to be widely accepted in clinical practice. At the same time, the research community is already debating whether “DTI is ready for retirement” as we know that DTI has limitations in the case of crossing fibers or non-Gaussian effects. An incredible amount of alternative models and representations are being ‘prepared’ to take its place (**Chapter 2** and Fig. 10.2). However, the ever growing amount of options to model the diffusion signal comes at the expense of being able to extensively compare them. Of the available methods, DKI and SD are particularly appealing because 1) their acquisition protocol is not as demanding as for some other dMRI techniques (e.g. DSI, Wedeen et al. (2005)), 2) various works have investigated how to reliably fit these models, 3) they have no or few assumptions on the underlying tissue structure, 4) they allow for the extraction of scalar measures beyond DTI measures and fiber direction estimates in the case of crossing fibers, 5) many software packages support them. Several studies have related changes in DKI or SD features to clinical parameters. The AFD derived from SD highlighted changes in the case of Motor Neurone Disease (Raffelt et al., 2012), while DKI has shown to provide valuable information in case of stroke (Jensen et al., 2011), attention-deficit

hyperactivity disorder (ADHD) (Helpern et al., 2011), gliomas (Van Cauter et al., 2012), epilepsy (Gao et al., 2012), and multiple sclerosis (Raz et al., 2013), amongst others. The increasing interest in these techniques is furthermore confirmed by the recent efforts to relate their measures to microstructural specific features (De Santis et al., 2012b; Fieremans et al., 2011; Hui et al., 2015), and to obtain them in a clinically feasible scan time (Hansen et al., 2013; Hansen et al., 2015).

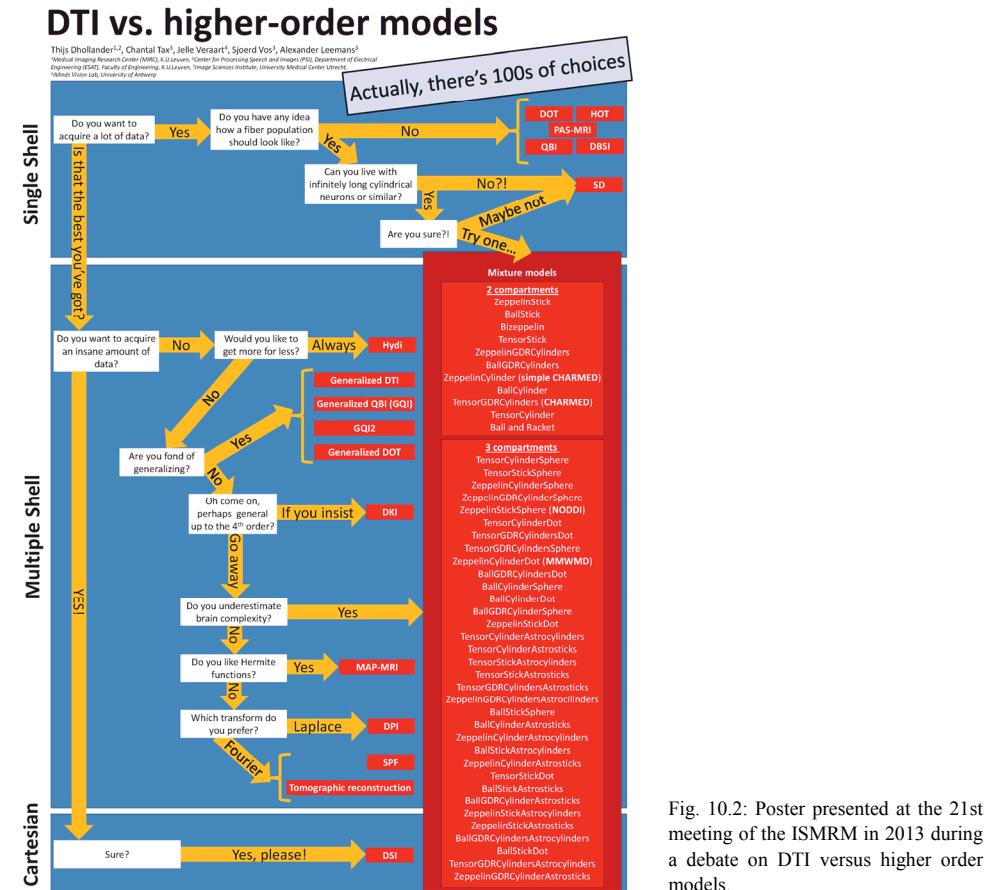


Fig. 10.2: Poster presented at the 21st meeting of the ISMRM in 2013 during a debate on DTI versus higher order models.

In **Chapters 4** and **5** we have proposed methodological improvements in the estimation of DKI (also discussed in the previous section) and SD. **Chapter 6** further investigates which voxels are selected as SFP voxels with the recursive calibration from Chapter 5. An important question that remains unanswered is how the RF should be calibrated in group studies. For quantitative analysis of the AFD the RF is generally chosen constant across groups, but it is unclear if this is optimal for tractography studies. An extensive investigation of the influence of RF shape on angular deviations of the fODF is therefore subject of ongoing work (Guo et al., 2016). Estimation in the context of SD is a fast moving field. In addition to the techniques described in Chapter 2 that extended the estimation fODF to multi-tissue and multi-shell by adapting the RF accordingly, recent works try to calibrate the RF in each voxel or per tissue type automatically. Anderson (2005) already estimated a different RF for each voxel by equating the MD of the

SFP to the MD of the whole voxel. However, it is known that the MD is lower in the case of crossing fibers (Vos et al., 2012), and this approach showed a reduced ability to solve crossing fibers. Schultz and Groeschel (2013) select the axially symmetric tensor kernel that reduces spurious peaks through an extra sparsity constraint. Such methods have the potential to reflect spatial variability of SFP properties in the RF, thereby extending ideas of Chapter 6. However, by design it is not possible to match the RF to properties of individual populations in a crossing fiber configuration, as only one kernel per voxel can be defined. Christiaens et al. (2015) and Jeurissen et al. (2015) use non-negative matrix factorization, which is a popular data-driven approach to find the hidden constituent parts of the data, to segment tissue into WM, GM, and CSF based on dMRI data only, and to compute a tissue-specific RF.

Other exciting new opportunities and challenges in the area of modelling and estimation include the characterization of time-dependent diffusion, the usage of other diffusion sequences beyond the ‘conventional’ single diffusion encoding, and the combination of diffusion measurements with MRI techniques beyond dMRI. It becomes increasingly apparent that a standard DTI sequence – which samples q -space on a single shell and thus is 5D (three dimensions for position and two for orientation) – gives an incomplete part of the picture. Multi-shell data (6D), generally obtained by varying the gradient strength while keeping the diffusion time constant, also provides information on the signal decay and can be used to derive the EAP and its features. When also the diffusion time is varied this adds yet another dimension. Studying time-dependent diffusion can provide new information on the fiber geometry at the micrometer scale (e.g. Baron and Beaulieu, 2014; Burcaw et al., 2015; De Santis et al., 2016b; Fick et al., 2015; Fieremans et al., 2016; Novikov et al., 2014). Advanced diffusion encoding methods, including multiple pairs of pulsed field gradients (e.g. Mitra, 1995; Ozarslan and Basser, 2008), oscillating gradients (e.g. Does et al., 2003), and q -space trajectory imaging (Westin et al., 2016), can potentially improve tissue characterization. In addition, combining dMRI with other microstructural techniques such as susceptibility imaging (Kleinnijenhuis et al., 2012), and relaxometry (De Santis et al., 2016a), might reveal additional and complementary information of the tissue microstructure. With these new measurements, the development of new modelling techniques and targeted estimation procedures is necessary.

10.4 INVESTIGATING THE BRAIN’S GEOMETRICAL ORGANIZATION WITH DIFFUSION MRI AND TRACTOGRAPHY

In addition to describing a structure or tract in terms of various diffusion measures as done in **Chapter 6**, its geometry may hold important additional information. Structures have a characteristic shape that varies across the healthy population and can be indicative of neurodevelopmental or neurodegenerative disorders. Examples of geometrical properties of tracts include length, curvature, torsion, dispersion, and orientation (Batchelor et al., 2006; Leemans et al., 2006; Savadjiev et al., 2012; Tax et al., 2012). Such features have for example been shown to correlate with clinical markers and genetic risk for schizophrenia (Savadjiev et al., 2016). Parker et al. (2016) furthermore showed that white matter tract shapes uniquely identifies individuals, also called structural fingerprinting. Shape features can also be used for the clustering of tracts.

In **Chapters 7 and 8** we investigate another geometrical feature of the brain: the occurrence sheet structure. Such structures only occur if the condition of zero Lie bracket component is fulfilled. Features such as the extent, orientation, and curvature of sheet structures can serve as new geometrical features to ‘fingerprint’ the brain. Qualitatively, the same high SPI areas could be found across individuals. In **Chapter 9** we also use geometrical features of tracts for visualization: their (local and global) orientation, and a measure for global tract dispersion. These features could for example be included in tract clustering or fingerprinting approaches.

Clear visualization of brain structures is highly challenging. In **Chapters 7 and 8** we propose the sheet tensor as a means to visualize the SPI. SPI-scalar maps are less informative as only the SPI of one pair of peaks can be shown (or a summary statistic of all SPI’s in a voxel), and the information on the orientation of the local surface is lost. While the sheet tensor provides useful information on a voxel-wise basis, the global 3D structure is hard to derive from such images. A surface propagation process – similar to the propagation of local peaks in tractography – will improve the investigation of such sheet-like structures in future work. In **Chapter 9** we have proposed a visualization method for improved exploration of tissue configurations that would otherwise be largely covered by other pathways. This approach was used to efficiently extract fiber bundles and provide a less cluttered visualization in a neurosurgical planning case. In our own experience, the method is very useful to quickly and interactively highlight different structures without the need for manual removal of streamlines that obscure others. This approach further facilitates the investigation of a potentially new fiber pathway (Heemskerk et al., 2015). In future work, the usefulness of this approach should be more extensively evaluated in a clinical context.

With dMRI we do not measure fibers or fiber bundles directly, but instead infer information on the maximum directions of diffusion at discrete positions in the brain that are assumed to coincide with the underlying fiber directions. As such, ‘vector fields’ in the brain do not actually exist (Chapters 7 and 8), but are derived from the data that is measured. dMRI tractography and methods based on tractography rely on directional data, and can be severely affected by factors such as noise and erroneous tract propagation. One famous problem is the crossing vs. kissing configuration: based on orientations alone such situations cannot be disentangled and the algorithm will generally reconstruct a crossing configuration by design. Tractography thus can easily generate false-positive tracts that do not correspond to histology. The extent of this problem was recently emphasized by Descoteaux et al. (2016). Microstructure-informed tractography methods, which rely on the assumption that microstructural features such as axon diameter are relatively constant along a tract, can provide necessary information to propagate or evaluate tracts beyond their orientation alone (Daducci et al., 2015b; Girard et al., 2015). Such methods could also be integrated into the computation of the SPI, where the clustering of peaks is currently only based on orientation. Microstructure-informed tractography is, however, not without challenges: it is for example hard to reliably fit microstructural models in crossing fiber configurations (Daducci et al., 2016; Chapter 6). In any case, validation of tractography and other derivatives of local peak directions (e.g. sheet structures) is crucial.

10.5 CONCLUSIONS

Overall, diffusion MRI is a rapidly evolving field with a lot of potential in clinical and biomedical research and applications. However, the dMRI processing pipeline is lengthy and complex, with numerous approaches available to tackle every step. While the development of new processing approaches is of paramount importance to improve the reliability of dMRI, we should not head in a direction where dMRI is only understandable and approachable to the methods developers themselves. Publically available evaluation frameworks, well-substantiated and evaluated algorithms, and intuitive guidelines are essential to aid the translation of dMRI processing methods into useful tools for applications.

REFERENCES

- Ackerman, J.J.H., Neil, J.J., 2010. Biophysics of Diffusion in Cells. In: Jones, D.K. (Ed.), *Diffusion MRI: Theory, Methods, and Applications*. Oxford University Press.
- Aganj, I., Lenglet, C., Sapiro, G., Yacoub, E., Ugurbil, K., Harel, N., 2010. Reconstruction of the orientation distribution function in single- and multiple-shell q-ball imaging within constant solid angle. *Magn.Reson.Med.* 64, 554-566.
- Aja-Fernandez, S., Alberola-Lopez, C., Westin, C.F., 2008. Noise and signal estimation in magnitude MRI and Rician distributed images: a LMMSE approach. *IEEE Trans.Image Process.* 17, 1383-1398.
- Aja-Fernandez, S., Pieciak, T., Vegas-Sanchez-Ferrero, G., 2015. Spatially variant noise estimation in MRI: a homomorphic approach. *Med.Image.Anal.* 20, 184-197.
- Aja-Fernandez, S., Tristan-Vega, A., 2012. Influence of noise correlation in multiple-coil statistical models with sum of squares reconstruction. *Magn.Reson.Med.* 67, 580-585.
- Aja-Fernandez, S., Tristan-Vega, A., Alberola-Lopez, C., 2009. Noise estimation in single- and multiple-coil magnetic resonance data based on statistical models. *Magn Reson.Imaging* 27, 1397-1409.
- Aja-Fernández, S., Tristán-Vega, A., 2013. A review on statistical noise models for Magnetic Resonance Imaging. LPI, ETSI Telecomunicacion, Universidad de Valladolid, Spain, Tech.Rep.
- Alexander, A.L., Hasan, K.M., Lazar, M., Tsuruda, J.S., Parker, D.L., 2001. Analysis of partial volume effects in diffusion-tensor MRI. *Magn.Reson.Med.* 45, 770-780.
- Alexander, D.C., 2005. Maximum entropy spherical deconvolution for diffusion MRI. *Inf.Process Med Imaging* 19, 76-87.
- Alexander, D.C., 2008. A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features. *Magn.Reson.Med.* 60, 439-448.
- Alexander, D.C., Hubbard, P.L., Hall, M.G., Moore, E.A., Ptito, M., Parker, G.J., Dyrby, T.B., 2010. Orientationally invariant indices of axon diameter and density from diffusion MRI. *Neuroimage* 52, 1374-1389.
- Anderson, A.W., 2001. Theoretical analysis of the effects of noise on diffusion tensor imaging. *Magn.Reson.Med.* 46, 1174-1188.
- Anderson, A.W., 2005. Measurement of fiber orientation distributions using high angular resolution diffusion imaging. *Magn.Reson.Med.* 54, 1194-1206.
- Andersson, J.L., 2008. Maximum a posteriori estimation of diffusion tensor parameters using a Rician noise model: why, how and but. *NeuroImage* 42, 1340-1356.
- Andersson, J.L., Richter, M., Richter, W., Skare, S., Nunes, R.G., Robson, M.D., Behrens, T.E., 2004. Effects of susceptibility distortions on tractography. In Proceedings of the 12th Annual Meeting of the ISMRM
- Andersson, J.L., Skare, S., 2002. A model-based method for retrospective correction of geometric distortions in diffusion-weighted EPI. *NeuroImage* 16, 177-199.
- Andersson, J.L., Skare, S., Ashburner, J., 2003. How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *NeuroImage* 20, 870-888.
- Andersson, J.L., Sotropoulos, S.N., 2016. An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging. *NeuroImage* 125, 1063-1078.
- Andersson, J.L.R., Skare, S., 2010. Image Distortion and Its Correction in Diffusion MRI. In: Jones, D.K. (Ed.), *Diffusion MRI: Theory, Methods, and Applications*. Oxford University Press.
- Ankele, M., Schultz, T., 2015. Quantifying Microstructure in Fiber Crossings with Diffusional Kurtosis. *Med.Image Comput.Comput.Assist.Interv.* 18, 150-157.
- Archibald, R., Gelb, A., 2002. A method to reduce the Gibbs ringing artifact in MRI scans while keeping tissue boundary integrity. *IEEE Trans.Med.Imaging* 21, 305-319.
- Ashburner, J., Friston, K.J., 2005. Unified segmentation. *NeuroImage* 26, 839-851.
- Assaf, Y., Basser, P.J., 2005. Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain. *NeuroImage* 27, 48-58.
- Assaf, Y., Blumenfeld-Katzir, T., Yovel, Y., Basser, P.J., 2008. AxCaliber: a method for measuring axon diameter distribution from diffusion MRI. *Magn.Reson.Med.* 59, 1347-1354.
- Assaf, Y., Freidlin, R.Z., Rohde, G.K., Basser, P.J., 2004. New modeling and experimental framework to characterize hindered and restricted water diffusion in brain white matter. *Magn.Reson.Med.* 52, 965-978.
- Assaf, Y., Mayk, A., Cohen, Y., 2000. Displacement imaging of spinal cord using q-space diffusion-weighted MRI. *Magn.Reson.Med.* 44, 713-722.
- Assaf, Y., Pasternak, O., 2008. Diffusion tensor imaging (DTI)-based white matter mapping in brain research: a review. *J.Mol.Neurosci.* 34, 51-61.

References

- Assemlal, H.E., Tschumperle, D., Brun, L., 2009. Efficient and robust computation of PDF features from diffusion MR signal. *Med.Image Anal.* 13, 715-729.
- Assemlal, H.E., Tschumperle, D., Brun, L., Siddiqi, K., 2011. Recent advances in diffusion MRI modeling: Angular and radial reconstruction. *Med.Image Anal.* 15, 369-396.
- Astola, L., Fuster, A., Florack, L., 2011. A Riemannian scalar measure for diffusion tensor images. *Pattern Recognition* 44, 1885-1891.
- Atkinson, D., Batchelor, P.G., Clark, C.A., 2008. Track ribbons--visualising structural information in diffusion tensor axial asymmetry. In Proceedings of the 15th Annual Meeting of the ISMRM 79.
- Auria, A., Daducci, A., Thiran, J.P., Wiaux, Y., 2015. Structured sparsity for spatially coherent fibre orientation estimation in diffusion MRI. *NeuroImage*. 115, 245-255.
- Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C., 2008. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Med.Image Anal.* 12, 26-41.
- Avram, A.V., Sarlls, J.E., Barnett, A.S., Özarslan, E., Thomas, C., Irfanoglu, M.O., Hutchinson, E., Pierpaoli, C., Basser, P.J., 2016. Clinical feasibility of using mean apparent propagator (MAP) MRI to characterize brain tissue microstructure. *NeuroImage*. 127, 422-434.
- Axer, H., 2010. Invasive methods for tracing white matter architecture. In: Jones, D.K. (Ed.), *Diffusion MRI: Theory, Methods, and Applications*. Oxford University Press.
- Axer, H., Axer, M., Krings, T., Keyserlingk, D.G., 2001. Quantitative estimation of 3-D fiber course in gross histological sections of the human brain using polarized light. *J.Neurosci.Methods* 105, 121-131.
- Bach, M., Fritzsche, K.H., Stieltjes, B., Laun, F.B., 2014. Investigation of resolution effects using a specialized diffusion tensor phantom. *Magn.Reson.Med.* 71, 1108-1116.
- Baete, S.H., Yutzy, S., Boada, F.E., 2015. Radial q-space sampling for DSI. *Magn.Reson.Med.* doi: 10.1002/mrm.25917.
- Bai, Y., Alexander, D.C., 2008. Model-based registration to correct for motion between acquisitions in diffusion MR imaging. *IEEE Int.Symp.Biomed.Imaging*, 947-950.
- Bakir, T., Reeves, S.J., 2000. A filter design method for minimizing ringing in a region of interest in MR spectroscopic images. *IEEE Trans.Med.Imaging* 19, 585-600.
- Bao, L., Robini, M., Liu, W., Zhu, Y., 2013. Structure-adaptive sparse denoising for diffusion-tensor MRI. *Med.Image Anal.* 17, 442-457.
- Barker, G.J., Parker, G.J.M., Wheeler-Kingshott, C.A., 2001. Gibbs ringing and negative ADC values. In Proceedings of the 9th Annual Meeting of the ISMRM, 1546.
- Barmoutis, A., Vemuri, B.C., Howland, D., Forder, J.R., 2008. Extracting tractosemas from a displacement probability field for tractography in DW-MRI. *Med.Image Comput.Comput.Assist.Interv.* 11, 9-16.
- Barnea-Goraly, N., Kwon, H., Menon, V., Eliez, S., Lotspeich, L., Reiss, A.L., 2004. White matter structure in autism: preliminary evidence from diffusion tensor imaging. *Biol.Psychiatry* 55, 323-326.
- Barnett, A.S., Hutchinson, E., Irfanoglu, M.O., Pierpaoli, C., 2014. Higher order correction of eddy current distortion in diffusion weighted echo planar images. *Proceedings of the 22nd Annual Meeting of the ISMRM*, 4463.
- Baron, C.A., Beaulieu, C., 2014. Oscillating gradient spin-echo (OGSE) diffusion tensor imaging of the human brain. *Magn.Reson.Med.* 72, 726-736.
- Baron, C.A., Lebel, R.M., Wilman, A.H., Beaulieu, C., 2012. The effect of concomitant gradient fields on diffusion tensor imaging. *Magn.Reson.Med.* 68, 1190-1201.
- Basser, P.J., 1995. Inferring microstructural features and the physiological state of tissues from diffusion-weighted images. *NMR Biomed.* 8, 333-344.
- Basser, P.J., Jones, D.K., 2002. Diffusion-tensor MRI: theory, experimental design and data analysis - a technical review. *NMR Biomed.* 15, 456-467.
- Basser, P.J., Mattiello, J., Le Bihan, D., 1994. Estimation of the effective self-diffusion tensor from the NMR spin echo. *J.Magn Reson.B* 103, 247-254.
- Basser, P.J., Pajevic, S., 2000. Statistical artifacts in diffusion tensor MRI (DT-MRI) caused by background noise. *Magn.Reson.Med.* 44, 41-50.
- Basser, P.J., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, A., 2000. In vivo fiber tractography using DT-MRI data. *Magn.Reson.Med.* 44, 625-632.
- Bastiani, M., Cottaar, M., Dikranian, K., Ghosh, A., Zhang, H., Alexander, D.C., Behrens, T., Jbabdi, S., Sotiroopoulos, S., 2016. Improved tractography by modelling sub-voxel fibre patterns using asymmetric fibre orientation distributions. In Proceedings of the 24th Annual Meeting of the ISMRM, 0010.
- Bastin, M.E., 2001. On the use of the FLAIR technique to improve the correction of eddy current induced artefacts in MR diffusion tensor imaging. *Magn Reson.Imaging* 19, 937-950.
- Batchelor, P.G., Calamante, F., Tournier, J.D., Atkinson, D., Hill, D.L., Connelly, A., 2006. Quantification of the shape of fiber tracts. *Magn.Reson.Med.* 55, 894-903.

Less Confusion in Diffusion MRI

- Beaulieu, C., 2002. The basis of anisotropic water diffusion in the nervous system - a technical review. *NMR Biomed.* 15, 435-455.
- Becker, S.M., Tabelow, K., Mohammadi, S., Weiskopf, N., Polzehl, J., 2014. Adaptive smoothing of multi-shell diffusion weighted magnetic resonance data by msPOAS. *NeuroImage*. 95, 90-105.
- Behrens, T.E., Woolrich, M.W., Jenkinson, M., Johansen-Berg, H., Nunes, R.G., Clare, S., Matthews, P.M., Brady, J.M., Smith, S.M., 2003. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn.Reson.Med.* 50, 1077-1088.
- Ben-Amitay, S., Jones, D.K., Assaf, Y., 2012. Motion correction and registration of high b-value diffusion weighted images. *Magn.Reson.Med.* 67, 1694-1702.
- Benner, T., van der Kouwe, A.J., Kirsch, J.E., Sorensen, A.G., 2006. Real-time RF pulse adjustment for B0 drift correction. *Magn.Reson.Med.* 56, 204-209.
- Bennett, K.M., Schmainka, K.M., Bennett, R.T., Rowe, D.B., Lu, H., Hyde, J.S., 2003. Characterization of continuously distributed cortical water diffusion rates with a stretched-exponential model. *Magn.Reson.Med.* 50, 727-734.
- Berl, M.M., Walker, L., Modi, P., Irfanoglu, M.O., Sarlls, J.E., Nayak, A., Pierpaoli, C., 2015. Investigation of vibration-induced artifact in clinical diffusion-weighted imaging of pediatric subjects. *Hum.Brain Mapp.* 36, 4745-4757.
- Bernier, M., Chamberland, M., Houde, J.C., Descoteaux, M., Whittingstall, K., 2014. Using fMRI non-local means denoising to uncover activation in sub-cortical structures at 1.5 T for guided HARDI tractography. *Front Hum.Neurosci.* 8, 715.
- Bickel, P.J., 1978. Using Residuals Robustly I: Tests for Heteroscedasticity, Nonlinearity. *The Annals of Statistics* 6, 266-291.
- Bodammer, N., Kaufmann, J., Kanowski, M., Tempelmann, C., 2004. Eddy current correction in diffusion-weighted imaging using pairs of images acquired with opposite diffusion gradient polarity. *Magn.Reson.Med.* 51, 188-193.
- Boujraf, S., Luypaert, R., Eisendrath, H., Osteaux, M., 2001. Echo planar magnetic resonance imaging of anisotropic diffusion in asparagus stems. *MAGMA*. 13, 82-90.
- Bowtell, R.W., McIntrye, D.J.O., Commandre, M.J., Glover, P.M., Mansfield, P., 1994. Correction of geometric distortions in echo planar images. In Proceedings of the 2nd Annual Meeting of the SMR, 411.
- Brion, V., Poupon, C., Riff, O., Aja-Fernandez, S., Tristan-Vega, A., Mangin, J.F., Le Bihan, D., Poupon, F., 2013. Noise correction for HARDI and HYDI data obtained with multi-channel coils and sum of squares reconstruction: an anisotropic extension of the LMMSE. *Magn.Reson.Imaging* 31, 1360-1371.
- Buades, A., Coll, B., Morel, J.M., 2005. A review of image denoising algorithms, with a new one. *Multiscale Modeling & Simulation* 4, 490-530.
- Buonocore, M.H., Gao, L., 1997. Ghost artifact reduction for echo planar imaging using image phase correction. *Magn.Reson.Med.* 38, 89-100.
- Burcaw, L.M., Fieremans, E., Novikov, D.S., 2015. Mesoscopic structure of neuronal tracts from time-dependent diffusion. *NeuroImage*. 114, 18-37.
- Calabrese, E., Badea, A., Coe, C.L., Lubach, G.R., Shi, Y., Styner, M.A., Johnson, G.A., 2015a. A diffusion tensor MRI atlas of the postmortem rhesus macaque brain. *NeuroImage*. 117, 408-416.
- Calabrese, E., Badea, A., Coe, C.L., Lubach, G.R., Styner, M.A., Johnson, G.A., 2014. Investigating the tradeoffs between spatial resolution and diffusion sampling for brain mapping with diffusion tractography: time well spent? *Hum.Brain Mapp.* 35, 5667-5685.
- Calabrese, E., Badea, A., Cofer, G., Qi, Y., Johnson, G.A., 2015b. A Diffusion MRI Tractography Connectome of the Mouse Brain and Comparison with Neuronal Tracer Data. *Cereb.Cortex*. 25, 4628-4637.
- Calabrese, E., Hickey, P., Hulette, C., Zhang, J., Parente, B., Lad, S.P., Johnson, G.A., 2015c. Postmortem diffusion MRI of the human brainstem and thalamus for deep brain stimulator electrode localization. *Hum.Brain Mapp.* 36, 3167-3178.
- Calamante, F., Tournier, J.D., Jackson, G.D., Connelly, A., 2010. Track-density imaging (TDI): super-resolution white matter imaging using whole-brain track-density mapping. *NeuroImage*. 53, 1233-1243.
- Callaghan, P.T., MacGowan, D., Packer, K.J., Zelaya, F.O., 1990. High-resolution q-space imaging in porous structures. *Journal of Magnetic Resonance (1969)* 90, 177-182.
- Campbell, J.S., Siddiqi, K., Rymar, V.V., Sadikot, A.F., Pike, G.B., 2005. Flow-based fiber tracking with diffusion tensor and q-ball data: validation and comparison to principal diffusion direction techniques. *NeuroImage*. 27, 725-736.
- Canales-Rodriguez, E.J., Daducci, A., Sotiroopoulos, S.N., Caruyer, E., Aja-Fernandez, S., Radua, J., Yurramendi Mendizabal, J.M., Iturria-Medina, Y., Melie-Garcia, L., Aleman-Gomez, Y., Thiran, J.P., Sarro, S., Pomarol-Clootet, E., Salvador, R., 2015. Spherical Deconvolution of Multichannel Diffusion MRI Data with Non-Gaussian Noise Models and Spatial Regularization. *PLoS.One*. 10, e0138910.
- Canales-Rodriguez, E.J., Iturria-Medina, Y., Aleman-Gomez, Y., Melie-Garcia, L., 2010a. Deconvolution in diffusion spectrum imaging. *NeuroImage*. 50, 136-149.

References

- Canales-Rodriguez, E.J., Lin, C.P., Iturria-Medina, Y., Yeh, C.H., Cho, K.H., Melie-Garcia, L., 2010b. Diffusion orientation transform revisited. *NeuroImage*. 49, 1326-1339.
- Carballido, A., Amico, F., Ugwu, I., Fagan, A.J., Fahey, C., Morris, D., Meaney, J.F., Leemans, A., Frodl, T., 2012. Reduced fractional anisotropy in the uncinate fasciculus in patients with major depression carrying the met-allele of the Val66Met brain-derived neurotrophic factor genotype. *Am.J.Med.Genet.B Neuropsychiatr.Genet.* 159B, 537-548.
- Carroll, R.J., Ruppert, D., 1982. Robust Estimation in Heteroscedastic Linear Models. *The Annals of Statistics* 10, 429-441.
- Caruyer, E., Deriche, R., 2012. Diffusion MRI signal reconstruction with continuity constraint and optimal regularization. *Med.Image Anal.* 16, 1113-1120.
- Caruyer, E., Lenglet, C., Sapiro, G., Deriche, R., 2013. Design of multishell sampling schemes with uniform coverage in diffusion MRI. *Magn.Reson.Med.* 69, 1534-1540.
- Catani, M., Bodi, I., Dell'acqua, F., 2012. Comment on "The geometric structure of the brain fiber pathways". *Science* 337, 1605.
- Cetin, S., Özarslan, E., Ünal, G.B., 2015. Elucidating Intravoxel Geometry in Diffusion-MRI: Asymmetric Orientation Distribution Functions (AODFs) Revealed by a Cone Model. *Med.Image Comput.Comput.Assist.Interv.* 18, 231-238.
- Chamberland, M., Whittingstall, K., Fortin, D., Mathieu, D., Descoteaux, M., 2014. Real-time multi-peak tractography for instantaneous connectivity display. *Front Neuroinform.* 8, 59.
- Chamberland, M., Bernier, M.I., Fortin, D., Whittingstall, K., Descoteaux, M., 2015. 3D interactive tractography-informed resting-state fMRI connectivity. *Frontiers in Neuroscience* 9.
- Chang, H., Fitzpatrick, J.M., 1992. A technique for accurate magnetic resonance imaging in the presence of field inhomogeneities. *IEEE Trans.Med.Imaging* 11, 319-329.
- Chang, L.C., Jones, D.K., Pierpaoli, C., 2005. RESTORE: robust estimation of tensors by outlier rejection. *Magn.Reson.Med.* 53, 1088-1095.
- Chang, L.C., Walker, L., Pierpaoli, C., 2012. Informed RESTORE: A method for robust estimation of diffusion tensor from low redundancy datasets in the presence of physiological noise artifacts. *Magn.Reson.Med.* 68, 1654-1663.
- Chen, Y., Tymofiyeva, O., Hess, C.P., Xu, D., 2015a. Effects of rejecting diffusion directions on tensor-derived parameters. *NeuroImage*. 109, 160-170.
- Chen, Y.H., Wei, D., Newstadt, G., DeGraef, M., Simmons, J., Hero, A., 2015b. Parameter estimation in spherical symmetry groups. *IEEE Signal Process.Lett.* 22, 1152-1155.
- Cheng, J., Deriche, R., Jiang, T., Shen, D., Yap, P.T., 2014. Non-Negative Spherical Deconvolution (NNSD) for estimation of fiber Orientation Distribution Function in single-/multi-shell diffusion MRI. *NeuroImage*. 101, 750-764.
- Cheng, J., Ghosh, A., Jiang, T., Deriche, R., 2010. Model-free and analytical EAP reconstruction via spherical polar Fourier diffusion MRI. *Med.Image Comput.Comput.Assist.Interv.* 13, 590-597.
- Cheng, J., Jiang, T., Deriche, R., 2011. Theoretical analysis and practical insights on EAP estimation via a unified hardi framework.
- Cheung, M.M., Hui, E.S., Chan, K.C., Helpner, J.A., Qi, L., Wu, E.X., 2009. Does diffusion kurtosis imaging lead to better neural tissue characterization? A rodent brain maturation study. *NeuroImage*. 45, 386-392.
- Christiaens, D., Maes, F., Sunaert, S., Suetens, P., 2015. Convex Non-negative Spherical Factorization of Multi-Shell Diffusion-Weighted Images. *Med.Image Comput.Comput.Assist.Interv.* 18, 166-173.
- Chung, K., Deisseroth, K., 2013. CLARITY for mapping the nervous system. *Nat.Methods*. 10, 508-513.
- Chung, K., Wallace, J., Kim, S.Y., Kalyanasundaram, S., Andelman, A.S., Davidson, T.J., Mirzabekov, J.J., Zalocousky, K.A., Mattis, J., Denisin, A.K., Pak, S., Bernstein, H., Ramakrishnan, C., Grosenick, L., Gradinaru, V., Deisseroth, K., 2013. Structural and molecular interrogation of intact biological systems. *Nature* 497, 332-337.
- Clark, C.A., Le Bihan, D., 2000. Water diffusion compartmentation and anisotropy at high b values in the human brain. *Magn.Reson.Med.* 44, 852-859.
- Close, T.G., Tournier, J.D., Calamante, F., Johnston, L.A., Mareels, I., Connelly, A., 2009. A software tool to generate simulated white matter structures for the assessment of fibre-tracking algorithms. *NeuroImage*. 47, 1288-1300.
- Collier, Q., Veraart, J., Jeurissen, B., den Dekker, A.J., Sijbers, J., 2015. Iterative reweighted linear least squares for accurate, fast, and robust estimation of diffusion magnetic resonance parameters. *Magn.Reson.Med.* 73, 2174-2184.
- Concha, L., Gross, D.W., Beaulieu, C., 2005. Diffusion tensor tractography of the limbic system. *AJNR.Am.J.Neuroradiol.* 26, 2267-2274.
- Constantinides, C.D., Atalar, E., McVeigh, E.R., 1997. Signal-to-noise measurements in magnitude images from NMR phased arrays. *Magn.Reson.Med.* 38, 852-857.

Less Confusion in Diffusion MRI

- Conturo, T.E., Lori, N.F., Cull, T.S., Akbudak, E., Snyder, A.Z., Shimony, J.S., McKinstry, R.C., Burton, H., Raichle, M.E., 1999. Tracking neuronal fiber pathways in the living human brain. *Proc.Natl.Acad.Sci.U.S.A.* 96, 10422-10427.
- Cook, P.A., Symms, M., Boulby, P.A., Alexander, D.C., 2007. Optimal acquisition orders of diffusion-weighted MRI measurements. *J.Magn.Reson.Imaging* 25, 1051-1058.
- Côté, M.A., Girard, G., Boré, A., Garyfallidis, E., Houde, J.C., Descoteaux, M., 2013. Tractometer: towards validation of tractography pipelines. *Med.Image Anal.* 17, 844-857.
- Coulon, O., Alexander, D.C., Arridge, S., 2004. Diffusion tensor magnetic resonance image regularization. *Med.Image Anal.* 8, 47-67.
- Coupé, P., Manjon, J.V., Gedamu, E., Arnold, D., Robles, M., Collins, D.L., 2010. Robust Rician noise estimation for MR images. *Med.Image Anal.* 14, 483-493.
- Daducci, A., Canales-Rodriguez, E.J., Descoteaux, M., Garyfallidis, E., Gur, Y., Lin, Y.C., Mani, M., Merlet, S., Paquette, M., Ramirez-Manzanares, A., Reisert, M., Reis, R.P., Sepehrband, F., Caruyer, E., Choupan, J., Deriche, R., Jacob, M., Menegaz, G., Prckovska, V., Rivera, M., Wiaux, Y., Thiran, J.P., 2014a. Quantitative comparison of reconstruction methods for intra-voxel fiber recovery from diffusion MRI. *IEEE Trans.Med.Imaging* 33, 384-399.
- Daducci, A., Canales-Rodriguez, E.J., Zhang, H., Dyrby, T.B., Alexander, D.C., Thiran, J.P., 2015a. Accelerated Microstructure Imaging via Convex Optimization (AMICO) from diffusion MRI data. *NeuroImage*. 105, 32-44.
- Daducci, A., Dal Palù, A., Lemkadem, A., Thiran, J.P., 2015b. COMMIT: Convex optimization modeling for microstructure informed tractography. *IEEE Trans.Med.Imaging* 34, 246-257.
- Daducci, A., Van De Ville, D., Thiran, J.P., Wiaux, Y., 2014b. Sparse regularization for fiber ODF reconstruction: from the suboptimality of l2 and l1 priors to l0. *Med.Image.Anal.* 18, 820-833.
- Daducci, A., Dal Palù, A., Descoteaux, M., Thiran, J.P., 2016. Microstructure informed tractography: pitfalls and open challenges. *Frontiers in Neuroscience* 10.
- De Bresser, J., Portegies, M.P., Leemans, A., Biessels, G.J., Kappelle, L.J., Viergever, M.A., 2011. A comparison of MR based segmentation methods for measuring brain atrophy progression. *NeuroImage*. 54, 760-768.
- De Groot, G., Verhoye, M., Poirier, C., Leemans, A., Eens, M., Darras, V.M., Van der Linden, A., 2009. Structural changes between seasons in the songbird auditory forebrain. *J.Neurosci.* 29, 13557-13565.
- De Groot, M., Vernooij, M.W., Klein, S., Ikram, M.A., Vos, F.M., Smith, S.M., Niessen, W.J., Andersson, J.L., 2013. Improving alignment in Tract-based spatial statistics: evaluation and optimization of image registration. *NeuroImage*. 76, 400-411.
- De Santis, S., Evans, C.J., Jones, D.K., 2012a. RAPID: A routine assurance pipeline for imaging of diffusion. *Magn.Reson.Med.*
- De Santis, S., Assaf, Y., Evans, C.J., Jones, D.K., 2014. Improved precision in CHARMED assessment of white matter through sampling scheme optimization and model parsimony testing. *Magn.Reson.Med.* 71, 661-671.
- De Santis, S., Assaf, Y., Jones, D.K., 2012b. Using the biophysical CHARMED model to elucidate the underpinnings of contrast in diffusional kurtosis analysis of diffusion-weighted MRI. *MAGMA*. 25, 267-276.
- De Santis, S., Barazany, D., Jones, D.K., Assaf, Y., 2016a. Resolving relaxometry and diffusion properties within the same voxel in the presence of crossing fibres by combining inversion recovery and diffusion-weighted acquisitions. *Magn.Reson.Med.* 75, 372-380.
- De Santis, S., Jones, D.K., Roebroeck, A., 2016b. Including diffusion time dependence in the extra-axonal space improves in vivo estimates of axonal diameter and density in human white matter. *NeuroImage*. 130, 91-103.
- Dela Haine, T.C.J., Duits, R., Tax, C.M.W., 2014. Sharpening Fibers in Diffusion Weighted MRI via Erosion. In: Westin, C.F., Vilanova, A., Burgeth, B. (Eds.), *Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data*. Springer Berlin Heidelberg, Berlin, Heidelberg, 97-126.
- Dell'Acqua, F., Rizzo, G., Scifo, P., Clarke, R.A., Scotti, G., Fazio, F., 2007. A model-based deconvolution approach to solve fiber crossing in diffusion-weighted MR imaging. *IEEE Trans.Biomed.Eng* 54, 462-472.
- Dell'Acqua, F., Scifo, P., Rizzo, G., Catani, M., Simmons, A., Scotti, G., Fazio, F., 2010. A modified damped Richardson-Lucy algorithm to reduce isotropic background effects in spherical deconvolution. *NeuroImage*. 49, 1446-1458.
- Dell'Acqua, F., Lacerda, L., Catani, M., Simmons, A., 2014. Anisotropic Power Maps: A diffusion contrast to reveal low anisotropy tissues from HARDI data. *Proceedings of the 22nd Annual Meeting of the ISMRM*, 0730.
- Dell'Acqua, F., Simmons, A., Williams, S.C., Catani, M., 2013. Can spherical deconvolution provide more information than fiber orientations? Hindrance modulated orientational anisotropy, a true-tract specific index to characterize white matter diffusion. *Hum.Brain Mapp.* 34, 2464-2483.
- Deprez, S., Billiet, T., Sunaert, S., Leemans, A., 2013. Diffusion tensor MRI of chemotherapy-induced cognitive impairment in non-CNS cancer patients: a review. *Brain Imaging Behav.* 7, 409-435.
- Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R., 2007. Regularized, fast, and robust analytical Q-ball imaging. *Magn.Reson.Med.* 58, 497-510.

References

- Descoteaux, M., Deriche, R., Knosche, T.R., Anwander, A., 2009. Deterministic and probabilistic tractography based on complex fibre orientation distributions. *IEEE Trans.Med.Imaging* 28, 269-286.
- Descoteaux, M., Deriche, R., Le Bihan, D., Mangin, J.F., Poupon, C., 2011. Multiple q-shell diffusion propagator imaging. *Med.Image.Anal.* 15, 603-621.
- Descoteaux, M., Sidhu, J., Garyfallidis, E., Houde, J.C., Neher, P.F., Stieltjes, B., Maier-Hein, K.H., 2016. False positive bundles in tractography. In Proceedings of the 24th Annual Meeting of the ISMRM, 0790.
- Dietrich, O., Raya, J.G., Reeder, S.B., Reiser, M.F., Schoenberg, S.O., 2007. Measurement of signal-to-noise ratios in MR images: influence of multichannel coils, parallel imaging, and reconstruction filters. *J.Magn.Reson.Imaging* 26, 375-385.
- Does, M.D., Parsons, E.C., Gore, J.C., 2003. Oscillating gradient measurements of water diffusion in normal and globally ischemic rat brain. *Magn.Reson.Med.* 49, 206-215.
- Donoho, D.L., 2006. Compressed sensing. *IEEE Trans.Inf.Theory* 52, 1289-1306.
- Du, Y.P., Joe, Z., X, Bernstein, M.A., 2002. Correction of concomitant magnetic field-induced image artifacts in nonaxial echo-planar imaging. *Magn.Reson.Med.* 48, 509-515.
- Duits, R., Franken, E., 2011. 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* 92, 231-264.
- Duning, T., Kloska, S., Steinruter, O., Kugel, H., Heindel, W., Knecht, S., 2005. Dehydration confounds the assessment of brain atrophy. *Neurology*. 64, 548-550.
- Dyrby, T.B., Baare, W.F., Alexander, D.C., Jelsing, J., Garde, E., Sogaard, L.V., 2011. An ex vivo imaging pipeline for producing high-quality and high-resolution diffusion-weighted imaging datasets. *Hum.Brain Mapp.* 32, 544-563.
- Eichelbaum, S., Hlawitschka, M., Scheuermann, G., 2013a. LineAO--improved three-dimensional line rendering. *IEEE Trans.Vis.Comput.Graph.* 19, 433-445.
- Eichelbaum, S., Hlawitschka, M., Scheuermann, G., 2013b. Openwalnut: An Open-Source Tool for Visualization of Medical and Bio-Signal Data. *Biomed.Tech.(Berl)*.
- Eickhoff, S., Nichols, T.E., Van Horn, J.D., Turner, J.A., 2016. Sharing the wealth: Neuroimaging data repositories. *NeuroImage*. 124, 1065-1068.
- Elhabian, S., Gur, Y., Vachet, C., Piven, J., Styner, M., Leppert, I.R., Pike, G.B., Gerig, G., 2014. Subject-Motion Correction in HARDI Acquisitions: Choices and Consequences. *Front Neurol.* 5, 240.
- Emself, L., Leemans, A., Langan, C., Van, H.W., Barker, G.J., McCarthy, P., Jeurissen, B., Sijbers, J., Sunaert, S., Cannon, D.M., McDonald, C., 2013. Limbic and callosal white matter changes in euthymic bipolar I disorder: an advanced diffusion magnetic resonance imaging tractography study. *Biol.Psychiatry* 73, 194-201.
- Evans, A.C., 2006. The NIH MRI study of normal brain development. *NeuroImage*. 30, 184-202.
- Everts, M.H., Bekker, H., Roerdink, J.B., Isenberg, T., 2009. Depth-dependent halos: Illustrative rendering of dense line data. *IEEE Trans.Vis.Comput.Graphics* 15, 1299-1306.
- Fan, Q., Witzel, T., Nummenmaa, A., Van Dijk, K.R., Van Horn, J.D., Drews, M.K., Somerville, L.H., Sheridan, M.A., Santillana, R.M., Snyder, J., Hedden, T., Shaw, E.E., Hollinshead, M.O., Renvall, V., Zanzonico, R., Keil, B., Cauley, S., Polimeni, J.R., Tisdall, D., Buckner, R.L., Wedeen, V.J., Wald, L.L., Toga, A.W., Rosen, B.R., 2016. MGH-USC Human Connectome Project datasets with ultra-high b-value diffusion MRI. *NeuroImage*. 124, 1108-1114.
- Farrell, J.A., Landman, B.A., Jones, C.K., Smith, S.A., Prince, J.L., van Zijl, P.C., Mori, S., 2007. Effects of signal-to-noise ratio on the accuracy and reproducibility of diffusion tensor imaging-derived fractional anisotropy, mean diffusivity, and principal eigenvector measurements at 1.5 T. *J.Magn.Reson.Imaging* 26, 756-767.
- Farzinfar, M., Oguz, I., Smith, R.G., Verde, A.R., Dietrich, C., Gupta, A., Escolar, M.L., Piven, J., Pujol, S., Vachet, C., Gouttard, S., Gerig, G., Dager, S., McKinstry, R.C., Paterson, S., Evans, A.C., Styner, M.A., 2013. Diffusion imaging quality control via entropy of principal direction distribution. *NeuroImage*. 82, 1-12.
- Fedorov, A., Beichel, R., Kalpathy-Cramer, J., Finet, J., Fillion-Robin, J.C., Pujol, S., Bauer, C., Jennings, D., Fennessy, F., Sonka, M., Buatti, J., Aylward, S., Miller, J.V., Pieper, S., Kikinis, R., 2012. 3D Slicer as an image computing platform for the Quantitative Imaging Network. *Magn.Reson.Imaging* 30, 1323-1341.
- Feng, Y., Wu, Y., Rathi, Y., Westin, C.F., 2015. Sparse deconvolution of higher order tensor for fiber orientation distribution estimation. *Artif.Intell.Med.* 65, 229-238.
- Ferizi, U., Schneider, T., Panagiotaki, E., Nedjati-Gilani, G., Zhang, H., Wheeler-Kingshott, C.A., Alexander, D.C., 2014. A ranking of diffusion MRI compartment models with in vivo human brain data. *Magn.Reson.Med.* 72, 1785-1792.
- Ferizi, U., Schneider, T., Tariq, M., Wheeler-Kingshott, C.A., Zhang, H., Alexander, D.C., 2013. The importance of being dispersed: A ranking of diffusion MRI models for fibre dispersion using in vivo human brain data. *Med.Image Comput.Comput.Assist.Interv.* 16, 74-81.
- Ferizi, U., Schneider, T., Witzel, T., Wald, L.L., Zhang, H., Wheeler-Kingshott, C.A., Alexander, D.C., 2015. White matter compartment models for in vivo diffusion MRI at 300mT/m. *NeuroImage*. 118, 468-483.

Less Confusion in Diffusion MRI

- Fick, R., Wassermann, D., Pizzolato, M., Deriche, R., 2015. A Unifying Framework for Spatial and Temporal Diffusion in Diffusion MRI. *Inf.Process Med.Imaging* 24, 167-178.
- Fick, R., Wassermann, D., Sanguinetti, G., Deriche, R., 2014. An analytical 3D Laplacian regularized SHORE basis and its impact on EAP reconstruction and microstructure recovery. In: O'Donnell, L., Nedjati-Gilani, G., Rathi, Y., Reisert, M., Schneider T. (Eds.), *Computational Diffusion MRI: Med.Image Comput.Comput.Assist.Interv. Workshop*, Boston, USA, 2014. Springer International Publishing, Cham, 151-165.
- Fick, R.H., Wassermann, D., Caruyer, E., Deriche, R., 2016. MAPL: Tissue microstructure estimation using Laplacian-regularized MAP-MRI and its application to HCP data. *NeuroImage*. 134, 365-385.
- Fieremans, E., Burcaw, L.M., Lee, H.H., Lemmerski, G., Veraart, J., Novikov, D.S., 2016. In vivo observation and biophysical interpretation of time-dependent diffusion in human white matter. *NeuroImage*. 129, 414-427.
- Fieremans, E., Jensen, J.H., Helpern, J.A., 2011. White matter characterization with diffusional kurtosis imaging. *NeuroImage* 58, 177-188.
- Fieremans, E., Novikov, D.S., Jensen, J.H., Helpern, J.A., 2010. Monte Carlo study of a two-compartment exchange model of diffusion. *NMR Biomed.* 23, 711-724.
- Fillard, P., Descoteaux, M., Goh, A., Gouttard, S., Jeurissen, B., Malcolm, J., Ramirez-Manzanares, A., Reisert, M., Sakaie, K., Tensaouti, F., Yo, T., Mangin, J.F., Poupon, C., 2011. Quantitative evaluation of 10 tractography algorithms on a realistic diffusion MR phantom. *NeuroImage*. 56, 220-234.
- Fillard, P., Pennec, X., Arsigny, V., Ayache, N., 2007. Clinical DT-MRI estimation, smoothing, and fiber tracking with log-Euclidean metrics. *IEEE Trans.Med.Imaging* 26, 1472-1482.
- Firbank, M.J., Coulthard, A., Harrison, R.M., Williams, E.D., 1999. A comparison of two methods for measuring the signal to noise ratio on MR images. *Phys.Med.Biol.* 44, N261-N264.
- Florack, L., 2008. Codomain scale space and regularization for high angular resolution diffusion imaging. *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW '08. IEEE Computer Society Conference on*, 1-6.
- Florack, L., 2013. Image structure, 10 ed. Springer Science & Business Media.
- Frank, L.R., 2002. Characterization of anisotropy in high angular resolution diffusion-weighted MRI. *Magn.Reson.Med.* 47, 1083-1099.
- Froeling, M., Tax, C.M.W., Vos, S.B., Leemans, A., 2014. MASSIVE: Multiple Acquisitions for Standardization of Structural Imaging Validation and Evaluation. Proceedings of the 22nd Annual Meeting of the ISMRM, 2582.
- Froeling, M., Tax, C.M.W., Vos, S.B., Luijten, P.R., Leemans, A., 2016. The "MASSIVE" brain dataset: Multiple Acquisitions for Standardization of Structural Imaging Validation and Evaluation. *Magn.Reson.Med.* doi: 10.1002/mrm.26259.
- Gallichan, D., Scholz, J., Bartsch, A., Behrens, T.E., Robson, M.D., Miller, K.L., 2010. Addressing a systematic vibration artifact in diffusion-weighted MRI. *Hum.Brain Mapp.* 31, 193-202.
- Gao, Y., Zhang, Y., Wong, C.S., Wu, P.M., Zhang, Z., Gao, J., Qiu, D., Huang, B., 2012. Diffusion abnormalities in temporal lobes of children with temporal lobe epilepsy: a preliminary diffusional kurtosis imaging study and comparison with diffusion tensor imaging. *NMR Biomed.* 25, 1369-1377.
- Garyfallidis, E., Brett, M., Amirbekian, B., Rokem, A., van der Walt, S., Descoteaux, M., Nimmo-Smith, I., 2014. Dipy, a library for the analysis of diffusion MRI data. *Front Neuroinform.* 8, 8.
- Garyfallidis, E., Ocegueda, O., Wassermann, D., Descoteaux, M., 2015. Robust and efficient linear registration of white-matter fascicles in the space of streamlines. *NeuroImage* 117, 124-140.
- Gholipour, A., Kehtarnavaz, N., Scherrer, B., Warfield, S.K., 2011. On the accuracy of unwarping techniques for the correction of susceptibility-induced geometric distortion in magnetic resonance Echo-planar images. *Conf.Proc.IEEE Eng.Med.Biol.Soc.* 2011, 6997-7000.
- Ghosh, A., Milne, T., Deriche, R., 2014. Constrained diffusion kurtosis imaging using ternary quartics & MLE. *Magn.Reson.Med.* 71, 1581-1591.
- Giltinan, D.M., Carroll, R.J., Ruppert, D., 1986. Some New Estimation Methods for Weighted Regression When There Are Possible Outliers. *Technometrics* 28, 219-230.
- Girard, G., Chamberland, M., Houde, J.C., Fortin, D., Descoteaux, M., 2012. Neurosurgical tracking at the Sherbrooke Connectivity Imaging Lab (SCIL). In *Med.Image Comput.Comput.Assist.Interv. DTI Challenge Workshop*, 55-73.
- Girard, G., Fick, R., Descoteaux, M., Deriche, R., Wassermann, D., 2015. AxTract: Microstructure-Driven Tractography Based on the Ensemble Average Propagator. *Inf.Process Med.Imaging* 24, 675-686.
- Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M., 2013. The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*. 80, 105-124.
- Glenn, G.R., Tabesh, A., Jensen, J.H., 2015. A simple noise correction scheme for diffusional kurtosis imaging. *Magn.Reson.Imaging* 33, 124-133.
- Glodek, D., Hesser, J., Zheng, L., 2016. Distortion correction of EPI data using multimodal nonrigid registration with an anisotropic regularization. *Magn Reson.Imaging* 34, 127-136.

References

- Gooijers, J., Leemans, A., Van Cauter, S., Sunaert, S., Swinnen, S.P., Caeyenberghs, K., 2014. White matter organization in relation to upper limb motor control in healthy subjects: exploring the added value of diffusion kurtosis imaging. *Brain Struct.Funct.* 219, 1627-1638.
- Graham, M.S., Drobniak, I., Zhang, H., 2015. A Simulation Framework for Quantitative Validation of Artefact Correction in Diffusion MRI. *Inf.Process Med.Imaging* 24, 638-649.
- Gramfort, A., Poupon, C., Descoteaux, M., 2014. Denoising and fast diffusion imaging with physically constrained sparse dictionary learning. *Med.Image Anal.* 18, 36-49.
- Griswold, M.A., Jakob, P.M., Heidemann, R.M., Nittka, M., Jellus, V., Wang, J., Kiefer, B., Haase, A., 2002. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn.Reson.Med.* 47, 1202-1210.
- Gudbjartsson, H., Patz, S., 1995. The Rician distribution of noisy MRI data. *Magn.Reson.Med.* 34, 910-914.
- Guo, F., Tax, C.M.W., Heemskerk, A.M., Leemans, A., 2016. Assessing angular characteristics of the fiber orientation distribution function in relation to the response function shape. In Proceedings of the 8th Annual Meeting of the ISMRM Benelux Chapter, O-001.
- Habib, J., Auer, D.P., Morgan, P.S., 2010. A quantitative analysis of the benefits of cardiac gating in practical diffusion tensor imaging of the brain. *Magn.Reson.Med.* 63, 1098-1103.
- Hagmann, P., Jonasson, L., Maeder, P., Thiran, J.P., Wedeen, V.J., Meuli, R., 2006. Understanding diffusion MR imaging techniques: from scalar diffusion-weighted imaging to diffusion tensor imaging and beyond. *Radiographics.* 26 Suppl 1, S205-S223.
- Hansen, B., Lund, T.E., Sangill, R., Jespersen, S.N., 2013. Experimentally and computationally fast method for estimation of a mean kurtosis. *Magn.Reson.Med.* 69, 1754-1760.
- Hansen, B., Lund, T.E., Sangill, R., Stubbé, E., Finsterbusch, J., Jespersen, S.N., 2015. Experimental considerations for fast kurtosis imaging. *Magn.Reson.Med.*
- Haselgrove, J.C., Moore, J.R., 1996. Correction for distortion of echo-planar images used to calculate the apparent diffusion coefficient. *Magn.Reson.Med.* 36, 960-964.
- Hau, J., Sarubbo, S., Perchey, G., Crivello, F., Joliot, M., Zago, L., Jobard, G., Mellet, E., Mazoyer, B., Tzourio-Mazoyer, N., Petit, L., 2014. Stem-based tractography to study the anatomical connectivity of human brain white matter pathways. Organization for Human Brain Mapping (OHB), 3930.
- Heemskerk, A.M., Leemans, A., Plaisier, A., Pieterman, K., Lequin, M.H., Dudink, J., 2013. Acquisition guidelines and quality assessment tools for analyzing neonatal diffusion tensor MRI data. *AJNR Am.J.Neuroradiol.* 34, 1496-1505.
- Heemskerk, A.M., Thiebaut de, S.M., Catani, M., Sarubbo, S., Petit, L., Viergever, M.A., Jones, D.K., Evans, J., Paus, T., Leemans, A., 2015. A new fiber bundle pathway identified with diffusion MRI fiber tractography: fact or fantasy? Proceedings of the 23rd Annual Meeting of the ISMRM, 2855.
- Hekimoglu, S., Berber, M., 2003. Effectiveness of robust methods in heterogeneous linear models. *Journal of Geodesy* 76, 706-713.
- Helpner, J.A., Adisetiyo, V., Falangola, M.F., Hu, C., Di, M.A., Williams, K., Castellanos, F.X., Jensen, J.H., 2011. Preliminary evidence of altered gray and white matter microstructural development in the frontal lobe of adolescents with attention-deficit hyperactivity disorder: a diffusional kurtosis imaging study. *J.Magn.Reson.Imaging* 33, 17-23.
- Henkelman, R.M., 1985. Measurement of signal intensities in the presence of noise in MR images. *Med Phys.* 12, 232-233.
- Hess, C.P., Mukherjee, P., Han, E.T., Xu, D., Vigneron, D.B., 2006. Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic basis. *Magn.Reson.Med.* 56, 104-117.
- Hiltunen, J., Hari, R., Jousmaki, V., Muller, K., Sepponen, R., Joensuu, R., 2006. Quantification of mechanical vibration during diffusion tensor imaging at 3 T. *NeuroImage.* 32, 93-103.
- Hodge, M.R., Horton, W., Brown, T., Herrick, R., Olsen, T., Hileman, M.E., McKay, M., Archie, K.A., Cler, E., Harms, M.P., Burgess, G.C., Glasser, M.F., Elam, J.S., Curtiss, S.W., Barch, D.M., Oostenveld, R., Larson-Prior, L.J., Ugurbil, K., Van Essen, D.C., Marcus, D.S., 2016. ConnectomeDB--Sharing human brain connectivity data. *NeuroImage.* 124, 1102-1107.
- Hogg, R.V., 1979. Statistical robustness: One view of its use in applications today. *The American Statistician* 33, 108-115.
- Holland, D., Kuperman, J.M., Dale, A.M., 2010. Efficient correction of inhomogeneous static magnetic field-induced distortion in Echo Planar Imaging. *NeuroImage.* 50, 175-183.
- Holland, P.W., Welsch, R.E., 1977. Robust regression using iteratively reweighted least-squares. *Communications in Statistics-Theory and Methods* 6, 813-827.
- Hosseini, A.P., Chung, M.K., Wu, Y.C., Alexander, A.L., 2013. Bessel Fourier Orientation Reconstruction (BFOR): an analytical diffusion propagator reconstruction for hybrid diffusion imaging and computation of q-space indices. *NeuroImage.* 64, 650-670.

Less Confusion in Diffusion MRI

- Hoy, A.R., Kecskemeti, S.R., Alexander, A.L., 2015. Free water elimination diffusion tractography: A comparison with conventional and fluid-attenuated inversion recovery, diffusion tensor imaging acquisitions. *J.Magn.Reson.Imaging* 42, 1572-1581.
- Hsu, Y.C., Lo, Y.C., Chen, Y.J., Wedeen, V.J., Isaac Tseng, W.Y., 2015. NTU-DSI-122: A diffusion spectrum imaging template with high anatomical matching to the ICBM-152 space. *Hum.Brain Mapp.* 36, 3528-3541.
- Huang, H., Ceritoglu, C., Li, X., Qiu, A., Miller, M.I., van Zijl, P.C., Mori, S., 2008. Correction of B0 susceptibility induced distortion in diffusion-weighted images using large-deformation diffeomorphic metric mapping. *Magn.Reson.Imaging* 26, 1294-1302.
- Hui, E.S., Russell Glenn, G., Helpern, J.A., Jensen, J.H., 2015. Kurtosis analysis of neural diffusion organization. *NeuroImage* 106, 391-403.
- Huizinga, W., Poot, D.H., Guyader, J.M., Klaassen, R., Coolen, B.F., van Kranenburg, M., van Geuns, R.J., Uitterdijk, A., Polfliet, M., Vandemeulebroucke, J., Leemans, A., Niessen, W.J., Klein, S., 2016. PCA-based groupwise image registration for quantitative MRI. *Med.Image Anal.* 29, 65-78.
- Irfanoglu, M.O., Modi, P., Nayak, A., Hutchinson, E.B., Sarlls, J., Pierpaoli, C., 2015. DR-BUDDI (Diffeomorphic Registration for Blip-Up blip-Down Diffusion Imaging) method for correcting echo planar imaging distortions. *NeuroImage*. 106, 284-299.
- Irfanoglu, M.O., Walker, L., Sammet, S., Pierpaoli, C., Machiraju, R., 2011. Susceptibility distortion correction for echo planar images with non-uniform B-spline grid sampling: a diffusion tensor image study. *Med.Image Comput.Comput.Assist.Interv.* 14, 174-181.
- Irfanoglu, M.O., Walker, L., Sarlls, J., Marenco, S., Pierpaoli, C., 2012. Effects of image distortions originating from susceptibility variations and concomitant fields on diffusion MRI tractography results. *NeuroImage*. 61, 275-288.
- Jbabdi, S., Behrens, T.E., Smith, S.M., 2010. Crossing fibres in tract-based spatial statistics. *NeuroImage*. 49, 249-256.
- Jelescu, I.O., Veraart, J., Fieremans, E., Novikov, D.S., 2016a. Degeneracy in model parameter estimation for multi-compartmental diffusion in neuronal tissue. *NMR Biomed.* 29, 33-47.
- Jelescu, I.O., Zurek, M., Winters, K.V., Veraart, J., Rajaratnam, A., Kim, N.S., Babb, J.S., Shepherd, T.M., Novikov, D.S., Kim, S.G., Fieremans, E., 2016b. In vivo quantification of demyelination and recovery using compartment-specific diffusion MRI metrics validated by electron microscopy. *NeuroImage*. 132, 104-114.
- Jensen, J.H., Falangola, M.F., Hu, C., Tabesh, A., Rapalino, O., Lo, C., Helpern, J.A., 2011. Preliminary observations of increased diffusional kurtosis in human brain following recent cerebral infarction. *NMR Biomed.* 24, 452-457.
- Jensen, J.H., Helpern, J.A., 2010. MRI quantification of non-Gaussian water diffusion by kurtosis analysis. *NMR Biomed.* 23, 698-710.
- Jensen, J.H., Helpern, J.A., Ramani, A., Lu, H., Kaczynski, K., 2005. Diffusional kurtosis imaging: the quantification of non-gaussian water diffusion by means of magnetic resonance imaging. *Magn.Reson.Med.* 53, 1432-1440.
- Jensen, J.H., Russell, G.G., Helpern, J.A., 2016. Fiber ball imaging. *NeuroImage*. 124, 824-833.
- Jeurissen, B., Leemans, A., Jones, D.K., Tournier, J.D., Sijbers, J., 2011. Probabilistic fiber tracking using the residual bootstrap with constrained spherical deconvolution. *Hum.Brain Mapp.* 32, 461-479.
- Jeurissen, B., Leemans, A., Tournier, J.D., Jones, D.K., Sijbers, J., 2013a. Investigating the prevalence of complex fiber configurations in white matter tissue with diffusion magnetic resonance imaging. *Hum.Brain Mapp.* 34, 2747-2766.
- Jeurissen, B., Tournier, J.D., Dhollander, T., Connelly, A., Sijbers, J., 2014. Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data. *NeuroImage*. 103, 411-426.
- Jeurissen, B., Tournier, J.D., Sijbers, J., 2015. Tissue-type segmentation using non-negative matrix factorization of multi-shell diffusion-weighted MRI images. *Proceedings of the 23rd Annual Meeting of the ISMRM*, 0349.
- Jezzard, P., Balaban, R.S., 1995. Correction for geometric distortion in echo planar images from B0 field variations. *Magn.Reson.Med.* 34, 65-73.
- Jezzard, P., Barnett, A.S., Pierpaoli, C., 1998. Characterization of and correction for eddy current artifacts in echo planar diffusion imaging. *Magn.Reson.Med.* 39, 801-812.
- Jian, B., Vemuri, B.C., 2007. A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted MRI. *IEEE Trans.Med.Imaging* 26, 1464-1471.
- Jian, B., Vemuri, B.C., Ozarslan, E., Carney, P.R., Mareci, T.H., 2007. A novel tensor distribution model for the diffusion-weighted MR signal. *NeuroImage*. 37, 164-176.
- Johansen-Berg, H., Behrens, T.E., 2009. Diffusion MRI: From quantitative measurement to in-vivo neuroanatomy. Academic Press.
- Jones, D.K., 2003. Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI. *Magn.Reson.Med.* 49, 7-12.
- Jones, D.K., 2004. The effect of gradient sampling schemes on measures derived from diffusion tensor MRI: a Monte Carlo study. *Magn.Reson.Med.* 51, 807-815.
- Jones, D.K., 2010a. Challenges and limitations of quantifying brain connectivity in vivo with diffusion MRI. *Imaging in Medicine* 2, 341-355.

References

- Jones, D.K., 2010b. Diffusion MRI: Theory, Methods, and Applications. Oxford University Press, USA.
- Jones, D.K., 2010c. The signal intensity must be modulated by the determinant of the jacobian when correcting for eddy currents in diffusion MRI. In Proceedings of the 18th Annual Meeting of the ISMRM, 1644.
- Jones, D.K., Basser, P.J., 2004. "Squashing peanuts and smashing pumpkins": how noise distorts diffusion-weighted MR data. *Magn.Reson.Med.* 52, 979-993.
- Jones, D.K., Cercignani, M., 2010. Twenty-five pitfalls in the analysis of diffusion MRI data. *NMR Biomed.* 23, 803-820.
- Jones, D.K., Horsfield, M.A., Simmons, A., 1999. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magn.Reson.Med.* 42, 515-525.
- Jones, D.K., Leemans, A., 2011. Diffusion tensor imaging. *Methods.Mol.Biol.* 711, 127-144.
- Jones, D.K., Travis, A.R., Eden, G., Pierpaoli, C., Basser, P.J., 2005. PASTA: pointwise assessment of streamline tractography attributes. *Magn.Reson.Med.* 53, 1462-1467.
- Karayannidis, F., Keuken, M.C., Wong, A., Rennie, J.L., de, H.G., Cooper, P.S., Ross, F.W., Lenroot, R., Parsons, M., Phillips, N., Michie, P.T., Forstmann, B.U., 2016. The Age-ility Project (Phase 1): Structural and functional imaging and electrophysiological data repository. *NeuroImage.* 124, 1137-1142.
- Kaufman, L., Kramer, D.M., Crooks, L.E., Ortendahl, D.A., 1989. Measuring signal-to-noise ratios in MR imaging. *Radiology.* 173, 265-267.
- Kellner, E., Dhital, B., Kiselev, V.G., Reisert, M., 2015. Gibbs-ringing artifact removal based on local subvoxel-shifts. *Magn.Reson.Med.* doi: 10.1002/mrm.26054.
- Kennis, M., van Rooij, S.J., Kahn, R.S., Geuze, E., Leemans, A., 2016. Choosing the polarity of the phase-encoding direction in diffusion MRI: Does it matter for group analysis? *NeuroImage.Clin.* 11, 539-547.
- Kindlmann, G., Tricoche, X., Westin, C.F., 2007. Delineating white matter structure in diffusion tensor MRI with anisotropy creases. *Medical Image Analysis* 11, 492-502.
- Kingsley, P.B., 2006. Introduction to diffusion tensor imaging mathematics: Part III. Tensor calculation, noise, simulations, and optimization. *Concepts in Magnetic Resonance Part A* 28A, 155-179.
- Kiselev, V.G., 2010. The cumulant expansion: An overarching mathematical framework for understanding diffusion MRI. In: Jones, D.K. (Ed.), *Diffusion MRI: Theory, Methods, and Applications*. Oxford University Press.
- Kiselev, V.G., Il'yasov, K.A., 2007. Is the "biexponential diffusion" biexponential? *Magn.Reson.Med.* 57, 464-469.
- Klein, S., Staring, M., Murphy, K., Viergever, M.A., Pluim, J.P., 2010. elastix: a toolbox for intensity-based medical image registration. *IEEE Trans.Med.Imaging* 29, 196-205.
- Kleinnijenhuis, M., Barth, M., Alexander, D.C., van Cappellen van Walsum AM, Norris, D.G., 2012. Structure Tensor Informed Fiber Tractography (STIFT) by combining gradient echo MRI and diffusion weighted imaging. *NeuroImage.* 59, 3941-3954.
- Kleinnijenhuis, M., Mollink, J., Kinchesh, P., Lam, W.W., Galinsky, V., Frank, L.R., Smart, S.C., Jbabdi, S., Miller, K.L., 2015. Monte Carlo diffusion simulations disambiguate the biophysical mechanisms of diffusion hinderance along tracts. *Proceedings of the 23rd Annual Meeting of the ISMRM,* 2756.
- Klingler, J., 1935. Erleichterung der makrokopischen Präparation des Gehirns durch den Gefrierprozess. *Schweiz Arch Neurol* 36, 247-256.
- Knutsson, H., Westin, C.F., 1993. Normalized and differential convolution. *IEEE CVPR,* 515-523.
- Koay, C.G., 2010. Least Squares Approaches to Diffusion Tensor Estimation. In: Jones, D.K. (Ed.), *Diffusion MRI: Theory, Methods, and Applications*. Oxford University Press.
- Koay, C.G., Basser, P.J., 2006. Analytically exact correction scheme for signal extraction from noisy magnitude MR signals. *J.Magn Reson.* 179, 317-322.
- Koay, C.G., Chang, L.C., Carew, J.D., Pierpaoli, C., Basser, P.J., 2006. A unifying theoretical and algorithmic framework for least squares methods of estimation in diffusion tensor imaging. *J.Magn Reson.* 182, 115-125.
- Koay, C.G., Ozarslan, E., Basser, P.J., 2009a. A signal transformational framework for breaking the noise floor and its applications in MRI. *J.Magn Reson.* 197, 108-119.
- Koay, C.G., Özarslan, E., Pierpaoli, C., 2009b. Probabilistic Identification and Estimation of Noise (PIESNO): a self-consistent approach and its applications in MRI. *J.Magn Reson.* 199, 94-103.
- Koenderink, J.J., Van Doorn, A.J., 1990. Receptive field families. *Biological cybernetics* 63, 291-297.
- Kuder, T.A., Stieltjes, B., Bachert, P., Semmler, W., Laun, F.B., 2012. Advanced fit of the diffusion kurtosis tensor by directional weighting and regularization. *Magn.Reson.Med.* 67, 1401-1411.
- Kuo, L.W., Chen, J.H., Wedeen, V.J., Tseng, W.Y., 2008. Optimization of diffusion spectrum imaging and q-ball imaging on clinical MRI system. *NeuroImage.* 41, 7-18.
- Kybic, J., Thevenaz, P., Nirko, A., Unser, M., 2000. Unwarping of unidirectionally distorted EPI images. *IEEE Trans.Med.Imaging* 19, 80-93.
- Lam, F., Babacan, S.D., Haldar, J.P., Weiner, M.W., Schuff, N., Liang, Z.P., 2014. Denoising diffusion-weighted magnitude MR images using rank and edge constraints. *Magn.Reson.Med.* 71, 1272-1284.

Less Confusion in Diffusion MRI

- Landman, B.A., Bazin, P.L., Prince, J.L., 2009. Estimation and application of spatially variable noise fields in diffusion tensor imaging. *Magn Reson Imaging* 27, 741-751.
- Landman, B.A., Bogovic, J.A., Wan, H., El Zahraa, E.F., Bazin, P.L., Prince, J.L., 2012. Resolution of crossing fibers with constrained compressed sensing using diffusion tensor MRI. *NeuroImage*. 59, 2175-2186.
- Landman, B.A., Farrell, J.A., Jones, C.K., Smith, S.A., Prince, J.L., Mori, S., 2007a. Effects of diffusion weighting schemes on the reproducibility of DTI-derived fractional anisotropy, mean diffusivity, and principal eigenvector measurements at 1.5T. *NeuroImage*. 36, 1123-1138.
- Landman, B.A., Mori, S., Prince, J.L., 2007b. Systematic evaluation of linear and nonlinear DTI estimation methods: an open framework. *Proceedings of the 15th Annual Meeting of the ISMRM* 1522.
- Lang, S., 1995. Differential and Riemannian Manifolds, Graduate Texts in Mathematics ed. Springer-Verlag New York.
- Langen, M., Leemans, A., Johnston, P., Ecker, C., Daly, E., Murphy, C.M., Dell'Acqua, F., Durston, S., Murphy, D.G., 2012. Fronto-striatal circuitry and inhibitory control in autism: findings from diffusion tensor imaging tractography. *Cortex* 48, 183-193.
- Lazar, M., Jensen, J.H., Xuan, L., Helpern, J.A., 2008. Estimation of the orientation distribution function from diffusional kurtosis imaging. *Magn.Reson.Med.* 60, 774-781.
- Le Bihan, D., 2010. Magnetic Resonance Diffusion Imaging: Introduction and Concepts. In: Jones, D.K. (Ed.), *Diffusion MRI: Theory, Methods, and Applications*. Oxford University Press.
- Le Bihan, D., Breton, E., Lallemand, D., Grenier, P., Cabanis, E., Laval-Jeantet, M., 1986. MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology* 161, 401-407.
- Le Bihan, D., Johansen-Berg, H., 2012. Diffusion MRI at 25: exploring brain tissue structure and function. *NeuroImage*. 61, 324-341.
- Le Bihan, D., Mangin, J.F., Poupon, C., Clark, C.A., Pappata, S., Molko, N., Chabriat, H., 2001. Diffusion tensor imaging: concepts and applications. *J.Magn Reson.Imaging* 13, 534-546.
- Lebel, C., Walker, L., Leemans, A., Phillips, L., Beaulieu, C., 2008. Microstructural maturation of the human brain from childhood to adulthood. *NeuroImage*. 40, 1044-1055.
- Leemans, A., 2010. Visualization of Diffusion MRI Data. In: Jones, D.K. (Ed.), *Diffusion MRI: Theory, Methods, and Applications*. Oxford University Press.
- Leemans, A., Jeurissen, B., Sijbers, J., Jones, D.K., 2009. ExploreDTI: a graphical toolbox for processing, analyzing, and visualizing diffusion MR data. In *Proceedings of the 17nd Annual Meeting of the ISMRM* 3537.
- Leemans, A., Jones, D.K., 2009. The B-matrix must be rotated when correcting for subject motion in DTI data. *Magn.Reson.Med.* 61, 1336-1349.
- Leemans, A., Sijbers, J., De, B.S., Vandervliet, E., Parizel, P., 2006. Multiscale white matter fiber tract coregistration: a new feature-based approach to align diffusion tensor data. *Magn.Reson.Med.* 55, 1414-1423.
- Leemans, A., Sijbers, J., Verhoye, M., Van der Linden, A., Van Dyck, D., 2005. Mathematical framework for simulating diffusion tensor MR neural fiber bundles. *Magn.Reson.Med.* 53, 944-953.
- Lim, K.O., Hedefus, M., Moseley, M., de, C.A., Sullivan, E.V., Pfefferbaum, A., 1999. Compromised white matter tract integrity in schizophrenia inferred from diffusion tensor imaging. *Arch.Gen.Psychiatry* 56, 367-374.
- Lin, C.P., Wedeen, V.J., Chen, J.H., Yao, C., Tseng, W.Y., 2003. Validation of diffusion spectrum magnetic resonance imaging with manganese-enhanced rat optic tracts and ex vivo phantoms. *NeuroImage*. 19, 482-495.
- Liu, M., Vemuri, B.C., Deriche, R., 2013. A robust variational approach for simultaneous smoothing and estimation of DTI. *NeuroImage*. 67, 33-41.
- Liu, X., Yuan, Z., Guo, Z., Xu, D., 2015. A localized Richardson-Lucy algorithm for fiber orientation estimation in high angular resolution diffusion imaging. *Med.Phys.* 42, 2524-2539.
- Liu, Z., Wang, Y., Gerig, G., Gouttard, S., Tao, R., Fletcher, T., Styner, M., 2010. Quality Control of Diffusion Weighted Images. *Proceedings of SPIE--the International Society for Optical Engineering* 7628, 76280J.
- Lorenz, R., Bellemann, M.E., Hennig, J., Il'yasov, K.A., 2008. Anisotropic phantoms for quantitative diffusion tensor imaging and fiber-tracking validation. *Applied Magnetic Resonance* 33, 419-429.
- Lu, H., Jensen, J.H., Ramani, A., Helpern, J.A., 2006. Three-dimensional characterization of non-gaussian water diffusion in humans using diffusion kurtosis imaging. *NMR Biomed.* 19, 236-247.
- Lundell, H., Nielsen, J.B., Ptito, M., Dyrby, T.B., 2011. Distribution of collateral fibers in the monkey cervical spinal cord detected with diffusion-weighted magnetic resonance imaging. *NeuroImage*. 56, 923-929.
- Maclarens, J., Han, Z., Vos, S.B., Fischbein, N., Bammer, R., 2014. Reliability of brain volume measurements: a test-retest dataset. *Sci.Data.* 1, 140037.
- Mangin, J.F., Poupon, C., Clark, C., Le Bihan, D., Bloch, I., 2002. Distortion correction and robust tensor estimation for MR diffusion imaging. *Med.Image Anal.* 6, 191-198.
- Manjon, J.V., Coupe, P., Concha, L., Buades, A., Collins, D.L., Robles, M., 2013. Diffusion weighted image denoising using overcomplete local PCA. *PLoS.One.* 8, e73021.

References

- Manjón, J.V., Coupe, P., Martí-Bonmatí, L., Collins, D.L., Robles, M., 2010. Adaptive non-local means denoising of MR images with spatially varying noise levels. *J.Magn Reson.Imaging* 31, 192-203.
- Mardia, K.V., Jupp, P.E., 2009. Directional statistics, 494 ed. Wiley. com.
- Margulies, D.S., Bottger, J., Watanabe, A., Gorgolewski, K.J., 2013. Visualizing the human connectome. *NeuroImage*. 80, 445-461.
- Maronna, R.A., Martin, R.D., Yohai, V.J., 2006. Robust statistics. J. Wiley.
- Mattiello, J., Basser, P.J., Le Bihan, D., 1997. The b matrix in diffusion tensor echo-planar imaging. *Magn.Reson.Med.* 37, 292-300.
- Mattiello, J., Basser, P.J., Le Bihan, D., 1994. Analytical expressions for the b matrix in NMR diffusion imaging and spectroscopy. *Journal of magnetic resonance, Series A* 108, 131-141.
- Maximov, I.I., Farrher, E., Grinberg, F., Shah, N.J., 2012. Spatially variable Rician noise in magnetic resonance imaging. *Med.Image Anal.* 16, 536-548.
- McGrath, J., Johnson, K., O'Hanlon, E., Garavan, H., Gallagher, L., Leemans, A., 2013. White Matter and Visuospatial Processing in Autism: A Constrained Spherical Deconvolution Tractography Study. *Autism Res.* 6, 307-319.
- Meer, P., Mintz, D., Rosenfeld, A., Kim, D.Y., 1991. Robust regression methods for computer vision: A review. *International journal of computer vision* 6, 59-70.
- Merhof, D., Sonntag, M., Enders, F., Nimsky, C., Hastreiter, P., Greiner, G., 2006. Hybrid visualization for white matter tracts using triangle strips and point sprites. *IEEE Trans.Vis.Comput.Graphics* 12, 1181-1188.
- Merlet, S.L., Deriche, R., 2013. Continuous diffusion signal, EAP and ODF estimation via Compressive Sensing in diffusion MRI. *Med.Image Anal.* 17, 556-572.
- Metzler-Baddeley, C., Hunt, S., Jones, D.K., Leemans, A., Aggleton, J.P., O'Sullivan, M.J., 2012a. Temporal association tracts and the breakdown of episodic memory in mild cognitive impairment. *Neurology* 79, 2233-2240.
- Metzler-Baddeley, C., O'Sullivan, M.J., Bells, S., Pasternak, O., Jones, D.K., 2012b. How and how not to correct for CSF-contamination in diffusion MRI. *NeuroImage*. 59, 1394-1403.
- Mikula, S., Binding, J., Denk, W., 2012. Staining and embedding the whole mouse brain for electron microscopy. *Nat.Methods*. 9, 1198-1201.
- Miller, K.L., Stagg, C.J., Douaud, G., Jbabdi, S., Smith, S.M., Behrens, T.E., Jenkinson, M., Chance, S.A., Esiri, M.M., Voets, N.L., Jenkinson, N., Aziz, T.Z., Turner, M.R., Johansen-Berg, H., McNab, J.A., 2011. Diffusion imaging of whole, post-mortem human brains on a clinical MRI scanner. *NeuroImage*. 57, 167-181.
- Misner, C.W., Thorne, K.S., Wheeler, J.A., 1973. Gravitation. W. H. Freeman.
- Mitra, P.P., 1995. Multiple wave-vector extensions of the NMR pulsed-field-gradient spin-echo diffusion measurement. *Phys.Rev.B.Condens.Matter* 51, 15074-15078.
- Mohammadi, S., Moller, H.E., Kugel, H., Muller, D.K., Deppe, M., 2010. Correcting eddy current and motion effects by affine whole-brain registrations: evaluation of three-dimensional distortions and comparison with slicewise correction. *Magn.Reson.Med.* 64, 1047-1056.
- Mohammadi, S., Tabelow, K., Ruthotto, L., Feiweier, T., Polzehl, J., Weiskopf, N., 2014. High-resolution diffusion kurtosis imaging at 3T enabled by advanced post-processing. *Front Neurosci.* 8, 427.
- Montgomery, D.C., Runger, G.C., 2010. Applied statistics and probability for engineers. Wiley.
- Morgan, P.S., Bowtell, R.W., McIntyre, D.J., Worthington, B.S., 2004. Correction of spatial distortion in EPI due to inhomogeneous static magnetic fields using the reversed gradient method. *J.Magn Reson.Imaging* 19, 499-507.
- Mori, S., Crain, B.J., Chacko, V.P., van Zijl, P.C., 1999. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Ann.Neurol.* 45, 265-269.
- Mori, S., van Zijl, P.C., 2002. Fiber tracking: principles and strategies - a technical review. *NMR Biomed.* 15, 468-480.
- Mori, S., Zhang, J., 2006. Principles of diffusion tensor imaging and its applications to basic neuroscience research. *Neuron* 51, 527-539.
- Mori, S., 2007. Introduction to diffusion tensor imaging. Elsevier.
- Mori, S., Wakana, S., Van Zijl, P.C., Nagae-Poetscher, L.M., 2005. MRI atlas of human white matter, 16 ed. Am Soc Neuroradiology.
- Moseley, M.E., Cohen, Y., Kucharczyk, J., Mintorovitch, J., Asgari, H.S., Wendland, M.F., Tsuruda, J., Norman, D., 1990. Diffusion-weighted MR imaging of anisotropic water diffusion in cat central nervous system. *Radiology* 176, 439-445.
- Mulkern, R.V., Haker, S.J., Maier, S.E., 2009. On high b diffusion imaging in the human brain: ruminations and experimental insights. *Magn.Reson.Imaging* 27, 1151-1162.
- Neher, P., Houde, J.C., Caruyer, E., Daducci, A., Dyrby, T., Maier-Hein, K.H., Stieltjes, B., Descoteaux, M., 2016. http://www.tractometer.org/ismrm_2015_challenge/.
- Neher, P.F., Laun, F.B., Stieltjes, B., Maier-Hein, K.H., 2014. Fiberfox: facilitating the creation of realistic white matter software phantoms. *Magn.Reson.Med.* 72, 1460-1470.

Less Confusion in Diffusion MRI

- Neto Henriques, R., Correia, M.M., Nunes, R.G., Ferreira, H.A., 2015. Exploring the 3D geometry of the diffusion kurtosis tensor--impact on the development of robust tractography procedures and novel biomarkers. *NeuroImage*. 111, 85-99.
- Niendorf, T., Dijkhuizen, R.M., Norris, D.G., van Lookeren, C.M., Nicolay, K., 1996. Biexponential diffusion attenuation in various states of brain tissue: implications for diffusion-weighted imaging. *Magn.Reson.Med.* 36, 847-857.
- Nilsson, M., Szczepankiewicz, F., van, W.D., Hansson, O., 2015. Extrapolation-Based References Improve Motion and Eddy-Current Correction of High B-Value DWI Data: Application in Parkinson's Disease Dementia. *PLoS One* 10, e0141825.
- Ning, L., Laun, F., Gur, Y., DiBella, E.V., Deslauriers-Gauthier, S., Megherbi, T., Ghosh, A., Zucchelli, M., Menegaz, G., Fick, R., St-Jean, S., Paquette, M., Aranda, R., Descoteaux, M., Deriche, R., O'Donnell, L., Rathi, Y., 2015. Sparse Reconstruction Challenge for diffusion MRI: Validation on a physical phantom to determine which acquisition scheme and analysis method to use? *Med.Image Anal.* 26, 316-331.
- Novikov, D.S., Jensen, J.H., Helpern, J.A., Fieremans, E., 2014. Revealing mesoscopic structural universality with diffusion. *Proc.Natl.Acad.Sci.U.S.A.* 111, 5088-5093.
- Nunes, R.G., Jezzard, P., Clare, S., 2005. Investigations on the efficiency of cardiac-gated methods for the acquisition of diffusion-weighted images. *J.Magn.Reson.* 177, 102-110.
- O'Donnell, L.J., Wells, W.M., III, Golby, A.J., Westin, C.F., 2012. Unbiased groupwise registration of white matter tractography. *Med.Image Comput.Comput.Assist.Interv.* 15, 123-130.
- O'Donnell, L.J., Westin, C.F., Golby, A.J., 2009. Tract-based morphometry for white matter group analysis. *NeuroImage*. 45, 832-844.
- O'Donnell, L.J., Westin, C.F., 2007. Automatic tractography segmentation using a high-dimensional white matter atlas. *IEEE Trans.Med.Imaging* 26, 1562-1575.
- Odish, O.F., Caeyenberghs, K., Hosseini, H., van den Bogaard, S.J., Roos, R.A., Leemans, A., 2015a. Dynamics of the connectome in Huntington's disease: A longitudinal diffusion MRI study. *NeuroImage.Clin.* 9, 32-43.
- Odish, O.F., Leemans, A., Reijntjes, R.H., van den Bogaard, S.J., Dumas, E.M., Wolterbeek, R., Tax, C.M., Kuijf, H.J., Vincken, K.L., van der Grond, J., Roos, R.A., 2015b. Microstructural brain abnormalities in Huntington's disease: A two-year follow-up. *Hum.Brain Mapp.* 36, 2061-2074.
- Oguz, I., Farzinfar, M., Matsui, J., Budin, F., Liu, Z., Gerig, G., Johnson, H.J., Styner, M., 2014. DTIPrep: quality control of diffusion-weighted images. *Front Neuroinform.* 8, 4.
- Olson, D., Arpinar, V., Muftuler, L.T., 2015. Comparison of diffusion kurtosis modeling algorithms: accuracy and application. *Proceedings of the 23rd Annual Meeting of the ISMRM*, 2986.
- Özarslan, E., Basser, P.J., 2008. Microscopic anisotropy revealed by NMR double pulsed field gradient experiments with arbitrary timing parameters. *J.Chem.Phys.* 128, 154511.
- Özarslan, E., Koay, C.G., Shepherd, T.M., Blackband, S., Basser, P.J., 2009. Simple harmonic oscillator based reconstruction and estimation for three-dimensional q-space MRI. *Proceedings of the 17th Annual Meeting of the ISMRM*, 1396.
- Özarslan, E., Koay, C.G., Shepherd, T.M., Komlosh, M.E., Irfanoglu, M.O., Pierpaoli, C., Basser, P.J., 2013. Mean apparent propagator (MAP) MRI: a novel diffusion imaging method for mapping tissue microstructure. *NeuroImage*. 78, 16-32.
- Özarslan, E., Shepherd, T.M., Vemuri, B.C., Blackband, S.J., Mareci, T.H., 2006. Resolution of complex tissue microarchitecture using the diffusion orientation transform (DOT). *NeuroImage*. 31, 1086-1103.
- Özarslan, E., Vemuri, B.C., Mareci, T.H., 2005. Generalized scalar measures for diffusion MRI using trace, variance, and entropy. *Magn.Reson.Med.* 53, 866-876.
- Panagiotaki, E., Schneider, T., Siow, B., Hall, M.G., Lythgoe, M.F., Alexander, D.C., 2012. Compartment models of the diffusion MR signal in brain white matter: a taxonomy and comparison. *NeuroImage*. 59, 2241-2254.
- Pannek, K., Raffelt, D., Bell, C., Mathias, J.L., Rose, S.E., 2012. HOMOR: higher order model outlier rejection for high b-value MR diffusion data. *NeuroImage*. 63, 835-842.
- Paquette, M., Merlet, S., Deriche, R., Descoteaux, M., 2014. DSI 101: Better ODFs for free! In *Proceedings of the 22nd Annual Meeting of the ISMRM*, 7018.
- Parker, G.D., Evans, G.J.A., Jones, D.K., 2016. Structural Fingerprinting of the Human Brain: How unique is tract shape to the individual? In *Proceedings of the 24th Annual Meeting of the ISMRM*, 0122.
- Parker, G.D., Marshall, D., Rosin, P.L., Drage, N., Richmond, S., Jones, D.K., 2013a. A pitfall in the reconstruction of fibre ODFs using spherical deconvolution of diffusion MRI data. *NeuroImage*. 65, 433-448.
- Parker, G.D., Marshall, D., Rosin, P.L., Drage, N., Richmond, S., Jones, D.K., 2013b. RESDORE: Robust Estimation in Spherical Deconvolution by Outlier Rejection. In *Proceedings of the 21st Annual Meeting of the ISMRM*, 3148.
- Pasternak, O., Rathi, Y., Shenton, M., Westin, C.F., 2012. Estimation of the Angle Between Crossing Fibers as a Novel Structural Quantity. In *Proceedings of the 20th Annual Meeting of the ISMRM*, 1915.

References

- Pautler, R.G., Silva, A.C., Koretsky, A.P., 1998. In vivo neuronal tract tracing using manganese-enhanced magnetic resonance imaging. *Magn.Reson.Med.* 40, 740-748.
- Peeters, T.H.J.M., Vilanova, A., Haar Romenij, B.M., 2006. Visualization of DTI fibers using hair-rendering techniques. *ASCI*, 66-73.
- Perrone, D., Aelterman, J., Pizurica, A., Jeurissen, B., Philips, W., Leemans, A., 2015. The effect of Gibbs ringing artifacts on measures derived from diffusion MRI. *NeuroImage* 120, 441-455.
- Perrone, D., Jeurissen, B., Aelterman, J., Roine, T., Sijbers, J., Pizurica, A., Leemans, A., Philips, W., 2016. D-BRAIN: Anatomically Accurate Simulated Diffusion MRI Brain Data. *PLoS.One.* 11, e0149778.
- Peterson, E., Bammer, R., 2016. Survivor's Guide to DTI Acquisition. In: Van Hecke, W., Emsell, L., Sunaert, S. (Eds.), *Diffusion Tensor Imaging*. Springer New York, 89-126.
- Petrovic, V., Fallon, J., Kuester, F., 2007. Visualizing whole-brain DTI tractography with GPU-based Tuboids and LoD management. *IEEE Trans.Vis.Comput.Graph.* 13, 1488-1495.
- Pierpaoli, C., 2010. Artifacts in Diffusion MRI. In: Jones, D.K. (Ed.), *Diffusion MRI: Theory, Methods, and Applications*. Oxford University Press.
- Pierpaoli, C., Walker, L., Irfanoglu, M.O., Barnett, A., Basser, P., Chang, L.C., Koay, C.G., Pajevic, S., Rohde, G., Sarlls, J., Wu, M., 2010. TORTOISE: an integrated software package for processing of diffusion MRI data. In: *Proceedings of the 18th Annual Meeting of the ISMRM*, 1597.
- Pintjens, W., Poot, D.H.J., Verhoye, M., Van der Linden, A., Sijbers, J., 2008. Susceptibility correction for improved tractography using high field DT-EPI. In *Proceedings of SPIE Medical Imaging* 6914, 69142I-9.
- Polders, D.L., Leemans, A., Hendrikse, J., Donahue, M.J., Luijten, P.R., Hoogduin, J.M., 2011. Signal to noise ratio and uncertainty in diffusion tensor imaging at 1.5, 3.0, and 7.0 Tesla. *J.Magn.Reson.Imaging* 33, 1456-1463.
- Poldrack, R.A., Laumann, T.O., Koyejo, O., Gregory, B., Hover, A., Chen, M.Y., Gorgolewski, K.J., Luci, J., Joo, S.J., Boyd, R.L., Hunnicke-Smith, S., Simpson, Z.B., Caven, T., Sochat, V., Shine, J.M., Gordon, E., Snyder, A.Z., Adeyemo, B., Petersen, S.E., Glahn, D.C., Reese, M.D., Curran, J.E., Goring, H.H., Carless, M.A., Blangero, J., Dougherty, R., Leemans, A., Handwerker, D.A., Frick, L., Marcotte, E.M., Mumford, J.A., 2015. Long-term neural and physiological phenotyping of a single human. *Nat.Commun.* 6, 8885.
- Poot, D.H., den Dekker, A.J., Achten, E., Verhoye, M., Sijbers, J., 2010. Optimal experimental design for diffusion kurtosis imaging. *IEEE Trans.Med.Imaging* 29, 819-829.
- Portegies, J.M., Fick, R.H., Sanguineti, G.R., Meesters, S.P., Girard, G., Duits, R., 2015. Improving Fiber Alignment in HARDI by Combining Contextual PDE Flow with Constrained Spherical Deconvolution. *PLoS.One.* 10, e0138122.
- Poupon, C., Rieul, B., Kezele, I., Perrin, M., Poupon, F., Mangin, J.F., 2008. New diffusion phantoms dedicated to the study and validation of high-angular-resolution diffusion imaging (HARDI) models. *Magn.Reson.Med.* 60, 1276-1283.
- Prckovska, V., Andorrá, M., Villoslada, P., Martinez-Heras, E., Duits, R., Fortin, D., Rodrigues, P., Descoteaux, M., 2015. Contextual Diffusion Image Post-processing Aids Clinical Applications. In: Hotz, I., Schultz, T. (Eds.), *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*. Springer International Publishing, Cham, 353-377.
- Presseau, C., Jodoin, P.M., Houde, J.C., Descoteaux, M., 2015. A new compression format for fiber tracking datasets. *NeuroImage* 109, 73-83.
- Pruessmann, K.P., Weiger, M., Scheidegger, M.B., Boesiger, P., 1999. SENSE: sensitivity encoding for fast MRI. *Magn.Reson.Med.* 42, 952-962.
- Pujol, S., Wells, W., Pierpaoli, C., Brun, C., Gee, J., Cheng, G., Vemuri, B., Commowick, O., Prima, S., Stamm, A., Goubran, M., Khan, A., Peters, T., Neher, P., Maier-Hein, K.H., Shi, Y., Tristan-Vega, A., Veni, G., Whitaker, R., Styner, M., Westin, C.F., Gouttard, S., Norton, I., Chauvin, L., Mamata, H., Gerig, G., Nabavi, A., Golby, A., Kikinis, R., 2015. The DTI Challenge: Toward Standardized Evaluation of Diffusion Tensor Imaging Tractography for Neurosurgery. *J.Neuroimaging* 25, 875-882.
- Pullens, P., Roebroeck, A., Goebel, R., 2010. Ground truth hardware phantoms for validation of diffusion-weighted MRI applications. *J.Magn.Reson.Imaging* 32, 482-488.
- Raffelt, D., Tournier, J.D., Fripp, J., Crozier, S., Connelly, A., Salvado, O., 2011. Symmetric diffeomorphic registration of fibre orientation distributions. *NeuroImage* 56, 1171-1180.
- Raffelt, D., Tournier, J.D., Rose, S., Ridgway, G.R., Henderson, R., Crozier, S., Salvado, O., Connelly, A., 2012. Apparent Fibre Density: a novel measure for the analysis of diffusion-weighted magnetic resonance images. *NeuroImage* 59, 3976-3994.
- Ramirez-Manzanares, A., Rivera, M., Vemuri, B.C., Carney, P., Mareci, T., 2007. Diffusion basis functions decomposition for estimating white matter intravoxel fiber geometry. *IEEE Trans.Med.Imaging* 26, 1091-1102.
- Raz, E., Bester, M., Sigmund, E.E., Tabesh, A., Babb, J.S., Jaggi, H., Helpern, J., Mitnick, R.J., Ingles, M., 2013. A better characterization of spinal cord damage in multiple sclerosis: a diffusional kurtosis imaging study. *AJNR.Am.J.Neuroradiol.* 34, 1846-1852.

Less Confusion in Diffusion MRI

- Reber, P.J., Wong, E.C., Buxton, R.B., Frank, L.R., 1998. Correction of off resonance-related distortion in echo-planar imaging using EPI-based field maps. *Magn.Reson.Med.* 39, 328-330.
- Reeder, S.B., Wintersperger, B.J., Dietrich, O., Lanz, T., Greiser, A., Reiser, M.F., Glazer, G.M., Schoenberg, S.O., 2005. Practical approaches to the evaluation of signal-to-noise ratio performance with parallel imaging: application with cardiac imaging and a 32-channel cardiac coil. *Magn.Reson.Med.* 54, 748-754.
- Reese, T.G., Heid, O., Weisskoff, R.M., Wedeen, V.J., 2003. Reduction of eddy-current-induced distortion in diffusion MRI using a twice-refocused spin echo. *Magn.Reson.Med.* 49, 177-182.
- Reijmer, Y.D., Brundel, M., de Bresser, J., Kappelle, L.J., Leemans, A., Biessels, G.J., 2013a. Microstructural white matter abnormalities and cognitive functioning in type 2 diabetes: a diffusion tensor imaging study. *Diabetes Care* 36, 137-144.
- Reijmer, Y.D., Leemans, A., Brundel, M., Kappelle, L.J., Biessels, G.J., 2013b. Disruption of the cerebral white matter network is related to slowing of information processing speed in patients with type 2 diabetes. *Diabetes* 62, 2112-2115.
- Reijmer, Y.D., Leemans, A., Heringa, S.M., Wielaard, I., Jeurissen, B., Koek, H.L., Biessels, G.J., 2012. Improved sensitivity to cerebral white matter abnormalities in Alzheimer's disease with spherical deconvolution based tractography. *PLoS One.* 7, e44074.
- Reina, G., Bidmon, K., Enders, F., Hastreiter, P., Ertl, T., 2006. GPU-Based Hyperstreamlines for Diffusion Tensor Imaging. *EuroVis* 6, 35-42.
- Reisert, M., Kiselev, V.G., Dihtal, B., Kellner, E., Novikov, D.S., 2014. MesoFT: unifying diffusion modelling and fiber tracking. *Med.Image Comput.Comput.Assist.Interv.* 17, 201-208.
- Reisert, M., Skibbe, H., 2016. Left-invariant diffusion on the motion group in terms of the irreducible representations of SO(3).
- Roalf, D.R., Quarckley, M., Elliott, M.A., Satterthwaite, T.D., Vandekar, S.N., Ruparel, K., Gennatas, E.D., Calkins, M.E., Moore, T.M., Hopson, R., Prabhakaran, K., Jackson, C.T., Verma, R., Hakonarson, H., Gur, R.C., Gur, R.E., 2016. The impact of quality assurance assessment on diffusion tensor imaging outcomes in a large-scale population-based cohort. *NeuroImage.* 125, 903-919.
- Rohde, G.K., Barnett, A.S., Bassier, P.J., Marenco, S., Pierpaoli, C., 2004. Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI. *Magn.Reson.Med.* 51, 103-114.
- Roine, T., Jeurissen, B., Perrone, D., Aelterman, J., Leemans, A., Philips, W., Sijbers, J., 2014. Isotropic non-white matter partial volume effects in constrained spherical deconvolution. *Front Neuroinform.* 8, 28.
- Roine, T., Jeurissen, B., Perrone, D., Aelterman, J., Philips, W., Leemans, A., Sijbers, J., 2015. Informed constrained spherical deconvolution (iCSD). *Med.Image.Anal.* 24, 269-281.
- Rokem, A., Yeatman, J.D., Pestilli, F., Kay, K.N., Mezer, A., van der Walt, S., Wandell, B.A., 2015. Evaluating the accuracy of diffusion MRI models in white matter. *PLoS One.* 10, e0123272.
- Ronen, I., Budde, M., Ercan, E., Annese, J., Techawiboonwong, A., Webb, A., 2014. Microstructural organization of axons in the human corpus callosum quantified by diffusion-weighted magnetic resonance spectroscopy of N-acetylaspartate and post-mortem histology. *Brain Struct.Funct.* 219, 1773-1785.
- Rousseeuw, P.J., Croux, C., 1993. Alternatives to the Median Absolute Deviation. *Journal of the American Statistical Association* 88, 1273-1283.
- Roze, E., Binders, M.J., Kersbergen, K.J., van der Aa, N.E., Groenendaal, F., van Haastert, I.C., Leemans, A., de Vries, L.S., 2015. Neonatal DTI early after birth predicts motor outcome in preterm infants with periventricular hemorrhagic infarction. *Pediatr.Res.* 78, 298-303.
- Ruthotto, L., Kugel, H., Olesch, J., Fischer, B., Modersitzki, J., Burger, M., Wolters, C.H., 2012. Diffeomorphic susceptibility artifact correction of diffusion-weighted magnetic resonance images. *Phys.Med.Biol.* 57, 5715-5731.
- Salvador, R., Pena, A., Menon, D.K., Carpenter, T.A., Pickard, J.D., Bullmore, E.T., 2005. Formal characterization and extension of the linearized diffusion tensor model. *Hum.Brain Mapp.* 24, 144-155.
- Sarlls, J.E., Pierpaoli, C., Talagala, S.L., Luh, W.M., 2011. Robust fat suppression at 3T in high-resolution diffusion-weighted single-shot echo-planar imaging of human brain. *Magn.Reson.Med.* 66, 1658-1665.
- Sarra, S.A., 2006. Digital total variation filtering as postprocessing for Chebyshev pseudospectral methods for conservation laws. *Numerical Algorithms* 41, 17-33.
- Savadjiev, P., Rathi, Y., Bouix, S., Verma, R., Westin, C.F., 2012. Multi-scale characterization of white matter tract geometry. *Med.Image Comput.Comput.Assist.Interv.* 15, 34-41.
- Savadjiev, P., Seidman, L.J., Thermenos, H., Keshavan, M., Whitfield-Gabrieli, S., Crow, T.J., Kubicki, M., 2016. Sexual dimorphic abnormalities in white matter geometry common to schizophrenia and non-psychotic high-risk subjects: Evidence for a neurodevelopmental risk marker? *Hum.Brain Mapp.* 37, 254-261.
- Savitzky, A., Golay, M.J.E., 1964. Smoothing and Differentiation of Data by Simplified Least Squares Procedures. *Analytical Chemistry* 36, 1627-1639.

References

- Scherrer, B., Schwartzman, A., Taquet, M., Sahin, Prabhu, S.P., Warfield, S.K., 2015. Characterizing brain tissue by assessment of the distribution of anisotropic microstructural environments in diffusion-compartment imaging (DIAMOND). *Magn.Res.Imag.* doi: 10.1002/mrm.25912.
- Scherrer, B., Warfield, S.K., 2012. Retrospective local artefacts detection in diffusion-weighted images using the random sample consensus (RANSAC) paradigm. *IEEE*, 546-549.
- Scholz, J., Klein, M.C., Behrens, T.E., Johansen-Berg, H., 2009. Training induces changes in white-matter architecture. *Nat.Neurosci.* 12, 1370-1371.
- Schultz, T., Groeschel, S., 2013. Auto-calibrating spherical deconvolution based on ODF sparsity. *Med.Image Comput.Comput.Assist.Interv.* 16, 663-670.
- Schultz, T., Theisel, H., Seidel, H.P., 2010. Crease surfaces: from theory to extraction and application to diffusion tensor MRI. *IEEE Trans.Vis.Comput.Graph.* 16, 109-119.
- Schurade, R., Hlawitschka, M., Hamann, B., Scheuermann, G., Knösche, T.R., Anwander, A., 2010. Visualizing white matter fiber tracts with optimally fitted curved dissection surfaces. In Proceedings of the 2nd Eurographics conference on Visual Computing for Biology and Medicine, 41-48.
- Setsompop, K., Kimmelingen, R., Eberlein, E., Witzel, T., Cohen-Adad, J., McNab, J.A., Keil, B., Tisdall, M.D., Hoecht, P., Dietz, P., Cauley, S.F., Tourneccheva, V., Matschl, V., Lenz, V.H., Heberlein, K., Pothast, A., Thein, H., Van, H.J., Toga, A., Schmitt, F., Lehne, D., Rosen, B.R., Wedeen, V., Wald, L.L., 2013. Pushing the limits of in vivo diffusion MRI for the Human Connectome Project. *NeuroImage* 80, 220-233.
- Sharman, M.A., Cohen-Adad, J., Descoteaux, M., Messe, A., Benali, H., Lehéricy, S., 2011. Impact of outliers on diffusion tensor and Q-ball imaging: clinical implications and correction strategies. *J.Magn.Reson.Imaging* 33, 1491-1502.
- Sherbondy, A.J., Dougherty, R.F., Ben-Shachar, M., Napel, S., Wandell, B.A., 2008. ConTrack: finding the most likely pathways between brain regions using diffusion tractography. *J.Vis.* 8, 15-16.
- Sijbers, J., den Dekker, A.J., 2004. Maximum likelihood estimation of signal amplitude and noise variance from MR data. *Magn.Reson.Med.* 51, 586-594.
- Sijbers, J., den Dekker, A.J., Van, A.J., Verhoye, M., Van, D.D., 1998. Estimation of the noise in magnitude MR images. *Magn Reson.Imaging* 16, 87-90.
- Sijbers, J., Poot, D., den Dekker, A.J., Pintjens, W., 2007. Automatic estimation of the noise variance from the histogram of a magnetic resonance image. *Phys.Med Biol.* 52, 1335-1348.
- Smith, R.E., 2016. http://mrtrix.readthedocs.io/en/latest/concepts/response_function_estimation.html.
- Smith, R.E., Tournier, J.D., Calamante, F., Connelly, A., 2013. SIFT: Spherical-deconvolution informed filtering of tractograms. *NeuroImage*. 67, 298-312.
- Smith, S.M., Jenkinson, M., Johansen-Berg, H., Rueckert, D., Nichols, T.E., Mackay, C.E., Watkins, K.E., Ciccarelli, O., Cader, M.Z., Matthews, P.M., Behrens, T.E., 2006. Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. *NeuroImage*. 31, 1487-1505.
- Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobniak, I., Flitney, D.E., Niazy, R.K., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J.M., Matthews, P.M., 2004. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*. 23 Suppl 1, S208-S219.
- Sotiroopoulos, S.N., Behrens, T.E., Jbabdi, S., 2012. Ball and rackets: Inferring fiber fanning from diffusion-weighted MRI. *NeuroImage* 60, 1412-1425.
- Sotiroopoulos, S.N., Jbabdi, S., Xu, J., Andersson, J.L., Moeller, S., Auerbach, E.J., Glasser, M.F., Hernandez, M., Sapiro, G., Jenkinson, M., Feinberg, D.A., Yacoub, E., Lenglet, C., Van Essen, D.C., Ugurbil, K., Behrens, T.E., 2013. Advances in diffusion MRI acquisition and processing in the Human Connectome Project. *NeuroImage*. 80, 125-143.
- Spivak, M., 1979. A comprehensive introduction to differential geometry. Publish or Perish, inc.
- St-Jean, S., Coupé, P., Descoteaux, M., 2016. Non Local Spatial and Angular Matching: Enabling higher spatial resolution diffusion MRI datasets through adaptive denoising. *Med.Image Anal.* 32, 115-130.
- Stanisz, G.J., Szafer, A., Wright, G.A., Henkelman, R.M., 1997. An analytical model of restricted diffusion in bovine optic nerve. *Magn.Reson.Med.* 37, 103-111.
- Stocker, T., Vahedipour, K., Pflugfelder, D., Shah, N.J., 2010. High-performance computing MRI simulations. *Magn.Reson.Med.* 64, 186-193.
- Street, J., Carroll, R.J., Ruppert, D., 1988. A note on computing robust regression estimates via iteratively reweighted least squares. *The American Statistician* 42, 152-154.
- Studholme, C., Hill, D.L.G., Hawkes, D.J., 1999. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition* 32, 71-86.
- Sundgren, P.C., Dong, Q., Gomez-Hassan, D., Mukherji, S.K., Maly, P., Welsh, R., 2004. Diffusion tensor imaging of the brain: review of clinical applications. *Neuroradiology* 46, 339-350.

Less Confusion in Diffusion MRI

- Szczepankiewicz, F., Latt, J., Wirestam, R., Leemans, A., Sundgren, P., van, W.D., Stahlberg, F., Nilsson, M., 2013. Variability in diffusion kurtosis imaging: impact on study design, statistical power and interpretation. *NeuroImage*. 76, 145-154.
- Tabelow, K., Voss, H.U., Polzehl, J., 2015. Local estimation of the noise level in MRI using structural adaptation. *Med.Image.Anal.* 20, 76-86.
- Tabesh, A., Jensen, J.H., Ardekani, B.A., Helpern, J.A., 2011a. Estimation of tensors and tensor-derived measures in diffusional kurtosis imaging. *Magn.Reson.Med.* 65, 823-836.
- Tabesh, A., Jensen, J.H., Fieremans, E., Helpern, J.A., 2011b. Q-space undersample diffusional kurtosis imaging. In Proceedings of the 19th Annual Meeting of the ISMRM, 1990.
- Tao, R., Fletcher, P.T., Gerber, S., Whitaker, R.T., 2009. A variational image-based approach to the correction of susceptibility artifacts in the alignment of diffusion weighted and structural MRI. *Inf.Process Med.Imaging* 21, 664-675.
- Taquet, M., Scherrer, B., Commowick, O., Peters, J.M., Sahin, M., Macq, B., Warfield, S.K., 2014. A mathematical framework for the registration and analysis of multi-fascicle models for population studies of the brain microstructure. *IEEE Trans.Med.Imaging* 33, 504-517.
- Tariq, M., Schneider, T., Alexander, D.C., Gandini Wheeler-Kingshott, C.A., Zhang, H., 2016. Bingham-NODDI: Mapping anisotropic orientation dispersion of neurites using diffusion MRI. *NeuroImage*. 133, 207-223.
- Tax, C.M.W., Dela Haije, T., Fuster, A., Duits, R., Viergever, M.A., Calabrese, E., Johnson, G.A., Florack, L.M.J., Leemans, A., 2015a. Towards quantification of the brain's sheet structure in diffusion MRI data. International BASP Frontiers Workshop, 74.
- Tax, C.M.W., Dela Haije, T., Fuster, A., Duits, R., Viergever, M.A., Florack, L.M.J., Leemans, A., 2014a. Towards quantification of the brain's sheet structure: Evaluation of the discrete Lie bracket. In Proceedings of the 22nd Annual Meeting of the ISMRM, 0975.
- Tax, C.M.W., Dela Haije, T., Fuster, A., Westin, C.F., Viergever, M.A., Florack, L.M.J., Leemans, A., 2016a. Mapping the brain's "Sheet Probability Index" (SPI) with diffusion MRI: Sheet happens?! In Proceedings of the 24th Annual Meeting of the ISMRM, 0791.
- Tax, C.M.W., Duits, R., Romeny, B.M., Vilanova, A., Ossenblok, P., 2012. Tractography of the Optic Radiation for Vision Sparing Epilepsy Surgery. In Proceedings of the IEEE ICIA, 441-445.
- Tax, C.M.W., Duits, R., Vilanova, A., Romeny, B.M., Hofman, P., Wagner, L., Leemans, A., Ossenblok, P., 2014b. Evaluating Contextual Processing in Diffusion MRI: Application to Optic Radiation Reconstruction for Epilepsy Surgery. *PLoS One* 9, e101524.
- Tax, C.M.W., Jeurissen, B., Vos, S.B., Viergever, M.A., Leemans, A., 2014c. Recursive calibration of the fiber response function for spherical deconvolution of diffusion MRI data. *NeuroImage* 86, 67-80.
- Tax, C.M.W., Novikov, D.S., Garyfallidis, E., Viergever, M.A., Descoteaux, M., Leemans, A., 2015b. Localizing and characterizing single fiber populations throughout the brain. Proceedings of the 23rd Annual Meeting of the ISMRM, 0473.
- Tax, C.M.W., Otte, W.M., Viergever, M.A., Dijkhuizen, R.M., Leemans, A., 2013. REKINDLE: Robust Extraction of Kurtosis INDices with Linear Estimation. In Proceedings of the 21st Annual Meeting of the ISMRM 0700.
- Tax, C.M.W., Otte, W.M., Viergever, M.A., Dijkhuizen, R.M., Leemans, A., 2015c. REKINDLE: Robust Extraction of Kurtosis INDices with Linear Estimation. *Magn.Reson.Med.* 73, 794-808.
- Tax, C.M.W., Stralen, M., Viergever, M.A., Ramsey, N., Leemans, A., 2014d. Orientation-dependent rendering of diffusion fiber tractography streamlines for improved visualization of complex tissue organization. In Proceedings of the 21st Annual Meeting of the ISMRM, 4529.
- Tax, C.M.W., Vos, S.B., Viergever, M.A., Froeling, M., Leemans, A., 2014e. Transforming grids to shells and vice versa: an evaluation of interpolation methods in diffusion MRI. In Proceedings of the 22nd Annual Meeting of the ISMRM, 4485.
- Tax, C.M.W., Westin, C.F., Dela Haije, T., Fuster, A., Viergever, M.A., Florack, L.M.J., Leemans, A., 2016b. Robust assessment of the brain's sheet structure using normalized convolution. In Proceedings of the 24th Annual Meeting of the ISMRM, 2073.
- Tax, C.M.W., Vos, S.B., Leemans, A., 2016c. Checking and Correcting DTI Data. In: Van Hecke, W., Emsell, L., Sunaert, S. (Eds.), *Diffusion Tensor Imaging*. Springer New York, 127-150.
- Teipel, S.J., Reuter, S., Stieltjes, B., Acosta-Cabronero, J., Ernemann, U., Fellgiebel, A., Filippi, M., Frisoni, G., Hentschel, F., Jessen, F., Kloppel, S., Meindl, T., Pouwels, P.J., Hauenstein, K.H., Hampel, H., 2011. Multicenter stability of diffusion tensor imaging measures: a European clinical and physical phantom study. *Psychiatry Res.* 194, 363-371.
- Tian, Q., Rokem, A., Folkerth, R.D., Nummenmaa, A., Fan, Q., Edlow, B.L., McNab, J.A., 2016. Q-space truncation and sampling in diffusion spectrum imaging. *Magn.Reson.Med.*

References

- Tobisch, A., Stöcker, T., Gröschel, S., Schultz, T., 2016. Iteratively Reweighted L1-Fitting for Model-Independent Outlier Removal and Regularization in Diffusion MRI.
- Tournier, J., Calamante, F., Connelly, A., 2012. MRtrix: diffusion tractography in crossing fiber regions. International Journal of Imaging Systems and Technology 22, 53-66.
- Tournier, J.D., Calamante, F., Connelly, A., 2007. Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution. NeuroImage. 35, 1459-1472.
- Tournier, J.D., Calamante, F., Connelly, A., 2009. How many diffusion gradient directions are required for HARDI. Proceedings of the 17nd Annual Meeting of the ISMRM, 358.
- Tournier, J.D., Calamante, F., Connelly, A., 2013a. A Robust Spherical Deconvolution Method for the Analysis of Low SNR or Low Angular Resolution Diffusion Data. Proceedings of the 21st Annual Meeting of the ISMRM, 772.
- Tournier, J.D., Calamante, F., Connelly, A., 2013b. Determination of the appropriate b value and number of gradient directions for high-angular-resolution diffusion-weighted imaging. NMR Biomed. 26, 1775-1786.
- Tournier, J.D., Calamante, F., Gadian, D.G., Connelly, A., 2004. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. NeuroImage. 23, 1176-1185.
- Tournier, J.D., Mori, S., Leemans, A., 2011. Diffusion tensor imaging and beyond. Magn.Reson.Med. 65, 1532-1556.
- Tristan-Vega, A., Aja-Fernandez, S., 2010. DWI filtering using joint information for DTI and HARDI. Med.Image Anal. 14, 205-218.
- Tristan-Vega, A., Aja-Fernandez, S., Westin, C.F., 2012. Least squares for diffusion tensor estimation revisited: propagation of uncertainty with Rician and non-Rician signals. NeuroImage. 59, 4032-4043.
- Tuch, D.S., 2004. Q-ball imaging. Magn.Reson.Med. 52, 1358-1372.
- Tuch, D.S., Reese, T.G., Wiegell, M.R., Makris, N., Belliveau, J.W., Wedeen, V.J., 2002. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. Magn.Reson.Med. 48, 577-582.
- Vaillancourt, O., Chamberland, M., Houde, J.C., Descoteaux, M., 2015. Visualization of Diffusion Propagator and Multiple Parameter Diffusion Signal. In: Schultz, T. (Ed.), Visualization and Processing of Higher Order Descriptors for Multi-Valued Data. Springer.
- Van Cauter, S., Veraart, J., Sijbers, J., Peeters, R.R., Himmelreich, U., De Keyzer, F., Van Gool, S.W., Van Calenbergh, F., De Vleeschouwer, S., Van Hecke, W., Sunaert, S., 2012. Gliomas: diffusion kurtosis MR imaging in grading. Radiology 263, 492-501.
- Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., 2013. The WU-Minn Human Connectome Project: an overview. NeuroImage 80, 62-79.
- Van Gelderen, P., de Vleeschouwer, M.H., DesPres, D., Pekar, J., van Zijl, P.C., Moonen, C.T., 1994. Water diffusion and acute stroke. Magn.Reson.Med. 31, 154-163.
- Van Hecke, W., Emsell, L., Sunaert, S., 2016. Diffusion Tensor Imaging: A Practical Handbook. Springer.
- Varentsova, A., Zhang, S., Arfanakis, K., 2014. Development of a high angular resolution diffusion imaging human brain template. NeuroImage. 91, 177-186.
- Veraart, J., Fieremans, E., Jelescu, I.O., Knoll, F., Novikov, D.S., 2016. Gibbs ringing in diffusion MRI. Magn.Reson.Med. 76, 301-314.
- Veraart, J., Fieremans, E., Novikov, D.S., 2015. Diffusion MRI noise mapping using random matrix theory. Magn.Reson.Med. doi: 10.1002/mrm.26059.
- Veraart, J., Poot, D.H., Van, H.W., Blockx, I., Van der Linden, A., Verhoye, M., Sijbers, J., 2011a. More accurate estimation of diffusion tensor parameters using diffusion Kurtosis imaging. Magn.Reson.Med. 65, 138-145.
- Veraart, J., Rajan, J., Peeters, R.R., Leemans, A., Sunaert, S., Sijbers, J., 2013a. Comprehensive framework for accurate diffusion MRI parameter estimation. Magn.Reson.Med. 70, 972-984.
- Veraart, J., Sijbers, J., Sunaert, S., Leemans, A., Jeurissen, B., 2013b. Weighted linear least squares estimation of diffusion MRI parameters: strengths, limitations, and pitfalls. NeuroImage. 81, 335-346.
- Veraart, J., Van, H.W., Sijbers, J., 2011b. Constrained maximum likelihood estimation of the diffusion kurtosis tensor using a Rician noise model. Magn.Reson.Med. 66, 678-686.
- Verhoeven, J.S., Rommel, N., Prodi, E., Leemans, A., Zink, I., Vandewalle, E., Noens, I., Wagemans, J., Steyaert, J., Boets, B., Van de Winckel, A., De, C.P., Lagae, L., Sunaert, S., 2012. Is There a Common Neuroanatomical Substrate of Language Deficit between Autism Spectrum Disorder and Specific Language Impairment? Cereb.Cortex 22, 2263-2271.
- Vilanova, A., Berenschot, G., Van Pul, C., 2004. DTI visualization with streamsurfaces and evenly-spaced volume seeding. Eurographics Association, 173-182.
- Vos, S.B., Aksoy, M., Han, Z., Holdsworth, S.J., Maclarens, J., Viergever, M.A., Leemans, A., Bammer, R., 2016a. Trade-off between angular and spatial resolutions in in vivo fiber tractography. NeuroImage. 129, 117-132.
- Vos, S.B., Jones, D.K., Jeurissen, B., Viergever, M.A., Leemans, A., 2012. The influence of complex white matter architecture on the mean diffusivity in diffusion tensor MRI of the human brain. NeuroImage. 59, 2208-2216.

Less Confusion in Diffusion MRI

- Vos, S.B., Tax, C.M., Luijten, P.R., Ourselin, S., Leemans, A., Froeling, M., 2016b. The importance of correcting for signal drift in diffusion MRI. *Magn.Reson.Med.* doi: 10.1002/mrm.26124.
- Vos, S.B., Tax, C.M.W., Froeling, M., Leemans, A., 2014. In vivo investigation of accuracy and precision of fiber orientations in crossing fibers in spherical deconvolution-based HARDI methods. In Proceedings of the 22nd Annual Meeting of the ISMRM, 0976.
- Vos, S.B., Viergever, M.A., Leemans, A., 2013. Multi-fiber tractography visualizations for diffusion MRI data. *PLoS.One*, 8, e81453.
- Wakana, S., Caprihan, A., Panzenboeck, M.M., Fallon, J.H., Perry, M., Gollub, R.L., Hua, K., Zhang, J., Jiang, H., Dubey, P., Blitz, A., van, Z.P., Mori, S., 2007. Reproducibility of quantitative tractography methods applied to cerebral white matter. *NeuroImage*. 36, 630-644.
- Walker, L., Chang, L.C., Koay, C.G., Sharma, N., Cohen, L., Verma, R., Pierpaoli, C., 2011. Effects of physiological noise in population analysis of diffusion tensor MRI data. *NeuroImage*. 54, 1168-1177.
- Walker, L., Chang, L.C., Nayak, A., Irfanoglu, M.O., Botteron, K.N., McCracken, J., McKinstry, R.C., Rivkin, M.J., Wang, D.J., Rumsey, J., Pierpaoli, C., 2016. The diffusion tensor imaging (DTI) component of the NIH MRI study of normal brain development (PedDTI). *NeuroImage*. 124, 1125-1130.
- Walker, L., Curry, M., Nayak, A., Lange, N., Pierpaoli, C., 2012. A framework for the analysis of phantom data in multicenter diffusion tensor imaging studies. *Hum.Brain Mapp.*
- Wang, H.C., Hsu, J.L., Leemans, A., 2012. Diffusion tensor imaging of vascular parkinsonism: structural changes in cerebral white matter and the association with clinical severity. *Arch.Neurol.* 69, 1340-1348.
- Wang, R., Benner, T., Sorensen, A.G., Wedeen, V.J., 2007. Diffusion Toolkit: A Software Package for Diffusion Imaging Data Processing and Tractography. In Proceedings of the 15th Annual Meeting of the ISMRM, 3720.
- Wedeen, V.J., Hagmann, P., Tseng, W.Y., Reese, T.G., Weisskoff, R.M., 2005. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magn.Reson.Med.* 54, 1377-1386.
- Wedeen, V.J., Mortazavi, F., Wang, R., Isaac Tseng, W., Witzel, T., Nummenmaa, A., Morrison, W., Stanley, H.E., Wald, L., Rosene, D., 2014. Grid structure of brain pathways - Validation and the character of turns. In Proceedings of the 22nd Annual Meeting of the ISMRM, 0803.
- Wedeen, V.J., Rosene, D.L., Wang, R., Dai, G., Mortazavi, F., Hagmann, P., Kaas, J.H., Tseng, W.Y., 2012a. Response to comment on "the geometric structure of the brain fiber pathways". *Science* 337, 1605.
- Wedeen, V.J., Rosene, D.L., Wang, R., Dai, G., Mortazavi, F., Hagmann, P., Kaas, J.H., Tseng, W.Y., 2012b. The geometric structure of the brain fiber pathways. *Science* 335, 1628-1634.
- Weldeselassie, Y.T., Barmpoutis, A., Atkins, M.S., 2012. Symmetric positive semi-definite Cartesian Tensor fiber orientation distributions (CT-FOD). *Med.Image.Anal.* 16, 1121-1129.
- Werring, D.J., Clark, C.A., Barker, G.J., Thompson, A.J., Miller, D.H., 1999. Diffusion tensor imaging of lesions and normal-appearing white matter in multiple sclerosis. *Neurology* 52, 1626-1632.
- Westin, C.F., Knutsson, H., Pasternak, O., Szczepankiewicz, F., Ozarslan, E., van, W.D., Mattison, C., Bogren, M., O'Donnell, L., Kubicki, M., Topgaard, D., Nilsson, M., 2016. q-space trajectory imaging for multidimensional diffusion MRI of the human brain. *NeuroImage*.
- Westin, C.F., Maier, S.E., Mamata, H., Nabavi, A., Jolesz, F.A., Kikinis, R., 2002. Processing and visualization for diffusion tensor MRI. *Med.Image.Anal.* 6, 93-108.
- Westin, C.F., Martin-Fernandez, M., Alberola-Lopez, C., Ruiz-Alzola, J., Knutsson, H., 2006. Tensor field regularization using normalized convolution and markov random fields in a Bayesian Framework. In: Weickert, J., Hagen, H. (Eds.), *Visualization and Image Processing of Tensor Fields*. Springer, 381-398.
- White, T., Nelson, M., Lim, K.O., 2008. Diffusion tensor imaging in psychiatric disorders. *Top.Magn Reson.Imaging* 19, 97-109.
- Whittingstall, K., Bernier, M., Houde, J.C., Fortin, D., Descoteaux, M., 2014. Structural network underlying visuospatial imagery in humans. *Cortex* 56, 85-98.
- Wu, M., Chang, L.C., Walker, L., Lemaitre, H., Barnett, A.S., Marenco, S., Pierpaoli, C., 2008a. Comparison of EPI distortion correction methods in diffusion tensor MRI using a novel framework. *Med.Image Comput.Comput.Assist.Interv.* 11, 321-329.
- Wu, Y.C., Alexander, A.L., 2007. Hybrid diffusion imaging. *NeuroImage*. 36, 617-629.
- Wu, Y.C., Field, A.S., Alexander, A.L., 2008b. Computation of diffusion function measures in q-space using magnetic resonance hybrid diffusion imaging. *IEEE Trans.Med.Imaging* 27, 858-865.
- Ye, C., Zhuo, J., Gullapalli, R.P., Prince, J.L., 2016. Estimation of Fiber Orientations Using Neighborhood Information. *Medical Image Analysis*.
- Yeh, F.C., Wedeen, V.J., Tseng, W.Y., 2010. Generalized q-sampling imaging. *IEEE Trans.Med.Imaging* 29, 1626-1635.
- Yendiki, A., Koldewyn, K., Kakunoori, S., Kanwisher, N., Fischl, B., 2013. Spurious group differences due to head motion in a diffusion MRI study. *NeuroImage* 88C, 79-90.

References

- Yushkevich, P.A., Zhang, H., Simon, T.J., Gee, J.C., 2008. Structure-specific statistical mapping of white matter tracts. *NeuroImage*. 41, 448-461.
- Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C., 2012. NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage*. 61, 1000-1016.
- Zhang, H., Yushkevich, P.A., Alexander, D.C., Gee, J.C., 2006a. Deformable registration of diffusion tensor MR images with explicit orientation optimization. *Med.Image Anal.* 10, 764-785.
- Zhang, J., van Zijl, P.C., Mori, S., 2006b. Image contrast using the secondary and tertiary eigenvectors in diffusion tensor imaging. *Magn.Reson.Med.* 55, 439-449.
- Zhang, M., Fletcher, P.T., 2015. Finite-Dimensional Lie Algebras for Fast Diffeomorphic Image Registration. *Inf.Process Med.Imaging* 24, 249-259.
- Zhang, S., Peng, H., Dawe, R.J., Arfanakis, K., 2011. Enhanced ICBM diffusion tensor template of the human brain. *NeuroImage*. 54, 974-984.
- Zhang, S., Demiralp, C., Laidlaw, D.H., 2003. Visualizing diffusion tensor MR images using streamtubes and streamsurfaces. *Visualization and Computer Graphics, IEEE Trans.Vis.Comput.Graphics* 9, 454-462.
- Zhang, Y., Brady, M., Smith, S., 2001. Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans.Med.Imaging* 20, 45-57.
- Zhang, Y., Wehrli, F.W., 2004. Reference-scan-free method for automated correction of Nyquist ghost artifacts in echoplanar brain images. *Magn.Reson.Med.* 51, 621-624.
- Zhou, M.X., Yan, X., Xie, H.B., Zheng, H., Xu, D., Yang, G., 2015. Evaluation of non-local means based denoising filters for diffusion kurtosis imaging using a new phantom. *PLoS.One.* 10, e0116986.
- Zhou, Z., Liu, W., Cui, J., Wang, X., Arias, D., Wen, Y., Bansal, R., Hao, X., Wang, Z., Peterson, B.S., Xu, D., 2011. Automated artifact detection and removal for improved tensor estimation in motion-corrupted DTI data sets using the combination of local binary patterns and 2D partial least squares. *Magn Reson.Imaging* 29, 230-242.
- Zhuang, J., Lu, Z.L., Vidal, C.B., Damasio, H., 2013. Correction of eddy current distortions in high angular resolution diffusion imaging. *J.Magn Reson.Imaging* 37, 1460-1467.
- Zucchelli, M., Brusini, L., Méndez, C.A.s., Menegaz, G., 2016. Multi-Tensor MAPMRI: How to Estimate Microstructural Information from Crossing Fibers. In: Fuster, A., Ghosh, A., Kaden, E., Rathi, Y., Reisert, M. (Eds.), *Computational Diffusion MRI: Med.Image Comput.Comput.Assist.Interv. Workshop*, Munich, Germany, 2015. Springer International Publishing, Cham, 65-74.
- Zwiers, M.P., 2010. Patching cardiac and head motion artefacts in diffusion-weighted images. *NeuroImage*. 53, 565-575.

Less Confusion in Diffusion MRI

SAMENVATTING

ACHTERGROND

Met magnetische kernspinresonantiebeeldvorming (magnetic resonance imaging, MRI) kunnen zachte weefsels in het lichaam afgebeeld worden op een niet-invasieve manier. MRI is daardoor onmisbaar bij de diagnose van veel ziektes. Veel zachte lichaamsweefsels zijn opgebouwd uit vezelachtige structuren, denk daarbij bijvoorbeeld aan zenuwvezels in de witte stof van de hersenen en spieren. Conventionele MRI technieken kunnen deze typen weefsels wel afbeelden, maar geven geen informatie over de onderliggende *architectuur* van de weefsels. De witte stof in de hersenen, bijvoorbeeld, wordt door deze technieken afgebeeld als één massa, waarbij elke informatie over de loop van de zenuwvezels ontbreekt. Diffusie MRI (dMRI) kan in tegenstelling tot conventionele MRI technieken ook informatie verschaffen over de configuratie van onderliggende vezelstructuren. Hiertoe wordt de MRI sequentie gevoelig gemaakt voor de willekeurige beweging van deeltjes (diffusie). dMRI is bijvoorbeeld bijzonder bruikbaar om gebieden getroffen door een acute beroerte te herkennen die niet zonder meer gedetecteerd kunnen worden op conventionele MRI scans.

Diffusie of Brownse beweging is de willekeurige beweging van moleculen ten gevolge van de thermische energie die ze bezitten. De relatie tussen de tijd die moleculen hebben om te diffunderen en hun gemiddelde verplaatsing kan beschreven worden met behulp van de diffusiecoëfficiënt. Deze is onder andere afhankelijk van de temperatuur, de grootte van de moleculen en het medium waarin diffusie plaatsvindt. De structuren die aanwezig zijn in weefsels, zoals celmembranen en myeline rond zenuwvezels, hebben invloed op de grootte van diffusie. De watermoleculen bewegen nog steeds op willekeurige wijze, maar worden gehinderd door verschillende microstructuren. Moleculen kunnen door deze hindering bijvoorbeeld meer bewegen in de richting parallel aan vezelstructuren dan loodrecht daarop. Deze *anisotropie* diffusie kan niet meer beschreven worden met één enkele diffusiecoëfficiënt. Door te kijken naar de grootte van diffusie in verschillende richtingen kan informatie verkregen worden over de onderliggende vezelstructuur.

Om MRI gevoelig te maken voor diffusie, wordt een extra magnetisch gradiëntveld toegepast. Door de richting van deze gradiënt te variëren, kan informatie verkregen worden over de richtingsafhankelijkheid van de diffusie. Om het MRI signaal te relateren aan het diffusieproces of de weefsel configuratie worden vaak modellen gebruikt. Het meest eenvoudige model dat het drie-dimensionale signaal relateert aan diffusie is een directe extensie van de diffusiecoëfficiënt naar drie dimensies, namelijk de *diffusietensor*. In diffusietensor imaging (DTI) wordt een dergelijke diffusietensor geschat op elke locatie, en vaak afgebeeld als een ellipsoïde. Belangrijke maten die uit de diffusietensor afgeleid kunnen worden zijn de fractionele anisotropie (FA) en de gemiddelde diffusiviteit (mean diffusivity, MD). Een FA dicht bij 1 beschrijft een ellipsoïde met een sigaarvorm, terwijl een lage FA overeenkomt met een bolvormige ellipsoïde. De FA en MD zijn gevoelige markers voor veranderingen in het weefsel als resultaat van verschillende ziektes.

Samenvatting

Naast het in kaart brengen van microstructurele eigenschappen, biedt dMRI de unieke mogelijkheid om de structurele connectiviteit van weefsel te bestuderen. De reconstructie en visualisatie van dergelijke vezelstructuren met behulp van dMRI wordt *tractografie* genoemd. Tractografie wordt het meest toegepast in de hersenen, maar kan ook in andere delen van het lichaam toegepast worden.

CONFUSIE IN DIFFUSIE MRI

De potentie van dMRI in een klinische en wetenschappelijke context werd al in een vroeg stadium erkend, en vandaag de dag groeit de interesse in het gebruik van diffusie MRI in klinisch en neurowetenschappelijk onderzoek nog steeds exponentieel. Dit heeft onderzoekers gemotiveerd om steeds nieuwe methodes te ontwikkelen om dMRI data te analyseren. Deze nieuwe ontwikkelingen richten zich bijvoorbeeld op het beter modelleren van het diffusieproces. Een limitatie van DTI is dat de diffusietensor geen informatie kan geven over kruisende of splitsende vezels. In deze situaties wordt de ellipsoïde bol- of pannekoek-vormig, en is het voor een tractografie-algoritme moeilijk om de richting van de onderliggende vezel te kunnen bepalen. *Sferische deconvolutie* (SD) is één van de methoden die dergelijke meer ingewikkelde constructies beter kunnen oplossen door het schatten van de vezel oriëntatie-distributie functie (fiber orientation distribution function, fODF). Een andere limitatie van DTI is de onderliggende Gaussische benadering van het diffusie proces, waaraan niet altijd voldaan wordt. *Diffusie kurtosis imaging* (DKI) is voorgesteld als een extensie van DTI waarmee ook niet-Gaussische diffusiematen berekend kunnen worden. Naast deze twee modellen, waarvan al is aangetoond dat ze bruikbaar zijn in verschillende toepassingen, bestaan er wel meer dan honderd alternatieve dMRI modellen.

Geavanceerde diffusie modellen maken vaak gebruik van meer en veeleisendere MRI acquisitions, wat de kans op artefacten vergroot. Het onderzoek naar nieuwe diffusie modellen gaat dan ook hand in hand met nieuwe ontwikkelingen in correctie strategieën voor onder andere beweging van de proefpersoon en geometrische vervormingen van het beeld. De toenemende complexiteit van deze modellen vraagt ook om nieuwe methoden op het gebied van tractografie, kwantitatieve vergelijking in groepsstudies, en visualisatie.

Door het alsmaar toenemende aanbod aan analyse-strategieën is het niet altijd duidelijk voor eindgebruikers hoe dMRI optimaal gebruikt kan worden om een vraagstuk te bestuderen, en welke conclusies er nu wel en niet betrouwbaar getrokken kunnen worden uit dMRI data. Dit leidt soms zelfs tot controverses in het onderzoeksgebied, zoals een recent debat in het wetenschappelijke tijdschrift Science over het wel of niet bestaan van zogenaamde ‘sheet-structuren’ in de hersenen. Zulke onduidelijkheden complicerden de translatie van nieuw ontwikkelde methodologie naar bruikbare instrumenten in klinische en biomedische toepassingen.

In dit proefschrift worden verschillende dMRI analyse stappen onder de loep genomen, met als doel om de confusie in diffusie MRI niet verder te vergroten en de vertaalslag naar neurowetenschappelijke en klinische toepassingen te faciliteren. De hoofdstukken in dit proefschrift focussen op het meer toegankelijk maken van dMRI analyse, en beschrijven nieuwe methodologie om intuitivere en datagestuurdere keuzes te kunnen maken en om

interpretatie en visualisatie van dMRI data te faciliteren. De methodologie ontwikkeld in dit proefschrift wordt verder gebruikt om fundamentele vraagstukken te onderzoeken, zoals de variabiliteit en karakteristieken van het dMRI signaal en de geometrische organisatie van de hersenen. Computer simulaties en evaluaties op echte data spelen een belangrijke rol.

EVALUATIE VAN DIFFUSIE MRI METHODOLOGIE

De beeldverwerkingspipeline voor dMRI data is vaak lang en gecompliceerd. Voor betrouwbaar gebruik van de data zijn veel verschillende stappen nodig, zoals 1) correctie voor artefacten, 2) het selecteren van een model, en 3) parameter schatting. **Hoofdstuk 2** beschrijft deze stappen en geeft een overzicht van de meest gebruikte en een aantal nieuwe dMRI technieken, met een focus op de hersenen. Dit overzicht kan helpen bij het kiezen van een optimale dMRI analyse strategie voor een bepaalde applicatie. Om de inhoud toegankelijk te maken voor een breder publiek, wordt er minimaal gebruikt gemaakt van vergelijkingen en technisch jargon.

Ondanks de ontwikkelingen op het gebied van dMRI beeldverwerking beschreven in hoofdstuk 2 blijft het optimaliseren van de beeldverwerkingspipeline nog steeds een actief onderzoeksgebied. Een omvangrijke evaluatie speelt daarom een belangrijke rol in het bereiken van een consensus over bepaalde keuzes voor de beeldverwerking. **Hoofdstuk 3** beschrijft de Multiple Acquisitions for Standardization of Structural Imaging Validation and Evaluation (MASSIVE) database voor het evalueren en vergelijken van nieuwe en bestaande MRI beeldverwerkingstechnieken. Deze unieke database bevat 8000 verschillende dMRI scans, het grootste aantal *in vivo* diffusie MRI scans van het brein van één persoon. Daarnaast bevat de database 10 T1-gewogen, 10 T2-gewogen, 10 fluid-attenuated inversion recovery (FLAIR) MRI scans, herhaalde ruis metingen, en herhaalde magnetische inhomogeniteitsmetingen. Door het samenstellen van deelverzamelingen uit de uitgebreide MASSIVE database kan bijvoorbeeld de invloed van het diffusie MRI protocol getest worden, de rol van andere MRI modaliteiten zoals T1 en T2-gewogen scans bestudeerd worden, en data-artefact correctiestrategieën ontwikkeld en vergeleken worden. Hoewel er andere databases beschikbaar zijn voor dergelijke doelen, zijn deze vaak gelimiteerd in het aantal diffusie MRI scans van eenzelfde persoon. In tegenstelling tot andere databases kan MASSIVE gebruikt worden om methoden uit elke mogelijk gewenste stap van de processing pipeline te evalueren. Bovendien is de MASSIVE database zo ontworpen dat resultaten op deelverzamelingen van de data goed te vertalen zijn naar acquisities die haalbaar zijn in een klinische setting. De data wordt gebruikt in verschillende hoofdstukken in dit proefschrift en is vrij beschikbaar op <http://massive-data.org/>

MODEL- EN PARAMETERSCHATTING IN DIFFUSIE MRI

Er wordt nog altijd veel onderzoek gedaan naar het nauwkeurig schatten van parameters in dMRI. Een betrouwbare schatting wordt vaak bemoeilijkt door ruis en andere problemen zoals uitbijters. Uitbijters kunnen ervoor zorgen dat de geschatte waardes ernstig afwijken, waardoor analyses die hierop gebaseerd zijn onbetrouwbaar worden. De extra informatie die DKI geeft

Samenvatting

ten opzichte van DTI kan bijvoorbeeld onbruikbaar worden in een klinische setting, omdat de schatting van DKI gevoeliger is voor uitbijters vanwege de hogere acquisitie-eisen en model complexiteit. In **Hoofdstuk 4** wordt een methode gepresenteerd voor robuuste parameter schatting in DKI, REKINDLE genaamd (Robust Extraction of Kurtosis INDices with Linear Estimation). REKINDLE heeft als doel uitbijters te identificeren en excluderen, en gebruikt daarvoor recursieve snelle lineaire schatting en een intuïtief detectie criterium voor uitbijters. REKINDLE is direct toe te passen op DTI en andere dMRI modellen die gelineariseerd kunnen worden.

Hoofdstuk 5 focust op de schatting van de fODF met SD. SD technieken zijn populair omdat ze een betrouwbare fODF kunnen schatten zonder al te hoge eisen te stellen aan de acquisitie. De juistheid van de resulterende fODF is afhankelijk van de keuze van een geschikte respons functie (RF), die het diffusiegewogen signaal profiel van een éénvezelpopulatie (single fiber population, SFP) representeert. Een onnauwkeurige calibratie van de RF kan sterke invloed hebben op de geschatte fODF en tractografie. Huidige methoden modelleren de RF met een diffusietensor of baseren de RF op voxels met een hoge FA na een DTI parameter schatting. Echter, het DTI model is onnauwkeurig bij typische acquisities gebruikt voor SD omdat ook niet-Gaussische effecten optreden. In Hoofdstuk 5 wordt een methode voorgesteld om de RF te calibreren uit de data, door recursief voxels met kruisende vezels te excluderen. Het algoritme localiseert zo voxels met éénvezelpopulaties, waarbij een alternatief intuïtiever selectiecriterium wordt voorgesteld dan de FA.

In **Hoofdstuk 6** wordt de recursieve calibratie uit hoofdstuk 5 gebruikt voor het localiseren van SFPs in dMRI data van verschillende proefpersonen. Er is in de laatste jaren veel aandacht besteed aan het oplossen van kruisende vezels, terwijl er nog veel onbekend is over de fundamentele eigenschappen van SFPs. In hoofdstuk 6 worden de signaal karakteristieken van SFPs verder bestudeerd door het schatten van verschillende dMRI modellen (waaronder DKI met REKINDLE en SD), en het uitrekenen van maten gerelateerd aan deze modellen. Ook wordt gekeken naar de variabiliteit van SFP eigenschappen tussen verschillende datasets, proefpersonen, en vezelbundels. Hierbij wordt gevonden dat de variabiliteit in signaal karakteristieken tussen verschillende bundels vaak groter is dan de variabiliteit tussen verschillende proefpersonen per bundel. De kennis over karakteristieken in simpelere configuraties zoals SFPs kan belangrijk zijn voor het beter begrijpen van de weefsel microstructuur en het ontwerpen van realistische simulatie experimenten.

ONDERZOEK NAAR DE GEOMETRISCHE ORGANISATIE VAN DE HERSENEN MET DIFFUSIE MRI EN TRACTOGRAFIE

Door de opkomst van dMRI technieken die kruisende vezels nauwkeuriger kunnen modelleren, kan de geometrische organisatie van banen in de witte materie beter bestudeerd worden. In een recente publicatie analyseren Wedeen et al. (2012) visueel de kruising, nabijheid, en ligging van vezelbanen gereconstrueerd met dMRI en tractografie. De auteurs ontdekten dat vezelbanen elkaar orthogonaal kruisen op oppervlakken. Ze vonden geen hersenbanen zonder een dergelijke ‘sheet structuur’, en suggereerden dat deze structuren een relatie kunnen hebben met embryogenese en de groei van axonen. In een reactie hierop stellen Catani et al. (2012) dat de bevinding zeer waarschijnlijk een artefact is, gecreëerd door de methodologie die gebruikt is

voor het reconstrueren van de vezelpaden. Deze methode kan minder goed vezels reconstrueren die kruisen met kleinere hoeken dan bijvoorbeeld de SD technieken. Volgens hen draagt deze lage angulaire resolutie bij aan het reconstrueren van vezels die kruisen met alleen maar orthogonale hoeken. Dit maakt op zijn beurt een geweven sheet-structuur een voor de hand liggende maar foutieve configuratie. Wedeen et al. (2012a) weerleggen deze kritieken en stellen dat er geen mechanismen bekend zijn die artificieel een dergelijke mathematisch zeer uitzonderlijke configuratie zouden creëren. Catani et al. (2012) geven ook aan dat Wedeen et al. (2012b) vooral qualitatieve resultaten laten zien, en dat paden gereconstrueerd met tractografie niet zonder meer gelijkgesteld kunnen worden aan echte axonen.

In **Hoofdstuk 7** van dit proefschrift wordt de confusie over de conditie voor het bestaan van sheet structuren weggenomen door de mathematische theorie die hieraan ten grondslag ligt te resumeren. Er wordt een methode voorgesteld om de sheet waarschijnlijkheidsindex (sheet probability index, SPI) te berekenen, die kwantificeert hoezeer de data het bestaan van sheet structuren onderbouwt. Deze methode is gebaseerd op het reconstrueren van vele paden met tractografie, en wordt veelvuldig geëvalueerd in simulaties en echte data, waarbij gebruikt wordt gemaakt van het werk gepresenteerd in eerdere hoofdstukken. **Hoofdstuk 8** beschrijft een alternatieve methode voor het berekenen van de SPI die niet afhangt van tractografie-resultaten en minder rekenintensief is.

Een andere consequentie van de opkomst van modellen die kruisende vezels kunnen modelleren is de toenemende complexiteit van het vezelnetwerk dat gereconstrueerd kan worden met tractografie. Door de grote overlap van paden is het een uitdaging om een duidelijke 3D visualisatie te maken voor het bestuderen van de organisatie van het netwerk. In **Hoofdstuk 9** wordt een nieuwe visualisatiemethode voorgesteld voor tractografie resultaten, waarbij de transparantie van paden bepaald wordt door diens lokale of globale oriëntatie. Dit leidt tot een verbeterde 3D weergave, waarbij structuren verkend kunnen worden die anders bedekt zouden zijn door meer oppervlakkige paden. Deze methode wordt gebruikt voor het efficiënt localiseren van vezelbundels. Verder wordt de visualisatie geëvalueerd in een neurochirurgische context.

TOEKOMSTPERSPECTIEF

Diffusie MRI is een snel ontwikkelend onderzoeksgebied met veel potentie in klinisch en biomedisch onderzoek en applicaties. De dMRI beeldverwerkingspipeline is echter lang en complex, en veel verschillende strategieën zijn beschikbaar voor elke stap in het proces. Het ontwikkelen van nieuwe methoden voor dMRI analyse is van groot belang om de betrouwbaarheid van dMRI te verbeteren. Tegelijkertijd is het essentieel om te waarborgen dat dMRI niet alleen te gebruiken en begrijpen is voor onderzoekers die de methoden ontwerpen. Openbaar beschikbare evaluatie databases, goed onderbouwde en geëvalueerde algoritmes, en zo intuïtief mogelijke richtlijnen zijn van belang om de translatie van dMRI methoden naar bruikbare instrumenten voor applicaties te faciliteren.

ACKNOWLEDGEMENTS

Onderzoeken doen en daar een proefschrift over schrijven klinkt misschien als een ietwat eenzame bezigheid. Niets is minder waar: ongelofelijk veel mensen hebben direct en indirect hun steentje bijgedragen aan dit boekje. Hoewel het bijna onmogelijk is om niemand te vergeten (sorry alvast!), wil ik toch een poging doen om een aantal belangrijke mensen te bedanken.

Alexander, ik had me simpelweg geen betere begeleider kunnen wensen tijdens de afgelopen jaren. Je weet precies wat er speelt in het veld, kent iedereen die ook maar iets met diffusie MRI te maken heeft en hebt ideeën genoeg om een heel leger van PhD studenten aan het werk te houden. Ik heb zelden iemand gezien die met zo veel enthousiasme de wetenschap beoefent, dat werkt ontzettend motiverend! Je weet altijd feilloos de wetenschappelijke problemen waarmee ik zit om te toveren in iets positiefs. Bedankt voor het vertrouwen dat je altijd hebt gehad in mij en de ideeën die ik inbracht, en de kansen en die je me hebt gegeven. Bedankt dat je me altijd in de goede richting hebt (bij)gestuurd wanneer dat nodig was, en voor je praktische tips en goed advies over mijn werk, de wetenschap en alles daaromheen, en het leven in het algemeen. Bedankt dat er ondanks je drukke schema altijd tijd was voor een inhoudelijke besprekking, een goed gesprek, en de nodige humor. Ik heb enorm veel van je geleerd, en met heel veel plezier mijn PhD bij je gedaan. Bedankt voor alles!

Max, ik ben erg dankbaar dat ik mijn PhD heb kunnen doen bij het ISI onder jouw begeleiding. Je enorme kennis over beeldverwerking en de onderzoeksWereld is al van onschabare waarde geweest voor velen die mij voorgingen, en zeker ook voor mij. Je inbreng op het werk in mijn proefschrift was altijd zeer waardevol en gedetailleerd. Ik ben van mening dat veel van je briljante zinsconstructies en fantastisch mooie verwoordingen eigenlijk wel hun eigen boekje verdienen. Ik bewonder hoe je van het ISI een fijne ‘thuisbasis’ maakt, waar er altijd ruimte en tijd is voor een inhoudelijk of persoonlijk gesprek en natuurlijk de nodige gezelligheid tijdens Sinterklaasavonden, het maken van stukjes, en groepsactiviteiten. Bedankt voor de fijne begeleiding, de betrokkenheid, en alle goede adviezen, ik heb daar al erg veel aan gehad en weet zeker dat ik er nog veel aan zal hebben!

Maxime D, I have greatly enjoyed the time working at the SCIL! I admire the way you are heading your lab, always taking time for your students, coming up with great and diverse projects that fit with the strengths and interests of each student, and making sure they get the most out of their PhD. This was no different when I was visiting your lab, I appreciate how you really took the time to support me in exactly the same way. I value your great ideas on every aspect of the diffusion MRI pipeline, and how these lead to useful tools that can be used by the community. Thanks for giving me this great opportunity, for engaging me in all the activities with the lab, for the amazing wakeboard-experience, for the nice discussions about work and life as a researcher, and for making me feel welcome in Sherbrooke! *Marci bin!*

C-F, I am extremely grateful for the opportunity you have given me to join your lab for six months. It is amazing how you could instantly come up with brilliant ideas that helped me to solve my scientific questions every time. In fact, you found the solution to our problem already in 1993! Thank you for all the opportunities you have given me during my stay, for introducing

me to the wonderful group of scientists you have gathered around you, for hosting us, and for the nice chats around a beer. I truly had a fantastic time in Boston!

Luc en Andrea, bedankt voor de fijne samenwerking, ik kwam met veel plezier weer ‘terug’ op de TU/e! Ik vond het geweldig om met jullie de wiskundige kant van ons project te bespreken, daar heb ik enorm veel van opgestoken!

I would like to thank all my co-authors for their great contributions to the papers and abstracts. Your feedback and ideas were of great value!

Tom en Floor, ik ben super blij dat jullie mij bij willen staan in mijn verdediging als paranimfen. Tom, al sinds onze studie op de TU/e ben je een partner-in-crime voor het aanpakken van wiskundige problemen waar we niets van snappen (maar waar jij toch altijd nog meer van snapt dan ik), carnaval en andere feestjes, het bespreken van vage levensvragen, en natuurlijk voor het in de wacht slepen van die toekomstige Nobelprijs. Ik vind het super cool dat we dat hebben kunnen voortzetten tijdens onze promotie. Als *keigoed* team Limbo’s hebben we een mooi project afgerond! Dankjewel voor de altijd-even-kritische vragen, flauwe grappen, medewerking als ik je weer eens tot laat op de avond lastigviel met sheet-vragen, koffiepauzes, en diepgaande wetenschappelijke discussies. Dit alles heeft mijn promotie een stuk waardevoller en vooral leuker gemaakt! Floor, als sinds dag 1 zijn we kamergenoten, eerst in het WKZ, toen op de OIOsteeg, en daarna zijn we definitief *geupgrade* naar de stafgang. Ook al was ik er niet altijd door het verre reizen (eerst naar de andere kant van de rivier en daarna naar de andere kant van de oceaan) was onze kamer altijd even gezellig. Die gezelligheid hebben we regelmatig voortgezet tijdens de borrels, congressen, etentjes, en natuurlijk de onvergetelijke *90ies-Now* parties! Het was fijn om altijd de eureka-momenten, frustraties, praktische zaken rondom promoveren, en een hoop zaken buiten het promoveren met iemand te kunnen delen. Hoewel onze onderwerpen ver uit elkaar liggen heb ik elke keer ontzettend veel gehad aan je input en eerlijke mening. Ik bewonder hoe je altijd voor alles en iedereen klaar staat, bedankt voor alles!

Thanks to my colleagues of the PROVIDI Lab! Sjoerd, Martijn, Anneriet, Wieke, Sam, Fenghua, Szabolcs, Hamed, Richard, and others outside the UMCU, I really enjoyed our discussions during (group) meetings and drinks during conferences and workshops. Sjoerd, ik heb super veel geleerd van onze discussies en de projecten waar we samen aan gewerkt hebben. Gelukkig kunnen we de discussies en gezelligheid voortzetten op jaarlijkse conferenties, en wie weet nu wel vaker in de UK! Martijn, bedankt voor je massive bijdrage aan dit proefschrift ;). All the other old and new ISI-PhD-student-colleagues (and here is where I will likely forget someone), Bruno, Floor, Miekee, Yolanda, Lisette, Pim, Hugo, Sjoerd, Nynke, Hendrik, Gerrit, Alexander, Sjoerd, Job, Yinghe, Jurica, Majd, Frank, Jelmer, Bob, Bas, Mike, Harriet, Thessa, Paul, Maaike, Sascha, Remco, Ieva, Hui Shan, Marielle, Nadieh, Britt, Nikolas, Jaap, Sandra, Erik, Majd, and others from the 7T, thanks for making work so much more enjoyable during the times that I was not abroad ;). Floor, Lisette, and Bruno, thanks for being great roommates. Hugo, bedankt dat ik altijd je kantoor kon binnenvallen voor welke PhD-gerelateerde of niet-PhD-gerelateerde vraag dan ook, en de gezellige diner-avonden! Miekee, bedankt voor de gezellige ISMRMs en de goede gesprekken. Mede-musISI Robin, Miekee, Pim, en Sjoerd, het muziek maken en pizza eten met jullie was altijd een aangename afleiding van het werk! A big thanks to all the other colleagues of the ISI. I enjoyed all the cake-moments, after-work drinks, stukjes, Ardennen weekends, courses, and other meetings! Jacqueline, Marjan, en Reneé, heel

Less Confusion in Diffusion MRI

erg bedankt dat ik altijd bij jullie terecht kon met mijn vragen rondom het promoveren of voor een goed gesprek. Gerard, bedankt voor alle steun op ICT gebied!

My research visits abroad wouldn't have been so amazing without great colleagues. Dear colleagues from the SCIL/SNAIL/Sherbrooke University, Elef, Sam, Mic, Maxime, Emmanuelle, JC, Sam, Gab, Mike, David, Alex, Kevin, David, Olivier, Etienne, Marc, Janice, Francois, LP, Jasmeen, Alessandro, and everyone from Hugo's lab, a big thanks for making Sherbrooke my second home during my stay. I think I laughed the most per time unit during my time in Sherbrooke ;)! Thanks for the unforgettable visits to cool places in Quebec, the 5-to-8s, the campings, the beers at Siboire and beer festivals, supporting Netherlands during the soccer games, the barbeques, pool parties, and everything else! Dear colleagues from LMI/Harvard Medical School, Lauren, Yogesh, Ofer, Peter, Sylvain, Raúl, Isaiah, Angela, Pegah, Fan, Hengameh, Ron, Lipeng, Orjan, and everyone from the Golby Lab, SPL, and PNL, thanks for making my internship in Boston such a great experience! I greatly enjoyed the discussions, meetings, and drinks after work. Thanks for all your support and for the opportunities you have given me, I learned incredibly much during my stay. Thanks to everyone from Polina's group at CSAIR/MIT for the wonderful discussions about Lie groups and machine learning, and the nice dinners!

Thanks to the great people in the diffusion MRI field, doing research and going to conferences becomes so much more fun. I'm grateful for all the nice discussions! Bedankt aan de Nederlands/Belgische diffusie collega's: Sjoerd, Ben, Jelle, Daan, Thijis, Timo, Quinten, Thibo, Timo, Gwendolyn, bedankt voor de gezelligheid en inhoudelijke discussies tijdens ISMRMs en ISMRM Beneluxen. Derek, thank you for the nice discussions we had about future projects and my future as a researcher. You always took the time to listen to my ideas and I am very much looking forward to start as a postdoc in CUBRIC! Remco, Anna, Pauly en Bart, bedankt voor de goede basis in diffusie MRI en onderzoek doen die jullie me hebben meegegeven tijdens mijn master.

Geachte leden van de leescommissie, prof. dr. Kahn, prof. dr. Hendrikse, prof. dr. Dumoulin, prof. dr. Niessen, en prof. dr. Geurts, hartelijk dank voor het lezen en beoordelen van mijn proefschrift.

Beste mevrouw van Damme en Universiteitsfonds Eindhoven, hartelijk dank voor het toekennen van de Marina van Damme beurs! Ik ben erg dankbaar dat ik de kans heb gekregen een deel van mijn PhD aan Harvard Medical School te doen, waarvan ik het resultaat heb kunnen opnemen in dit proefschrift. Bedankt aan de leden van het Marina van Damme netwerk voor de inspirerende discussies en gezelligheid tijdens de bijeenkomsten. Paulo and Vesna, during my Master I learned a lot about diffusion MRI by reading your theses, thanks for now supporting the printing of my thesis with Mint Labs! Hartelijk dank aan Philips voor de bijdrage aan het drukken van dit proefschrift.

Naast onderzoek doen heb ik de kans gehad om me bezig te houden met een aantal leuke en uitdagende nevenactiviteiten. Dank aan de KNAW, De Jonge Akademie, Faces of Science en NEMO Kennislink dat ik de kans heb gekregen om over mijn werk te bloggen. Dit heb ik altijd met erg veel plezier gedaan! Het geeft je als onderzoeker de kans om je onderzoek te communiceren naar een breed publiek en om eens vanuit een ander perspectief naar je werk te

kijken. Bedankt aan de mede-promovendi waarmee ik samen de MISP 2013 en de ISMRM Benelux 2015 heb mogen organiseren!

Ik mag mezelf gelukkig prijzen met een aantal geweldige mensen om me heen die me veel hebben gesteund, tijdens mijn promotie maar eigenlijk vooral daarbuiten. Zonder de aangename afleiding tijdens de vele uitjes waren de afgelopen 4 jaar absoluut een stuk minder leuk geweest. Bedankt voor alle gezelligheid, het samen lachen en huilen, de flauwekul maar ook serieuze gesprekken, jullie interesse en steun! Femke, al sinds we kunnen lopen (of schommelen ;)) zijn we vriendinnen en ik ben ontzettend blij dat onze paden elkaar nog altijd blijven kruisen, ook al zat er soms een oceaan tussen. Ik weet zeker dat we dat relatief klein stukje water tussen Nederland en Wales ook wel aankunnen! Lore, je zou bijna denken dat onze Normaalkromme-sjablonen tijdens de wiskundelessen op de middelbare school uiteindelijk hun vruchten hebben afgeworpen ;). Na een gezellige tijd op de TU/e en tijdens de snowboardvakanties weten we ook nu zonder moeite die gezelligheid voort te zetten! Dunja, ik weet nog goed hoe we samen als brugpieper begonnen zijn, inclusief grote tassen vol met boeken en super-geavanceerde ezelsbruggetjes. Gelukkig is er ook nu nog steeds regelmatig tijd voor de nodige humor, een lekker etentje of dagje shoppen ;). Jolijn, Micha en Simone, de gezelligheid tijdens de danstrainingen werd al snel uitgebreid tot gezelligheid tijdens feestjes, high teas, shop-dagjes, en vakanties. Het maakt eigenlijk niet uit wie waar woont, we vinden altijd wel een tijd en plek om bij te kletsen! Zelfs al moest ik ver weg op conferentie, dan gingen jullie gewoon mee om de week daarvoor of daarna lekker vakantie te vieren ;) Simone, als introzusjes bij BMT hadden we elkaar al snel gevonden. Studeren was een stuk leuker met jou als bovenbuurvrouw en mede-intromama! De diners en weekendjes samen met Rik, Jurre, Mies, Tom en May zijn altijd supergezellig en de perfecte afleiding tussen alle drukte door, thanks guys! Rachel, tussen ons vele reizen door vinden we gelukkig toch soms de tijd om gezellig een drankje te doen, ik vind het super om alle verhalen en ervaringen met je uit te kunnen wisselen! Karlijn, echt leuk dat je nu ook de wetenschappelijke kant uit gaat, dat geeft ons vast nog meer stof om over bij te praten! Bas en Coen, hoewel we allemaal een compleet andere richting uit zijn gegaan ben ik blij dat we elkaar nog regelmatig tegenkomen. De etentjes met Tanja en Tinka zijn altijd even gezellig!

Opa en oma Meulen, bedankt voor al jullie liefde en steun. Jullie hebben het hele promotietraject met veel belangstelling gevolgd, en bij jullie kon ik altijd terecht om even tot rust te komen met *un tas koffie* en *un kukske*. Mijn verre reizen weerhielden ons er niet van om die gezelligheid voort te zetten via Skype, al moest ik de Limburgse vlaai dan helaas wel missen! Opa en oma Tax, wat had ik graag gewild dat we samen nog veel uitstapjes, feestdagen, kaart-rondjes, Tax-dagen, enzovoort hadden kunnen meemaken. Ik weet zeker dat jullie meekeken! Tante Rina, bedankt voor de interesse die u altijd getoond heeft en de gezellige bezoekjes waarbij een lekkere Bossche bol niet mocht ontbreken. Ook de rest van de familie bedankt voor alle steun en interesse. Johane, Daniel, Jinny et Sophie, je vous remercie de m'avoir reçu avec tant d'amour et accueilli dans votre famille!

Marc en Richard, kleine *brotties*, ik denk dat jullie meer bijgedragen hebben aan dit proefschrift dan jullie misschien erg in hebben. Marc, jij bent mijn grote motivatie om de techniek voor hersenonderzoek een handje vooruit te willen helpen, als kind wilde ik er al achter komen wat er in je gedachten speelde. Gelukkig zegt die grote glimlach van jou vaak meer dan 1000 woorden. Met je vrolijke muziekjes is thuis komen altijd net een stukje gezelliger, en na een

wandeling met jou heb ik weer energie voor 10! Richard, als kleine uk zei je al tegen mij “als je het wilt dan kun je het”, en geloof me, die wijsheid heb ik vaak in mijn hoofd moeten herhalen tijdens een moeilijke klus. Super leuk dat we samen kunnen *nerden* over techniek, gitaren, superhelden-films, flauwe grappen, en meer. Ik hoop dat ik net zo’n grote steun mag zijn voor jou in de dingen die je nog gaat meemaken als jij voor mij was!

Maxime C, it is hard to think of a passion that we do *not* have in common. Good food, traveling, camping, awkward humour, movies and series, snowboarding, good food... The enthusiasm we share for science is just one out of many things. The things we do not always have in common are exactly the things that I admire the most about you. Thanks for always being there for me, for making everything so much fun, and for your amazing support in difficult or busy times. Thank you for all your direct and indirect contributions to this thesis! Our life is quite an adventure with all the traveling, but I’m extremely grateful that we can do this together, always with a big smile! *Je t'aime Maxime, t'es toute ma vie :).*

Pap en mam, waar was ik nou helemaal geweest zonder jullie. Jullie hebben mij altijd onvoorwaardelijk gesteund in al mijn beslissingen, waar die ook heen leidden. Verhuizen van hot naar her, studeren en nog verder studeren en promoveren, dagvullende dans-uitjes in gymzalen, typisch volgeplande dagen met 1 minuut spelning tussen verschillende afspraken, verre reizen met alle bijbehorende praktische zaken, het was jullie niet snel te gek. Tussen alle drukte door was het enorm fijn dat er altijd een thuis was waar ik tot rust kon komen en 100% mezelf kon zijn. Dat we straks een stukje verder weg gaan wonen zal daar niets aan veranderen! We hebben al veel lol beleefd aan het plannen van de promotiedag. Als ik alleen al zie hoe jullie ernaar toe leven, dan is al het werk van de afgelopen 4 jaar meer dan de moeite waard geweest. Pap en mam, bedankt voor alles, deze is voor jullie!

PUBLICATIONS

JOURNAL PUBLICATIONS

1. **C.M.W. Tax**, T.C.J. Dela Haije, A. Fuster, C.F. Westin, M.A. Viergever, L.M.J. Florack, A. Leemans, "Sheet Probability Index (SPI): Characterizing the geometrical organization of the white matter with diffusion MRI", *NeuroImage* (in press)
2. M. Froeling, **C.M.W. Tax**, S.B. Vos, P.R. Luijten, A. Leemans, The "MASSIVE" brain dataset: Multiple Acquisitions for Standardization of Structural Imaging Validation and Evaluation, *Magnetic Resonance in Medicine*, (in press).
3. S.B. Vos, **C.M.W. Tax**, P.R. Luijten, S. Ourselin, A. Leemans, M. Froeling, The importance of correcting for signal drift in diffusion MRI, *Magnetic Resonance in Medicine*, (in press)
4. **C.M.W. Tax**, M. Chamberland, M. van Stralen, M.A. Viergever, K. Whittingstall, D. Fortin, M. Descoteaux, A. Leemans, "Seeing more by showing less: Orientation-dependent transparency rendering for fiber tractography visualization", *PLoS One*, 2015, nr. 10, vol. 10, p. e0139434
5. **C.M.W. Tax**, W.M. Otte, M.A. Viergever, R.M. Dijkhuizen, A. Leemans, REKINDLE: Robust Extraction of Kurtosis INDices with Linear Estimation, *Magnetic Resonance in Medicine*, 2015, nr. 2, vol. 73, pp. 794-808
6. O.F.F. Odish, A. Leemans, R.H.A.M. Reijntjes, S.J.A. van den Bogaard, E.M. Dumas, R. Wolterbeek, **C.M.W. Tax**, H.J. Kuijf, K.L. Vincken, J. van der Grond, R.A.C. Roos, Microstructural brain abnormalities in Huntington's disease: a two year follow-up, *Human Brain Mapping*, 2015, nr. 6, vol. 36, pp. 2061-2074
7. M. Bach, F.B. Laun, A. Leemans, **C.M.W. Tax**, G.J. Biessels, B. Stieltjes, K.H. Maier-Hein, Methodological considerations on Tract-Based Spatial Statistics (TBSS), *NeuroImage*, 2014, vol. 100, pp. 358-369
8. **C.M.W. Tax**, A. Leemans, 2014. Hoe magnetisme en diffusie kunnen bijdragen aan hersenonderzoek: Toepassingen van diffusie MRI tractografie, *Neuropraxis*, 2014, nr. 3, vol. 18, pp. 92-105
9. **C.M.W. Tax**, R. Duits, A. Vilanova, B.M. ter Haar Romeny, P. Hofman, L. Wagner, A. Leemans, and P. Ossenblok, 2014. Evaluating Contextual Processing in Diffusion MRI: Application to Optic Radiation Reconstruction for Epilepsy Surgery, *PLOS ONE*, nr. 7, vol. 9, p. e101524
10. **C.M.W. Tax**, B. Jeurissen, S.B. Vos, M.A. Viergever, A. Leemans, 2014. Recursive calibration of the fiber response function for spherical deconvolution of diffusion MRI data, *NeuroImage* vol. 86, pp. 67-80
11. **C.M.W. Tax**, A.P. Bom, R. Fitzpatrick, R.L. Taylor, N. Todd, M. Welgampola, 2013. The galvanic whole-body sway response in health and disease, *Clinical Neurophysiology* nr. 10, vol. 124, pp. 2036-2045

CONFERENCE PUBLICATIONS

1. H.J. Kuijf, **C.M.W. Tax**, L.K. Zaanen, W.H. Bouvy, J. de Bresser, A. Leemans, M.A. Viergever, G.J. Biessels, K.L. Vincken, The added value of diffusion tensor imaging for automated white matter hyperintensity segmentation, in: Computational Diffusion MRI, Editor(s): L. O'Donnell, G. Nedjati-Gilani, Y. Rathi, M. Reisert, T. Schneider, Springer International Publishing, 2014, Mathematics and Visualization, pp. 45-53.
2. **C.M.W. Tax**, R. Duits, B.M. ter Haar Romeny, A. Vilanova and P. Ossenblok, 2012, Tractography of the Optic Radiation for Vision Sparing Epilepsy Surgery In: Proceedings IEEE (ICIA) International Conference on Information and Automation in Shenyang, China, pp.441-445
3. **C.M.W. Tax**, R. Duits, B.M. ter Haar Romeny, A. Vilanova and P. Ossenblok, 2012 Tractography of the Optic Radiation for Vision Sparing Epilepsy Surgery In: ASCI proceedings ICT.OPEN conference in Rotterdam, The Netherlands

CONFERENCE ABSTRACTS (FIRST AUTHOR)

1. **C.M.W. Tax**, C.F. Westin, T.C.J. Dela Haije, A. Fuster, M.A. Viergever, L.M.J. Florack, A. Leemans, Robust assessment of the brain's sheet structure using normalized convolution, International Society Magnetic Resonance in Medicine, 2016, 2073
2. **C.M.W. Tax**, T.C.J. Dela Haije, A. Fuster, C.F. Westin, M.A. Viergever, L.M.J. Florack, A. Leemans, Mapping the brain's "Sheet Probability Index" (SPI) with diffusion MRI: Sheet happens?!, International Society Magnetic Resonance in Medicine, 2016, 0791
3. **C.M.W. Tax**, D.S. Novikov, E. Garyfallidis, M.A. Viergever, M. Descoteaux, A. Leemans, Localizing and characterizing single fiber populations throughout the brain, International Society Magnetic Resonance in Medicine, 2015, 0473.
4. **C.M.W. Tax**, T.C.J. Dela Haije, A. Fuster, R. Duits, M.A. Viergever, E. Calabrese, G.A. Johnson, L.M.J. Florack, A. Leemans, Towards Quantification of the Brain's Sheet Structure in Diffusion MRI Data, International BASP Frontiers Workshop, 2015, 74.
5. **C.M.W. Tax**, M. van Stralen, M.A. Viergever, N. Ramsey, A. Leemans, Orientation-dependent rendering of diffusion fiber tractography streamlines for improved visualization of complex tissue organization, International Society Magnetic Resonance in Medicine, 2014, 4529.
6. **C.M.W. Tax**, S.B. Vos, M.A. Viergever, M. Froeling, A. Leemans, Transforming grids to shells and vice versa: an evaluation of interpolation methods in diffusion MRI q- and b-space, International Society Magnetic Resonance in Medicine, 2014, 4485.
7. **C.M.W. Tax**, T.C.J. Dela Haije, A. Fuster, R. Duits, M.A. Viergever, L.M.J. Florack, A. Leemans, Towards quantification of the brain's sheet structure: Evaluation of the discrete Lie bracket, International Society Magnetic Resonance in Medicine, 2014, 0975.
8. **C.M.W. Tax**, B. Jeurissen, S.B. Vos, M.A. Viergever, A. Leemans, 2013, Recursive calibration of the fiber response function for spherical deconvolution diffusion ODF sharpening, ISMRM Workshop on Diffusion as a Probe of Neural Tissue Microstructure in Podstrana, Croatia

9. **C.M.W. Tax**, T.C.J. Dela Haije, A. Fuster, M.A. Viergever, L.M.J. Florack, A. Leemans, 2013, Considerations on the theory of sheet structure of cerebral pathways, ISMRM Workshop on Diffusion as a Probe of Neural Tissue Microstructure in Podstrana, Croatia
10. **C.M.W. Tax**, W.M. Otte, M.A. Viergever, R.M. Dijkhuizen, and A. Leemans, 2013, REKINDLE: Robust Extraction of Kurtosis Indices with Linear Estimation International Society for Magnetic Resonance in Medicine, 21th Scientific Meeting in Salt Lake City, USA, 0700
11. **C.M.W. Tax**, B. Jeurissen, M.A. Viergever, A. Leemans, 2013, Robust fiber response function estimation for deconvolution based diffusion MRI methods International Society for Magnetic Resonance in Medicine, 21th Scientific Meeting in Salt Lake City, USA, 3149
12. **C.M.W. Tax**, B. Jeurissen, M.A. Viergever, A. Leemans, 2013, Robust fiber response function estimation for deconvolution based diffusion MRI methods International Society for Magnetic Resonance in Medicine Benelux, 5th Scientific Meeting in Rotterdam, The Netherlands, PRES8-05
13. **C.M.W. Tax**, S.B. Vos, J. Veraart, J. Sijbers, M.A. Viergever, and A. Leemans, 2012, The effect of the kurtosis on the accuracy of diffusion tensor based fiber tractography International Society for Magnetic Resonance in Medicine, 20th Scientific Meeting in Melbourne, Australia, 3629
14. **C.M.W. Tax**, A.P. Bom, R. Fitzpatrick, R.L. Taylor, N. Todd, M. Welgampola, 2012, The galvanic whole-body sway response in health and disease Bárány Society Meeting in Uppsala, Sweden, O. 146
15. **C.M.W. Tax**, R. Duits, A. Vilanova, B.M. ter Haar Romeny, A. Leemans, P. Ossenblok, 2012, Localization and visualization of the optic radiation for anterior temporal lobe resection NVPHBV Spring Meeting in The Hague, The Netherlands.
16. **C.M.W. Tax**, P. Ossenblok, L. Wagner, R. Duits, R. Berting, A. Vilanova, 2011, Is fMRI-seeded DTI tractography feasible in case of a cystic encephalomalacia? 13th edition of the annual international clinical symposium 'Update@kempenhaeghe.nl' in Heeze, The Netherlands

BOOK CHAPTERS

1. **C.M.W. Tax**, S.B. Vos, A. Leemans, Checking and Correcting DTI Data, Editor(s): W. Van Hecke, L. Emsell, S. Sunaert, Springer New York, 2016, Diffusion Tensor Imaging: A Practical Handbook, p. 127-150.
2. S.B. Vos, **C.M.W. Tax**, A. Leemans, Diffusion MRI and fiber tractography: The State-Of-The-Art and its Potential Impact on Patient Management, Editor(s): T.C. Kwee, H. Zaidi, Elsevier, 2013, vol. 8, Evolving Medical Imaging Techniques - PET Clinics, p. 279–293.
3. T.C.J. Dela Haije, R. Duits and **C.M.W. Tax**, Sharpening Fibers in Diffusion Weighted MRI via Erosion, Editors: C.F. Westin, A. Vilanova, B. Burgeth, Springer Berlin Heidelberg, 2014, Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data Mathematics and Visualization, p. 97–126

ABOUT THE AUTHOR

Chantal Tax was born on September 10 1988 in Weert, The Netherlands. She completed her Biomedical Engineering bachelor at Eindhoven University of Technology. During her master Medical Engineering at the same university she focused on biomedical modelling and image processing and analysis. She did a four-month research project at the Royal Prince Alfred Hospital in Sydney, Australia, where she investigated a new clinical test for balance disorders in the lab of Michael Halmagyi and Miriam Welgampola. Chantal completed her master thesis in the group of Bart ter Haar Romeny in collaboration with epilepsy center Kempenhaeghe, under the supervision of Remco Duits, Anna Vilanova, and Pauly Ossenblok. The focus of this project was to improve the reconstruction of the optic radiation for epilepsy surgery using functional MRI and diffusion MRI tractography. After graduation in 2012 she started as a PhD candidate at the PROVIDI Lab, part of the Image Sciences Institute, UMC Utrecht. Under the supervision of Alexander Leemans and Max Viergever she focused on improving the analysis of diffusion MRI data. During her PhD she spent three months at the Sherbrooke Connectivity Imaging Lab (SCIL), Université de Sherbrooke, Canada, working with Maxime Descoteaux. After receiving the Marina van Damme grant she spent six months at the Laboratory of Mathematics in Imaging (LMI), Harvard Medical School, USA, working with Carl-Fredrik Westin. She blogs for ‘Faces of Science’ (an initiative of the Royal Netherlands Academy of Arts and Sciences) about her work and life as a researcher. In 2016 she became a Junior Fellow of the ISMRM. The results of the work during her PhD are described in this thesis.

Chantal Tax werd geboren op 10 september 1988 in Weert. Na de bachelor Biomedische Technologie aan de Technische Universiteit Eindhoven heeft ze zich tijdens haar master Medical Engineering gespecialiseerd in biomedische modellering en beeldverwerking. Gedurende een externe stage in het lab van Michael Halmagyi en Miriam Welgampola in het Royal Prince Alfred Hospital in Sydney (Australië) heeft ze gewerkt aan een nieuwe klinische test voor evenwichtstoornissen. Chantal heeft haar master onderzoek gedaan in de groep van Bart ter Haar Romeny in samenwerking met epilepsiecentrum Kempenhaeghe. Hier heeft ze onder de supervisie van Remco Duits, Anna Vilanova, en Pauly Ossenblok onderzoek gedaan naar het verbeteren van de reconstructie van de optische radiatie voor epilepsie chirurgie met functionele en diffusie MRI. Na haar afstuderen in 2012 begon ze haar promotieonderzoek bij het PROVIDI Lab, Image Sciences Institute binnen het UMC Utrecht. Onder de supervisie van Alexander Leemans en Max Viergever heeft ze gewerkt aan het verbeteren van de beeldanalyse van diffusie MRI data. Tijdens haar promotieonderzoek heeft ze gedurende drie maanden gewerkt in het Sherbrooke Connectivity Imaging Lab (SCIL), Université de Sherbrooke (Canada) onder supervisie van Maxime Descoteaux. Na het ontvangen van de Marina van Damme beurs heeft ze gedurende zes maanden onderzoek gedaan in het Laboratory of Mathematics and Imaging (LMI), Harvard Medical School (USA) onder supervisie van Carl-Fredrik Westin. Ze blogt voor ‘Faces of Science’ (een initiatief van de Koninklijke Nederlandse Akademie van Wetenschappen) over haar werk en leven als onderzoeker. In 2016 werd ze een Junior Fellow van de ISMRM. Het werk van haar promotie is gebundeld in dit proefschrift.