Finsler geometry and diffusion MRI

Document Version:
Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

• A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
• The final author version and the galley proof are versions of the publication after peer review.
• The final published version features the final layout of the paper including the volume, issue and page numbers.

Link to publication

Download date: 16. Nov. 2023

Finsler Geometry and Diffusion MRI

Finsler Geometry and Diffusion MRI

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven, op gezag van de
rector magnificus, prof. dr. ir. F.P.T. Baaijens, voor een
commissie aangewezen door het College voor
Promoties in het openbaar te verdedigen
op dinsdag 16 mei 2017 om 16.00 uur

door

Thomas Carolus Johannes Dela Haije

geboren te Landgraaf

Dit proefschrift is goedgekeurd door de promotoren en de samenstelling van de promotiecommissie is als volgt:

Voorzitter:	Prof. Dr. Ir. B. Koren
Promotor:	Prof. Dr. L.M.J. Florack
Copromotor:	Dr. A. Fuster
Leden:	Prof. Dr. C.-F. Westin (Harvard University)
	Prof. Dr. R. Deriche (Université Côte d'Azur)
	Prof. Dr. Ir. B.M. ter Haar Romeny
	Dr. E. Özarslan (Linköpings Universitet)
	Dr. A. Feragen (Københavns Universitet)

Het onderzoek dat in dit proefschrift wordt beschreven is uitgevoerd in overeenstemming met de TU/e Gedragscode Wetenschapsbeoefening

Summary

Finsler geometry and diffusion MRI

Diffusion-weighted magnetic resonance imaging (MRI) enables one to record localized measurements of ensemble-averaged molecular displacements in vivo, which provide valuable information on e.g. structural anisotropy in brain tissue. Even the most basic diffusion-weighted acquisition produces a vast trove of data, offering a glimpse of the full (diffusion) time- and position-dependent ensemble average propagator (EAP) that is probed with this modality. This thesis covers aspects of the various roles geometry plays in the organization and understanding of this data, which is apparent first and foremost in modeling. The EAP is de facto insensitive to any position-dependence within the localization window (the voxel), and it thus makes sense to assume a model for the underlying structure that has the same spatial scale as the acquired data. If we then reconstruct the diffusion taking place within a model voxel, we need only consider a spatially homogeneous process. This idea leads logically to tangent spaces as a basis for the description of the data. The voxel location is associated to a point in a manifold, and the local EAP resides in the tangent space at that point.

From there on the introduction of a geometrical framework proceeds as follows. We constrain the geometry of the manifold, find a relation between the geometry and the measured diffusion, and perform further analyses on the geometric structure of the manifold instead of on the unwieldy probability density functions. This has advantages both at a local and at a global level; the important local characteristics of a geometric structure are typically well-understood, while global characteristics such as (long range) geodesics can capture relationships in the data across space. The first part of this thesis addresses technicalities involved in the development of such frameworks for Riemannian and Finslerian geometry.

In the first chapter of Part I we reproduce and refine known relations between the Riemannian framework and the widely used diffusion tensor imaging model. Specifically, we extend the existing framework to include a position-dependent scaling of the Riemannian metric, which has advantages in practical applications. In the second

chapter we derive a fundamental relationship between the geometrical structure of a Finsler manifold and large displacements in a Brownian motion-like process on this manifold. This in effect establishes a Finslerian framework as a generalization of the Riemannian one, and preliminary results illustrate the added power that comes at the cost of a concomitant complexity increase. We find that the Finslerian framework is a natural fit for diffusion-weighted data acquired at high gradient strengths, meaning that the scanner is more sensitive to large displacements, whereas the Riemannian framework is mainly applicable to low gradient strength data.

Part II covers our development on the post-modeling processing steps in the standard diffusion MRI pipeline. In the first chapter we look at a basic sharpening strategy intended to address Gaussian blurring of the propagator caused by acquisition constraints. This very general idea is shown to lead to very simple algorithms for many classes of diffusion MRI models, some of which have been proposed as ad hoc enhancement methods in other works and different contexts. Basic diffusion tensor imaging experiments illustrate the potential of the strategy both for visualization purposes and for more advanced analysis techniques such as tractography. In the second chapter of Part II we discuss the role of Finsler geometry in interpolation. We describe the existing methods applicable to the Riemannian framework, show how these methods can be applied in the Finslerian framework, and explain how they could be formally generalized. Finally, we also explain how ideas from Finsler geometry could help improve analysis of low gradient strength data that is more suitable for analysis in the Riemannian framework.

In Part III of the thesis we study geometry-inspired analysis techniques. The three chapters in this part concern local tissue descriptors, reconstruction of long range connections in the manifold, and an analysis of large-scale higher-dimensional structures in the data. The local measures and connection applications are predominantly based on the proposed Finslerian framework, and use results from e.g. convex geometry to relate the geometrical structure to the structure of the underlying tissue. In the final chapter of Part III we investigate the sheet structure hypothesis, which poses that long-range connections in the brain, so-called fiber pathways, can be grouped together in a natural way to form two-dimensional surfaces. The main contribution in this chapter is the proposition and validation of an algorithm that can detect this structure, based on the Lie bracket of vector fields derived from diffusion-weighted data. Using this technique we show that the sheet structure hypothesis does not hold everywhere in the brain, and that its validity is at the very least restricted to specific scales.

Considerations on the implementation of the Finslerian framework have led to a number of practical advancements, such as a reformulation of the essential constraints in (generalized) diffusion tensor imaging using semi-definite programming. Most of the code used for the results in this thesis is bundled in the open source Mathematica package Diffusica` that is currently in development.

Contents

List of Figures

List of Nomenclature

Abbreviations

ADE asymmetric diffusion encoding
AP anterior-posterior
ASD autism spectrum disorder
bODF barrier orientation distribution function
CC corpus callosum
CSD constrained spherical deconvolution
CSF cerebrospinal fluid
CST corticospinal tract
DDE double diffusion encoding
DKI diffusional kurtosis imaging
dODF diffusion orientation distribution function
DSI diffusion spectrum imaging
DTI diffusion tensor imaging
EAP ensemble average propagator
fODF fiber orientation density function
FTM Finsler tractography module
GDE generalized diffusion encoding
HCP Human Connectome Project
IS inferior-superior
LR left-right
MANOVA multivariate analysis of variance
MDE multi-diffusion encoding
MRI magnetic resonance imaging
ODF orientation distribution/density function
OGSE oscillating gradient spin echo
PDE partial differential equation
PGSE pulsed gradient spin echo

RGB red-green-blue
RTM Riemannian tractography module
SDE single diffusion encoding
SLF superior longitudinal fasciculus
SNR signal-to-noise ratio
SPI sheet probability index
TDC typically developing controls
WM white matter

Parameters and physical constants

α diffusion coefficient
b b-value
B seed region
$\boldsymbol{B}$ static magnetic field
δ pulsed gradient spin echo pulse width or voxel size
Δ diffusion time
Δh step size
ϵ sharpening parameter upper bound
f sharpening parameter
γ (proton) gyromagnetic ratio
G gradient strength
$\boldsymbol{G}$ pulsed gradient spin echo diffusion encoding gradient
h arc length
h_{max} sheet resolution parameter
k relative signal-to-noise ratio
κ gradient scaling parameter or curvature
$\bar{\kappa}$ dimensionless gradient scaling parameter
λ collision rate or SPI sensitivity parameter
n_{b} number of bootstraps
n_{n} number of noise iterations
ν connectivity scaling parameter
$\boldsymbol{q}$ gradient wave vector
ρ reciprocal curvature
s power sharpening parameter
t threshold
τ effective diffusion time
TE echo time
TI inversion time
TR repetition time

T_1 spin-lattice relaxation time
T_2 spin-spin relaxation time
T_2^* spin-spin and static field non-uniformity-induced relaxation time
v velocity

MRI-related objects

AD axial diffusivity
ADC apparent diffusion coefficient
b_{max} maximum b-value employed in reconstruction
c_l linearity coefficient
c_p planarity coefficient
c_s sphericity coefficient
C connectivity matrix
C_{avg} connectivity path measure based on average cost
C_{max} connectivity path measure based on maximum cost
D diffusion tensor
F_{new} revised connectivity Finsler function
F_{old} previous connectivity Finsler function
FA fractional anisotropy
Ψ_b barrier ODF
Ψ_d normalized diffusion ODF
Ψ_f fiber ODF
Ψ_l linear barrier ODF
Ψ_p planar barrier ODF
Ψ_{sa} solid angle diffusion ODF
λ_{max} first (largest) eigenvalue
λ_{med} second eigenvalue
λ_{min} third (smallest) eigenvalue
Λ_{avg} average estimated maximum displacement
Λ_{max} largest estimated displacement
Λ_{min} smallest estimated maximum displacement
Λ_{rel} difference between largest and smallest estimated maximum displacements
m average region-wise connectivity
μ mean
MD (generalized) mean diffusivity
NG non-Gaussianity
P probability function
P_λ sheet probability index
Q Gaussian deconvolved ensemble average propagator

R closure of quadrilaterals
R_κ κ level set
RA radial anisotropy
RD radial diffusivity
σ standard deviation
S normalized diffusion MRI signal
$\boldsymbol{S}$ sheet tensor
$\boldsymbol{T}$ tensorial representation of the direction-dependent apparent diffusion coefficient

Geometrical objects

α integral curve
C (geodesic) curve
d general distance function
δ delta function
$\boldsymbol{\delta}$ Kronecker delta
dσ Lebesgue measure on the sphere
$\partial\Omega$ boundary of Ω
Δ Laplace–Beltrami operator
η induced measure on indicatrix
E Euclidean norm
ϕ azimuthal angle
F Finsler function, or, Finslerian norm
Φ flow operator
$\boldsymbol{g}$ metric or fundamental tensor
$\tilde{\boldsymbol{g}}$ adjugate-based Riemannian metric tensor
γ^i_{jk} formal Christoffel symbols of the second kind
G Riemannian norm
Γ^i_{jk} Christoffel symbols of the second kind
$\boldsymbol{h}$ dual fundamental tensor
H dual Finsler function
$\boldsymbol{I}$ identity
IM indicatrix bundle of M
$\mathcal{L}$ length of curve
μ induced probability measure on indicatrix
M manifold, position space of spins
N subset of M
Ω support of displacement probability distribution
$\boldsymbol{p}$ fiducial point on M

r radial distance

$\boldsymbol{r}$ displacement vector

$\boldsymbol{R}$ Ricci curvature tensor

$\mathbb{R}^n$ n-dimensional real vector space

S sheet structure

S^n n-dimensional sphere

SM scaled indicatrix bundle

θ polar angle

ϑ angle

TM tangent bundle of M

$\overset{\circ}{T}M$ slit tangent bundle of M

V family of voxels of M

$\boldsymbol{\xi}$ induced coordinates on TM

Z geodesic spray

Transforms

C cosine transform

$\mathcal{F}$ Fourier transform

ρ_σ Gaussian deblurring transform

$\mathcal{R}$ Funk–Radon transform

$\mathcal{S}$ sine transform

Functions

f_{W} Watson distribution

Γ Gamma function

M Kummer function

Y_l^m spherical harmonic (seismological convention)

$\tilde{Y}_l^m$ real-valued spherical harmonic

Brackets

$\|\cdot\|$ (dual) norm

$\|\cdot\|_{\mathrm{op}}$ operator norm

$\langle .,. \rangle$ Kronecker tensor

$[\cdot,\cdot]$ Lie bracket

$[\ldots]$ ordered list

$\{\ldots\}$ unordered list

Modifiers

$\cdot_x$ localized at x

$\cdot_F$ relative to norm F

$\cdot_\kappa$ estimated using parameter κ

$\hat{\cdot}$ normalized quantity

$\widehat{\cdot}$ estimated quantity

$\bar{\cdot}$ complex conjugate

$\cdot^\perp$ normal component

$\cdot^*$ convex conjugate, dual, or optimal

Preface

Chapter 1

Introduction

Neurogeometry and Diffusion MRI

One of the major challenges faced by scientists today is understanding the human brain. Questions about consciousness, dreams, and memory have been slowly moving from the realm of philosophy to the multidisciplinary field of neuroscience, which investigates the anatomy and function of the brain at all scales. Fueled by major developments in a.o. imaging and genetics [240], interest in this field has increased tremendously. As a result, there are now large initiatives underway [77, 186, 351, 358] to acquire and combine results from a large number of subjects using different modalities [351], and "to integrate this data in a unified picture of the brain as a single multi-level system" [358].

In this thesis we look at the brain with one such modality; diffusion-weighted magnetic resonance imaging, or diffusion MRI—a non-invasive, in vivo imaging technique that is uniquely suitable for investigations of neuronal fiber pathway anatomy [158]. Diffusion MRI is sensitive to structure at the micrometer scale of neurons, averaged over a millimeter scale window called the voxel. Properties of incoherent structures—structures that are sufficiently random at a sub-voxel level—are mostly lost in this voxel-averaging, and mainly characteristics of large coherent structures such as fiber pathways remain. The recovery of tissue characteristics that are perceivable through diffusion MRI poses a non-trivial inverse problem, and the work presented in this manuscript focuses on the various roles that geometry plays in this and related problems.

This introductory chapter first covers the relevant background of diffusion MRI in Section 1.1, and we then briefly discuss the structure of the thesis in Section 1.2. The remaining sections are used to declare some of the notations, conventions, and resources used throughout the work.

1.1 The basics of diffusion magnetic resonance imaging

1.1.1 Nuclear magnetic resonance

In the presence of a strong, static, uniform magnetic field B [T], the intrinsic magnetic moments of spin-carrying particles (or 'spins') in a specimen align to produce a net magnetization vector parallel to the direction of the static field. These aligned spins can be perturbed in a controlled manner through the application of an electromagnetic pulse with a specific resonance frequency called the Larmor frequency, determined by the product of the nucleus' gyromagnetic ratio γ $\left[\frac{\text{rad}}{\text{s T}}\right]$ and the *local* strength of the magnetic field. A perturbation pulse results in an excess of excited spins—spins that have a component perpendicular (or 'transversal') to the static field B—which return to thermal equilibrium under the remittance of another set of electromagnetic waves. This phenomenon, called nuclear magnetic resonance, lies at the

basis of magnetic resonance imaging, which uses this interaction between electro-magnetic fields and spin-carrying particles to measure various characteristics of these particles and their surroundings [37, 157, 314]. Because the resonance frequency of spins depends on the strength of the local magnetic field, we can add magnetic field gradients that make the effective field strength encountered by spins dependent on position, allowing us to e.g. submit localized pulses of electromagnetic radiation.

There are a number of basic parameters that describe the dynamics of perturbed spins in a standard MRI experiment. Firstly, the relaxation time T_1 describes the recovery of the net magnetization component along the static field, caused by thermal interactions between spins and their surroundings. Secondly, excited spins precess around B at the Larmor frequency. Small dynamic fluctuations in the local magnetic field due to interactions between spins cause small variations in their respective precession frequencies, and the resulting decoherence—spins acquire accumulative phase differences relative to the expected precession—results in a decay of the rotating transversal component of the net magnetization. This decay is characterized by the T_2 relaxation time. Finally there are small but constant non-uniformities in the static field that speed up transversal decay, which are combined with the T_2 effects in the T_2^* characteristic time. Additional details can be found in the references [49, 157, 314].

1.1.2 The pulsed gradient spin echo sequence

Because the additional decoherence due to T_2^* effects is caused by *constant* non-uniformities, it is to a degree reversible: an inversion pulse can be used to reverse the acquired dephasing, causing a temporary refocusing in the transversal plane that produces a measurable echo [53, 159]. In diffusion MRI this idea of spin echoes is combined with gradients to induce a resonance signal that is encoded with spin motion. These specific motion-sensitive sequences of gradients and electromagnetic pulses are called diffusion encoding pulse sequences. The sequences of interest in this thesis are variations on the classical Stejskal–Tanner sequence [323], or pulsed gradient spin echo sequence (PGSE). A PGSE sequence starts with a 90 degree pulse that moves the net magnetization vector into the transversal plane perpendicular to B. As stated, the magnetization vector then precesses around the orientation of the static field, and a set of canceling gradients are applied to produce an additional displacement-dependent dephasing in the spins that make up the net magnetization vector. This causes an additional suppression in the produced echo, as illustrated in Fig. 1.1.

1.1.3 The diffusion-weighted signal

Consider now the compact, embedded, three-dimensional differentiable submanifold $M \subset \mathbb{R}^3$ that represents the position space of (ensembles of) nuclear spins in the
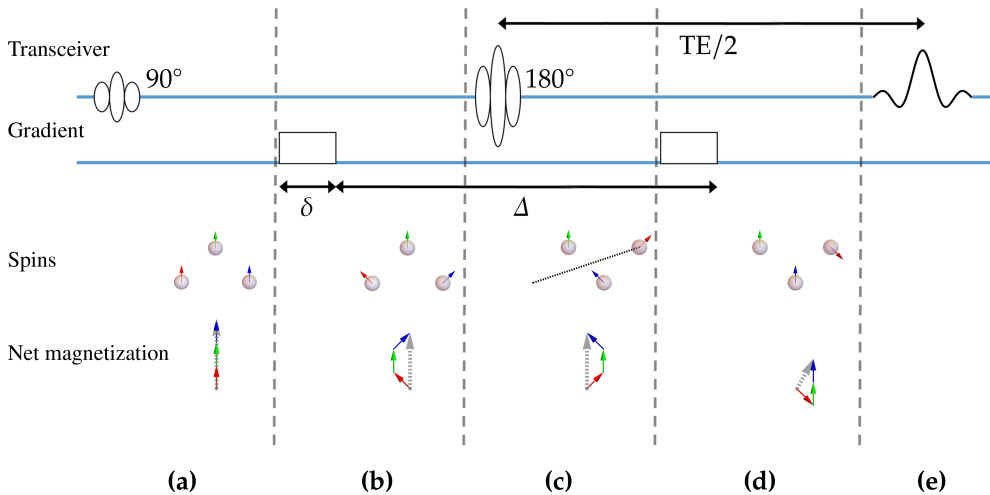
Figure 1.1: A schematic representation of the standard pulsed gradient spin echo sequence, with time on the horizontal axis. In this idealized scenario, the relevant acquisition parameters are the pulse width δ, the pulse separation width Δ, and the echo time TE. The first row shows the electromagnetic pulses that are applied, as well as the echo that is measured. The second row shows the diffusion encoding gradients. The last two rows show the effect of the diffusion encoding gradients on a representative selection of spins in the transversal plane at four different times, and the corresponding net magnetization vector. For the purpose of this illustration we do not include T_2 relaxation and other non-diffusion related effects, and we only show the phase of the spins relative to a reference that rotates with the Larmor frequency determined by B. (a) After the initial 90° pulse, excited spins are perfectly aligned in the transversal plane, where they precess at the Larmor frequency. They immediately start to incur phase shifts due to small static fluctuations in the magnetic field (T_2^* effects, not illustrated). (b) During the application of the first gradient pulse, the precession frequency of spins is temporarily shifted by an amount depending on their position, which results in a position-dependent phase shift. (c) A 180° degree pulse effectively reverses the acquired dephasing; spins that lagged behind before the refocusing pulse, now run ahead of the expected precession. (d) A second, identical gradient pulse causes stationary spins (green and blue) to return to their original state. For spins that have moved during or between the gradient pulses—in this illustration the red spin moved along the black dashed line—this refocusing is imperfect. (e) At this point in time the spins rephase—the static non-uniformities in the magnetic field have caused the spins to realign, resulting in a spin echo at the echo time TE. The echo is smaller in magnitude compared to that of the standard spin echo sequence (without diffusion encoding gradients), because the gradients introduced an additional displacement-dependent phase shift. The non-diffusion-weighted attenuation can be factored out, so that a normalized signal remains that only reflects spin displacement.

scanner. Elements of the tangent space $T_x M$ at $x \in M$ typically represent (ensemble) velocity or displacement vectors of spin-carrying particles at x, and the union of tangent spaces for all $x \in M$ forms the tangent bundle TM. An asterisk denotes the algebraic dual of a space, e.g. T^*M denotes the set of linear functionals (covectors) acting on elements of TM. The Kronecker tensor $\langle .,. \rangle : V^* \times V \to \mathbb{K}$ is the bilinear mapping acting on a covector $f \in V^*$ and a vector $v \in V$ such that $\langle f, v \rangle = f(v)$, for some vector space V (typically $T_x M$ for some $x \in M$) over a field $\mathbb{K}$ (typically $\mathbb{R}$).

A typical diffusion-weighted magnetic resonance experiment now consists of a series of (voxel-wise) measurements such as shown in Fig. 1.1 parameterized by a diffusion time $\Delta > 0$, a pulse width $\delta > 0$, and an applied gradient $G \in (\mathbb{R}^3)^*$ with magnitude G. We will write $P(r) := P_{\Delta,\delta}(p, r)$ for the probability that spins initially at p undergo a displacement $r \in T_p M$ during the diffusion time Δ, and we write the transversal component of the net magnetization vector in complex exponential form. Both Δ and δ are assumed constant, and for simplicity we will only consider $\delta \ll \Delta$.[1] With the definition of the wave vector[2]

$$q := \gamma \delta G_p \in T_p^* M, \tag{1.1}$$

where γ is the gyromagnetic ratio and G_p is the (induced[3]) gradient at p, it follows that a spin that undergoes a translation r between the two gradient pulses has its magnetic moment phase-shifted by $\exp(\mathrm{i}\langle q, r \rangle)$ [49, 50]. The normalized signal $S(q) := S_{\Delta,\delta}(p, q)$, which represents the total contribution from all spins to the measured echo (at TE) relative to the magnitude of the non-diffusion-weighted spin echo, is thus related to P through an inverse Fourier transform [48–50, 251]:

$$S(q) = \int_{T_p M} \mathrm{e}^{\mathrm{i}\langle q, r \rangle} P(r) \, \mathrm{d}r. \tag{1.2}$$

We will assume that the probability density function P is bounded, Lebesgue integrable, and that it has a well-defined moment generating function, which implies that the domain of S can be extended to complex-valued arguments [165]. For finite diffusion times Δ, P is compactly supported on the set $\Omega := \Omega_p(\Delta, \delta) \subset T_p M$. Detailed balance arguments can be invoked to argue that $P(r) = P(-r)$, so that the normalized

---

[1] The case when $\delta \ll \Delta$ is referred to as the narrow pulse regime, which does not apply in typical diffusion MRI experiments. The concepts introduced in this section still apply if the narrow pulse approximation does not hold, though with P instead viewed as a center-of-mass propagator [49, 200, 234, 283]. Model-dependent strategies to address this issue [222, 312, 379], like the correction $\Delta \to \Delta - \frac{\delta}{3}$ used for purely Gaussian diffusion, will generally be disregarded in the theoretical developments presented in this thesis. Additional corrective terms that account for e.g. finite slew rates are similarly excluded.

[2] We opt for the convention $q = \gamma \delta G_p$ over $q' = (2\pi)^{-1} \gamma \delta G_p$, as this simplifies many of the computations in later chapters.

[3] M inherits the Euclidean structure from $\mathbb{R}^3$, which allows us to trivially identify $(\mathbb{R}^3)^*$ with $T_p^* M$ for any $p \in M$.

signal in Eq. (1.2) would be real-valued[1] [153, 177, 222, 258, 352], where we assume that (active) flow is negligible in our considered experiments [244].

Remark 1. *The gradient strength G and time parameters δ and Δ are typically combined into a single parameter called the b-value:*

$$b = \gamma^2 \delta^2 G^2 \Delta. \tag{1.3}$$

We assume γ, δ, and Δ to be constant for a single experiment.

Remark 2. *The fact that P has compact support Ω is a physical reality [194], but many useful models (like diffusion tensor imaging [24, 25]) do not account for this explicitly. In the current manuscript compactness of Ω forms an essential part of our theoretical analysis of P.*

In practice we cannot measure the motion of spins at a single point $\boldsymbol{p}$. We are limited to measuring the ensemble dynamics of spins in a small neighborhood of $\boldsymbol{p}$: the voxel $V_{\boldsymbol{p}}$. We thus measure the voxel-averaged displacement probability, called the ensemble average propagator (EAP), and the corresponding normalized average signal. We retain the form and interpretation of Eq. (1.2) with the introduction of the following assumption:

Assumption 1. *At the voxel level we invariably assume that structures that impact spin dynamics are distributed approximately uniformly over $V_{\boldsymbol{p}}$, in the sense that spin dynamics over the diffusion time Δ can be taken to be similar at different points $\boldsymbol{x} \in V_{\boldsymbol{p}}$. Furthermore, we assume that the spatial variation of non-diffusion related parameters such as relaxation times can be considered negligible within $V_{\boldsymbol{p}}$.*

The validity of this assumption is dependent on the spatial resolution (voxel size) and on the diffusion time. In practice a voxel may fail to be homogeneous, and instead consist of two or more distinct partitions that separately satisfy Assumption 1. In this case we speak of partial volume effects, which we will discuss at several points in the thesis. Assumption 1 also justifies our treatment of non-diffusion-related effects such as T_2 relaxation, see e.g. the work by Novikov et al. [250, 251].

1.2   Thesis overview

In the following chapters we will focus on three different domains of diffusion MRI research: modeling diffusion in complex environments, practical processing of the

---

[1]In reality there are distortions that can introduce an additional phase in the signal, and usually one considers the magnitude of the signal to account for this. This has a significant practical impact on the statistical properties of the signal [156], and there have recently been proposed alternatives that aim to improve on this [116, 277]. It also destroys any sign information present in the complex-valued signal, though the signal is typically considered to be positive at practically attainable b-values [367].

data, and the inference of (structural) information. These three domains are each covered in a separate part, and the general ideas behind each part are briefly introduced here.

1.2.1 Metrics as models

MRI scans can be made sensitive to molecular displacements, be they the results of diffusion, convection, or otherwise. Even the simplest such experiments, e.g. the original diffusion-weighted acquisitions of fluid volumes, are difficult to interpret without the use of models. In the case of a simple liquid, one for example assumes that the relevant bulk behavior of the fluid can be described by the diffusion equation [117, 120], in which case the EAP is simply the Green's function

$$P(\boldsymbol{r}) = \frac{1}{\sqrt{(4\pi\alpha\Delta)^3}} e^{-\frac{\|\boldsymbol{r}\|^2}{4\alpha\Delta}} \tag{1.4}$$

with α the diffusion coefficient of the fluid and $\|\cdot\| : T_xM \to [0, \infty)$ the Euclidean norm[1]. The diffusion-weighted data can then be used to measure the diffusion coefficient intrinsic to different fluids [349] using

$$S(\boldsymbol{q}) = e^{-\alpha\Delta\|\boldsymbol{q}\|^2}, \tag{1.5}$$

which follows by substituting Eq. (1.4) in Eq. (1.2).

In the practical and clinical settings where diffusion MRI is mostly applied nowadays [6, 154, 166, 183], models play an even bigger role. Instead of acquiring data for a pure liquid, one generally observes the interaction between the liquid and some surrounding structure, and it is often information about the liquid's surroundings that one is after in an experiment. Not surprisingly, the importance and complexity of appropriate models grows with the complexity of the sample in the scanner, and without accurate models it is practically impossible to learn anything from an acquisition.

In modern diffusion MRI, measurements are thus sensitized to *structure* at the typical diffusion length scale of microns, and the obvious challenge is to find the relation between parameters that describe the diffusion, and quantities that can describe some aspects of the local structure. The archetypical example—that is also the main testing ground in this thesis—is the human brain, wherein diffusion MRI is used to answer questions about the neuronal composition of tissue. Common considerations are the densities, diameters, and local orientations of neurons in a voxel, and how these factors influence the effective diffusion process.

---

[1]The dual norm associated with $\|\cdot\|$ is simply the standard Euclidean norm on T_x^*M, and we denote this dual norm by $\|\cdot\|$ as well.

(a) Schematic of the structure in a voxel   **(b)** Geometrical model corresponding to (a)
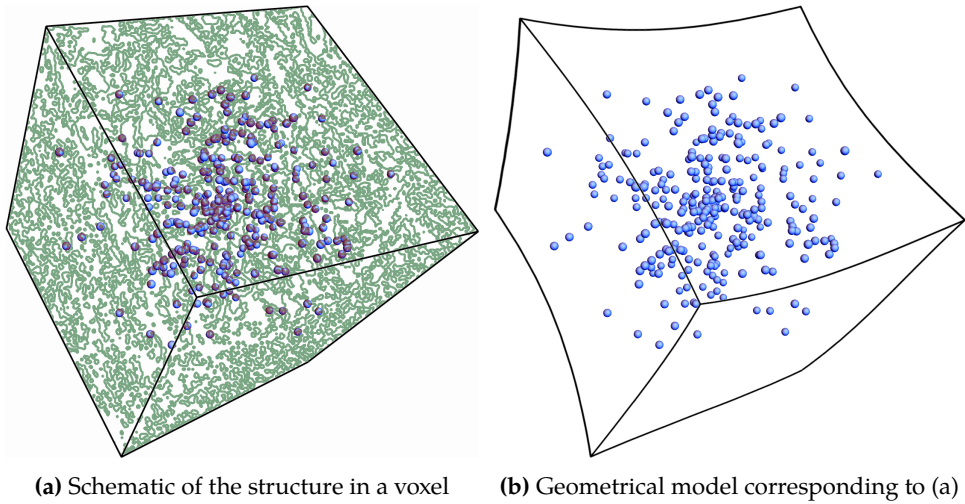
Figure 1.2: An illustration of the idea behind geometrical modeling. The black box in (a) represents the space in which diffusing spins (blue, originating from the center) are measured, i.e., the voxel. Tissue/structural complexity is represented in (a) by the green lines that interact in different ways with the spins, and this structure is replaced in (b) by a deformation of the space (deformed black box). The ultimate goal of geometrical modeling is to come up with the simplest geometry (or deformation) that reproduces the same effective diffusion of spins as in the original space, in which case all measurable information about the structural complexity in the imaged sample is encoded in the geometry of the deformed space.

With the wide array of potential applications of diffusion imaging, an enormous variety in models has been proposed in the literature, see the references for overviews [16, 119, 250, 265]. In Part I of this thesis we introduce the class of geometrical models, which attempt to model arbitrary spatial stochastic processes in Euclidean space as simple Brownian, isotropic diffusion processes in a more complicated space. The principal idea behind this is that we translate information in the stochastic process (which represents the observed diffusion) to a space, in such a way that the local geometry of the space describes some or all structural properties that are of import to the acquisition. This transition is illustrated in Fig. 1.2.

More specifically, we will assume in Part I that the observed diffusion can be viewed as if taking place in a (geometric) manifold, obtained through controlled deformations of Euclidean space. If one allows arbitrary deformations, then such a manifold can reproduce any diffusion process, and so geometrical models can in principle model every possible observation. But the degree (and type) of deformation that is allowed must be limited in reality, because one needs to be able to derive the structure of the manifold from the acquired measurements. Diffusion MRI modelling is thus reduced to finding descriptive geometries that can be related one-to-one with the measured stochastic process, where more complex geometries lead to more general descriptors of the diffusion and ambient structure.

The two chapters in Part I cover the first non-trivial geometrical models, using Rie-

mannian (Chapter 2) and Finsler (Chapter 3) manifolds. The metrics that describe the local structure of these manifolds are shown to be inferrable from diffusion-weighted measurements under the assumptions discussed in this chapter, and in general these metrics capture complementary information about the underlying structure. The theories of Riemannian and Finslerian geometry are well-established [21, 54, 187, 293], and particularly relevant for diffusion MRI is that spatial (inter-voxel) relations/interactions can be handled naturally (see also Section 1.2.3).

The resulting Riemannian framework was introduced over a decade ago by O'Donnell et al [254], and corresponds directly to one of the foundational diffusion MRI models, diffusion tensor imaging [24, 25, 218]. Our recent contributions to this framework can be found in the references [86, 138, 140–142]. Finsler geometry was first applied in diffusion MRI by Pichon et al. [271] and Melonakos et al. [227, 228], and a Finslerian framework was developed over several years in the works leading up to this thesis [86–90]. The Finslerian framework applies for example to generalized diffusion tensor imaging [222] and diffusional kurtosis imaging [178] (DKI). Methods based on the Finslerian framework are presented in Chapters 6 and 7, where we show for example that the local Finsler geometry can in principle capture details about the orientations of crossing elongated structures, which is impossible in the Riemannian framework. The application of Finsler geometry in diffusion MRI led to the development of a constrained reconstruction algorithm that is also useful outside of geometrical modeling applications.

1.2.2   Data processing

Once we have a local description of the data in terms of physically relevant model parameters, we can further analyze and present the data attuned to a specific application. This typically involves some data processing—registration [147, 155, 176, 201, 216, 253, 284, 326, 381] and resampling [35, 74, 115, 133, 198, 209, 270] to facilitate comparisons between different data sets, enhancements so that features of interest stand out [31, 93, 103, 112, 195, 211], normalization [3, 232] to handle different acquisition protocols, differences between scanners, or different gradient coil configurations, and so forth. Some of these processing steps translate well from standard image processing, but specialized algorithms arise e.g. from cursory knowledge about what kind of local structures one may expect, or by exploiting potential correlations between voxels [31, 93, 112–114]. Some insights on this part of the diffusion MRI pipeline are discussed in the two chapters of Part II.

In Chapter 4 we consider an enhancement algorithm involving controlled deblurring of the EAP [94]. The motivation for this approach is rooted in the notion that extraneous interactions between spins and most unsought structures have no long-range spatial correlations. This implies that the effect of these interactions on the EAP can be modeled approximately as a Gaussian convolution, which we combine with

a position-dependent heuristic to define a locally adaptive enhancement algorithm. Preliminary results show that this method can be used e.g. to drastically improve contrast in visualizations of commonly used scalar measures.

Next we discuss interpolation in Chapter 5. Interpolation concerns the inference of model parameters at positions where we do not have any direct measurements. A dependable interpolation scheme is crucial in many practical applications—it allows a data set to be upsampled as a simple way to improve the accuracy of many algorithms, downsampled to speed up an analysis, or resampled to define a point-wise correspondence with a different data set. The simplest way to do interpolation in diffusion MRI is by applying standard (e.g. polynomial or spline) interpolation on the raw measurements before reconstructing the model parameters, but this approach does not incorporate non-trivial spatial correlations that are often inherent to the data. Starting from the geometrical frameworks of Part I, we briefly review some of the existing interpolation schemes in the Riemannian setting, and discuss how they might fit in the Finslerian framework. We also propose a way to apply Finsler-based interpolation schemes to Riemannian data, by embedding the Riemannian structure in a Finsler manifold [133].

1.2.3   Information from geometry

In the final part of this thesis we treat various geometry-inspired means by which information can be extracted from diffusion-weighted MRI data, and how these can be used in neuroimaging applications. We consider not only features that follow naturally from the frameworks proposed in Part I, but also take a closer look at the large-scale geometrical organization of the human brain. Part III is divided in three chapters: 'Point', 'Curve', and 'Surface'.

Chapter 6 – 'Point', covers a broad range of local features—properties of the geometry or the measured diffusion that are defined at the level of the voxel. The key contribution in this chapter is the proposal of the barrier orientation distribution function (barrier ODF or bODF) [84, 89, 90], which describes the orientations of correlations in structures encountered by diffusing spins. This distribution is inspired by related concepts in convex and stochastic geometry [298, 324], and it can be computed from the local structure of either a Riemannian or Finsler manifold. By assuming that the majority of these correlations in neuroimaging are the result of axonal fiber pathways [30], the barrier ODF is hypothesized to capture the orientation distribution of axons in a voxel, placing it in the same realm as e.g. the widely used fiber ODF [96, 97, 342].

We further look at inter-voxel characteristics of a sample captured by our geometrical models. As a rule, these are described by the differential structure of the manifold, i.e., in terms of derivatives of the local geometrical structure. The most important of these characteristics is likely the geodesic—the analog to a straight line

in curved space. Geodesics can be used to compute curves of minimal length relative to the geometry of the manifold, and as such describe the evolution of the diffusion propagation front [29, 254]. This leads to promising applications in e.g. the simulation of tumor growth in fibrous tissues [243]. Geodesics have also been used for tractography—the pursuit to create a map of the structural connections in the human brain [26, 33, 99, 242, 287]—and more abstractly as a tool to describe the *degree* of structural connectivity between two regions [29, 130, 218, 269]. These last two applications are the focus of Chapter 7 – 'Curve' [90, 91].

In 'Surface', Chapter 8, we study the geometrical structure of the brain at the level of fiber pathways. A remarkable hypothesis about the organization of these pathways was formulated and argued for in a widely publicized study by Wedeen et al. [365]. This hypothesis concerns the observation that "pathways formed parallel sheets of interwoven paths [throughout the brain]" [365], and has since become known as the 'sheet structure hypothesis'. The presented evidence in favor of a basal sheet structure was mainly qualitative, but its validity would impose significant and *verifiable* integrability constraints on the planes formed by vectors tangent to fiber pathways.

The sheet structure hypothesis proved a catalyst for debate in the diffusion MRI and neuroscience communities [57, 95, 327, 331, 363–365], and so far there has been no decisive answer regarding the hypothesis' validity. Our contributions to this ongoing discussion are presented in the last chapter of this thesis, where we formalize the integrability conditions imposed by the presence of sheet structure, and introduce a quantitative measure of the local integrability called the sheet probability index (SPI) [328–335].

1.3   Notational conventions

Effort has been made to adhere to the following notational conventions. Vectors, tensors, and elements of M are denoted by boldface italics. Scalar variables, including tensorial components, are written as light italics. Covectors are generally viewed as tensors and follow the same notational conventions. Unless specified otherwise, summation indices will run from 1 through the number of dimensions, which is typically 3. A coordinate chart $(x^1, x^2, \dots)$ on a neighborhood of $x \in M$ induces a basis $\left\{ \frac{\partial}{\partial x^1}, \frac{\partial}{\partial x^2}, \dots \right\}$ on $T_x M$, so that we can express $y \in T_x M$ locally as $y = \sum_i y^i \frac{\partial}{\partial x^i}$. Hence, a tangent vector y_x at $x \in M$ can be identified with the directional derivative of any suitably differentiable function $f : M \to \mathbb{R}$, and we may write $y_x f$ to denote the derivative of f at point x along y.

A few symbols are doubly defined, usually in cases where their meaning is so ingrained in a certain context, that changing them would only be more confusing.

1.4 Algorithms and data

Most of the code used to generate the results in this thesis is included in a Wolfram Language package called Diffusica`, which is built on top of the Classes` package. Classes` is already publicly available (`github.com/tomdelahaije`), and an alpha release of Diffusica` is currently in preparation. Both packages will be released under the Apache 2.0 open source license, and some details on the implementation can be found in Appendix A. Other software used to generate results is referenced explicitly throughout the text.

The experiments presented in this thesis are performed on synthetic and real data sets, which are described in the relevant sections. Data were provided in part by the Human Connectome Project (HCP), WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. Data collection and sharing for this project was provided in part by the MGH-USC Human Connectome Project (Principal Investigators: Bruce Rosen, M.D., Ph.D., Arthur W. Toga, Ph.D., Van J. Wedeen, M.D.). HCP funding was provided by the National Institute of Dental and Craniofacial Research (NIDCR), the National Institute of Mental Health (NIMH), and the National Institute of Neurological Disorders and Stroke (NINDS). HCP data are disseminated by the Laboratory of Neuro Imaging at the University of California, Los Angeles.

Chapter 1

Part I

Modeling

Chapter 2

Riemann

The Riemannian Framework

Based in part on:

T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. "Finslerian diffusion and the Bloch-Torrey equation". In: *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*. Ed. by I. Hotz and T. Schultz. Cham: Springer International Publishing, 2015, pp. 21–35. ISBN: 978-3-319-15089-5 978-3-319-15090-1.

A. Fuster, T. Dela Haije, A. Tristán-Vega, B. Plantinga, C.-F. Westin, and L. Florack. "Adjugate diffusion tensors for geodesic tractography in white matter". In: *Journal of Mathematical Imaging and Vision* 54.1 (2016), pp. 1–14. ISSN: 0924-9907, 1573-7683. DOI: 10.1007/s10851-015-0586-8.

In a diffusion magnetic resonance imaging (MRI) experiment the scanner is sensitized to the micro-scale Brownian movement of water molecules, and the resulting measurements reflect the voxel-aggregated molecular displacements along a number of predetermined directions, recall Section 1.1. In diffusion tensor imaging (DTI) [24, 25] specifically it is assumed that this averaged behavior can be described by an anisotropic diffusion equation. The Riemannian framework for diffusion MRI pioneered by O'Donnell et al. [254], provides an intuitive geometrical interpretation of this model. Instead of assuming anisotropic diffusion, one considers the water molecules to be subject to a simple isotropic diffusion process in a Riemannian space; it has been known since the work of Lenglet et al. [218] that the resulting model is equivalent to DTI. Hence the measured sample is effectively modeled as a Riemannian manifold, on which the ensemble movement of water molecules is governed by a simple Brownian stochastic process [208], Section 1.2.1. The geometrical viewpoint in DTI has led to elegant solutions to problems such as interpolation, tractography, and so forth [13, 128, 198, 237, 238, 254, 270, 305].

The goal of this chapter is not only the introduction of the Riemannian framework as a first example of geometrical frameworks, but also its preceding practical counterpart, diffusion tensor imaging. DTI initiated the explosive growth in diffusion MRI research of the last decade, and much of the basic information available in diffusion MRI data is based on DTI. DTI results are often used as a reference in this thesis, and this chapter summarizes the relevant implementation and analysis details that we take for common knowledge in the remaining chapters.

In Section 1.2.1 we have outlined the overarching idea of geometrical frameworks for diffusion MRI, and for reasons that will become clear in this part, Riemannian geometry plays an integral part in these frameworks. We first summarize the relevant aspects of Riemannian geometry in Section 2.1. In Section 2.2 we next derive the diffusion tensor imaging model from the ansatz that the underlying tissue can be modeled as a Riemannian manifold, by which we establish a simple, invertible relation between the manifold metric structure and the second order moments of the ensemble average propagator (EAP). The presented derivation is based on a stochastic particle process called the isotropic (scattering) transport process, which reproduces the DTI model in the central limit. With this relation we can then define the Riemannian framework, slightly extending the original works by a.o. O'Donnell et al. [254]. In this section we also cover how we can define a rescaled Riemannian manifold, which is convenient for the geodesic computations that we look at in Chapter 7. In Section 2.3 we then discuss some aspects regarding the implementation of the Riemannian framework, and in Section 2.4 we present a basic illustrative example of reconstructed DTI data. The chapter is concluded with a brief discussion on the Riemannian framework, and its (potential) value in the analysis of diffusion MRI data. Parts of this chapter have been presented at the CDMRI and BASP workshops [138, 140, 142].

2.1 Background

2.1.1 Manifolds and norms

Riemannian geometry is a generalization of (standard) Euclidean geometry, and in many regards it can be understood in terms of the same familiar concepts—lengths of curves and angles between lines are for example well-defined notions in a Riemannian manifold as well. A Riemannian manifold can be viewed as a curved space, and we will refer to the local deformation as the (locally specified) geometry (recall Fig. 1.2). The local geometry of the manifold is described in full by the Riemannian metric tensor: "Riemannian geometry: [the] metric as foundation of all" [233, Chapter 13]. For a general introduction to the theory of Riemannian geometry, see e.g. the work of Jost [187].

Let M be a three-dimensional Riemannian manifold, equipped with a (non-degenerate) metric tensor g [187, 233]. Angles and distances are defined in terms of an inner product induced by the metric tensor, and the most important notion in this work is the corresponding norm on T_xM defined as

$$G(x, y) := \sqrt{\sum_{i,j} g_{ij}(x) y^i y^j},$$

(2.1)

for all $y \in T_xM$. This norm varies smoothly on M, and has all the usual properties (e.g. the metric tensor is symmetric; $g_{ij} = g_{ji}$, and positive-definite; $G(x, y) \geq 0$ for all $(x, y) \in TM$). If the metric tensor is independent of the position $x \in M$, where M contains some fiducial point p, we may identify all tangent spaces in TM with T_pM and thus make the identification

$$G(x, y) \equiv G(p, y) =: G(y)$$

(2.2)

for all $x \in M$. We did something similar (implicitly) in Eq. (1.2), where the Euclidean vector space structure justifies integration over T_pM, with $M \subset \mathbb{R}^3$ the set of points in the scanner.

2.1.2 The indicatrix

The indicatrix is the Riemannian analog of a unit sphere in T_xM defined as

$$I_xM = \{y \in T_xM \mid G(x, y) = 1\},$$

(2.3)

which carries a natural volume measure[1] η_x induced by the restricted metric. The associated probability measure μ is defined as

$$\mu_x := \frac{\eta_x}{\eta_x(I_x M)}.\tag{2.4}$$

The Riemannian indicatrix is always ellipsoidal, as it is a level set of a positive semi-definite quadratic polynomial, and its semi-principal axes are defined by the eigensystem of g. The eigensystem can be computed in the usual way from a matrix representation of the tensor [204]. Note that the Riemannian indicatrix and the Riemannian metric tensor represent the same exact information.

2.1.3 Geodesics

Geodesics extend the concept of a 'straight line' to more general geometries. A Riemannian geodesic between two points $x_1, x_2 \in M$ is a curve $C : [0, T] \to M$ that (locally) minimizes the length integral

$$\mathcal{L}_G(C) := \int_0^T G\left(C(t), \frac{\mathrm{d}C}{\mathrm{d}t}(t)\right) \mathrm{d}t,\tag{2.5}$$

where $C(0) = x_1$ and $C(T) = x_2$. Uniqueness and existence criteria can be found in the references [187], and we will provide some additional details in the next chapters.

2.1.4 Duality

The set of all linear functionals acting on $T_x M$ forms the dual vector space $T_x^* M$, which is equipped with the natural norm

$$G^*(x, w) := \sup_{y \in I_x M} \langle w, y \rangle = \sum_{i,j} g^{ij}(x) w_i w_j,\tag{2.6}$$

where $w \in T_x^* M$ and $g^{ij}(x)$ are the components of the inverse metric tensor $g^{-1}(x)$. The unit level set of the dual norm is referred to as the figuratrix.

2.2 Theory

The purpose of this chapter is to introduce the Riemannian framework for diffusion MRI, where Riemannian manifolds are used to model characteristics of the underlying (tissue) structure that have an influence on local diffusion properties. Reversing the chronological development of this framework, we start in Section 2.2.1 from the

---

[1] The exact form of η_x is not relevant for the discussion in this chapter.

assumption that the acquired measurements are actually obtained from a Brownian motion in a Riemannian manifold. In the next section, Section 2.2.2, we show how this leads to a model of the diffusion MRI signal from which we can reconstruct the local metric structure of the presupposed Riemannian manifold. This relation slightly extends the typical Riemannian framework as introduced in the references [218, 254]. In Section 2.2.3 finally we introduce an auxiliary conformal factor that narrows the gap between large-scale anisotropic diffusion in Euclidean space, and isotropic diffusion in Riemannian space [141].

2.2.1 Diffusion in a Riemannian manifold

There are many ways in which one can derive the bulk behavior of particles subject to a Brownian motion, which can be understood intuitively as oft-colliding particles that move along straight lines in-between collisions [43, 117]. The derivation of the Brownian motion equations that is most relevant in this work remains true to this intuitive notion of colliding particles, formalized as a Markov process (X_t, Y_t) called the isotropic (scattering) transport process. It is instructive to consider this process first in Euclidean space, defined as follows in the works of Monin [239] and Watanabe et al. [361, 362], where Brownian motion is generated by the scaled Laplacian $\Delta_E :=$ $\alpha \sum_i \frac{\partial^2}{\partial x^i \partial x^i}$ through the diffusion equation

$$\frac{\partial}{\partial t} = \Delta_E, \tag{2.7}$$

with α the diffusion coefficient as in Chapter 1.

Consider a particle initially at a position $x_0 \in \mathbb{R}^3$ moving in a straight line, with unit speed along the direction $y_0 \in S^2$, and denote its path by $\xi(x_0, y_0)(t)$ with $0 \leq t \leq \tau_1$ and τ_1 the time at which the particle first collides. Next assume that the time between collisions $e_i = \tau_i - \tau_{i-1}$ (where we set $\tau_0 = 0$) satisfies $P\{e_i > t\} = \exp(-t)$ with $t \geq 0$ [220]. After the first collision the particle continues from the point of collision $x_1 = \xi(x_0, y_0)(\tau_1)$ along a new direction y_1 chosen according to the uniform probability measure on the sphere S^2. A continuation of this construction produces the transport process (X_t, Y_t):

$$\left. \begin{array}{l} X_t = \xi(x_{i-1}, y_{i-1})(t) \\ Y_t = \frac{dX_t}{dt} \end{array} \right|_{\tau_{i-1} \leq t \leq \tau_i} \tag{2.8}$$

for all positive integers i. This process is illustrated in Fig. 2.1.

In order to show (weak) convergence of X_t to a Brownian motion we introduce a parameter $\beta > 0$ such that the mean distance the particle traverses between collisions scales linearly with β, while its speed scales with β^{-1}. This results in a fixed (β-independent) mean squared displacement for a fixed diffusion time, and we set the
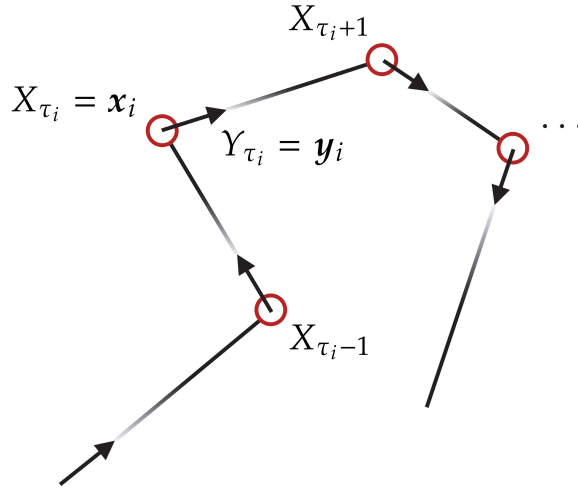
Figure 2.1: An illustration of the isotropic transport process. A particle (represented by a red circle) moves along a straight line along y_i from the point x_i, until at a time $t = \tau_{i+1}$ a (virtual) collision causes the particle to change direction. The collision times satisfy $P\{\tau_i - \tau_{i-1} > t\} = \exp(-t)$, where $\tau_0 = 0$. The transport process is defined by these paths X_t (black lines), and the direction of movement Y_t at each point, cf. Eq. (2.8).

ratio between the collision rate and the velocity squared to $\frac{2}{3}\alpha$, α as in Eq. (2.7). In the limit $\beta \downarrow 0$, we then recover Δ_E as the generator of this process (of which we will forego the technicalities discussed in the references [361, 362]) such that the density evolution of isotropic transport particles is dictated by the expected standard diffusion equation, Eq. (2.7).

Following the works by Watanabe et al. [361, 362], Pinsky [275] showed that an analogous stochastic process can be defined on the tangent bundle of a Riemannian manifold. His adaption requires a number of significant changes (see Pinsky [275] and the references therein for more details):

1. The trajectories that particles follow are now Riemannian geodesics, cf. Section 2.1.3, rather than straight lines (Euclidean geodesics).

2. The new direction of a particle after collision is now based on the rotationally invariant probability measure μ_x on the locally defined Riemannian unit sphere $I_x M$, cf. Eqs. (2.3) and (2.4), rather than a probability measure on the globally defined Euclidean unit sphere.

The remainder of the derivation is similar to the Euclidean case, and again a limit theorem on a parameter β is invoked to surmise a diffusion generator of the form

$$\Delta_R := \alpha \frac{1}{\sqrt{\det g(x)}} \sum_{i,j} \frac{\partial}{\partial x^i} \left(\sqrt{\det g(x)} \, g^{ij}(x) \frac{\partial}{\partial x^j} \right), \tag{2.9}$$

where det g is the determinant of the (matrix representation of the) metric tensor g. The generator Δ_R is the unique generalization of the scaled Laplacian Δ_E to Riemannian manifolds, and coincides for $\alpha = 1$ with the usual expression for the generator of Brownian motion in Riemannian manifolds, the Laplace–Beltrami operator [187]. As mentioned at the beginning of this section, this generator can be derived by various means, e.g. both Δ_E and Δ_R can be defined through the divergence of the gradient, concepts that are well-defined in both Riemannian manifolds and Euclidean space.

In the end, we thus consider the bulk behavior of diffusing spins in a sample as given by the diffusion equation

$$\frac{\partial}{\partial t} = \Delta_R. \tag{2.10}$$

Note that we tacitly assume that α is constant throughout M—though this means we ignore temperature differences, spatial dependencies in the chemical composition of the diffusing substance, and other effects that may impact the free diffusion coefficient. For now, we instead group these factors with the tissue components that determine the geometrical make-up of the manifold. When relevant, we will similarly assume that the collision rate and average velocity of particles in a voxel are constant.

2.2.2 The Riemannian framework for diffusion MRI

With the assumed voxel-wise homogeneity (Assumption 1), we can take $g(x) \equiv g(p) =: g$ for $x \in V_p$, recall Section 1.1, so that Eq. (2.10) simplifies to

$$\frac{\partial}{\partial t} = \alpha \sum_{i,j} g^{ij}(p) \frac{\partial^2}{\partial x^i \partial x^j}. \tag{2.11}$$

We obtain the ensemble average propagator P for the voxel at the implicit point p as the Green's function of this equation at time $t = \Delta$,

$$P(r) = \sqrt{\frac{\det g}{(4\pi\alpha\Delta)^3}} e^{-\sum_{i,j} \frac{g_{ij} r^i r^j}{4\alpha\Delta}}, \tag{2.12}$$

and the relation between g and the normalized signal S follows via Eq. (1.2):

$$S(q) = e^{-\alpha\Delta \sum_{i,j} g^{ij} q_i q_j}. \tag{2.13}$$

By taking the logarithm on both sides we obtain an equation that is linear in the components of the metric tensor g, from which we can reconstruct g up to the factor α if we have measurements of S for at least six non-collinear q. We can thus conclude that it is possible to reconstruct the local structure of a Riemannian manifold from observations of particles in Brownian motion using diffusion MRI: Eq. (2.13) defines

a Riemannian (geometrical) framework.

The link between Riemannian geometry and diffusion tensor imaging follows straightforwardly. The diffusion tensor imaging model assumes that the signal can be written as

$$S(\boldsymbol{q}) = \mathrm{e}^{-\Delta \sum_{i,j} D^{ij} q_i q_j}, \tag{2.14}$$

with D^{ij} the so-called diffusion tensor, which leads (by comparing with Eq. (2.13)) to the duality relation

$$\boldsymbol{D}^{-1} = \frac{1}{\alpha} \boldsymbol{g}. \tag{2.15}$$

The diffusion coefficient of the spin-carrying medium, α, was not explicitly included in earlier works on the Riemannian framework. We include it here because it serves important practical, theoretical, but also pedagogical purposes. As explained in a.o. Section 1.2.1, our goal with the introduction of these geometrical frameworks is to extract information about the structure surrounding the diffusing substance. This structure is independent of the diffusion process used to probe it[1], and this is made explicit by separating the diffusion coefficient α, with typical units mm^2/s, and the intrinsically dimensionless metric tensor $\boldsymbol{g}$. At the same time this separation allows us to compare measurements obtained using different liquid media, and to correct for external factors like temperature and pressure. If there is a spatial dependency in α, this can be included through a minor modification of Eq. (2.9) that we do not consider here [141].

2.2.3 A scaled Riemannian manifold

So far we have assumed that the metric tensor $\boldsymbol{g}$ extracted from $\boldsymbol{D}$ after factorization of α is equivalent to the geometry of interest. However, we can consider alternative manifold definitions by adapting Eq. (2.15), where we could add additional properties of the sample (known by other means) to the geometry, or remove unwanted properties from the structure. Perhaps the simplest example of such an adaptation is a position-dependent scaling of $\boldsymbol{D}$, corresponding to a conformal scaling of the metric. Such a conformal factor was proposed by e.g. Hao et al. [162] to improve geodesic tractography, which we will discuss in Chapter 7.

The use of a voxel-dependent conformal factor that we propose in this chapter is related to the following discussion. In the previous sections we derived the Riemannian framework for DTI by considering a Brownian motion on a Riemannian manifold, culminating in Eq. (2.15). This relation followed the requisite assumption that within a voxel the Riemannian structure was constant, allowing us to 'zoom in'

---

[1]With the one exception that different media (or different diffusion times) can be used to probe a structure at different *scales*, which will generally result in different effective geometries. This scale-dependence was briefly mentioned in Chapter 1, but here and in the remainder of the thesis we consider this scale to be fixed.

on the local geometry. Classically, derivations of DTI were not based on Δ_R, but on the anisotropic diffusion equation

$$\frac{\partial}{\partial t} = \sum_{i,j} \frac{\partial}{\partial x^i}\left(D^{ij}(x)\frac{\partial}{\partial x^j}\right), \qquad (2.16)$$

where D^{ij} is the diffusion tensor as before. Under Assumption 1, this equation obviously simplifies to the form of Eq. (2.11), so the local, observable behavior of the isotropic Riemannian and anisotropic Euclidean processes are indistinguishable. However, when looking at large-scale (inter-voxel) diffusion these equations generally prescribe qualitatively different behaviors [29, 254]. This is not a fundamental problem, as the anisotropic diffusion in Euclidean space still induces the Riemannian structure $g \propto D^{-1}$ (see e.g. de Lara [208]), but it does suggest the following D-dependent conformal factor [141].

Recall the identity in Eq. (2.15), and apply the substitution $g \rightarrow \alpha f(x)\, D^{-1}$ in Eq. (2.9) with $f \in C^\infty(M, (0, \infty))$. We then find

$$\begin{aligned}
\Delta_R &= \sqrt{\frac{\det D(x)}{f(x)^3}} \sum_{i,j} \frac{\partial}{\partial x^i}\left(\sqrt{\frac{f(x)}{\det D(x)}}D^{ij}(x)\frac{\partial}{\partial x^j}\right) \\
&= \sqrt{\frac{\det D(x)}{f(x)^3}} \sum_{i,j}\left(\frac{\partial}{\partial x^i}\sqrt{\frac{f(x)}{\det D(x)}}\right)D^{ij}(x)\frac{\partial}{\partial x^j} + \frac{1}{f(x)}\sum_{i,j}\left(\frac{\partial}{\partial x^i}D^{ij}(x)\frac{\partial}{\partial x^j}\right),
\end{aligned}$$
$$(2.17)$$

giving

$$\sum_{i,j}\left(\frac{\partial}{\partial x^i}D^{ij}(x)\frac{\partial}{\partial x^j}\right) = f(x)\Delta_R - \sqrt{\frac{\det D(x)}{f(x)}} \sum_{i,j}\left(\frac{\partial}{\partial x^i}\sqrt{\frac{f(x)}{\det D(x)}}\right)D^{ij}(x)\frac{\partial}{\partial x^j}. \quad (2.18)$$

It follows that by taking $f \propto \det D$, the anisotropic diffusion process scaled by a factor $f(x)^{-1}$ is equivalent to a Brownian motion in a Riemannian manifold defined by the metric $(\alpha \det D)\, D^{-1}$ [141, 142]. When there are reasons (e.g. of a phenomenological nature) to assume that Eq. (2.16) is valid at the voxel scale in the measured sample, then it can be argued that it can be advantageous to include the corrective factor $\det D$ in the definition of the metric tensor, as the resulting geometry will more closely reproduce the actual spin dynamics. Note that this advantage is null in regions where the determinant is constant, which can alternatively be concluded by comparing Eq. (2.10) and Eq. (2.16).

2.3 Methods

The Riemannian framework [218, 254], defined in Eq. (2.13), provides an interpretation of diffusion tensor imaging data in terms of Riemannian manifolds [187] through the remarkable observation that Brownian motion in a Riemannian manifold captures the complete geometric structure. In this section we describe the implementation used to perform the basic DTI experiments in this chapter and the comparative experiments in the remainder of the thesis. We also discuss basic means to gather information from reconstructed DTI data, and introduce some visualization concepts that are used throughout this work. Reconstruction of the DTI model, and all subsequent computations and visualizations shown in this chapter were performed with Diffusica`, see Appendix A.2.

2.3.1 Scalar measures

Reconstruction of D produces a tensor field—a positive-definite matrix at each $p \in M$. These tensors describe the directionality and general amount of diffusion, which are typically visualized using different scalar maps derived from D. Of course we could create maps of the components of D, but these are typically not very informative, as they lack invariance under a change of coordinate basis. Practically useful diffusion descriptors do satisfy this basic invariance [25], and can thus be expressed in terms of tensor invariants, in this case e.g. the eigenvalues of the diffusion tensor or the invariants defined in the original work by Basser et al. [25]. By positive-definiteness of D all its eigenvalues are positive, and sorted from small to large they are denoted $\lambda_{\min}$, λ_{med}, and $\lambda_{\max}$. The largest eigenvalue $\lambda_{\max}$ is also called the axial diffusivity (AD), while the average of the other two eigenvalues $\frac{\lambda_{\mathrm{med}} + \lambda_{\min}}{2}$ is called the radial diffusivity (RD).

The first important scalar measures that are commonly used in practice convey information about the overall amount of diffusion in a voxel. Of these, the most relevant are the mean diffusivity (MD), defined as

$$\mathrm{MD}(D) := \frac{1}{3}\operatorname{tr} D = \frac{\lambda_{\min} + \lambda_{\mathrm{med}} + \lambda_{\max}}{3}, \qquad (2.19)$$

and the determinant

$$\det D = \lambda_{\min}\lambda_{\mathrm{med}}\lambda_{\max}. \qquad (2.20)$$

A higher MD or determinant implies a greater amount of diffusion, and especially the mean diffusivity is also used extensively in clinical applications [41, 42, 60, 71, 303, 311].

The second set of scalar measures that are often used are the geometrical measures (or shape measures) of the diffusion tensor [372, 373]. These do not tell anything about

the amount of diffusion, but instead provide a quantification or classification of the orientation dependence. The most important shape measures we use here are the fractional anisotropy (FA) [27]

$$\mathrm{FA}(\boldsymbol{D}) := \sqrt{\frac{1}{2}} \frac{\sqrt{(\lambda_{\max} - \lambda_{\mathrm{med}})^2 + (\lambda_{\max} - \lambda_{\min})^2 + (\lambda_{\mathrm{med}} - \lambda_{\min})^2}}{\sqrt{\lambda_{\min}^2 + \lambda_{\mathrm{med}}^2 + \lambda_{\max}^2}} \qquad (2.21)$$

and the linearity, planarity, and sphericity coefficients defined respectively as

$$c_{\mathrm{l}}(\boldsymbol{D}) := \frac{\lambda_{\max} - \lambda_{\mathrm{med}}}{\lambda_{\min} + \lambda_{\mathrm{med}} + \lambda_{\max}}, \qquad (2.22)$$

$$c_{\mathrm{p}}(\boldsymbol{D}) := \frac{2\,(\lambda_{\mathrm{med}} - \lambda_{\min})}{\lambda_{\min} + \lambda_{\mathrm{med}} + \lambda_{\max}}, \qquad (2.23)$$

$$c_{\mathrm{s}}(\boldsymbol{D}) := \frac{3\lambda_{\min}}{\lambda_{\min} + \lambda_{\mathrm{med}} + \lambda_{\max}}. \qquad (2.24)$$

All of these measures lie in the range $[0, 1]$, and the latter three are normalized such that they sum to 1 for a given tensor. The FA is high if there is a large variation between the diffusivities along different orientations, and the linearity, planarity, and sphericity coefficients are high if the diffusion is close to linear, planar, or spherical respectively. Like MD, FA has been considered in a large number of clinical studies [42, 71, 107, 109, 230, 290, 311].

2.3.2 Glyph visualizations

If we want to show the full information contained in the tensor field, the scalar measures discussed in the previous section are not sufficient. Instead, we occasionally rely on glyph visualizations. A glyph in this context is a local surface rendering of a spherical function derived from the data at a given position, created by deforming a sphere such that its radius for a given direction is proportional to the corresponding function value. To visualize a DTI tensor, we typically define glyphs as (scaled) level sets of $\sum_{i,j} D_{ij}(\boldsymbol{x})y^i y^j$ [274] (Section 2.1.2), with D_{ij} the components of $\boldsymbol{D}^{-1}$, though there are alternatives [197, 356]. The DTI tensor is thus visualized as an ellipsoid, with radii that are representative of the root mean squared displacement of the spin ensemble along an orientation. The semi-principal axes of the ellipsoid have lengths that are proportional to the square roots of the eigenvalues, and the axes are aligned with the eigenvectors of $\boldsymbol{D}$. The geometrical measures introduced in the previous section describe the shape of the DTI ellipsoids. In terms of the Riemannian framework, the DTI ellipsoids are conformal to the indicatrices in the Riemannian manifold.

Unless otherwise specified, we encode the approximate orientation of the diffusion tensor in the color of the DTI ellipsoid, where the primary colors red, green, and blue are associated with the ellipsoid's alignment along three orthogonal axes [264]. In neuroimaging experiments, these are the standard subject space axes: red is associated to the left-right (LR) axis, green to the anterior-posterior axis (AP), and blue to the inferior-superior axis (IS). We determine the weights of each color either by normalizing the $\lambda_{\max}$ eigenvector (expressed in subject space coordinates), or by looking at the relative values on the diagonal of the diffusion tensor (again expressed in subject space coordinates). The latter of these is in some cases preferred, as the resulting color map is guaranteed to be smooth if the tensor field is smooth.

We will rely on glyph visualizations for other types of data in later chapters as well.

2.3.3 Data

By their abstract nature, most methods presented in this thesis are quite generally applicable, though they are developed with neuroimaging applications in mind. The original Riemannian framework was proposed in this context, and there are high quality public data sets available that are helpful in method prototyping and testing, recall Section 1.4. Many of the figures in this thesis intended as proof-of-principle or as illustrations, including those in this chapter, are generated from such data.

The preliminary real data experiments presented in the next section are based on brain images and parcellations of a subject of the WU-Minn Human Connectome Project (HCP) [351]. The diffusion-weighted data has an isotropic voxel size of 1.25 mm, and consists of three shells with 90 approximately uniformly distributed orientations each, $b = \{1000, 2000, 3000\}$ s/mm^2 (Eq. (1.3)), and 18 $b = 0$ s/mm^2 baseline images.

The Riemannian framework imposes the same data requirements as the DTI model—at least six non-collinear diffusion-weighted images, and a baseline image to compute the normalized signal. The optimal implementation of DTI has been extensively investigated [24, 203, 294, 354], and based on these works we use a weighted linear least squares approach using the iterative reweighting strategy proposed by Salvador et al. [294], using the baseline images and the $b = 1000$ s/mm^2 shell [184, 375]. The local Riemannian metric tensors are computed through a matrix inversion of the reconstructed diffusion tensors based on Eq. (2.15), where we arbitrarily fix $\alpha = 1$ mm^2/s. The scaled metric is computed by multiplying the metric tensors with $\det D$, cf. Section 2.2.3. It was not necessary to impose a positive-definite constraint on the diffusion tensors [202, 360], and we do not use any outlier correction [62]. None of the experimental data sets require additional pre-processing—we used the pre-processed HCP data which is already corrected for distortions due to eddy currents and motion [217, 289].
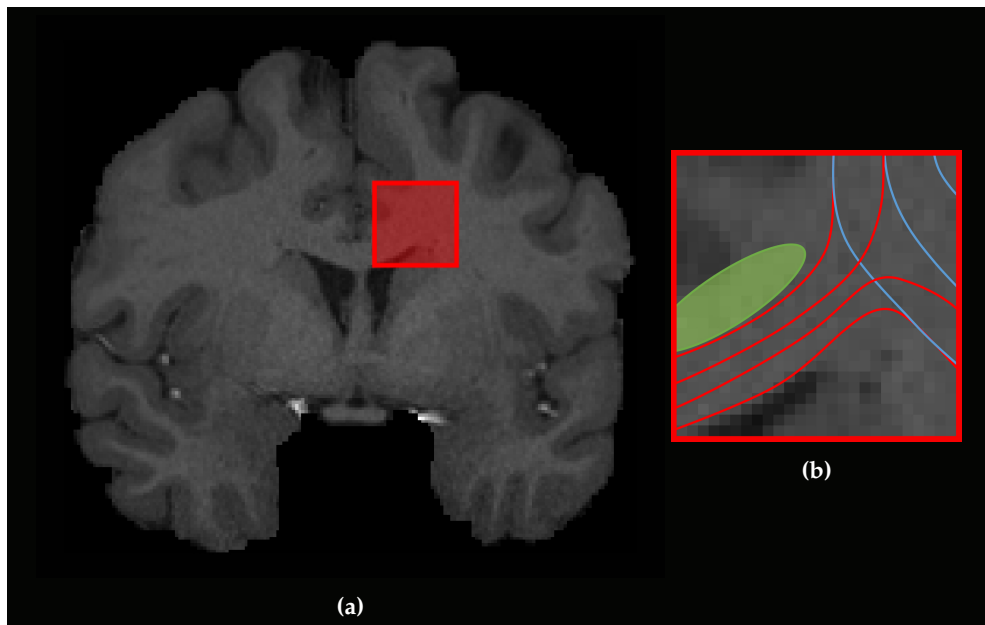
Figure 2.2: (a) An anatomical (T1-weighted) image of a coronal slice of the HCP data set used in this work. The indicated region marks a segment of the slice used in a number of experiments. (b) Part of the coronal slice shown in (a) with a schematic drawing that highlights the structure of some large bundles that are expected to be present in this segment based on anatomical references [59]. The red lines indicate fibers originating from the corpus callosum, which consists of strongly aligned fiber tracts that connect the left and right hemispheres. They cross and merge in the centrum semiovale with the corticospinal tracts indicated in blue. The darker voxels below the corpus callosum show a small part of a ventricle. The green region indicates tracts in the cingulum, which lie perpendicular to the visualized coronal slice.

The experiments in this chapter are focussed on a representative coronal slice in the HCP data. In Fig. 2.2 we plot this slice in an anatomical (T1-weighted) image, where we can recognize most of the qualitatively different structures that we expect to find in the data. The slice contains white matter regions, consisting of strongly aligned, long-range myelinated axons in single and crossing fiber regions (light gray); gray matter regions containing unmyelinated neurons and cell bodies (dark gray); and cerebrospinal fluid (CSF) in the ventricles and around the brain (black). A segmentation of these different regions is shown in Fig. 2.3. In Fig. 2.2(b) we sketched the expected fiber pathways near the centrum semiovale (indicated by the red box in Fig. 2.2(a)), which is an often used test bed in diffusion MRI because of the complex configuration of crossing fibers that can be found there.
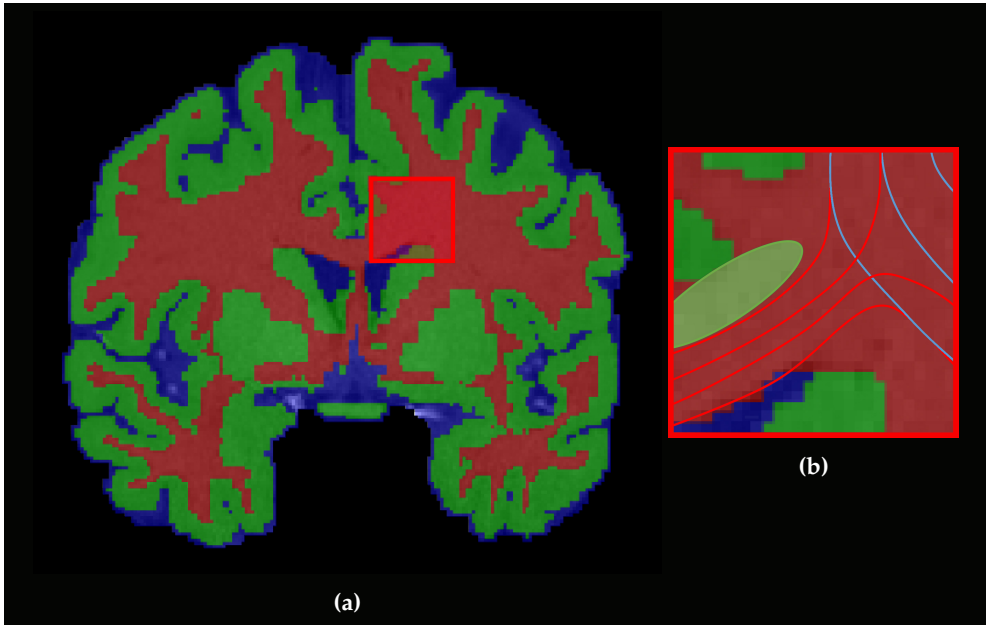
Figure 2.3: The coronal slice shown in Fig. 2.2 with an overlay identifying the white matter (red), gray matter (green), and other voxels including those containing significant volumes of cerebrospinal fluid (blue).

2.4   Results

In this section we present some examples of the scalar measures and glyph visualization presented in the previous section, as well as explain how the proposed scaled metric compares to the normal metric. An evaluation of the scaled metric in a practical setting is done in Chapter 7. The goal of this section, as of the chapter in general, is to acquaint one with the diffusion tensor imaging model and its geometrical interpretation.

In Figs. 2.4, 2.5, and 2.6 we plot the scalar measures introduced in Section 2.3.1 for the coronal slice of the HCP data shown in Fig. 2.2(a), and in Fig. 2.7(a) we show the corresponding ellipsoidal level sets in the small region of interest near the centrum semiovale.

The different tissues observable in the HCP data (Fig. 2.3) have very distinct diffusion characteristics. In white matter, axons act as barriers to diffusing particles, resulting in smaller values for the mean diffusivity and the determinant, see Fig. 2.4. Because these axons are strongly aligned and relatively densely packed, spins can move more freely along the fiber pathways than perpendicular to them, resulting in a relatively high linearity coefficient and fractional anisotropy, Figs. 2.5 and 2.6(a). The

(a) Mean diffusivity   (b) Determinant
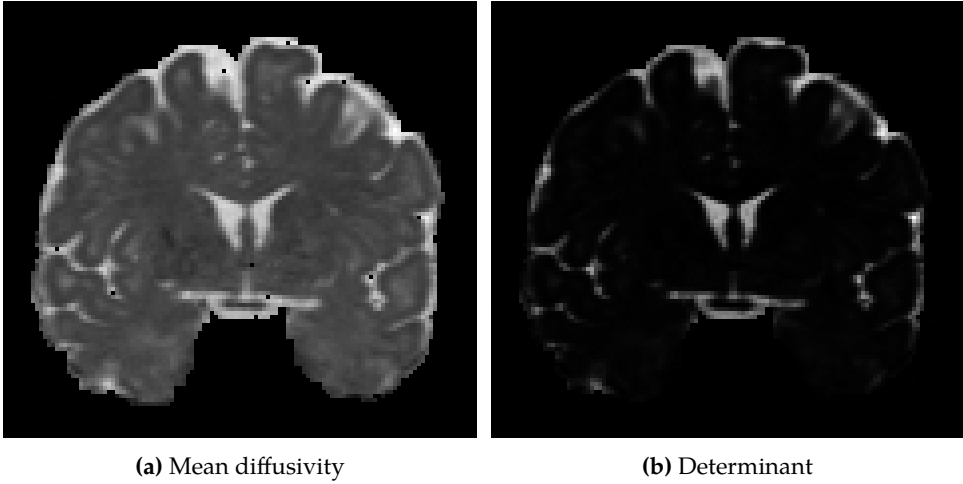
Figure 2.4: The mean diffusivity (a) and determinant (a) of the local diffusion tensors reconstructed for data in a coronal slice of the HCP data set described in Section 2.3.3. The mean diffusivity represents the average rate of diffusion in a voxel, and the determinant can be interpreted as the (Euclidean) volume of the DTI ellipsoid. These measures are comparable in white and gray matter (cf. Fig. 2.3)—where diffusing spins typically encounter a large number of barriers—and much higher in e.g. the cerebrospinal fluid, where spins move relatively unobstructed.

corresponding level sets—the glyph visualization in Fig. 2.7(a)—are less spherical, and the sphericity coefficient is relatively low, Fig. 2.6(c). In white matter regions containing e.g. crossings of fiber pathways, anisotropy is a bit reduced comparatively, and the glyphs are more flat than elongated (Fig. 2.7(a)), which is reflected in a higher planarity coefficient (Fig. 2.6(b)) and a lower linearity coefficient (Fig. 2.6(a)). Without the presence of large-scale structural coherence, both gray matter and CSF regions are fairly isotropic, i.e., they have a high sphericity coefficient and a low anisotropy measure (Figs. 2.6(c) and 2.5). The mean diffusivity is comparable between white matter and gray matter, but much higher in CSF, cf. Fig. 2.4(a).

Through Eq. (2.15) we link the diffusion tensor to a metric tensor, providing an alternative interpretation of the diffusion-weighted measurements. While the diffusion tensor describes the mean squared displacements of spins at a given position and along a given orientation, the metric tensor describes the local, orientation-dependent deformation of the space. Consequently, diffusion scalar measures have an analogous interpretation in terms of this deformation. A small mean diffusivity or determinant implies space is compressed relative to the standard Euclidean space, which means a unit displacement relative to the Riemannian geometry is smaller than a unit displacement relative to the Euclidean geometry. A high fractional anisotropy means that the deformation of the space is strongly dependent on direction, and the geometric measures further characterize this anisotropy.

With this in mind, we can also interpret the scaled Riemannian manifold pro-
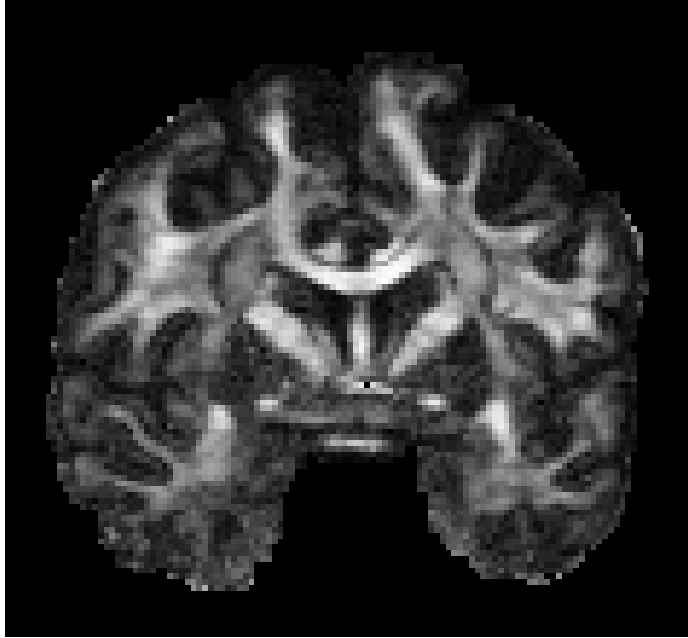
Figure 2.5: A map of the fractional anisotropy (Eq. (2.21)) shown in the same slice as shown in Figs. 2.2 and 2.4. FA is zero when diffusion is isotropic, i.e., when $\lambda_{\min} = \lambda_{\text{med}} = \lambda_{\max}$, which is the case in the ventricles (dark regions in the center of the brain). If there is more diffusion along one direction than others, FA increases, as can be seen in e.g. the corpus callosum (recognizable as the bright, slightly curved set of voxels above the ventricles that connect the left and right hemispheres).

(a) Linearity   **(b)** Planarity   **(c)** Sphericity
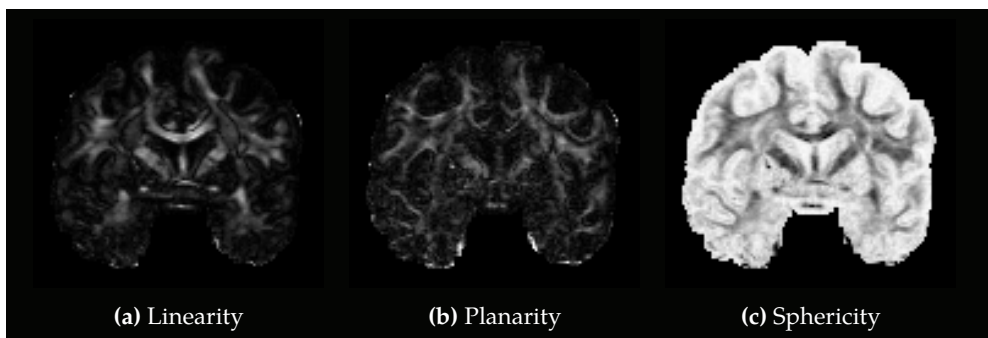
Figure 2.6: The geometrical measures proposed by Westin et al. [372, 373], that are used to characterize the shape of the DTI ellipsoids that represent the full local diffusion information. These measures are insensitive to variations in the overall size of the ellipsoids, and instead give a measure of elongation (a), planarity (b), and sphericity (c). A high sphericity implies low fractional anisotropy (i.e. a high degree of isotropy), and a high linearity implies high fractional anisotropy, cf. Fig. 2.5.
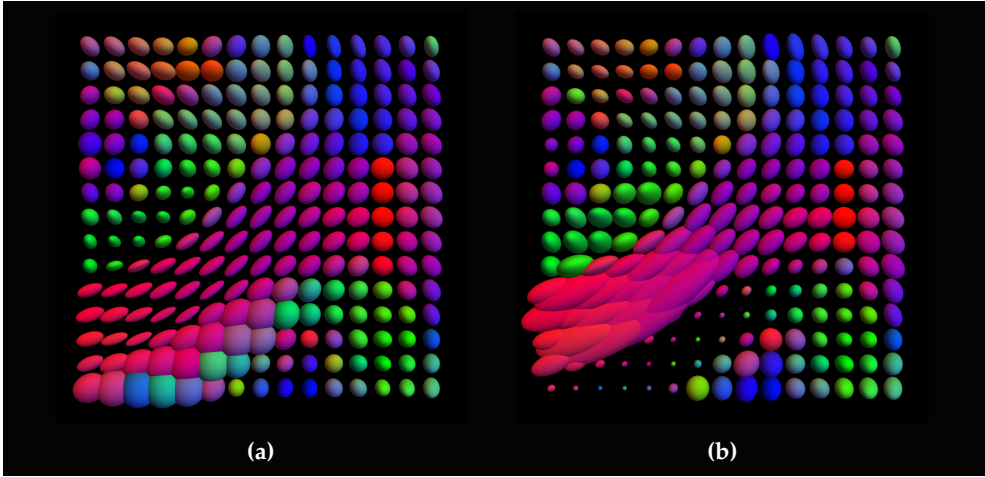
Figure 2.7: (a) DTI ellipsoids in a coronal slice of a Human Connectome Project data set, specifically in the region indicated by the red box in Fig. 2.2(a). The radii of a glyph represent the root mean squared displacement of spins along that orientation, and the peaks of these glyphs—given by the eigenvector corresponding to the largest eigenvalue—are approximately aligned with the main orientation of neuronal tracts in the underlying tissue [30, 247]. The glyphs are colored according to the main eigenvector of the diffusion tensor, with red indicating alignment with the left-right axis (LR), green with the anterior-posterior (AP), and blue with the inferior-superior (IS). From the viewpoint provided by the geometrical framework of Eq. (2.15), the displayed ellipsoids are proportional to the indicatrices of the metric tensors, which represent the compression/dilation of the model space relative to Euclidean space. We could alternatively plot the figuratrices associated to the metric tensors—which would illustrate how displacements along an orientation which has a large associated diffusivity, are assigned a short Riemannian distance [254]—but these are less intuitive and provide no additional information. (b) The same data as shown in (a), but locally rescaled in an attempt to unify the large-scale behavior described by the Riemannian Laplace–Beltrami (Eq. (2.9)) and the anisotropic diffusion equation (Eq. (2.16)) [254]. Scaling consists of a division of D by its determinant (Fig. 2.4(b)), which greatly suppresses the magnitude of diffusion tensors in voxels with CSF, while the data in gray and white matter is affected approximately uniformly.

posed in Section 2.2.3, and shown in Fig. 2.7(b). As could be seen in Fig. 2.4(b), the determinant used in the definition of the scaled metric tensor is similar in gray and white matter regions, and the effect of scaling is small in most regions outside the corpus callosum, where glyphs—and thus space—are somewhat dilated. Apart from this, the largest effect can be observed in voxels with significant CSF, where the scaled metric represents a significant compression in the model space relative to white/gray matter.

2.5 Discussion

2.5.1 The Riemannian framework

In this chapter we introduced the first geometrical framework for diffusion MRI, based on the idea that the diffusion characteristics of a complex specimen can be reproduced

by a deformed 'model space' in which we observe a simple Brownian motion, recall Section 1.2.1. We specifically considered the case of a Riemannian geometry, which restricts the deformations of the model space to smoothly varying scalings, along three smoothly varying orthogonal axes. If the model space is compressed along a given orientation, the distance that particles can travel relative to the Euclidean geometry in a given time is decreased accordingly. As explained in Section 2.4, this framework allows us to interpret diffusion measurements as a measurement of the degree of deformation of the model space relative to Euclidean space, and by virtue of Eq. (2.13) a small set of diffusion-weighted measurements is sufficient to recover the complete local geometry if the medium's diffusion coefficient is known. If we were to actually be observing a Brownian motion in a Riemannian manifold, this relationship would be one-to-one: the observable behavior of Brownian motion particles is (naturally) completely governed by the metric tensor, viz. Eq. (2.10), while the metric tensor is in turn fully determined by the observable behavior of the Brownian motion particles. We will see in the remainder of the thesis that this is quite a remarkable property—one that is generally lost if we for example allow more general deformations of the model space.

At first glance, the restriction to Riemannian geometry may be regarded as quite arbitrary. What makes Riemannian geometry in particular so interesting? The answer to this question lies with the aforementioned one-to-one relationship between Brownian motion and the Riemannian metric tensor. By definition, Brownian motion follows as in Section 2.2.1 through a central limit theorem, which in our setting of guaranteed finite displacement variances implies that the resulting propagator will tend to a Gaussian distribution. The implication is that the Gaussian propagator found in Eq. (2.12) can describe *any* Brownian motion in a subset $V_p \subset M$ that satisfies Assumption 1, while every Gaussian propagator can be linked in a unique way to a Riemannian metric tensor. As a result, the Riemannian framework is very broadly applicable, contingent only on the conditions for the central limit theorem: a single diffusion compartment wherein the large velocity – large collision rate approximation described in Section 2.2.1 is appropriate.[1] In the case of multiple relevant diffusion compartments, the framework can be applied to each compartment individually.

Simpler geometries whose structure can be deduced from observations of Brownian motion—think for example of a geometry where local scaling is constrained to be isotropic—will necessarily lack the generality of Riemannian geometry, while more complicated geometries cannot possibly be derived from basic observations of particles that satisfy the conditions of the central limit theorem. This leaves only the question of whether there could be non-Riemannian geometries that can be similarly

---

[1]One caveat that we must mention is that the Brownian motion could be degenerate—particles may in theory be restricted to what is essentially a one or two-dimensional motion, which would result in one or two eigenvalues of *D* becoming zero. In this case the central limit theorem applies to the lower-dimensional stochastic process, and Eq. (2.12) has to be modified accordingly.

related to Brownian motion in a unique manner. While we are unaware of examples, this seems entirely plausible; the local, observable behavior would be indistinguishable from a Riemannian manifold, but behavior at the voxel scale could be different. Identifying such geometries could potentially lead to interesting methodological advances. In light of this we intend to investigate e.g. sub-Riemannian geometry, which has already found applications in diffusion MRI [93, 112, 113, 121, 282].

2.5.2 Relation to Gaussian diffusion

The Riemannian framework leads to a Gaussian model for the diffusion MRI signal, recall Eq. (2.13). Aside from the explicit separation of medium- and structure-specific characteristics achieved by the inclusion of the diffusion coefficient α, this model is functionally the same as the standard Gaussian diffusion model—diffusion tensor imaging, or DTI [24, 25]. Given a DTI data set and a known (or assumed) value for α, the metric tensor g follows straightforwardly through a matrix inversion, viz. Eq. (2.15). Because of the simplicity of this relation, and because DTI is so widely used, we can view the Riemannian framework as an additional means *specifically* to interpret and analyze DTI data.

There are also historical reasons to view the Riemannian framework in tandem with DTI. In the foundational works of O'Donnell et al. [254] and Lenglet et al. [218], Riemannian geometry was not invoked to model the diffusion MRI signal, but as a tool to compute neuronal fiber pathways from diffusion tensor data. In this context, O'Donnell et al. [254] simply postulated a duality relation between g and D, based on the expectation that a valid metric should associate short distances to displacements along orientations with a high diffusivity. The work of Lenglet et al. [218] is closer to ours, but still somewhat different; the authors start with a field of diffusion tensors, and then argue that a Riemannian manifold that produces those tensors must satisfy a duality relation between the metric and diffusion tensors. Note that at inception of both approaches we have DTI—there is no Riemannian interpretation of diffusion MRI data without DTI. The presented 'stand-alone' derivation of the Riemannian framework—where the assumption of a Riemannian manifold leads to Gaussian diffusion, instead of the other way around—is perhaps the most valuable paradigmatic contribution of this chapter.

The DTI model can be considered the de facto standard in most applications—the assumption of Gaussian diffusion has become so ubiquitous that more general models are typically grouped under the umbrella term 'non-Gaussian diffusion'. These models either include several non-exchanging compartments—leading to multi-exponential and multi-compartment models [4, 14, 15, 32, 119, 212, 246, 265, 325]—or they assume that the remaining conditions of the central limit theorem do not apply, e.g. that there is anomalous diffusion [40, 160, 161, 202, 257] or that we need to take additional moments of the EAP into account [49, 178, 222, 259, 366].

Anomalous diffusion describes scenarios where e.g. the velocity tends to infinity independent from the collision rate, and where the microscopic definition of the diffusion constant α is thus not well-defined in terms of the average velocity and average collision rate. Anomalous diffusion has not been proven to occur in neuronal tissue [45, 250], and we do not consider it in this manuscript. Multi-compartment models describe diffusion in different compartments with no or minimal exchange, which in a geometrical framework should be modeled separately as a rule. Additional moments of the EAP become relevant when neither the collision rate nor the velocity can be considered sufficiently large, and models that assume a more general form for the EAP have been found to be relevant in neuroimaging applications [178, 222, 223, 250]. We will find that a consideration of this last deviation from Gaussian behavior is actually essential for the second, more complex geometrical framework we propose in the next chapter.

2.5.3   Information in the Riemannian structure

The incorporation of geometry into diffusion MRI is first and foremost intended to help analyze the data obtained from diffusion MRI scans. The basic idea behind this is geometrization: if available, it is generally easier to obtain relevant information using existing, generic tools from geometry—which have been developed and applied over more than a century—than it is to develop a new set of tools specifically tailored to a particular problem. The Riemannian framework serves, in a sense, as a means to translate problems in the language of diffusion MRI to the more general language of Riemannian geometry. In this section we look at applications of Riemannian geometry at the local level of a single voxel, as well as at the inter-voxel level.

2.5.3.1   Local information

The defining equations of the Riemannian framework, Eqs. (2.12) and (2.15), relate the diffusion tensor (the covariance matrix of a Gaussian diffusion) to a Riemannian metric tensor. Because this relation is (essentially) one-to-one, the information that can be extracted from the metric tensor is fundamentally the same as that in the diffusion tensor. The Riemannian framework does allow us to interpret common DTI descriptors differently, viz. in terms of the metric structure of an associated Riemannian manifold. We discussed some examples in Section 2.4; the mean diffusivity and DTI tensor determinant, fractional anisotropy, and the geometric shape measures, cf. Figs. 2.4, 2.6, and 2.7. A large number of scalar measures has been proposed for specific applications [5, 27, 273, 350, 372] in addition to the measures discussed in this chapter, and they likewise have geometrical interpretations.

Scalar measures in the Riemannian framework can be expressed in terms of tensor invariants like eigenvalues, which together with the eigenvectors characterize the

metric tensor completely. While scalar invariants capture local scalar information, eigenvectors capture information about orientation. In biological experiments for example, most notably neuroimaging and the imaging of skeletal and cardiac muscles, the main eigenvector orientation coincides (approximately) with the dominant orientation of fibrous structures in the tissue [30, 247]—provided this orientation is well-defined. The precise relation between local pathway orientations and diffusion is unknown, and as of this writing there seem to be several adequate solutions to this problem [50, 96, 97, 102, 167, 252, 259, 261, 342, 367]. In Chapter 6 we approach this question from a geometrical perspective based on the generalized framework that we propose in the next chapter.

Note that local descriptors in particular are strongly dependent on the diffusion time. A longer diffusion time means particles can move greater distances and probe larger structures, while they become less sensitive to properties of smaller structures [250, 251].

2.5.3.2 Global information

In samples where diffusion probes large, coherent structures that span multiple voxels, there is a question of integration. Is it possible to glean information about larger objects from local diffusion measurements? In neuroimaging these large structures of interest are typically fiber pathways—bundles of neuronal tracts that connect different parts of the nervous system. In this setting, the primary question is whether and how two regions are connected. Anatomical questions like this are addressed with tractography, which is the field of study concerned with the reconstruction of these fiber pathways.

The most common tractography algorithms work by computing integral curves in vector fields derived from the diffusion data [26, 69, 99, 242]. For DTI for example one can consider the vector field formed by the main eigenvectors of the diffusion tensors, which should be approximately tangent to the dominant orientation of the underlying fiber pathways. The existence of integral curves (or streamlines) is guaranteed if a vector field is Lipschitz continuous, so given an initial point in M one can trace the likely physical connections for as long as successive vectors are sufficiently similar. This approach works fairly well in regions where tissue orientation is well-defined, but is likely to fail if the pathways are organized in crossings or other complex configurations [134, 135, 348]. Modern methods solve this by going to more complex models that are capable of resolving these complex configurations (we will focus more on this in the next chapter), and/or by considering more involved tracking algorithms, e.g. global [68, 287] or probabilistic tractography methods [33, 99, 180, 299, 310]. These more complicated tracking algorithms have as a common trait that they tend to use all the information in the reconstructed model or the raw diffusion

data (instead of reducing the data to a set of vectors) and this trait is shared by the geometry-based geodesic tractography.

The general assumption in geodesic tractography is that there is a relation between physical connections and geodesics in M. This has some significant advantages, e.g. geodesic tractography will also work in regions where the dominant orientation of the tissue is not clearly or uniquely defined. The other side of the coin is that because not every set of points is connected by a physical fiber pathway, the relation between geodesics and fiber pathways is surely not one-to-one—we assumed that M is compact (Section 1.1.3), which implies geodesic completeness: there exist geodesics between any two points in M. This means that one needs to have reliable boundary conditions—knowledge about which regions are connected—or an additional criterion that can be used to determine retroactively whether a geodesic is likely to represent a valid connection. More problematic, is that there are also simple examples where (current) geodesic tractography is not able to reproduce known connections. Whether this issue can be resolved is an ongoing question that was the subject of a number of recent works [139–142, 163], including the works by Fuster et al. where the scaled manifold discussed in Section 2.2.3 was first proposed. We will look at geodesic tractography and connectivity in Chapter 7.

Where local tensor invariants are the building blocks of all useful descriptors of sub-voxel characteristics, larger-than-voxel-scale properties in the Riemannian framework are expressed in terms of the invariant *differential* structure of the manifold. Geodesics for example can be computed by solving the Euler–Lagrange equations [187, Lemma 1.4.4]

$$\frac{d^2 C^i}{dt^2}(t) + \sum_{j,k} \Gamma^i_{jk}(C(t)) \frac{dC^j}{dt}(t) \frac{dC^k}{dt}(t) = 0, \tag{2.25}$$

where C^i is the i-th component of the geodesic C in local coordinates, and Γ^i_{jk} are the Christoffel symbols given by

$$\Gamma^i_{jk}(x) = \frac{1}{2} \sum_l g^{il}(x) \left(\frac{\partial g_{kl}}{\partial x^j}(x) + \frac{\partial g_{jl}}{\partial x^k}(x) - \frac{\partial g_{jk}}{\partial x^l}(x) \right). \tag{2.26}$$

Another example is the curvature of the Riemannian space, which can be characterized by e.g. the Ricci curvature tensor

$$R_{ij}(x) = \sum_{k,l} \left[\frac{\partial \Gamma^k_{ji}}{\partial x^k}(x) - \frac{\partial \Gamma^k_{ki}}{\partial x^j}(x) + \Gamma^k_{kl}(x)\Gamma^l_{ji}(x) - \Gamma^k_{jl}(x)\Gamma^l_{ki}(x) \right], \tag{2.27}$$

based on which a number of scalar measures have been proposed [19, 137]. This tensor also tells us something about how contiguous geodesics behave; the geodesic

deviation at $(\boldsymbol{x}, \boldsymbol{y}) \in TM$, defined as

$$\operatorname{dev}(\boldsymbol{x}, \boldsymbol{y}) = \sum_{i,j} R_{ij}(\boldsymbol{x}) y^i y^j, \tag{2.28}$$

is positive only when geodesics tangent to $\boldsymbol{y}$ locally diverge. Geodesic deviation thus characterizes infinitesimal dilations of space along a geodesic, and has been suggested as an additional force in a modified geodesic tractography algorithm proposed by Sepasian et al. [306].

2.5.4 The scaled manifold

The first contribution of this chapter is the abstracted derivation of the Riemannian framework, where we explicitly separate the diffusivity of the probing medium (the diffusion coefficient α) from the surrounding structure (represented by the metric tensor $\boldsymbol{g}$). In this framework diffusion is modeled using Gaussian distributions, resulting in a model that is functionally similar to DTI. However, DTI can alternatively be derived—as in many early works—from the anisotropic diffusion equation, Eq. (2.16), which prescribes different behavior for spins at the (unobserved) millimeter scale [29]. This motivated the introduction of an optional conformal factor in a modified definition of the Riemannian manifold, in which Brownian motion more closely resembles an anisotropic diffusion of particles in Euclidean space. The scaled Riemannian manifold is obtained in practice by dividing the diffusion tensor by its determinant before computing the metric in the normal way using Eq. (2.15).

We plotted the determinant in a representative slice of a Human Connectome Project data set (Section 2.3.3) in Fig. 2.4(b). The determinant is relatively large in regions with cerebrospinal fluid, like the ventricles and near the cortex, and much smaller in regions of gray and white matter. Between different regions of gray and white matter there is little variation, and from its definition it is obvious that the modified metric will not be qualitatively different from the usual definition in those regions. Practically speaking, we thus expect the largest differences introduced by this scaling to be visible near the boundaries between CSF and gray/white matter, and we will indeed see this in the geodesic tractography experiments presented in Chapter 7. In that chapter we will also discuss a geometrical interpretation of the scaled manifold in terms of how geodesics are computed.

Chapter 3

Finsler

The Finslerian Framework

Based in part on:

T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. "Finslerian diffusion and the Bloch-Torrey equation". In: *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*. Ed. by I. Hotz and T. Schultz. Cham: Springer International Publishing, 2015, pp. 21–35. ISBN: 978-3-319-15089-5 978-3-319-15090-1.

T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. *The Finslerian character of the diffusion MRI signal*. In preparation, preprint available upon request.

The data necessary to reconstruct the structurally informative diffusion tensor imaging (DTI) model can be acquired in a matter of minutes, and in part because of this it is a popular technique. Unfortunately its applicability as a tissue model is limited to voxels with predominantly co-aligned fibers [30, 348], recall Section 2.5.3.2. With the advent of more complex models, propelled by the abundant research showing the limitations of DTI in many regions with complex fiber architectures [181], focus has likewise shifted away from the Riemannian framework. In its wake a number of authors have investigated the possible application of Finsler geometry [21, 55, 125, 288], a generalization of the theory of Riemannian geometry on which the DTI framework of the previous chapter was founded. A number of promising results based on the application of Finsler geometry in diffusion MRI have already been published [17, 130, 133, 227, 228, 307], showcasing significant improvement in known problem areas over similar DTI-based experiments.

What lacks in the existing works on Finsler geometry is a concrete relation to the physics of diffusion [129], a relation that proved so valuable in the Riemannian case: it is this relation that tells us how to describe the manifold's metric structure in terms of the MRI signal, and it is this relation that reveals the meaning of geometry-derived structures like the geodesic tracts that we discuss in Chapter 7. Ultimately it is the absence of this fundamental relation that explains why none of the recent papers reference a canonical definition for the Finslerian metric structure, and why results based on Finsler geometry are often so difficult to interpret.

The primary purpose of this chapter is to show how a natural Finsler metric structure can be derived, both in theory and in practice, from high gradient strength diffusion MRI measurements. With this relation, we effectively establish a Finslerian geometrical framework similar to the Riemannian framework of the previous chapter, each applicable to a specific gradient strength regime. Insights from this new interpretation of high gradient strength data quickly lead to various interesting quantifiers of local tissue properties, which we review in Chapter 6.

Most of the major theoretical results of this chapter are described in Section 3.2, where we first show that a straightforward application of the diffusion equation corresponding to Brownian motion in a Finsler manifold [10, App. A] does not result in a signal model from which one can derive the Finsler metric [86]. This result is based on the isotropic transport process that was already featured prominently in the previous chapter. Next we show that a Finslerian structure *can* in principle be extracted from diffusion measurements by considering an adapted isotropic transport process with finite velocity/collision rate. Though the resulting model is obviously inspired by Finslerian Brownian motions, it can be used to analyze any process with well-defined cumulants.

Practical issues pertaining to the reconstruction of the model are covered in Section 3.3. The first illustrative results can be found in Section 3.4, and the chapter

is concluded with a discussion in Section 3.5 on the strengths and weaknesses of the framework, along with a glance ahead to possible future work. The chapter is preluded by a brief introduction to some of the relevant concepts and objects in Finsler geometry, Section 3.1. Preliminary results of this work have been presented at Dagstuhl and the ISMRM [86–89]

3.1 Background

3.1.1 Manifolds and norms

Finsler geometry is a generalization of Riemannian geometry, where the squared norm is no longer restricted to a quadratic form [65]. As a result, the local geometry in a Finsler manifold can describe more elastic deformations applied along arbitrary orientations—not only along orthogonal ones (Section 2.5.1). For a general introduction to the theory of Finsler geometry, see e.g. the work of Bao et al. [21].

Let $M := (M, F)$ be a three-dimensional (reversible) Finsler manifold, meaning that M is endowed with an additional geometric structure captured by the (continuous) Finsler function $F : TM \to [0, \infty)$ that satisfies the following conditions.

- **Regularity** F is C^∞ on the slit tangent bundle $\overset{\circ}{T}M = \{(x, y) \in TM \mid F(x, y) \neq 0\}$.

- **Homogeneity** For all $\beta > 0$ and all $(x, y) \in TM$, it holds that

$$F(x, \beta y) = \beta F(x, y). \tag{3.1}$$

- **Subadditivity** For all $(x, y), (x, z) \in TM$, F satisfies the triangle inequality

$$F(x, y + z) \leq F(x, y) + F(x, z). \tag{3.2}$$

- **Reversibility** For all $(x, y) \in TM$,

$$F(x, -y) = F(x, y). \tag{3.3}$$

The Finsler function induces a norm $F =: \|.\|_F : T_x M \to [0, \infty)$ on the tangent space at each (implicit) point $x \in M$. It follows that

$$F(x, y) = \|y\|_F = \sqrt{\sum_{i,j} g_{ij}(x, y) y^i y^j}, \tag{3.4}$$

where g_{ij} is the coordinate representation of the fundamental tensor[1] $\mathbf{g} = \{g_{ij}\}$ defined as

$$g_{ij}(\mathbf{x}, \mathbf{y}) := \frac{1}{2} \frac{\partial^2 F^2(\mathbf{x}, \mathbf{y})}{\partial y^i \partial y^j} \tag{3.5}$$

relative to the local coordinates (y^1, y^2, y^3) on $T_x M$. This norm provides the machinery to compute speeds and distances relative to the Finsler geometry, much like the standard Euclidean norm $\|.\| : T_x M \to [0, \infty)$ does for Euclidean geometry.

A Finsler manifold is Riemannian when the fundamental tensor is independent of the vector $\mathbf{y}$, in which case $\mathbf{g}(\mathbf{x}) := \mathbf{g}(\mathbf{x}, \mathbf{y})$ corresponds to the Riemannian metric tensor for all $\mathbf{y}$, recall Eq. (2.1). Informally, a Finsler manifold is locally Minkowskian if the Finsler function is instead independent of the position $\mathbf{x} \in M$. For locally Minkowskian manifolds M (containing some fiducial point $\mathbf{p}$) we may identify all tangent spaces in TM with $T_p M$, and thus make the identification

$$F(\mathbf{x}, \mathbf{y}) \equiv F(\mathbf{p}, \mathbf{y}) =: F(\mathbf{y}) \tag{3.6}$$

for all $\mathbf{x} \in M$.

3.1.2   The indicatrix

The indicatrix[2] is the Finslerian analog of a unit sphere in $T_x M$ defined as

$$I_x M = \{\mathbf{y} \in T_x M \mid F(\mathbf{x}, \mathbf{y}) = 1\}, \tag{3.7}$$

and can itself be viewed as a (two-dimensional) Riemannian manifold by restricting F to $I_x M$. As a Riemannian manifold the indicatrix carries a natural volume measure η_x induced by the restricted metric, and the associated probability measure μ is defined as

$$\mu_x := \frac{\eta_x}{\eta_x(I_x M)}. \tag{3.8}$$

Explicit expressions for μ_x will be given when used.

The indicatrix is of particular interest in this chapter, because homogeneity (Eq. (3.1)) ensures that the indicatrix and the Finsler function represent essentially the same information. By the properties listed in Section 3.1.1 the indicatrix is a centrally symmetric convex set, i.e., it has antipodal symmetry.

---

[1]The fundamental tensor can be considered the Finslerian analogue of the metric tensor in Riemannian geometry, and we use the same symbol $\mathbf{g}$ for both. Where necessary they can be distinguished by their arguments.

[2]Like with the fundamental tensor, we use the same symbol for the Riemannian and the Finslerian indicatrix.

3.1.3 Geodesics

A geodesic between two points $x_1, x_2 \in M$ is a curve $C : [0, T] \to M$ that locally minimizes the length integral

$$\mathcal{L}_F(C) := \int_0^T F\left(C(t), \frac{\mathrm{d}C}{\mathrm{d}t}(t)\right) \mathrm{d}t, \tag{3.9}$$

where $C(0) = x_1$ and $C(T) = x_2$. Uniqueness and existence criteria can be found in the references [21]. For a geodesic $C(t)$, $t > 0$, with initial conditions $C(0) = x$ and $\frac{\mathrm{d}}{\mathrm{d}t} C(0) = y \in T_x M$, the exponential map is defined as

$$\exp(x, t y) := \begin{cases} C(t) & \text{if } y \neq \mathbf{0} \\ x & \text{if } y = \mathbf{0} \end{cases}. \tag{3.10}$$

The exponential map is C^∞ in the parameter t, and for any suitably differentiable function $t \mapsto u(\exp(x, t y))$ it satisfies the partial differential equation

$$\frac{\partial u}{\partial t} = Z u, \tag{3.11}$$

where the geodesic spray Z is given (for constant speed geodesics) by [21, Section 3.8]

$$Z := Z(x, y) := \sum_i y^i \frac{\partial}{\partial x^i} - \sum_{i,j,k} \gamma^i_{jk}(x, y) y^j y^k \frac{\partial}{\partial y^i}. \tag{3.12}$$

Here γ^i_{jk} are the formal Christoffel symbols of the second kind on $\overset{\circ}{T}M$, given by

$$\gamma^i_{jk}(x, y) = \frac{1}{2} \sum_l g^{il}(x, y) \left(\frac{\partial g_{kl}}{\partial x^j}(x, y) + \frac{\partial g_{jl}}{\partial x^k}(x, y) - \frac{\partial g_{jk}}{\partial x^l}(x, y) \right), \tag{3.13}$$

with g^{ij} the components of the inverse metric tensor as in Chapter 2. Note that on a locally Minkowskian manifold $\gamma^i_{jk}(x, y) = 0$, so geodesics are simply straight lines defined by the initial conditions (x, y), giving an isomorphism $T_p M \to M$ for any $p \in M$.

3.1.4 Duality

The set of all linear functionals acting on $T_x M$ forms the dual vector space $T^*_x M$, which is equipped with the natural norm

$$F^*(x, w) := \sup_{y \in I_x M} \langle w, y \rangle =: H(x, w), \tag{3.14}$$

where $w \in T_x^* M$. The duality relation between F and H is invertible, with $H^* = F^{**} = F$, and—like in the Riemannian setting—the unit level set of this norm is typically referred to as the figuratrix. If the fiducial point x is clear from context, we may omit it when writing H (or quantities derived from H).

The dual norm H has the same properties as F listed in Section 3.1.1, and in particular has an associated dual fundamental tensor

$$h^{ij}(x, w) := \frac{1}{2} \frac{\partial^2 H^2(x, w)}{\partial w_i \partial w_j}. \tag{3.15}$$

This tensor induces natural coordinates $\omega := \omega(x, w)$ on $T_x M$ [21, Eq. 14.8.1] with

$$\omega^i(x, w) := \sum_j h^{ij}(x, w) \, w_j, \tag{3.16}$$

which have the property that

$$H(x, w) = H^*(\omega(x, w)) = F(\omega(x, w)). \tag{3.17}$$

3.2   Theory

In this section we introduce the Finslerian framework for diffusion MRI through the same progressive steps as taken in the previous chapter. In Section 3.2.1 we first discuss the isotropic scattering transport process in a Finsler manifold, i.e., a stochastic process of colliding particles that reduces to a Brownian motion under certain conditions. We illustrate this process using simulations in a locally Minkowskian manifold, and discuss its central limit and large displacement behaviors in separate subsections. In Section 3.2.2 we then explain how the large displacement behavior of isotropic diffusion processes in a Finsler manifold can be related to the geometric structure of the manifold in a unique way. In principle, this allows us to reconstruct a Finsler geometry from high gradient strength diffusion measurements, thus establishing a geometrical framework for diffusion MRI based on Finsler geometry.

3.2.1   Diffusion in a Finsler manifold

As explained in the previous chapter, it is possible to derive a geometrical framework for diffusion tensor imaging by modeling the spin position space as a Riemannian manifold. In a first attempt to generalize this idea to more complex diffusion models, we investigated isotropic diffusion in a Finsler manifold as a model for the ensemble dynamics of the measured spins [86]. Sections 3.2.1.1 and 3.2.1.2 cover the Finslerian isotropic transport process and its central limit behavior. In Section 3.2.1.3 we briefly

discuss the large displacement behavior of the isotropic transport process, which becomes relevant when the central limit approximation does not hold.

3.2.1.1 Isotropic transport

Mimicking approaches employed in Riemannian manifolds [275], recall Section 2.2.1, one can obtain the Finslerian isotropic diffusion equation as a limiting case of the isotropic transport process, a stochastic process defined globally on any (compact) reversible Finsler manifold. The Finslerian isotropic transport process describes particles that travel with constant Finslerian speed $F(y) = v$ along geodesics until a collision event, after which movement continues in a new direction at the same speed. This scattering is uniform relative to the measure μ_x, Eq. (3.8), and the time between successive collision events again follows an exponential distribution with a rate parameter λ.

We illustrate the bulk behavior of spins subject to this process through a simple simulation. Let the characteristic speed of an isotropic transport process be v, let $v^2/\lambda = 1$ be the constant that defines the collision rate λ, and denote by Δ the total diffusion time in the simulation. We will assume the number of dimensions $n = 2$. For a particular evolution with initial position x_0 and initial velocity y_0 we draw k collision times τ_i from an exponential distribution, where k is the smallest integer that gives $\sum_{i=1}^{k} \tau_i \geq \Delta$. Given this collection of $\{\tau_i\}$ we define $\tau_0 := 0$ and truncate τ_k so as to obtain a sequence of times $\tau_0, \ldots, \tau_k$ that totals to Δ. Both the Finsler function F and the direction of the initial velocity vector are free, and we assume that the Finsler function is independent of position, recall Section 3.1.1 and Assumption 1. The random evolution $C(t)$ of a single test particle is then piecewise geodesic, meaning for these simple simulations that for each interval $\tau_i \leq t < \tau_{i+1}$, $i = 0, \ldots, k-1$ the curve $C(t)$ forms a straight line. Finally at each τ_i, $i = 1, \ldots, k-1$, a new unit velocity vector y_i is selected subject to

$$P\{y_i \in dy_{i-1}\} = \mu_{x_{i-1}}(dy_{i-1}) = \sqrt{\det g(y)}\,(y_{i-1}^1 dy_{i-1}^2 - y_{i-1}^2 dy_{i-1}^1)\,, \qquad (3.18)$$

such that $\lim_{t\downarrow\tau_i}(dC(t)/dt) = y_i$. Practical sampling is done numerically [106, Section II.2].

As a basic example, Fig. 3.1 shows convergence of the isotropic transport process to an isotropic diffusion by plotting approximate density functions obtained through a simulation with $F(y) = \|y\| = \sqrt{(y^1)^2 + (y^2)^2}$ and for increasing velocities v. The support of the propagator grows with velocity. As one might expect, simulations with more general F as shown in Fig. 3.2 appear to converge to an anisotropic (Gaussian) diffusion.

(a) $v = 0.75$

(b) $v = 1.00$

(c) $v = 1.25$

(d) $v = 1.50$
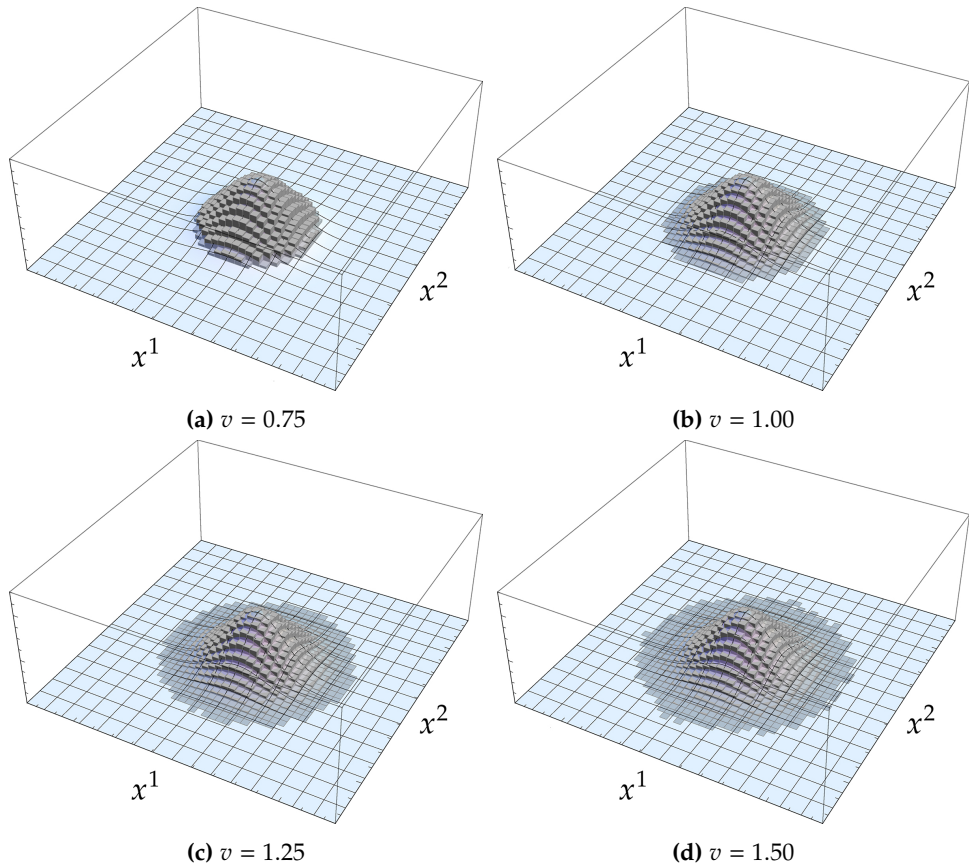
Figure 3.1: Approximate histograms (gray bars) corresponding to the isotropic transport process simulated with various values of v and with $\alpha = 1$ (recall Chapter 2). Simulations consist of 10^7 collisions, with a diffusion time $t = 7$. The vertical axis ranges from 0 to 0.04, while the planar axes range from -15 to 15. Simulation takes place in a Euclidean space, i.e., a locally Minkowskian manifold whose Finsler function is given by $F(\boldsymbol{y}) = \sqrt{(y^1)^2 + (y^2)^2}$. The corresponding asymptotic distribution function is computed by evaluating Eq. (3.19), and plotted in blue.
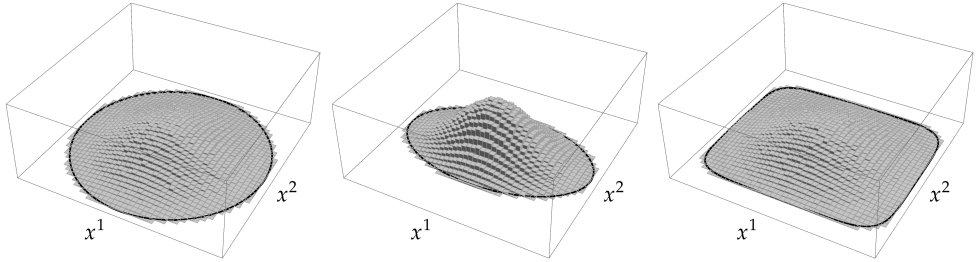
(a) $F(\boldsymbol{y}) = 3\left((y^1)^2 + (y^2)^2\right)^{\frac{1}{2}}$ **(b)** $F(\boldsymbol{y}) = \left((y^1)^2 + 3(y^2)^2\right)^{\frac{1}{2}}$ **(c)** $F(\boldsymbol{y}) = \left((y^1)^6 + 3(y^2)^6\right)^{\frac{1}{6}}$

Figure 3.2: Approximate histograms (gray bars) corresponding to the isotropic transport process simulated with various Finsler functions F, and with $v = \alpha = 1$. Simulations consist of 10^7 collisions, and take place on a locally Minkowskian manifold. The theoretical boundaries, indicating the farthest distance particles can travel, are plotted in black.

3.2.1.2   Central limit behavior

As suggested by the simulations in Fig. 3.2, and as was the case in Riemannian manifolds, the Finslerian isotropic transport process leads to a simple anisotropic diffusion equation in the Brownian motion limit. From the work of Antonelli and Zastawniak [10, App. A] it follows that in this limit, where both v and λ go to infinity subject to $0 < \frac{3}{2}\frac{v^2}{\lambda} =: \alpha$ constant, we obtain a general Laplace–Beltrami operator Δ_F given by an indicatrix average of $Z^2 = Z \circ Z$ at some anchor point $\boldsymbol{x} \in M$:

$$\Delta_F := 3\,\alpha \int_{I_x M} Z^2 \, \mu_{\boldsymbol{x}}(\mathrm{d}\boldsymbol{y}). \tag{3.19}$$

This operator generates a diffusion on M according to the diffusion equation

$$\frac{\partial}{\partial t} = \Delta_F. \tag{3.20}$$

We can solve this equation to obtain a displacement probability density function for the asymptotic (central limit) behavior of the particles [117], and note that this asymptotic probability distribution corresponds for Riemannian Finsler functions corresponds to the expression obtained in Chapter 2.

Remark 3. *For diffusions in Euclidean and Riemannian geometries, α is the diffusion coefficient, recall Sections 1.2.1 and 2.2.1. In principle α could thus be determined from measurements of free diffusion, and in contrast with the oft-used apparent diffusion coefficient it is an intrinsic property of the ambient medium independent of the manifold or tissue. As in Chapter 2, we assume throughout that α is independent of position.*

If we wish to apply the isotropic transport process to diffusion MRI modeling as in Section 2.2.1, we have to make use of the spatial homogeneity assumption postulated

in Section 1.1. This assumption states that the structure captured by F is independent of position within the fiducial voxel V_p, which means that V_p can be viewed as a locally Minkowskian manifold. The partial differential equation describing the diffusion in the voxel (Eq. (3.20)) reduces to

$$\frac{\partial}{\partial t} = 3\alpha \sum_{i,j} \int_{I_p M} y^i y^j \, \mu_p(\mathrm{d}y) \frac{\partial^2}{\partial x^i \partial x^j} =: \sum_{i,j} D^{ij}(p) \frac{\partial^2}{\partial x^i \partial x^j}, \qquad (3.21)$$

and we obtain what is essentially the premise of DTI, with a simple expression for the diffusion tensor of Eq. (2.14) in terms of F. As such we may conclude that Brownian motion in a Finsler manifold leads to a viable, but not exceptionally novel, model for diffusion MRI. This model coincides in form with diffusion tensor imaging, and in spirit with the Riemannian interpretation of DTI. Of course, the fact that the DTI model describes diffusion with the relatively simple symmetric 2-tensor D, immediately implies that in general most of the information about the underlying Finsler geometry is lost. While D can be described in terms of the Finsler structure, the relation is many-to-one and it is thus not feasible to e.g. obtain the metric by simply measuring D.

3.2.1.3   Large displacements

Let us now look closer at the isotropic transport process on V_p. As explained in the previous section, it is generally impossible to recover the Finsler function from diffusion that can be accurately modeled through a central limit approximation, as this approximation is necessarily Gaussian (see e.g. [66]). But at what stage in the limiting process $v, \lambda \to \infty$ does this loss of information occur? It is clear that the displacement of particles at $x \in V_p$ subject to the isotropic transport process traveling at a constant *finite* Finslerian speed v for a time Δ, is bounded by a scaled indicatrix

$$S_p M := \{ r \in T_x M \,|\, F(r) = \Delta v \} \subset T_p M, \qquad (3.22)$$

as a direct consequence of the triangle inequality [21, Section 1.2B], recall Eq. (3.6). This implies that the displacement probability density function P of the finite speed process is strictly nonzero inside this boundary and zero outside, i.e., that the boundary $\partial \Omega$ of the domain Ω of P is identical to the scaled indicatrix:

$$\partial \Omega = S_p M. \qquad (3.23)$$

Because of the relation between the indicatrix and the Finsler function (Section 3.1.2), it thus follows that we can in principle still recover F from the probability density function P whenever $v < \infty$. In other words, though we are (expectedly) unable to

infer the complete Finslerian geometry from observations of Gaussian diffusion, we may be able to recover it from observations of non-Gaussian diffusion.

3.2.2 The Finslerian framework for diffusion MRI

Independent of the assumptions specific to the isotropic transport process, asymptotic distributions obtained through central limit theorems are accurate for physically realistic systems only near the origin. This is a simple consequence of the particles' actual finite speed of movement. For high velocity particles near the origin, which are practically guaranteed to have incurred a large number of collisions, the asymptotic density function will be a reasonable approximation to the true density function. But particles that have travelled far from the origin tend to have suffered only a small number of collisions, with the most extreme case being particles that have travelled along a direct collision-free path from the origin. For these particles the Brownian motion limit is not appropriate, and ensemble behavior can deviate significantly from that predicted by the asymptotic distribution. This can already be seen in Fig. 3.2, where although the bulk behavior is well represented by a Gaussian distribution, the support of the displacement probability density functions is not necessarily ellipsoidal.

The fringe behavior of such large-displacement particles can typically be considered irrelevant, because the probabilities associated with these paths are small when the velocity and the collision rate are large. This is also observed in practice, where the diffusion tensor imaging model (which corresponds to the Brownian motion limit, recall Chapter 2) generally represents the data very well. A useful exception is the specific case of high gradient strength diffusion MRI, where spins that undergo large displacements have a relatively strong impact on the signal, viz. [374, I-4; Example 2].

In order to obtain an estimate of the scaled indicatrix S_pM and the Finsler function F, we investigate the boundary $\partial\Omega$ of the support Ω. It turns out that this boundary can be obtained from a large gradient strength asymptotic formula [374] for the cumulant generating function $q \mapsto \log S(-iq)$ as follows. Assume that Ω is a compact set, define $m(q) := \sup_{r\in\Omega} \exp(\langle q, r\rangle)$, and let $\Omega_\epsilon(q) := \{r \in \Omega \,|\, \exp(\langle q, r\rangle) \geq m(q) - \epsilon\}$ for $0 < \epsilon < m(q)$. Then

$$\lim_{\beta\to\infty} \frac{\log S(-i\beta q)}{\beta} \overset{(1.2)}{=} \lim_{\beta\to\infty} \frac{1}{\beta} \log \int_\Omega e^{\beta\langle q,r\rangle} P(r)\, dr$$

$$\geq \lim_{\beta\to\infty} \log\left[\left(\int_{\Omega_\epsilon(q)} (m(q)-\epsilon)^\beta P(r)\, dr\right)^{\frac{1}{\beta}}\right]$$

$$= \log\left[m(q) - \epsilon\right]. \tag{3.24}$$

Additionally we have that, independently of β,

$$\frac{\log S(-\mathrm{i}\beta q)}{\beta} \leq \log \left[\left(\int_\Omega m(q)^\beta P(r) \, \mathrm{d}r \right)^{\frac{1}{\beta}} \right] = \log m(q), \tag{3.25}$$

and so for any ϵ we get

$$\log\left[m(q) - \epsilon \right] \leq \lim_{\beta \to \infty} \frac{\log S(-\mathrm{i}\beta q)}{\beta} \leq \log m(q) = \sup_{r \in \Omega} \langle q, r \rangle. \tag{3.26}$$

We thus find the asymptotic relation

$$\log S(-\mathrm{i}q) \sim \sup_{r \in \partial\Omega} \langle q, r \rangle = \sup_{r \in S_p M} \langle q, r \rangle = v\,\Delta \sup_{r \in I_p M} \langle q, r \rangle \quad (\|q\| \to \infty), \tag{3.27}$$

wherein we recognize the definition of the dual Finsler function in the locally Minkowskian manifold V_p for which $\partial\Omega = S_p M \propto I_p M$:

$$\frac{1}{v\,\Delta} \log S(-\mathrm{i}q) \sim \sup_{r \in I_p M} \langle q, r \rangle = \sup_{F(r)=1} \langle q, r \rangle = F^*(q) \quad (\|q\| \to \infty). \tag{3.28}$$

$F^* = H$ is the algebraic dual introduced in Section 3.1.4, and by invertibility of the duality relation (Section 3.1.4) the Finsler function F is obtained up to a constant v (recall Section 2.2.1) from the large gradient strength signal using

$$F(r) \sim \left[\frac{1}{v\,\Delta} \log S(-\mathrm{i}\cdot) \right]^* (r) \quad (\|q\| \to \infty). \tag{3.29}$$

Note that the (finite speed) isotropic transport process in V_p produces an ensemble average propagator (EAP) with convex support. Sufficient conditions for arbitrary processes are similar to the conditions imposed by the central limit theorem of the previous chapter: the displacement probability distribution of a finite sequence of independent and identically (compactly) distributed random variables has convex support. To see this, let the Minkowski sum of two sets $S_1, S_2 \subset \mathbb{R}^3$ be defined as

$$S_1 + S_2 := \{ x_1 + x_2 \mid x_1 \in S_1, x_2 \in S_2 \}, \tag{3.30}$$

and let scalar multiplication for a set S be defined as $\beta S = \{ \beta x \mid x \in S \}$ for any $\beta \in \mathbb{R}$. Minkowski addition commutes with the taking of the convex hull, meaning $\operatorname{conv} S_1 + \operatorname{conv} S_2 = \operatorname{conv}(S_1 + S_2)$, where the convex hull $\operatorname{conv} S$ of a set S is defined as

$$\operatorname{conv} S := \left\{ \beta x_1 + (1 - \beta) x_2 \mid x_1, x_2 \in S, 0 \leq \beta \leq 1 \right\}. \tag{3.31}$$
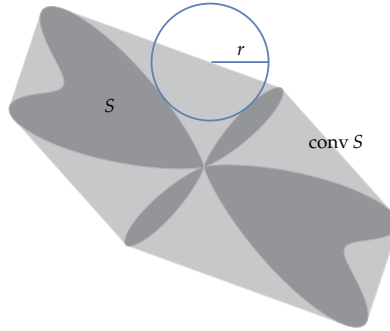
Figure 3.3: An illustration of the inner radius of a set S (displayed in dark gray). The inner radius $r(S)$ is the smallest number r for which circles with radius r centered anywhere within conv S (light gray), contain a subset of S.

The inner radius $r(S)$ of a set S is defined as the smallest radius r for which circles with radius r centered within conv S contain a subset of S, see Fig. 3.3. We then have the following theorem [322].

Theorem 1 (Shapley–Folkman–Starr Theorem). *Let $S_1, \ldots, S_m$ be compact sets in $\mathbb{R}^n$, and let $x \in \text{conv}(S_1 + \cdots + S_m)$. Then there are points $x_i \in S_i$, $i = 1, \ldots, m$, such that $\|x - \sum_{i=1}^m x_i\| \leq R$, where R is the sum of the $\min\{m, n\}$ largest $r(S_i)$.*

Assume now that the propagator of a diffusion process, for any given constant diffusion time smaller than or equal to some arbitrary but fixed Δ, is the same independent of position. Then consider a temporal partitioning $\{s_1, \ldots, s_n\}$, such that $s_i > 0$ and $s_1 + \cdots + s_n = \Delta$, and let S_i denote the (not necessarily convex) support of the diffusion propagator associated with the diffusion time s_i. Note that each of the points within S_i can be reached in the time s_i, and so by the homogeneity of the process all points in $S_i + S_j$ can be reached in time $s_i + s_j$. It then follows from Theorem 1 that for increasing n, the set $\sum_{i=1}^n S_i$ converges to conv S, showing that the support of the propagator with arbitrary diffusion time Δ is convex. Theorem 1 can similarly be employed to quantify the deviation from convexity if the diffusion propagator is known completely, and one of the few simple cases where this deviation can be significant is a spatially averaged propagator whose constituent diffusion processes cannot interact, e.g. as is the case in multi-compartment models [4].

Remark 4. *In voxels that contain more than one diffusion compartment, the support of the EAP may not be convex. In that case the Finsler function that can be computed using Eq. (3.29) will be representative of the convex hull of the actual support, and the apparent boundary and the Finsler function should be interpreted in terms of this upper bound on the largest displacements. A more extreme example in which convexity may not be guaranteed is anomalous diffusion [164], which is not considered here.*

3.3 Methods

In the previous section we derived a relation between high gradient strength diffusion MRI data and the metric structure of an associated Finsler manifold, Eq. (3.29). By virtue of this relation any high gradient strength data set can in principle be modeled using Finsler geometry, and in this section we will discuss some practical aspects of this approach. We address in particular the issue of infinite gradient strengths, and discuss how the techniques we develop based on the infinite gradient strength case are directly applicable in the case of finite gradient strengths.

In Sections 3.3.1 and 3.3.2 we look at a practical reconstruction of the dual Finsler function H based on Eq. (3.28), and on the reconstruction of F from H based on Eq. (3.17). Details about the implementation of the resulting reconstruction algorithms are given in Section 3.3.3. The proof-of-principle experiments presented in this chapter are based on the same Human Connectome Project (HCP) data used in Chapter 2, and we additionally consider the artificial data described in Section 3.3.4.

It now becomes convenient to introduce a global Cartesian coordinate system on $\mathbb{R}^3$, which induces local coordinates (q_1, q_2, q_3) on $T_p^* M$. We will additionally make use of spherical coordinates (r, θ, ϕ) on $T_p^* M$ defined by

$$
\begin{cases}
r(\boldsymbol{q}) := \sqrt{(q_1)^2 + (q_2)^2 + (q_3)^2} \in [0, \infty) \\[2mm]
\theta(\boldsymbol{q}) := \cos^{-1} \dfrac{q_3}{\sqrt{(q_1)^2 + (q_2)^2 + (q_3)^2}} \in [0, \pi] \\[2mm]
\phi(\boldsymbol{q}) := \tan^{-1} \dfrac{q_2}{q_1} \in [0, 2\pi)
\end{cases}
\tag{3.32}
$$

3.3.1 Reconstruction of the dual Finsler function

The first step in reconstructing the Finsler function F is the analytic continuation of the signal $S(\boldsymbol{q})$ to the cumulant generating function $S(-\mathrm{i}\boldsymbol{q})$. According to the Paley–Wiener theorem [81], $S(-\mathrm{i}\boldsymbol{q})$ follows once we have an analytical expression for $S(\boldsymbol{q})$, so we can consider for example the cumulant expansion of the EAP [131, 178, 222, 251] where

$$
\frac{1}{\Delta} \log S(\boldsymbol{q}) = \sum_{k=1}^{\infty} \frac{(-1)^k}{(2k)!} \sum_{i_1, \ldots, i_{2k}} D^{i_1 \cdots i_{2k}} q_{i_1} \cdots q_{i_{2k}}
\tag{3.33}
$$

and

$$
\frac{1}{\Delta} \log S(-\mathrm{i}\boldsymbol{q}) = \sum_{k=1}^{\infty} \frac{1}{(2k)!} \sum_{i_1, \ldots, i_{2k}} D^{i_1 \cdots i_{2k}} q_{i_1} \cdots q_{i_{2k}},
\tag{3.34}
$$

with $D^{i_1 \cdots i_{2k}}$ the coefficients of fully symmetric positive-definite $2k$-tensors. Many other approaches allow such a continuation as well.

Next we have to reconstruct the dual $F^* = H$, which follows from Eq. (3.28) in the large gradient strength limit $G \to \infty$. In practice the maximum attainable gradient strength G is of course limited by hardware constraints, so instead of directly reconstructing H we determine an approximation H_κ based on a level set R_κ of $\log S(-i\boldsymbol{q})$, defined in spherical coordinates by the equality

$$\log S(-iR_\kappa(\theta, \phi), \theta, \phi) = v\Delta\kappa =: \bar{\kappa} \tag{3.35}$$

for some $\kappa > 0$. As H is homogeneous of degree one, it makes sense to then define

$$H_\kappa(r, \theta, \phi) := \frac{r}{v\Delta} R_\kappa^{-1}(\theta, \phi) \log S\left(-iR_\kappa(\theta, \phi), \theta, \phi\right) = r\kappa R_\kappa^{-1}(\theta, \phi), \tag{3.36}$$

which can easily be shown to have all the properties associated with a (dual) Finsler function (Sections 3.1.1 and 3.1.4), and has the additional property that

$$\lim_{\kappa \to \infty} H_\kappa = H. \tag{3.37}$$

The subscript κ will henceforth be used to denote estimates of objects.

Remark 5. *The parameter κ, or its dimensionless counterpart $\bar{\kappa}$, provides a natural means to handle hardware gradient limitations. In practice one would choose κ to correspond to the highest gradient strength for which there is data of acceptable quality. $\bar{\kappa}$ is computed in our experiments by (arbitrarily) fixing $v = \frac{1\,\mathrm{mm}}{\Delta}$, as we fixed α in Chapter 2.*

Proceeding with the example of the cumulant expansion, R_κ can be obtained by solving for $\boldsymbol{q}$ in

$$\sum_{k=1}^{\infty} \frac{1}{(2k)!} \sum_{i_1, \ldots, i_{2k}} D^{i_1 \cdots i_{2k}} q_{i_1} \cdots q_{i_{2k}} = v\kappa, \tag{3.38}$$

which uniquely solves to a convex body because the left-hand side of this equation is by definition convex in $\boldsymbol{q}$. H_κ then follows from Eq. (3.36). The derivation for other models and expansions will be quite similar, and to keep the discussion general we will assume in the following that we have an expression for H_κ of the form

$$H_\kappa(r, \theta, \phi) = r \sum_{l=0}^{\infty} \sum_{m=-l}^{l} c_l^m(\kappa) Y_l^m(\theta, \phi), \tag{3.39}$$

with Y_l^m the spherical harmonics in the physics convention [12, Section 12.6] and c_l^m the expansion coefficients.

3.3.2 Reconstruction of the Finsler function

For Eq. (3.39) we can compute analytical expressions for the (estimated) fundamental dual tensor, defined in analogy with Eq. (3.15) as

$$h_\kappa^{ij}(q) := \frac{1}{2}\frac{\partial^2 H_\kappa^2(q)}{\partial q_i \partial q_j},\tag{3.40}$$

see Appendix B. The estimated dual fundamental tensor induces coordinates $\xi := \xi(q)$ on TM as in Eq. (3.16) with

$$\xi^i(q) := \sum_j h_\kappa^{ij}(q)\,q_j,\tag{3.41}$$

which have the property that

$$H_\kappa(q) = H_\kappa^*(\xi(q)) =: F_\kappa(\xi(q)).\tag{3.42}$$

Based on this relation we can evaluate the Finsler function estimate F_κ for a number of $\xi(q)$, which we use to estimate the coefficients d_l^m in the general expression

$$F_\kappa(r,\theta,\phi) = r\sum_{l=0}^{\infty}\sum_{m=-l}^{l} d_l^m(\kappa)Y_l^m(\theta,\phi)\tag{3.43}$$

where (r,θ,ϕ) are now the induced spherical coordinates on TM. As $F^* = H$, we naturally have that

$$\lim_{\kappa\to\infty} F_\kappa = F.\tag{3.44}$$

3.3.3 Implementation

The cumulant expansion, introduced as an example in Section 3.3.1, can be used to estimate the spherical harmonic coefficients d_l^m of the Finsler function (Eq. (3.43)) in a practical setting. Note that because we have a finite number of measurements, the summation in the cumulant expansion (Eq. (3.33)) has to be truncated. Generally we justify this with the assumption that the higher order coefficients, or their analogues in other models, are negligible in the neighborhood of R_κ. Similarly we assume that a finite number of coefficients are sufficient to describe H_κ and F_κ.

Reconstruction of the cumulant expansion coefficients $D^{i_1\cdots i_{2k}}$ (together with the baseline signal) is done using (weighted) linear least-squares, subject to the constraint that the left-hand side of Eq. (3.38) is convex. This convexity constraint is enforced by optimizing over the set of polynomials whose Hessian is sum-of-squares [67].

Specifically, given the model

$$f(q) := \sum_{k=0}^{K} \frac{(-1)^k}{(2k)!} \sum_{i_1,\dots,i_{2k}} D^{i_1 \cdots i_{2k}} q_{i_1} \cdots q_{i_{2k}} \qquad (3.45)$$

for a positive integer K, a sufficiently large set $\{q^i\}_{i=1}^N$, and the corresponding (linearized) measurements $y^i = \frac{1}{\Delta} \log S(q^i)$, we solve the following optimization problem:

$$\min_{D^{i_1 \cdots i_{2k}}} \sqrt{\sum_{i=1}^{N} (w^i)^2 \left(f(q^i) - y^i \right)^2}, \qquad (3.46)$$
$$\text{s.t. } G(f) \geq 0$$

with w^i the weights. Here $G(f) \geq 0$ denotes positive semi-definiteness of any Gram representation matrix [67] of the polynomial[1]

$$\sum_{i,j} \frac{\partial^2 f}{\partial q_i \partial q_j} (-iq) s_i s_j \qquad (3.47)$$

with $s \in T^*M$ an additional (dummy) variable.

Eq. (3.46) can be cast as a semi-definite programming problem [85, 224], which is solved using SDPA[2] [377]. Eq. (3.38) is then solved numerically for a set of 247 orientations[3] and a given value of κ, after which spherical harmonics are fitted with order $4K$ (real data) or $6K$ (artificial data). The weights are set to $w^i = 1$ in reconstructions using artificial data, and to $w^i = y^i$ in the case of real data [203]. This leaves us with an expression for H_κ as in Eq. (3.39). We take $K = 2$ for all results in this thesis, unless otherwise specified. At the moment we do not use iterative reweighting strategies [294, 354] or outlier rejection [337], though these are relatively straightforwardly incorporated.

When computing the dual of a Finsler function, we have to generate a dense sampling of dual orientations ξ. Because a uniform sampling of q can produce a highly non-uniform distribution of ξ, cf. Fig. 3.4, these ξ are generated dynamically. Starting from a uniform but sparse set of q, samples are added in the center of each polygon that is larger than a specified threshold in the dual space, Fig. 3.5. Once

---

[1]This constraint actually follows as a necessary condition from the positivity of P.
[2]SDPA is run with the adapted parameter settings $\epsilon^* = 1 \times 10^{-6}, \bar{\epsilon} = 1 \times 10^{-10}, \lambda^* = 1 \times 10^2, \omega^* = 2, \beta^* = 0.1, \bar{\beta} = 0.2$, and $\gamma^* = 0.8$. Some ad hoc preconditioning is applied to account for the large differences between the order of magnitudes of the different order D components. The algorithm typically converges to the optimal solution in few dozen iterations.
[3]The orientations are obtained by repeated tessellation of an icosahedron, and are roughly uniformly distributed on the hemisphere.

(a) *q* samples   **(b)** *ξ* samples
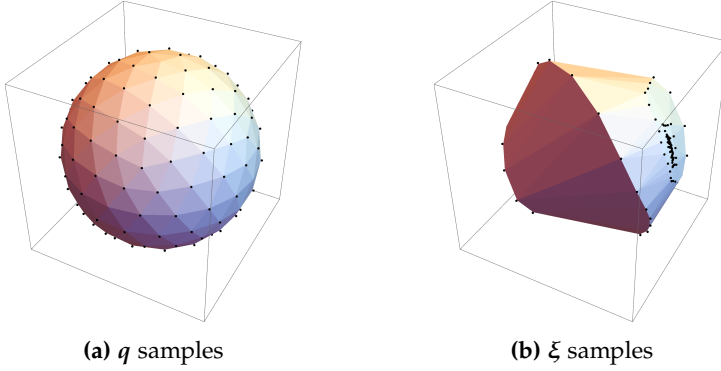
Figure 3.4: A uniformly distributed set of q, and the corresponding highly non-uniformly distributed ξ. This example is based on a single voxel in the corpus callosum of the HCP data set described in Section 2.3.3.

we obtain a sampling density similar to the icosahedron tessellation, we reconstruct the coefficients d_l^m (Eq. (3.43)) for F_κ from Eq. (3.42) (after division by $\|\xi\|$). These spherical harmonic reconstructions are Laplacian regularized [101] with $\lambda = 0.01$ to suppress ringing artifacts (see the results in Section 3.4.3). Relevant identities and explicit expressions can be found in Appendix B.

3.3.4   Data

In addition to the HCP data described in Section 2.3.3, both one- and three-dimensional artificial data are used to validate the implementation of the Finslerian framework and to illustrate the effect of changes in the maximum b-value and κ. The one-dimensional EAP is a family of distributions with bounded support defined by

$$P(\boldsymbol{r}) \equiv P(r) = \begin{cases} N \exp\left[\frac{c^4 + 2c^2\rho^2 - 3\rho^4}{2\rho^2(r^2 - c^2)}\right] & \text{if } -c < r < c \\ 0 & \text{otherwise} \end{cases}, \tag{3.48}$$

where r is the Cartesian coordinate, N is a (numerically computed) normalization factor, $c > 0$ is the boundary of the support of P, and $0 < \rho \leq c$ is a measure of the width of the distribution. This EAP is very similar to a Gaussian near the origin, but quickly tapers of to 0, so that the EAP is compactly supported. The EAP is intended to resemble a finite speed stochastic process such as discussed in Section 3.2.1.

We define three-dimensional EAPs based on the cumulant expansion given in Eq. (3.33), truncated to either order two or order four. For order two we define the tensor $A := 10^{-3} \operatorname{diag}(3, 1, 1)$ (in Cartesian coordinates), and in the fourth order case we take $\{D^{ij}\} = 10^{-3} \operatorname{diag}(2, 2, 1)$, while the non-zero elements of the fourth order tensor are given (up to permutation) by $D^{1111} = D^{2222} = 10^{-12}$ and $D^{1122} = -\frac{1}{2} \times 10^{-12}$. Both of these correspond to convex H.
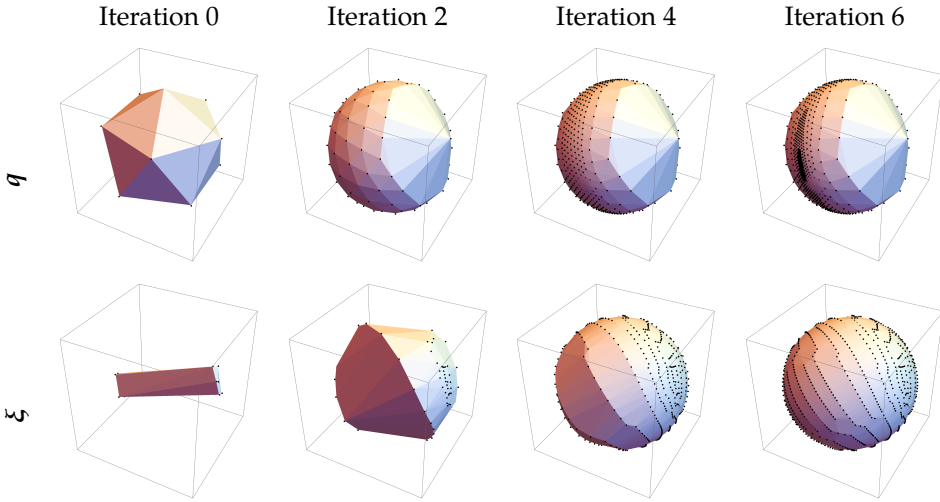
Figure 3.5: Dynamic generation of a uniform set of ξ, suitable for subsequent spherical harmonic reconstruction. The algorithm converged after 6 iterations.

Finally we take a three-dimensional EAP given by a multi-tensor model [4] defined by the average of two DTI models, A and $A(\vartheta)$, where $A(\vartheta)$ is simply a rotated version of A defined as

$$A(\vartheta) := 10^{-3} \begin{pmatrix} 2 + \cos 2\vartheta & -\sin 2\vartheta & 0 \\ -\sin 2\vartheta & 2 - \cos 2\vartheta & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{3.49}$$

The multi-tensor model assumes the presence of two separate compartments within a voxel, which means that the standard interpretation of the Finsler function does not apply. However, situations similar to this one can reasonably be expected to occur in real data, so it is important to check the behavior of our algorithms in this case as well.

The real data experiments are based on the $b = 1000\,\text{s/mm}^2$ and $b = 2000\,\text{s/mm}^2$ shells of the data set described in Section 2.3.3. We reconstruct the Finsler function and its dual from the second and fourth order cumulants of the EAP, using the steps described in Section 3.3.3.

3.4 Results

3.4.1 Finite gradient strengths

If we consider the physically realistic (one-dimensional) EAP, Eq. (3.48), and realistic gradient strengths, we expectedly find that $H_\kappa(q)$ underestimates $H(q)$ in the range of clinical MRI gradient strengths, cf. Fig. 3.6. This simple experiment suggests that
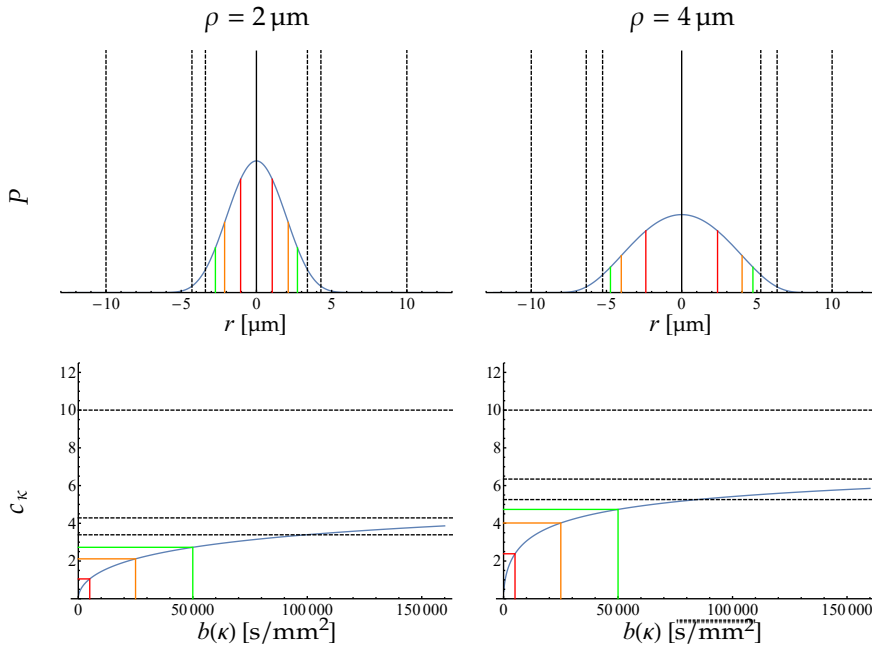
Figure 3.6: A simple one-dimensional diffusion model (Eq. (3.48) with $c = 10\,\mu\mathrm{m}$ and varying ρ). The EAPs are plotted in the top figures, and the bottom figures shows the estimated range $c_\kappa := H_\kappa(1) = \sup_{r \in \Omega} r$ measured with gradient strengths that correspond to given b-values (Eq. (1.3)). The black dashed lines represent (from outside inward) the true boundary, the boundary within which 99 percent of the particles remain, and the boundary within which 95 percent of particles remain. The red lines represent the estimates obtained with a gradient setting that is feasible with modern clinical scanners ($b = 5000\,\mathrm{s/mm}^2$), and the orange and green lines represent estimates with $b = 25000\,\mathrm{s/mm}^2$ and $b = 50000\,\mathrm{s/mm}^2$ respectively.

in order to get reliable estimates of H (and thus of the range of P), peak gradient strengths common in preclinical scanners are probably necessary. These results also show that estimates obtained with the low gradient strength data are more accurately interpreted as characteristic widths of the EAP.

3.4.2 Model reconstruction

3.4.2.1 Artificial data

Fig. 3.7 shows a very basic example of the (dual) Finsler function. In this figure we show a level set of F and H ($\bar{\kappa} = 1$) computed from the artificial DTI A data set, recall Section 3.3.4. The glyphs are colored in the standard red-green-blue (RGB) scheme, with the relative importance of each color given by the absolute values of the Cartesian orientation coordinates (cf. Section 2.3.2). As we are in the Riemannian setting, the shape of the level sets is ellipsoidal, independent of the choice for κ; the level set of F in this case is exactly the ellipsoid typically used in the visualization

(a) H_κ **(b)** F_κ
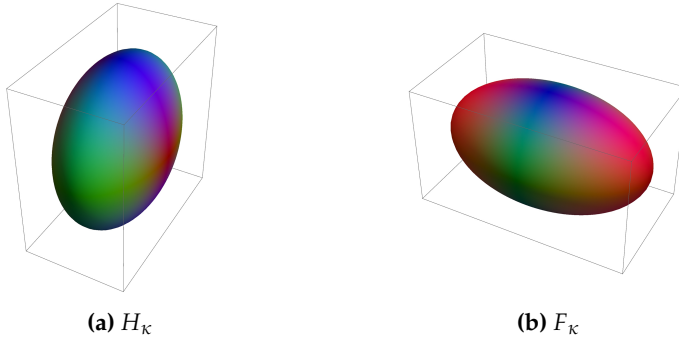
Figure 3.7: Level sets of H_κ (a) and F_κ (b) for $\bar{\kappa} = 1$, based on the artificial DTI data A described in Section 3.3.4. The level set of F corresponds to the ellipsoid typically used in DTI visualization.

of DTI (Section 2.3.2). The higher order tensor data is more interesting, cf. Fig. 3.8 (top row). We show once more a level set of F, but in this case the choice for κ has a small but noticeable influence on the shape. The level sets are no longer ellipsoidal, but they remain convex. The maximum b-values required to compute the level sets are listed in the figure, as estimated from the dual Finsler function through

$$b_{\max}(\kappa) := \max_{r(q)=1\,\mathrm{s/mm^2}} \Delta \left(\frac{\kappa}{H_\kappa(q)} \right)^2, \qquad (3.50)$$

which follows from solving Eq. (3.36) for R_κ and applying Eq. (1.3). If the maximum b-value is much larger than the largest b-value in an acquisition then the constrained reconstruction loses accuracy, which might make further analysis impossible. Similar level sets for the artificial multi-tensor data are plotted in Fig. 3.8 (bottom row), where convexity of the level sets is ensured by the constrained reconstruction (Section 3.3.3).

3.4.2.2 Real data

In Fig. 3.9 we show level sets of the estimated dual Finsler function H_κ and the estimated Finsler function F_κ, computed in part of the coronal slice of a HCP data set (Fig. 2.2(b)) with $\bar{\kappa} = 100$. By its relation to the Riemannian framework, the level sets of the Finsler function can be interpreted in roughly the same way is the standard ellipsoid visualization commonly used in DTI, while the level sets of H correspond approximately to an ellipsoid visualization of the inverse DTI tensor. For the relatively low κ that we could set for this data set, the level sets are, at least to the naked eye, still close to ellipsoidal. Taking even smaller κ brings us back to the Riemannian framework of Chapter 2, in which case we the level sets of F_κ and the DTI ellipsoids correspond exactly.

$\bar{\kappa} = 1$ $\bar{\kappa} = 50$ $\bar{\kappa} = 100$ $\bar{\kappa} = 150$

HT

$19\,\mathrm{s/mm^2}$ $935\,\mathrm{s/mm^2}$ $1865\,\mathrm{s/mm^2}$ $2790\,\mathrm{s/mm^2}$

MT

$18\,\mathrm{s/mm^2}$ $917\,\mathrm{s/mm^2}$ $1822\,\mathrm{s/mm^2}$ $2715\,\mathrm{s/mm^2}$
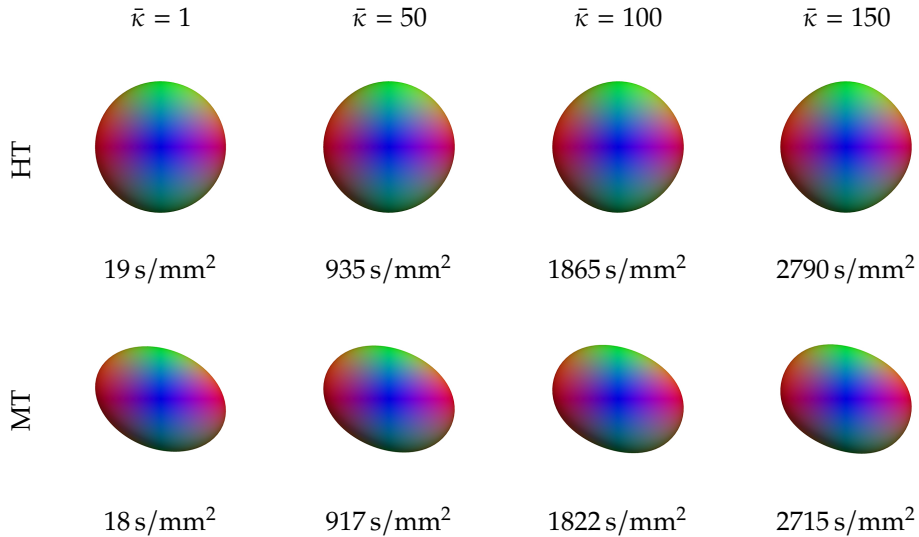
Figure 3.8: Level sets of F_κ for different values of κ, based on the artificial higher order data (HT, top row) and multi-tensor data (MT, bottom row, $\vartheta = \pi/4$ in $A(\vartheta)$) described in Section 3.3.4. The estimated maximum b-values that were necessary for the reconstruction are listed in the captions. The dominant diffusion directions become more apparent with increasing κ. Though the differences are visually minute, they have a considerable impact on e.g. the shape of the barrier orientation distribution function that will be introduced in Chapter 6.

(a) H_κ, $\bar{\kappa} = 100$.   **(b)** F_κ, $\bar{\kappa} = 100$.
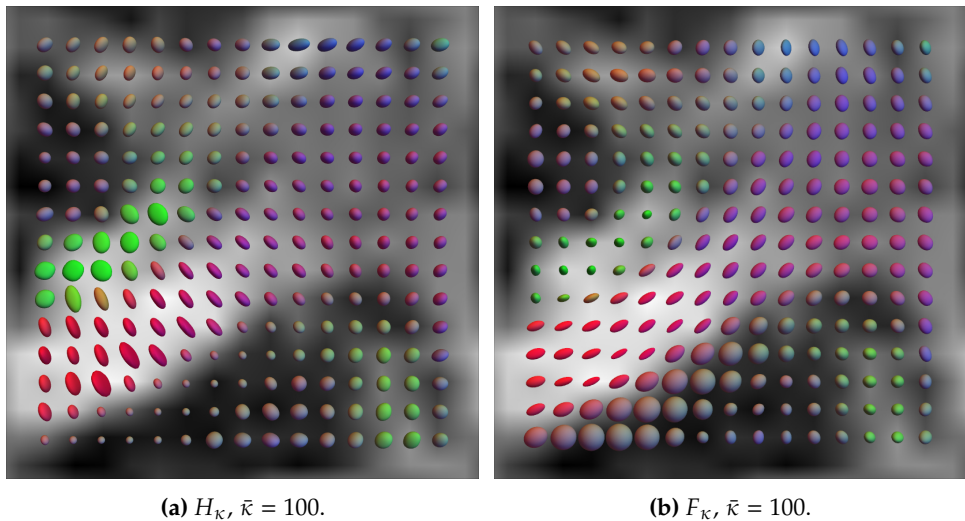
Figure 3.9: Part of a coronal slice of a HCP data set, showing scaled level sets of H_κ and F_κ. The colors are determined from the diagonal entries in the matrix representation of the second order diffusion tensor.
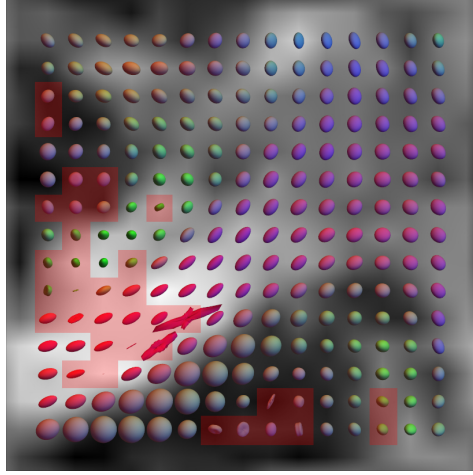
Figure 3.10: Part of a coronal slice of a HCP data set, showing level sets of F_κ obtained with unconstrained optimization corresponding to Fig. 3.9(b). Level sets with guaranteed convexity can be identified using semi-definite programming, producing the highlighted voxels that pinpoint positions that require constrained optimization.

3.4.3 Constrained and regularized reconstruction

Fig. 3.10 shows a level set of F_κ for $\bar{\kappa} = 100$ in a small region of the HCP data set, computed using unconstrained reconstruction. Though numerical problems can result in minor artifacts in the constrained reconstruction shown in the previous section, the artifacts in the unconstrained reconstruction are of a completely different order. Level sets with guaranteed convexity can be identified by solving a small semi-definite program, resulting in the overlay that highlights corrupted voxels. Overall the residuals do not change more than a few percentage points between the constrained and the unconstrained reconstructions.

Regularization suppresses oscillatory artifacts that are occasionally present in the spherical harmonic expansion of a Finsler function, see Fig. 3.11. We address the benefits of regularization in more detail Chapter 6, where we look at different ways to extract information from the Finsler function.

3.5 Discussion

3.5.1 The Finslerian vs. the Riemannian framework

In the context of diffusion magnetic resonance imaging, the global structure of a complex specimen like the brain can be modeled as a Riemannian manifold [218]; the metric structure of the manifold can be extracted directly from diffusion measurements,
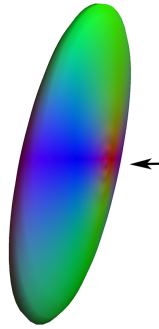
Figure 3.11: A level set of H_κ for a hand-picked voxel in the corpus callosum in the HCP data set. The oscillations indicated in the level set of the dual Finsler function produce high frequency noise in the barrier orientation distribution function that we will discuss in Chapter 6.

using the inverse relationship between the diffusion tensor and the Riemannian metric tensor given in Eq. (2.15). In this chapter we have shown how Finsler manifolds can function as a more complex model of global structure; the domain of definition of the ensemble average propagator (EAP), estimable from level sets of the corresponding cumulant generating function, uniquely relates diffusion measurements to a Finslerian metric per Eq. (3.29).

The Riemannian and Finslerian framework are primarily distinguished by the range of gradient strengths to which they are applicable. The Riemannian framework is generally constrained to low gradient strengths, and the Finslerian framework to high gradient strengths. In a practical setting the gradient strength is of course limited, but this is naturally solved by the introduction of the κ parameter, Eq. (3.35), that determines the specific level set of the cumulant generating function used in the reconstruction algorithm. A proper Finsler function can be computed for every value of $\kappa > 0$, where $\kappa \to \infty$ allows for an exact reconstruction. In the limit $\kappa \to 0$ the Finslerian framework reduces to the Riemannian framework, resulting in an equivalence between the Finsler function and the Riemannian norm. The maximum κ is limited by the accuracy with which level sets of the cumulant generating function can be determined, which is in turn limited by the availability and quality of data acquired at high gradient strengths. Note that the acquisition of high quality (high signal-to-noise ratio) data becomes more and more challenging at increasing gradient strengths, both because of technical (hardware) constraints, and because the signal itself tends to zero away from the origin.

In the experiments presented in Fig. 3.6 we applied the Finslerian framework to estimate the support of artificial one-dimensional displacement EAPs, which were non-zero only within the range $(-10\,\mu m, 10\,\mu m)$ and had varying widths. For every b-value (with $\Delta = 20\,ms$), we computed the maximum κ so that $b_{max}(\kappa) = b$. While estimates improve with the width of the distribution, in clinical data with b-values generally below $5000\,s/mm^2$ the estimated maximum displacements can be expected to be significant underestimations of the true maxima. Even with the large gradient strengths common in preclinical and molecular research we can only expect to measure a (very) high percentile that may still be far from the true boundary, as there are simply too few particles outside this region to induce a measurable signal. With $0 < \kappa < \infty$, the estimated shape of the boundary lies between the true shape ($\kappa \to \infty$) and the ellipsoidal approximation ($\kappa \to 0$, as in diffusion tensor imaging (DTI)), so choosing a smaller κ in a Finsler manifold can be seen as having a blurring effect on the boundary estimation (see e.g. Fig. 3.8). The estimated maximum displacements should thus be explicitly referred to as estimated or apparent maximum displacements, whose value should be understood to depend on the acquisition. In future work we will consider more dynamical choices for κ that e.g. produce level sets that correspond to a globally fixed average b-value, to ensure that there is no risk of a significant position-dependence in the blurring effect.

In principle the Finslerian and Riemannian approach capture different, but complementary information about the diffusion process. If however the assumption holds that the diffusion in a sample is truly Finslerian in nature, then the Finsler function is a complete representation of the process, similar to how a Riemannian metric fully characterizes a Gaussian process. In this case the Finslerian framework will encompass the Riemannian one, and the Riemannian metric tensor can then be recovered from the Finsler function through Eq. (3.21). A Riemannian metric tensor can also be computed when the diffusion process is not necessarily Finslerian, in which case its difference with a measured Riemannian metric tensor (from low b-value data, $\kappa \to 0$) gives an indication of the non-Finslerian character of the diffusion process. Significant non-Finslerianity can be due to severe inhomogeneities within the voxel (cf. Section 3.2.1.3 and Assumption 1), or because the displacement probabilities in the stochastic process depend on higher order derivatives of the particle paths. We formalize and compute this non-Finslerianity measure in Chapter 6.

If the measured stochastic process can indeed be described by a (finite speed) Finslerian isotropic transport process, then the Finsler function determines the full displacement probability density function. Because the Finsler function is in turn determined by a single one of its level sets, we should in principle be able to derive a single shell (fixed gradient strength) Finsler geometry model for diffusion MRI. In Chapter 6 we will present results that suggest that the Finslerian framework is

sufficient in large parts of the brain, which would make this an interesting avenue for future research.

The Finslerian approach we propose here is distinct from the work of Florack and Fuster [129], where the diffusion tensor D is postulated to depend on the orientation of the gradient wave vector q in the diffusion tensor imaging model Eq. (2.14). D can then be equated to the inverse of the fundamental tensor, which associates an alternative Finsler geometry to the diffusion MRI measurements. The downside of this approach is that the defining properties of the Finsler function, in particular homogeneity Eq. (3.1) and subadditivity Eq. (3.2), impose somewhat unnatural constraints on the modeled signal. These properties are naturally present in the Finslerian framework introduced in the current work.

3.5.2   Implementation of the Finslerian framework

The Finslerian framework can be applied to any compactly supported propagator with a well-defined cumulant generating function. It functions as an interpretation layer for higher order models like generalized diffusion tensor imaging [222], diffusional kurtosis imaging [178] (DKI), and functional basis expansions like 3D SHORE [259], just like the Riemannian framework does for diffusion tensor imaging [25]. Spherical deconvolution approaches [182, 302, 342, 344] and many other multi-compartment methods [4, 14, 15, 32] are notably excluded from this list as they specifically model voxels as spatially inhomogeneous, which complicates the relation between metric and large displacements, cf. Section 3.2.1.3. A comparison between the different models to which the Finslerian framework applies lies beyond the scope of this thesis.

If a model provides an analytical expression for the displacement probability density function, we can (in theory) obtain the cumulant generating function through analytical continuation [81]. For example, the preliminary experiments presented in this manuscript are based on a cumulant expansion of the propagator, which straightforwardly gives a polynomial expression for the cumulant generating function, see Eq. (3.38). Unfortunately the reconstruction of polynomial models like the cumulant expansion has to be constrained, which means that they generally require the solution of a computationally expensive semi-definite program for every voxel [85, 224]. Aside from optimizing the implementation, either by using different solvers [11, 39, 377] or different sets of parameters, we can save a significant amount of time by running the constrained reconstruction only on voxels where unconstrained reconstruction fails. To check whether an unconstrained reconstruction satisfies a convexity constraint (Fig. 3.10) also requires the solution of a semi-definite program, but because this semi-definite program is much smaller, and because we find that the percentage of voxels that requires constrained reconstruction is fairly low (< 10% in the considered Human Connectome Project (HCP) data with the described model), this extra effort

can pay off. In data sets or models where this percentage is much higher alternative models should be considered.

The local geometry in a Finsler manifold is captured by the Finsler function, whose dual is approximated from level sets of the cumulant generating function using Eq. (3.36). These level sets are expressed in terms of spherical harmonics, Eq. (3.39), which greatly simplifies visualization of glyphs, as well as a number of further analysis steps that will be discussed in Chapter 6. We are however unaware of a way to compute this expansion directly from the polynomial coefficients of the cumulant generating function, which means that we have to compute a large number of polynomial roots that can be used to fit the spherical harmonics. The roots of the low degree polynomials we used in the presented experiments have closed-form solutions, so the dual Finsler function can still be computed efficiently.

The Finsler function itself can be computed from the duality relation given in Eq. (3.42), which once more requires a fit to spherical harmonics. Preliminary results of the Finsler function computation are presented in Figs. 3.7, 3.8, and 3.9. Computing H and F from reconstructed cumulants can be done in a matter of minutes for an entire data set, whereas constrained reconstruction of the cumulants takes hours in comparison. Nevertheless, we expect that after optimizing the constrained reconstruction, also the dual computation could and should be improved. Room for improvement lies particularly in the adaptive sample generation discussed in Section 3.3.3, which is now by far the most time-consuming step in the dual computation algorithm.

Although the spherical harmonics-based computation of the Finsler function and its dual is accurate and relatively efficient, we lose the guaranteed convexity of the Finsler function that we had for the constrained polynomial expression. This can result in the minor artifacts that are (barely) noticeable in Fig. 3.9, but are more clearly shown in Fig. 3.11. Here we suppress these artifacts by using adaptive sample generation and regularized spherical harmonic reconstructions, though in future work this may be addressed by replacing the spherical harmonic expansion with a polyhedral approximation [145] or with splines [189]. Dual computations for these types of expressions are potentially simpler and more accurate, and might obviate the need for both regularization and adaptive sampling.

3.5.3   Prospects and limitations

On the practical side there are still a number of steps that can be improved. A faster and more stable reconstruction of the (dual) Finsler function, is at the top of that list. In particular with the computation of the dual there may also be some theoretical advances that can help to improve the implementation. Results may also improve with alternative acquisitions; at the moment most diffusion-weighted data is acquired on

densely populated shells, while a more uniform sampling of the measurement space is likely to be preferable for models like Eq. (3.45).

Forgetting for a moment the improvements that can still be made to the proposed framework, one is left with the question whether more general frameworks can be meaningfully defined. The Finslerian framework is suitable for the analysis of large gradient strength data, but it is fundamentally restricted to purely orientation-dependent properties. In terms of brain fiber configurations, this means that there is no way within this geometry to distinguish between identical fiber bundles in varying configurations, e.g. crossing, fanning, and kissing fibers [340]. In order to do this, we would have to be able to measure the bending of fibers, meaning aggregate particle dynamics would not only have to depend on position and orientation, but also on curvature. We will briefly return to this in Chapter 9.

A second limitation of the Finslerian framework is the postulate of a fixed diffusion time Δ. A lot of work has been done recently to harness the additional information that comes with varying Δ (e.g. [122, 251]), and it would be interesting to see what role geometry could play in this. The other fixed parameters α (or v) and κ also provide some flexibility. We assumed α constant, but this constraint could be relaxed e.g. by making α dependent on the local tissue type. Different values of α for white matter, gray matter, and cerebrospinal fluid could help in handling partial volume effects. Freedom in κ could be used to improve the angular resolution locally.

Finally, the general framework we introduced in this chapter does not trivially translate to a clinical setting. We have established a link between diffusion-weighted images and Finsler geometry, and the next step in the application of this framework is the evaluation of different tissue and diffusion properties that are captured by the local and global geometric structure of the manifold. Analogous to the Riemannian case, local properties can be defined using invariants of the local structure, e.g. of the fundamental tensor, while global properties can be expressed in terms of the differential structure of the manifold. In the last part of this thesis we will present a preliminary discussion on the potential of different such properties, though their evaluation in a more clinical context is yet to be done.

Chapter 3

Part II

Processing

Chapter 4

Enhancement

Propagator Deblurring

Based in part on:

T. Dela Haije, N. Sepasian, A. Fuster, and L. Florack. "Adaptive enhancement in diffusion MRI through propagator sharpening". In: *Computational Diffusion MRI*. Ed. by A. Fuster, A. Ghosh, E. Kaden, Y. Rathi, and M. Reisert. Cham: Springer International Publishing, 2016, pp. 131–143. ISBN: 978-3-319-28586-3 978-3-319-28588-7. Joint first authors Dela Haije and Sepasian.

Much of Chapters 2 and 3 focused on the relation between geometry and diffusion, which is but one of many approaches to handle the inference of structural information from diffusion-weighted MRI data [4, 14, 15, 27, 32, 100, 148, 182, 214, 259, 263, 302, 325, 342, 344, 382]. It stands to reason that this inference is generally imperfect in all but the simplest experiments. Taking once more the brain as an example, particles interact not only with axons—our structures of interest—but also with essentially randomly placed proteins, glial cells, and other small, confounding obstacles.

A simple way to potentially improve subsequent analysis steps is by noting that the mesoscopic effect of these microscopic random barriers is naturally isotropic, which implies that it can be modeled to a first approximation as a Gaussian blurring. We can thus attempt to remove such confounding effects through propagator *deblurring*, suggesting a general enhancement strategy that can be grouped with a number of existing approaches that have been proposed for specific models and in more specific practical settings.[1]

For example, many advanced tractography techniques allow tracts to deviate from the principal direction of diffusion. This can be advantageous in terms of robustness to noise, but if the diffusion profile is not sufficiently sharp the tracts can deviate significantly from the expected orientations. Issues like this inspired various pre-processing techniques, such as the log-Euclidean framework to handle noisy data [13], spherical deconvolution of the diffusion tensor to enhance the anisotropy [103, 130], sharpening by raising the diffusion tensors to a certain power [211], and morphological enhancements of orientation distribution functions [93, 111, 113].

The purpose of this chapter is the introduction of an enhancement method based on deblurring the ensemble average propagator. The theoretical basis of the method is described in Section 4.1, and explicit expressions for specific models are presented in Section 4.1.2. In Section 4.2 we describe some basic experiments on diffusion tensor imaging (DTI) data to illustrate the potential of the method, of which the results are presented in Section 4.3.

4.1   Theory

4.1.1   The enhanced diffusion signal

The ensemble average propagator (EAP) $P(r)$ represents the likelihood of a displacement r occurring in a voxel within an implicit diffusion time Δ, which is assumed to satisfy $P(-r) = P(r)$. Under the narrow-pulse approximation, i.e., when the duration δ of the applied[2] diffusion encoding gradients G is much smaller than Δ, P is related

---

[1]This in contrast to various *blurring* methods that have been proposed e.g. in the context of denoising, which are typically applied before model reconstruction [31, 46, 75, 111, 114, 121, 174, 276].
[2]We assume a standard pulsed gradient spin echo sequence.

to the normalized signal S through the Fourier transform $\mathcal{F}$ [49]:

$$S(\boldsymbol{q}) = \mathcal{F}^{-1}\{P\}(\boldsymbol{q}), \tag{4.1}$$

recall Section 1.1 and Eq. (1.2). Here $\boldsymbol{q} = \gamma\delta\boldsymbol{G}$ is the wave vector encoding information regarding the applied gradients, with γ the gyromagnetic ratio, cf. Eq. (1.1). In the typical situation where $\delta \approx \Delta$, a relation analogous to Eq. (4.1) holds [234, 251], recall the discussion in Section 1.1.3. One generally acquires a number of samples of S for various $\boldsymbol{q}$ in each voxel, though we will not state this voxel dependence explicitly until Section 4.3.

Interesting features of the propagator, like differences between angular diffusivities, can be difficult to detect. This difficulty is aggravated by e.g. sparse sampling of the Fourier space and interactions between spins and small, extraneous, and random obstacles, which have a blurring effect on P. To this end we consider what happens when the propagator P is blurred with a Gaussian $G_\sigma(\boldsymbol{r}) := \frac{1}{\left(\sqrt{2\pi}\sigma\right)^3}e^{-\frac{\|\boldsymbol{r}\|^2}{2\sigma^2}}$, $\sigma \geq 0$, and define Q as the deconvolution of P, i.e., $P = G_\sigma * Q$. Consequently,

$$S = \mathcal{F}^{-1}\{G_\sigma * Q\} = \mathcal{F}^{-1}\{G_\sigma\}\,\mathcal{F}^{-1}\{Q\}. \tag{4.2}$$

Using $\mathcal{F}^{-1}\{G_\sigma\}(\boldsymbol{q}) = e^{-\frac{1}{2}\sigma^2\|\boldsymbol{q}\|^2}$ we obtain

$$\rho_\sigma(S)(\boldsymbol{q}) := \mathcal{F}^{-1}\{Q\}(\boldsymbol{q}) = e^{\frac{1}{2}\sigma^2\|\boldsymbol{q}\|^2}S(\boldsymbol{q}), \tag{4.3}$$

where ρ_σ is the effective deblurring operator acting on the signal.

We note that σ is constrained by the fact that the signal S—and thus also the sharpened signal $\rho_\sigma(S)$—is required to be strictly less than 1 everywhere except in the origin.[1] This means that σ should be chosen such that

$$\forall_{\boldsymbol{q}\neq\boldsymbol{0}}\ e^{\frac{1}{2}\sigma^2\|\boldsymbol{q}\|^2}S(\boldsymbol{q}) < 1. \tag{4.4}$$

Eq. (4.3) can be used to process raw diffusion MRI data, but this should be avoided due to ill-posedness. Note for example that for large values of $\|\boldsymbol{q}\|$, the presence of noise means Eq. (4.4) will impose $\sigma^2 = 0$. In fact, it is desirable to apply the enhancement after S is expressed in terms of a specific model, the specific structure of which can be exploited to simplify Eq. (4.3). In the following section we present details for three commonly used models, including the DTI model used in Chapter 2 and the generalized DTI model (cumulant expansion) used in Chapter 3.

---

[1]This follows from Eq. (4.1) under some regularity conditions.

4.1.2 Simplified deconvolution for specific models

4.1.2.1 Diffusion tensor imaging

In diffusion tensor imaging [24, 25], already introduced in Eq. (2.14), the ensemble average propagator is assumed to be a multivariate normal distribution. The signal model can be written in matrix notation as

$$S(q) = e^{-\tau\, q^T D\, q},$$ (4.5)

with a constant τ defined as $\Delta - \frac{\delta}{3}$ and with D the positive-definite second order diffusion tensor. It follows that the signal corresponding to the sharpened propagator is given by

$$\rho_\sigma\,(S)\,(q) = e^{-\tau\, q^T D\, q + \frac{1}{2}\sigma^2\, q^T I\, q} = e^{-\tau\, q^T \left(D - \frac{\sigma^2}{2\tau} I\right) q}$$ (4.6)

with I the 3×3 identity matrix. In this case the enhancement consists therefore simply of subtracting a constant value (depending on the chosen σ) from the diagonal elements of the diffusion tensor D.

Note that the diffusion tensor is required to be positive-definite, which imposes a cap on the allowed values for σ:

$$\frac{\sigma^2}{2\tau} < \lambda_{\min}.$$ (4.7)

Here $\lambda_{\min}$ is the smallest eigenvalue, recall Section 2.3.1. This constraint corresponds exactly to Eq. (4.4).

4.1.2.2 Apparent diffusion coefficient models

For higher order models based on the apparent diffusion coefficient, the signal is described by [100, 129, 260]

$$S(q) = e^{-\tau\, \mathrm{ADC}(q)},$$ (4.8)

in which the apparent diffusion coefficient ADC is assumed to satisfy

$$\mathrm{ADC}(\beta\, q) = \beta^2 \mathrm{ADC}(q)$$ (4.9)

for $\beta \in \mathbb{R}$, i.e., the ADC is homogeneous of degree two. Applying Eq. (4.3) to this model gives

$$\rho_\sigma\,(S)\,(q) = e^{-\tau\, \mathrm{ADC}(q) + \frac{1}{2}\sigma^2 \|q\|^2},$$ (4.10)

a common 'ad hoc' enhancement principle already adopted in practice [262]. The ADC is typically expressed in terms of a fully symmetric higher even order (Cartesian) tensor or in terms of real-valued spherical harmonic functions, both of which transform in a straightforward manner under the proposed enhancement.

In the case of a tensor expansion the ADC is parameterized by fully symmetric coefficients $T^{i_1\cdots i_{2L}}$ of the tensor $\boldsymbol{T}$. L is called the order of the expansion. By defining $\hat{\boldsymbol{q}} := \boldsymbol{q}/\|\boldsymbol{q}\| \in S^2$ as the direction of the wave vector, we can write the tensor expansion of the ADC as

$$\text{ADC}(\boldsymbol{q}) = \|\boldsymbol{q}\|^2 \sum_{i_1,\ldots,i_{2L}=1}^{3} T^{i_1\cdots i_{2L}} \hat{q}_{i_1} \cdots \hat{q}_{i_{2L}}, \tag{4.11}$$

where $\hat{q}_i$ denotes the i-th component of $\hat{\boldsymbol{q}}$ in Cartesian coordinates. Plugging Eq. (4.11) into Eq. (4.10) gives

$$\rho_\sigma\,(S)\,(\boldsymbol{q}) = \mathrm{e}^{-\tau\,\|q\|^2\left[\Sigma_{i_1,\ldots,i_{2L}=1}^3\left(T^{i_1\cdots i_{2L}} - \frac{\sigma^2}{2\tau}I^{i_1\cdots i_{2L}}\right)\hat{q}_{i_1}\cdots\hat{q}_{i_{2L}}\right]}, \tag{4.12}$$

in which the tensor with components $I^{i_1\cdots i_{2L}}$ has the property that it produces one when contracted with any unit vector $\hat{\boldsymbol{q}}$, viz.

$$I^{i_1\cdots i_{2L}} = \delta^{(i_1 i_2} \cdots \delta^{i_{2L-1} i_{2L})}. \tag{4.13}$$

Here δ^{ij} is 1 for $i = j$ and 0 otherwise, and parentheses denote index symmetrization [233, p. 126]. For $L = 2$ for example, the only non-zero components (excluding symmetries) are given by [263]

$$I^{1111} = I^{2222} = I^{3333} = 3I^{1122} = 3I^{1133} = 3I^{2233} = 1. \tag{4.14}$$

The real-valued and symmetrical spherical harmonics $\tilde{Y}_l^m$ are defined [100] for $l = 0, 2, \ldots$ and $-l \le m \le l$ as in Appendix B. The ADC is then parameterized by the spherical harmonic coefficients c_l^m giving

$$\text{ADC}(\boldsymbol{q}) = \|\boldsymbol{q}\|^2 \sum_{l=0}^{L} \sum_{m=-2l}^{2l} c_{2l}^m \tilde{Y}_{2l}^m(\hat{\boldsymbol{q}}), \tag{4.15}$$

with $2L$ again the maximum order of the expansion. Since $\tilde{Y}_0^0(\hat{\boldsymbol{q}}) = (2\sqrt{\pi})^{-1}$ is the only constant basis function, enhancement in the spherical harmonic basis has the simple form

$$\rho_\sigma\,(S)\,(\boldsymbol{q}) = \mathrm{e}^{-\tau\,\|q\|^2\left[\left(c_0^0 - \frac{\sqrt{\pi}\sigma^2}{\tau}\right)\tilde{Y}_0^0(\hat{\boldsymbol{q}}) + \Sigma_{l=1}^L \Sigma_{m=-2l}^{2l} c_{2l}^m \tilde{Y}_{2l}^m(\hat{\boldsymbol{q}})\right]}, \tag{4.16}$$

and so enhancement boils down to subtracting the constant $\frac{\sqrt{\pi}\sigma^2}{\tau}$ from the coefficient c_0^0.

There are no known conditions on the coefficients (independent of the chosen basis) that impose the constraint in Eq. (4.4) exactly, so one would have to resort to numerical techniques when explicit constraints are needed. The constraint does

simplify significantly though:

$$\sigma^2 < \min_{\{q \,|\, \|q\|=1\}} \left[2\tau \text{ADC}(q) \right]. \tag{4.17}$$

Alternatively one could check if the coefficients satisfy a sufficient (but *not* necessary) condition for positive-definiteness [149].

4.1.2.3 Generalized diffusion tensor imaging

The last model we consider is another generalization of DTI proposed by Liu et al. [222], which is equivalent to the fixed-time cumulant expansion introduced in Eq. (3.33). In this case the signal is modeled as

$$S(q) = e^{-\sum_{i=2}^{L} \tau_i \sum_{j_1,\dots,j_i=1}^{3} D^{j_1 \cdots j_i} q_{j_1} \cdots q_{j_i}}, \tag{4.18}$$

with $\tau_i := \Delta - \frac{i-1}{i+1}\delta$ and $D^{j_1 \cdots j_i}$ the components of the generalized diffusion tensor, and now L the maximum order considered (in this case not forced to be even). In this case the analogy with DTI is immediately apparent, and we observe that the sharpening is achieved by subtracting the value $\frac{\sigma^2}{2\tau}$ from the coefficients D^{11}, D^{22}, and D^{33}. The same holds for specific models of this general form, like diffusional kurtosis imaging [178] (DKI).

Again there are no known conditions on the coefficients $D^{j_1 \cdots j_i}$ that enforce Eq. (4.4) generally, and unlike in the previous section Eq. (4.4) does not simplify. There do exist similar sufficient conditions that can be useful in this context [83, 150], cf. Section 3.3.3.

4.1.3 Adaptive enhancement

To apply the proposed enhancement the user would have to select a reasonable value for σ. We know from Eq. (4.4) that if we allow the parameter σ to surpass the given limit, then enhancement will produce incorrect results. In the case of DTI for example, a single globally optimal value for σ will—provided it exists—be based on a fraction of the globally smallest eigenvalue.

Let us define the theoretical bound on σ as

$$\epsilon := \min_{p \in M} \max_{\sigma \in \mathbb{R}} \left\{ \sigma \,\middle|\, \forall_{q \neq 0} \; \rho_\sigma(S)(p, q) < 1 \right\}, \tag{4.19}$$

where $p \in M \subset \mathbb{R}^3$ specifies a voxel in the region M, which will typically comprise all voxels in the brain, and where we include the p-dependence of the signal S. The

enhancement operator is then defined as

$$\rho_f^{\text{global}}(S)(\boldsymbol{p}, \boldsymbol{q}) := \rho_{f \cdot \epsilon}(S)(\boldsymbol{p}, \boldsymbol{q}), \tag{4.20}$$

where $f \in [0, 1)$ is a user-selected fraction of ϵ that determines the enhancement strength.

Though a single parameter value for an entire data set makes sense from a theoretical perspective, in practice it can be rather limiting. Therefore we propose to select a global fraction f of the *local* bound $\epsilon(\boldsymbol{p})$. This local adaptive enhancement can be defined as

$$\rho_f^{\text{local}}(S)(\boldsymbol{p}, \boldsymbol{q}) := \rho_{f \cdot \epsilon(\boldsymbol{p})}(S)(\boldsymbol{p}, \boldsymbol{q}). \tag{4.21}$$

In this case f is still a global parameter, but the bound ϵ is now a function of position:

$$\epsilon(\boldsymbol{p}) := \max_{\sigma \in \mathbb{R}} \left\{ \sigma \,\middle|\, \forall_{\boldsymbol{q} \neq \boldsymbol{0}} \; \rho_\sigma(S)(\boldsymbol{p}, \boldsymbol{q}) < 1 \right\}. \tag{4.22}$$

4.2 Methods

The first experiment we present in the next section illustrates the effect of the method on DTI scalar maps, where it provides a simple means to improve contrast. As a preliminary investigation into its use as a data pre-processing step, we additionally present results on DTI tractography. For this second experiment we select two deterministic tractography algorithms; a streamline method that traces the vector field defined by the main eigenvector of the diffusion tensor, and a multivalued geodesic tractography method.

Since the enhancement presented here does not change the main directions of diffusivity, its use in streamline tractography is limited to a modification of the boundary conditions (such as the stopping criterion). In geodesic tractography fiber pathways are reconstructed as local minimizers of the length functional in a given metric space, where the main assumption is that pathways resemble the path of the most efficient diffusion propagation [162], recall Section 2.5.3.2. Specifically, we use a recently introduced geodesic tractography technique that obtains these pathways as solutions of Euler-Lagrange equations in Riemannian or Finsler manifolds [305]. This approach can capture (multivalued) geodesics connecting two given points or regions by considering the geodesics as functions of position and orientation. Since in this case tractography relies on the entire diffusion profile instead of just the main directions of diffusion, enhancement has a more pronounced effect. Though there are alternatives [140, 141, 163, 299], we use the inverse of the diffusion tensor as the metric (cf. Chapter 2) as it is the most well-known definition.

Apart from the proposed enhancement we use a simple power transform as an alternative diffusion tensor imaging pre-processing step for a comparison [103, 211,

343]. The power sharpening depends on a positive parameter s, and is given by the matrix power of the diffusion tensor:

$$(\boldsymbol{D}, s) \mapsto \boldsymbol{D}^s. \tag{4.23}$$

We use two data sets to perform the experiments. The first data set (referred to as the Siemens data set) is acquired with a 3T Siemens scanner at a resolution of $1{\times}1{\times}1\,\text{mm}^3$ and a b-value of $1000\,\text{s/mm}^2$ (66 gradient directions). The second data set is provided (pre-processed) by the Human Connectome Project (HCP), and we work with a weighted linear least squares DTI reconstruction based on the $b = 1000\,\text{s/mm}^2$ shell.

The regions of interest are selected within the corpus callosum and derived from the expert-annotated Mori tract atlas [241]. Experiments were performed with Diffusica` and vIST/e (`bmia.bmt.tue.nl/software/viste`). In the geodesic tractography experiments we place two seed points in each voxel, and assign four random orientations within an elliptic cone around the main eigenvector to each seed point [305]. In all tracking experiments the algorithms terminate once the fractional anisotropy (FA, Eq. (2.21)) becomes less than 0.1.

Since there is no ground truth available to quantify the tractography, we look at the true positive percentile defined as the percentage of fibers which connect a seed region with given regions of interest selected using available white matter bundle atlases [59, 241].

4.3   Results

4.3.1   Adaptive enhancement of scalar maps

The effect of adaptive enhancement for a small region in the human brain is shown in Fig. 4.1, and in Fig. 4.2 we consider the effect of enhancement on the fractional anisotropy and mean diffusivity (MD, Eq. (2.19)) indices for diffusion tensor imaging [27], defined in Section 2.3.1. In both cases contrast improves markedly, but concomitant changes such as a strong decrease of the MD in areas of large and isotropic diffusivity may, depending on the application, be undesirable. Based on the results in Figs. 4.1 and 2.3.1, we can make the following observations about the effects of the local adaptive operator:
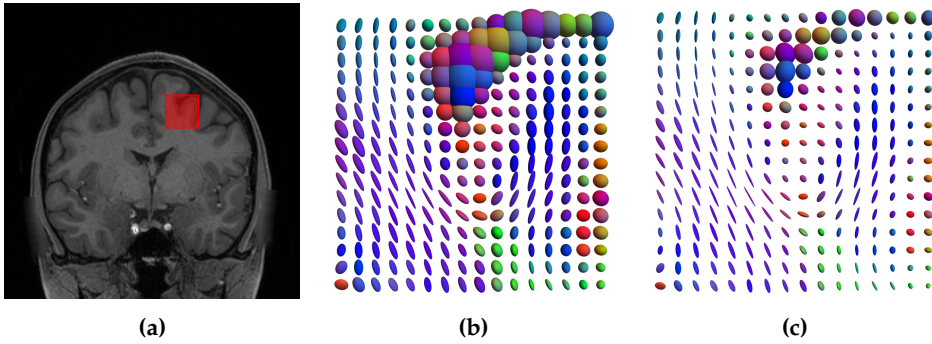
Figure 4.1: Illustration of adaptive enhancement for DTI in the cortex (Human Connectome Project data set, see Section 4.2). (a) T1 weighted data with the red box indicating the investigated region. (b) The tensors in the region indicated in (a). (c) The enhanced tensors of (b) after applying Eq. (4.21) with $f = 0.5$.

1. Voxels with large and relatively isotropic diffusion (e.g. in the ventricles) are affected the strongest; the amount of diffusion is decreased while the anisotropy is increased slightly.

2. In areas where the diffusion is anisotropic (white matter), the amount of diffusion decreases a little bit, while anisotropy increases far more significantly.

3. With small isotropic diffusion, both the (absolute) amount of diffusion and the anisotropy remain relatively untouched.

4.3.2 The effect of adaptive enhancement on DTI tractography

In Fig. 4.3 the seed region is placed in the postcentral gyri areas of the corpus callosum, Siemens data set (Section 4.2), and tractography is done using streamlining. The resulting tracts are known to correspond reasonably well to the anatomy, even in the case of DTI-based tractography. As such we would like the enhancement not to change the tracts too much, which we indeed see in the top row. Note also that the tracts are recovered consistently while varying f. Additionally we find that tracts continue a little bit farther into the gray matter compared to the original data due to the increased anisotropy. The second row shows that sharpening with a power transform produces a number of incoherent tracts even for low powers.

More interesting are the results of geodesic tractography, seen in Fig. 4.4, for the same data set and seed region. In this case we see that after enhancement, we obtain much denser, more coherent, and generally more cogent tracts. The number of spurious fibers increases as expected with increasing f, and the best results are seemingly obtained with $f \approx 0.25$. Of particular interest are the areas indicated with the orange arrows, where we recover tracts in the enhanced data (top row) that are
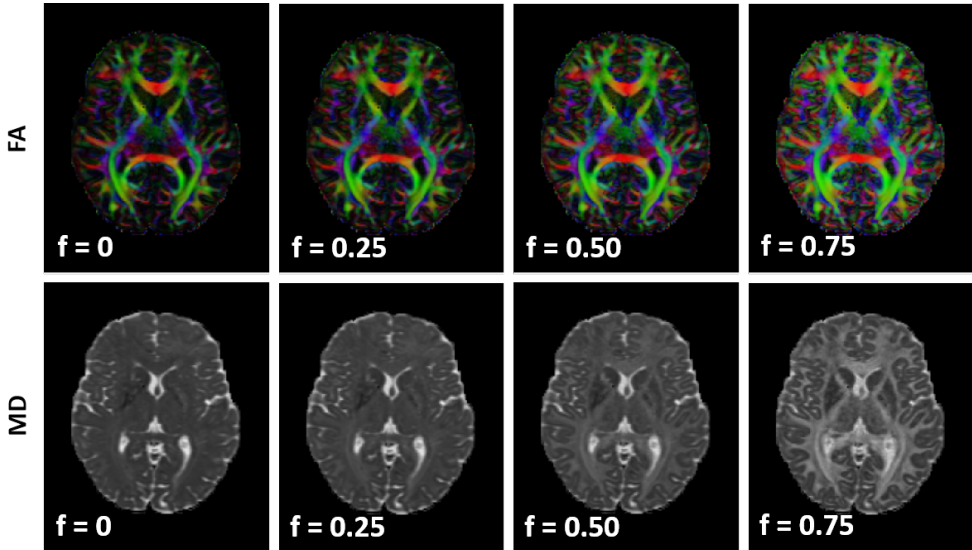
Figure 4.2: DTI-based FA and MD scalar maps of the Human Connectome Project data after adaptive enhancement, for increasing values of f, cf. Eq. (4.21).

missing in the tractography based on the original data. These fibers are expected to be present from known anatomy and from symmetry arguments. For reference we include the results obtained when applying a power transform (Eq. (4.23)) to the tensors, which introduces far more spurious fibers and does not produce any tracts in the indicated region. As a very simple quantification we show the true positive percentiles in Fig. 4.5, as explained in Section 4.2.

4.4 Discussion

In this chapter we proposed an adaptive enhancement technique for diffusion-weighted data, intended to address blurring effects caused by e.g. interactions between diffusing spins and randomly oriented obstacles that are of no practical interest. Because these effects can be modeled to a first approximation as a Gaussian blurring of the ensemble average propagator (EAP), we can try to suppress these effects by deblurring the propagator. As we have shown in Section 4.1.2, this leads to simple expressions of the deblurred signal for a variety of diffusion MRI models. We expect that for many models not considered here the resulting expressions will be equally uninvolved, e.g. for multi-compartment models in which the compartments are modeled by tensors. In other models (such as fiber orientation distribution models [344]) application of the enhancement might become too difficult to be useful in practice.
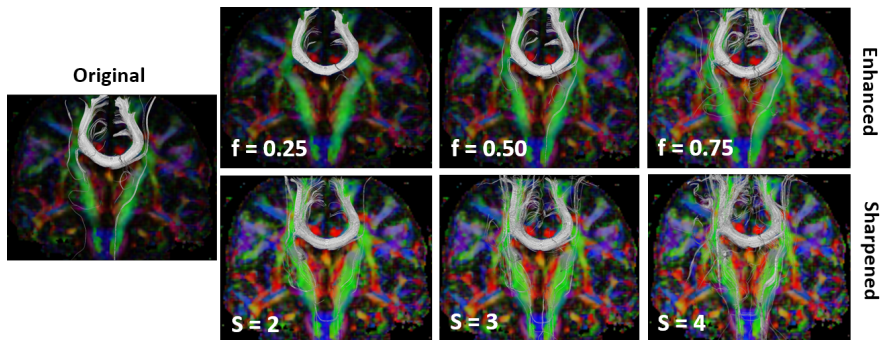
Figure 4.3: The results of streamline tractography seeded from the postcentral gyri areas of the corpus callosum (Siemens data), using no additional pre-processing (left), the presented deconvolution enhancement with $f = 0.25, 0.50, 0.75$ (top row, Eq. (4.21)), and sharpened diffusion tensors with powers 2, 3, and 4 (bottom row, Eq. (4.23)).
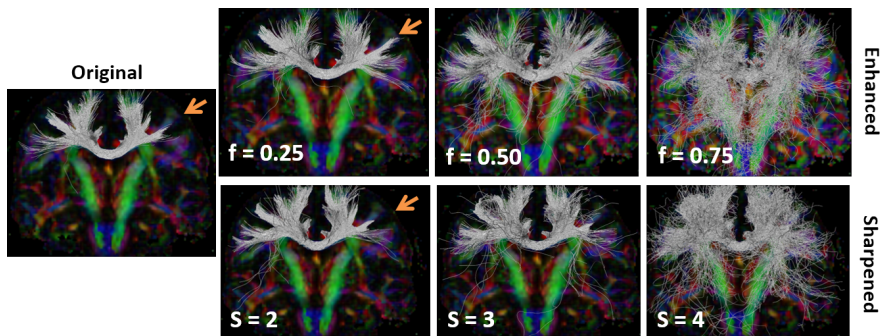
Figure 4.4: The results of geodesic tractography seeded from the postcentral gyri areas of the corpus callosum (Siemens data), using no additional pre-processing (left), the presented deconvolution enhancement with $f = 0.25, 0.50, 0.75$ (top row, Eq. (4.21)), and sharpened diffusion tensors with powers 2, 3, and 4 (bottom row, Eq. (4.23)). The orange arrows point to regions where symmetry of the tracts, lacking in the original and the power sharpened data, is partly recovered after enhancement.

(a) Streamline tractography   **(b)** Geodesic tractography
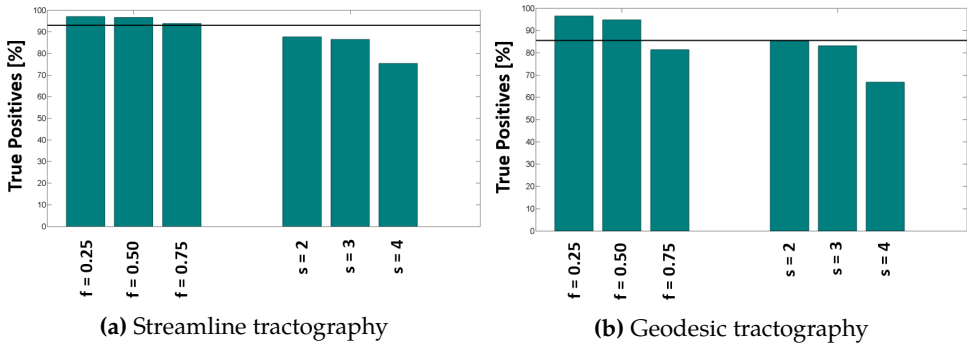
Figure 4.5: The true positive percentile of fibers connecting the given regions in postcentral gyri part of the corpus callosum in the Siemens data. The horizontal solid black lines indicate the true positive percentiles obtained when using the original data.

The enhancement method discussed here is similar to, but distinctly different from techniques like spherical deconvolution [103, 130]. A precise description of their relation is subject of future work. The proposed method also has obvious ties to the work of Canales-Rodríguez et al. [52], who looked at deconvolving a fully reconstructed propagator (based on diffusion spectrum imaging [50, 367]) with a sinc function.

Preliminary experiments illustrate that the method has some positive effects on the visualization of scalar measures and on tractography in the simple diffusion tensor imaging (DTI) case. Though the presented experiments are quite straightforward, combined with the fact that the method is fairly easy to apply even for more advanced models they do demonstrate potential. Generally speaking though, the method remains to be validated for each model in which it is to be applied, including DTI. In addition to this there are still a number of model-dependent open questions—like the proper range of the enhancement strength f—that would have to be considered.

We have also yet to investigate the optimal way to adaptively set the local enhancement strength. For now we take a fixed fraction of the maximum allowed value locally, but depending on the application there may be better choices for this. We should also note that the assumption of purely Gaussian diffusion for the isotropic contribution of random obstacles may be false if we move to higher b-values. As we discussed in Chapter 3, higher order moments of the EAP become relevant at high gradient strengths, which holds even for unrestricted/unhindered diffusion.

One final point we have not discussed in depth concerns the computation of the bound $\epsilon(p)$, Eq. (4.22). In the experiments presented here simple analytical formulae provided an unambiguous choice for this value, but for other models this value may have to be approximated numerically. This will complicate the implementation, and in some cases may even be a limiting factor.

Chapter 4

Chapter 5

Interpolation

Geometry-inspired Interpolation Techniques

Based in part on:

L. Florack, T. Dela Haije, and A. Fuster. "Direction-controlled DTI interpolation". In: *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*. Ed. by I. Hotz and T. Schultz. Cham: Springer International Publishing, 2015, pp. 149–162. ISBN: 978-3-319-15089-5 978-3-319-15090-1.

Interpolation is an important technique in signal and image processing, concerned with predicting data at positions where no measurements are available. The idea behind interpolation is that nearby measurements can be assumed to provide information about data at nearby points, motivated by the assumption that the data is continuous and its variations bounded.

In diffusion-weighted MRI, the complex interactions between spins and structure that underly the measurements have a non-trivial impact on the expected variations between data points, and the standard assumptions of interpolation have to be adapted accordingly. For example, we typically assume in diffusion MRI that smooth variations in the data are the *result* of smooth variations in the underlying structure, and this structure should thus be taken into account when developing interpolation methods. Not doing so can result in physically unrealistic interpolations, such as an artificially increased mean diffusivity or determinant—the so-called swelling effect described by e.g. Pennec et al. [270].

In this chapter we discuss interpolation as it applies to the Riemannian and Finslerian frameworks of Part I. By virtue of the relations exposed in Chapter 2, interpolation of Riemannian metric tensors is entirely equivalent to interpolation of diffusion tensor imaging (DTI) data, and we summarize some of the existing approaches in Section 5.1.1. In Section 5.1.2 we then use the fact that Riemannian geometry can be considered a constrained Finslerian geometry [65], recall Section 3.1.1, which allows us to embed the Riemannian structure in a Finsler manifold. Interpolation can then proceed using Finsler-specific schemes, and the resulting interpolants can be projected back onto the space of Riemannian metric tensors. There is an extra orientational degree of freedom in this projection, and interpolation schemes based on this embedding are hence dubbed direction-controlled schemes. In Sections 5.1.3 and 5.1.4 we next discuss how we could formally extend concepts from interpolation in the Riemannian case to the Finslerian one. The direction-controlled interpolation of ellipsoids is expanded on in Sections 5.2 and 5.3, where we discuss the implementation of this technique and present some preliminary results. The chapter is concluded in Section 5.4 with a brief discussion of the presented results.

Remark 6. *Because measurements are acquired in Euclidean space, we rely on the Euclidean vector space structure when discussing distances or when comparing data at different positions and along different orientations. The postulated geometry represented by the data is thus viewed as embedded in Euclidean space.*

5.1 Theory

5.1.1 Interpolation of ellipsoids

As stated in Chapter 1.2.2, interpolation can be implemented in a model-independent manner by using standard techniques on raw diffusion-weighted data [115]. This approach can be slow however, and does not make use of the structural information that is represented by the signal, which has led to the development of new interpolation methods that can be applied after model reconstruction. In the Riemannian framework this comes down to interpolation of metric tensor fields, which has been extensively investigated in diffusion tensor imaging literature. Note that a naive application of componentwise interpolation is not very suitable in this scenario, as those often result in physically unrealistic interpolated tensors [270]. The majority of proposed alternatives can be assigned to one of the following two categories.

The first category is that of (codomain) geodesic interpolators. Geodesic interpolators[1] are motivated by the observation that a proper interpolation scheme requires a distance on the codomain of the considered field—a (trivial) distance on $\mathbb{R}$ for scalar images, and a distance on the cone of positive semi-definite matrices in the case of a metric tensor field. One then considers geodesics relative to this distance, which provide smooth transitions between the objects defined at two (nearby) points. The Euclidean distances between the point of interpolation and the known data points are used to relate points in the domain to the interpolated objects in the codomain. The distance on the codomain is generally not uniquely defined, and a number of different distances have been proposed in the context of DTI [13, 23, 28, 128, 270].

The second category of solutions decomposes the tensor into a set of scalar invariants, recall Section 2.3.1, and a vector or set of vectors. The data is thus represented by a number of scalar images, which can be interpolated with traditional techniques, and one or more vector fields that have to be interpolated using more specialized techniques. The interpolated scalar images and vector fields are then recombined into an interpolated tensor field. Examples of this approach are the 'geodesic-loxodrome' method [198], 'dyadic-tensor based interpolation' [143], 'feature-based interpolation' [378], and the 'linear invariant tensor interpolation' scheme [144]. Note that these techniques are dependent on e.g. the set of invariants used, meaning they are typically designated for specific applications.

Remark 7. *We omit a few other useful and interesting approaches that do not fit well in either of these categories, such as the PDE-based tensor field interpolation of Weickert and Welk* [369].

---

[1]The term 'geodesic' in the name geodesic interpolator does not refer to the spatial geodesics discussed earlier in this thesis, but to geodesics in the codomain that are relevant to the interpolation.
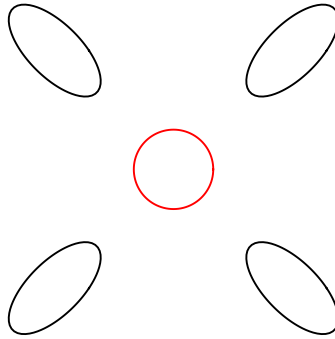
Figure 5.1: An example of destructive interference resulting from DTI interpolation in a 'crossing' region. The black ellipses represent the measured diffusion, and are assumed to accurately represent the underlying structure. The circular ellipse in red is obtained using log-Euclidean interpolation, and illustrates the resulting loss of orientation information in the center. Note that the specifics of the interpolation scheme are irrelevant; any proper (non-parameterized) scheme will have to respect the intrinsic symmetry in the depicted structure, and will therefore produce an isotropic interpolant.

5.1.2 Direction-controlled Riemannian interpolation

In Chapter 3 we discussed recent efforts to generalize DTI and the Riemannian framework, which have been primarily motivated by limitations of the diffusion/metric tensor in neuroimaging applications [181]. At a very basic level, these limitations arise due to restrictions of the quadratic form $g_{ij}(x)y^i y^j$ used to describe the local geometry, recall Section 2.1.1. The metric tensor $g(x)$ simply does not have enough degrees of freedom to accurately describe e.g. the structure of tissue containing crossing fiber pathways. In this section we argue that this is not only a problem at the modeling stage—where this problem can only be addressed by moving towards more complex models—but can also be problematic in interpolation. Even if the ellipsoids faithfully portray the underlying structure, interpolation schemes such as those described in the previous section can in some situations result in a quantifiable loss of information, as illustrated in Fig. 5.1.

The paradox of DTI interpolation without the destructive interference illustrated in Fig. 5.1 may be resolved if we (1) relax the quadratic assumption (and thus the restriction to six degrees of freedom) outside grid points, and (2) specify a projection back onto quadratic forms *a posteriori* given a context-dependent bias in the form of a preferred direction. We can then apply more general interpolation schemes that have the potential to preserve information that is otherwise lost, and make use of the specified preferred direction to reconstruct biased diffusion tensor fields that can be used for further processing. Geometrically speaking this solution entails an embedding of the data into a Finsler manifold (where we do not need to sacrifice any contextual orientation information), and an orientation-parameterized projection back to the space of Riemannian metric tensors.
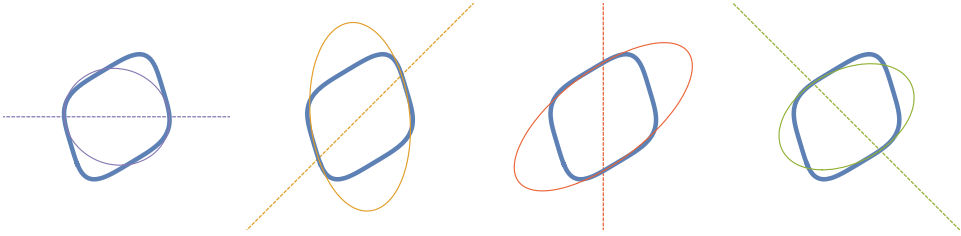
Figure 5.2: A schematic illustration of the relation between the Finsler function $F(x, y)$ and the fundamental tensor $g(x, y)$, cf. Eq. (3.5). The thick blue convex body is the indicatrix of the Finsler function $F(x, y) = \left(\left(y^1 + 0.3y^2 \right)^4 + \left(y^1 - 0.6y^2 \right)^4 \right)^{\frac{1}{4}}$, and the thin ellipses are level sets of the corresponding $g(x, y)$ for different y (indicated by the dashed lines). The ellipses describe the shape of the indicatrix at the intersections with the corresponding dashed lines.
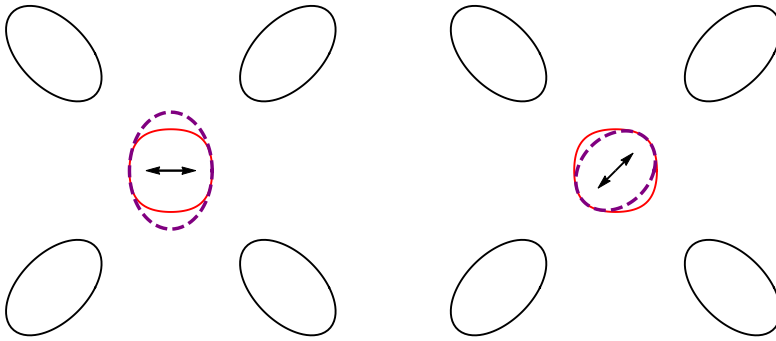
Figure 5.3: Direction-controlled interpolation of the data shown in Fig. 5.1. The newly proposed scheme produces the indicatrix shown in red, using equal weights in the computation of the geometric mean. The interpolated Finsler function can be projected onto the cone of positive semi-definite matrices, resulting in the dashed purple ellipses. The black arrows indicate the preferred orientation used in the fundamental tensor-based projection. The projected fundamental tensor by design produces ellipses that are biased towards ellipses aligned with the specified preferred orientation. The strongly anisotropic, vertically oriented ellipse obtained with the horizontal bias emphasizes that the preferred orientation lacks an expected vertical component.

As an example we consider the weighted geometric mean of quadratic forms, which is subadditive and homogeneous of degree two, making it a natural candidate for the interpolation of quadratic forms in a Finsler manifold. To obtain a physically meaningful quadratic form from the interpolated Finsler function we adopt the fundamental tensor defined in Eq. (3.5). The fundamental tensor represents the shape of the interpolated function's indicatrix near a specified orientation, see Fig. 5.2. The projected geometric mean thus balances the contribution of surrounding ellipsoids based on their Euclidean distance to the point of interpolation *and* on their radius along the specified orientation, which allows us to avoid potential ambiguities as shown in Fig. 5.3. We will show a few additional examples of this 'direction-controlled' interpolation of in Section 5.3.

5.1.3 Direction-controlled Finslerian interpolation

In the previous section we embedded the Riemannian metric tensor field in a Finsler manifold, which allowed us to use a basic Finslerian interpolation scheme that postpones the problematic restriction to quadratic forms. In this section we discuss the opposite—using the fundamental tensor to reduce data modeled by a Finsler function near a given position to a Riemannian structure, and then applying one of the Riemannian interpolation techniques described in Section 5.1.1. We can in principle do this for every orientation to effectively compute an interpolated fundamental tensor, from which we can compute an interpolated Finsler function using Eq. (3.4).

This ad hoc interpolation of projected structures necessarily disregards any correlations in the information associated with different orientations in the Finslerian structure, and thus makes very limited use of the advantages Finsler geometry has over Riemannian geometry. One implication is that the resulting interpolation scheme is simply equivalent to the selected Riemannian scheme for purely Riemannian Finsler functions, meaning for example that it cannot be used for direction-controlled interpolation in Riemannian manifolds as described in Section 5.1.2.

5.1.4 Interpolation of convex bodies

The weighted geometric mean interpolation suggested in Section 5.1.2 was used as a simple way to generate an interpolated Finsler function from a set of quadratic forms. The alternative approach in Section 5.1.3 is applicable to arbitrary Finsler functions, and perhaps more intuitive, but as noted there does not generally make use of the full scope of information present in the Finslerian structure. Ideally, we would like to formally extend the work done in the Riemannian framework to the Finslerian case. Because the Finsler function in principle has an infinite number of scalar invariants it is difficult to extend the decomposition approach—but it is theoretically possible to generalize the idea of geodesic interpolation to the Finslerian framework.

Geodesic interpolation of Finsler functions would involve the definition of a distance on the set of convex bodies, or alternatively, on the set of normed spaces. The most common example of such a distance is the Banach–Mazur distance, and the space of Finsler functions equipped with this metric defines a Banach–Mazur compactum [339]. If we define $GL(V_1, V_2)$ to be the set of all linear isomorphisms between an F_1-normed space V_1 and an F_2-normed space V_2, with F_1 and F_2 Finsler functions, and write $\|T\|_{\mathrm{op}}$ ($T \in GL(V_1, V_2)$) for the corresponding operator norm

$$\|T\|_{\mathrm{op}} = \inf_{\beta \geq 0} \left\{ \beta \,\middle|\, \forall_{v \in V_1} \, \|T(v)\|_{F_2} \leq \beta \|v\|_{F_2} \right\}, \tag{5.1}$$

then we can define the Banach–Mazur distance d by

$$d(F_1, F_2) = \log \left[\inf_{T \in GL(V_1, V_2)} \|T\|_{\mathrm{op}} \|T^{-1}\|_{\mathrm{op}} \right].\tag{5.2}$$

Practical implementations of algorithms based on this distance are not available to our knowledge [35, 339], and—although theoretically feasible—we do not further investigate geodesic interpolation of Finsler functions in this thesis.

5.2 Methods

In the remainder we will only look at direction-controlled interpolation in the Riemannian case—the development and evaluation of interpolation techniques for Finsler functions is a topic of ongoing work. The implementation of direction-controlled interpolation proceeds as follows. For $k = 1, 2, \ldots, K$, let $F_k^2(y) :=$ $g_{ij}(x_k) y^i y^j$ be a set of K spatially distributed quadratic forms, with g a metric tensor field derived from diffusion-weighted data acquired at nearby points $x_k \in M$ (e.g. as in Chapter 2). Given a set of weights $w_1(x), w_2(x), \ldots, w_K(x) > 0$ that satisfy $\sum_{k=1}^{K} w_k(x) = 1$, we define an interpolated Finsler function $F(x, y)$ as the weighted geometric mean given by

$$F(x, y) = \Pi_{k=1}^{K} [F_k(y)]^{w_k(x)}.\tag{5.3}$$

Explicit expressions for the corresponding fundamental tensor follow using Eq. (3.5), e.g. for $K = 2$ we find

$$g_{ij}(x, y) = F^2(x, y) \left[w_1(x)(g_2)_{ij} + w_2(x)(g_1)_{ij} - 2w_1(x)w_2(x)\xi_i\xi_j \right],\tag{5.4}$$

in which $(g_k)_{ij} := \frac{g_{ij}(x_k)}{F_k^2}$ and $\xi_i(y) := \sum_j \left[(g_1)_{ij} y^j - (g_2)_{ij} y^j \right]$. The relation between the weight $w_k(x)$ and the point of interpolation x can be defined using any suitable scheme, see e.g. Fig. 5.4 for the weighting scheme used in the presented (two-dimensional) experiments. For one-dimensional interpolation we take $w_1(x) = \frac{\|x_2 - x\|}{\|x_2 - x_1\|}$ and $w_2(x) = 1 - w_1(x)$.

5.3 Results

Fig. 5.5 shows the results of one-dimensional interpolations between two (two-dimensional) quadratic forms according to Eqs. (5.3) and (5.4), both for the unambiguous indicatrix (a) and for some of its projections (b). Looking at the vertical spoke in Fig. 5.5(a) we note that the interpolated indicatrix seems ellipsoidal if the interpolated tensors are proportional, and we can indeed conclude from Eq. (5.3) that in this

$$w_1(x) + w_2(x) + w_3(x) + w_4(x) = 1$$
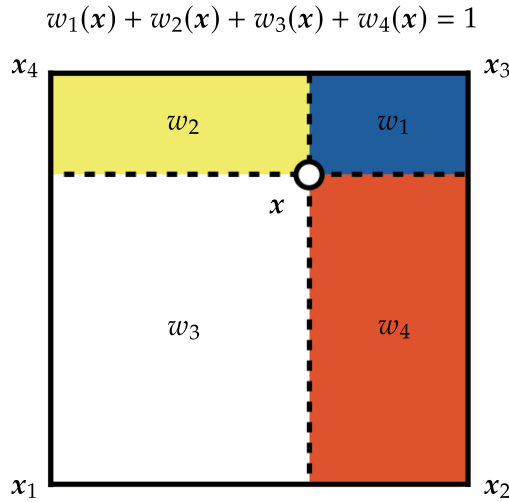
Figure 5.4: First order linear interpolation on a regular grid in two dimensions. The weight $w_k(x)$ for an interpolation point x is proportional to the indicated surface areas opposite of the associated data point x_k. The interpolated Finsler function $F(x, y)$ is defined as the weighted geometric mean $\Pi_{k=1}^4 [F_k(y)]^{w_k(x)}$, in which $F_k^2(y) = F^2(x_k, y)$ are quadratic forms in y given by the DTI data at the grid points x_k. The corresponding expression for the interpolated fundamental tensor $g(x, y)$ follows analogously to the derivation of Eq. (5.4).

case interpolation boils down to isotropic scaling. Generically, however, the interpolated indicatrix is non-ellipsoidal, which is most pronounced on the horizontal spoke where the main eigenvectors of the tensors at the ends are perpendicular. The corresponding projections on the other hand, see e.g. Fig. 5.5(b), are always ellipsoidal by construction. Given a pair of tensors and a fixed weight there are infinitely many such projections, and a unique one is only singled out after providing a preferred orientation. If this orientation happens to be aligned with either of the tensors, then this alignment biases the interpolation to the extent that the corresponding orientation tends to be emphasized.

To better appreciate this phenomenon, consider Fig. 5.6, which shows the interpolation along the horizontal spoke of Fig. 5.5 in more detail. Figs. 5.6(b) and (c) show ellipses obtained through projection based on horizontal, respectively vertical orientations. It is apparent that a horizontal (vertical) bias enhances the influence of the horizontally (vertically) oriented tensor, thus breaking the naively expected symmetry in this example where the tensors are 90°-rotated copies. Observe for example the occurrence of an isotropic transition strictly before the midpoint between the two outer ellipses in Fig. 5.6(b)—the horizontal orientation used in the projection results in a bias that emphasizes the horizontally oriented ellipse on the right. In the vertical projection shown in Fig. 5.6(c), which coincides with the orientation of the main

(a)   **(b)**
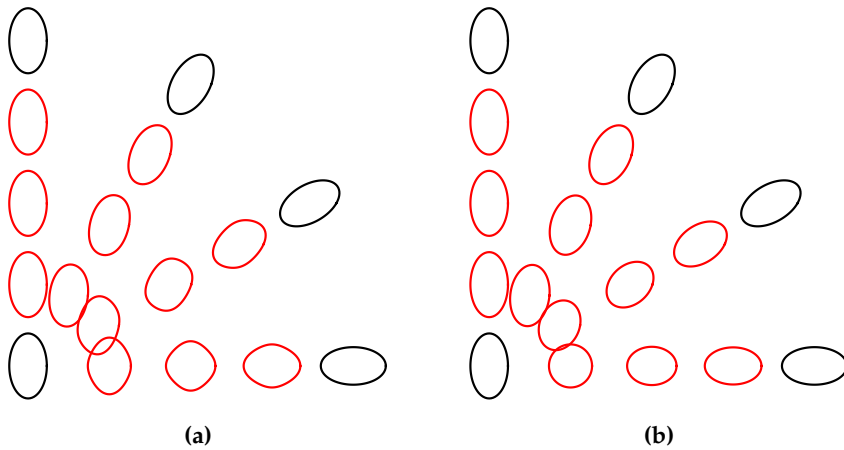
Figure 5.5: (a) Indicatrices of Finsler functions obtained with Eqs. (5.3) and (5.4). The acquired tensor data is represented by the black ellipses at either end of a spoke. The radial distance along a spoke reflects the relative weights given to the ellipses at the ends, recall Section 5.2. (b) Ellipses obtained after projection onto the cone of positive semi-definite matrices. Projections are computed using the orientation of the spoke as the preferred orientation.

axis of the ellipse on the left, the isotropic transition occurs after the midpoint. This bias may be exploited in a spatial interpolation scheme to express one's expectation about the orientation of the structure of interest, if available. In particular, crossing fiber pathways that are common in neuroimaging 'see' different ellipsoidal interpolations at the same crossing point depending on their tangents, as could already be seen in Fig. 5.3. In any reasonable non-parametric interpolation scheme the isotropic transition necessarily occurs exactly in the middle, cf. Fig 5.6(d).

Fig. 5.7 illustrates grid interpolation according to the bilinear interpolation scheme sketched in Fig. 5.4. The experiment shown in (a) shows a discrete tensor field sampled from a smooth 'single fiber orientation' region of interest, simulating a mildly curved underlying fiber bundle (black ellipses), together with its interpolation to a refined grid. Fig. 5.7(b) illustrates the case of a sampled, randomly oriented tensor field.

It can be observed in Fig. 5.7(a) that the interpolating gauge figures are almost elliptical, and thus have a fairly well-defined, though not entirely crisp, preferred orientation. Although orientation is (by construction) likewise unambiguous in each sample point in Fig. 5.7(b), it is much more fuzzy and less anisotropic at typical interpolation points, and the non-elliptical shape of the (convex) gauge figures is also more apparent. In both cases one needs a field of preferred orientations to single out an unambiguous member from the osculating family of ellipses so as to re-obtain a symmetric positive-definite second order tensor at each point. Such a field may be relatively straightforward and inspired by the data in the former case, but less trivial in the latter.

(a) Indicatrices of the direction-controlled interpolation

(b) Horizontal projection

(c) Vertical projection

(d) Log-Euclidean interpolation
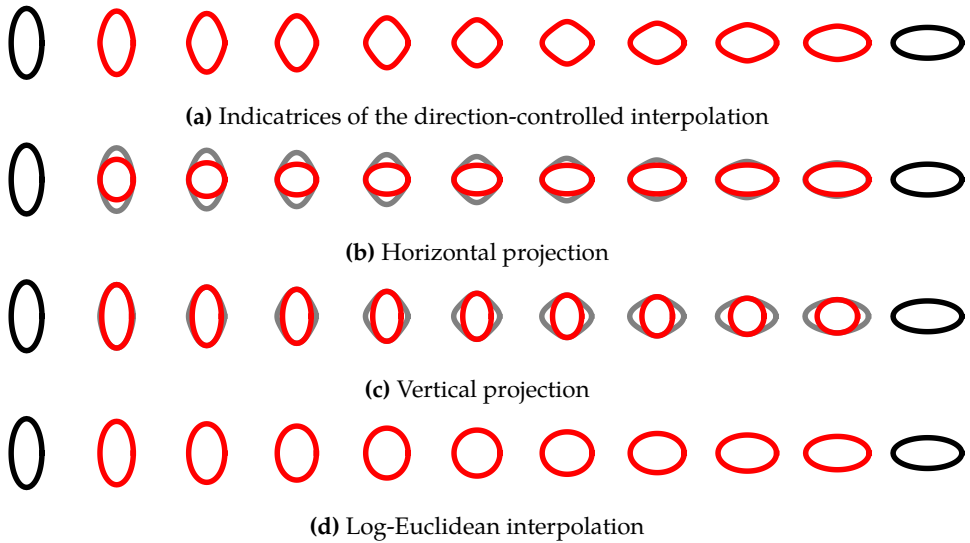
Figure 5.6: Comparison of linear interpolation schemes on two-dimensional tensors. The acquired data is represented by the black ellipses at the ends, and the interpolations are shown in red. (a) Indicatrices of the interpolated Finsler functions based on the scheme proposed in Eqs. (5.3) and (5.4). (b) Ellipses obtained from (a) by evaluating the fundamental tensor biased to the horizontal axis. (c) The vertical projection of (a). (d) Log-Euclidean interpolation [13]. In this case the interpolation is symmetric, and the ellipse in the middle is isotropic. In (b) and (c) the a posteriori imposed bias allows this symmetry to be broken in a controlled manner.

(a)

(b)
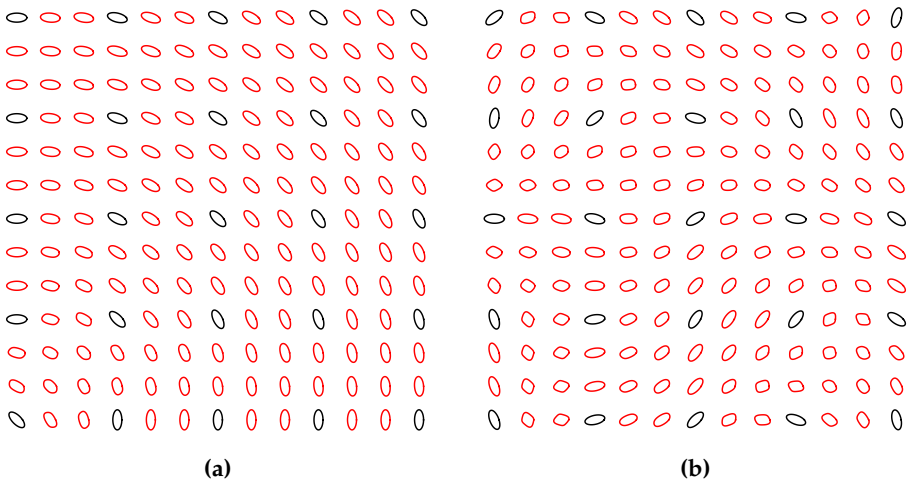
Figure 5.7: Two-dimensional grid refinement through interpolation according to the scheme of Fig. 5.4. The black ellipses are grid samples, the red ellipses are interpolations. (a) Interpolation of a smoothly varying discrete tensor field. (b) Interpolation of a randomly varying discrete tensor field.
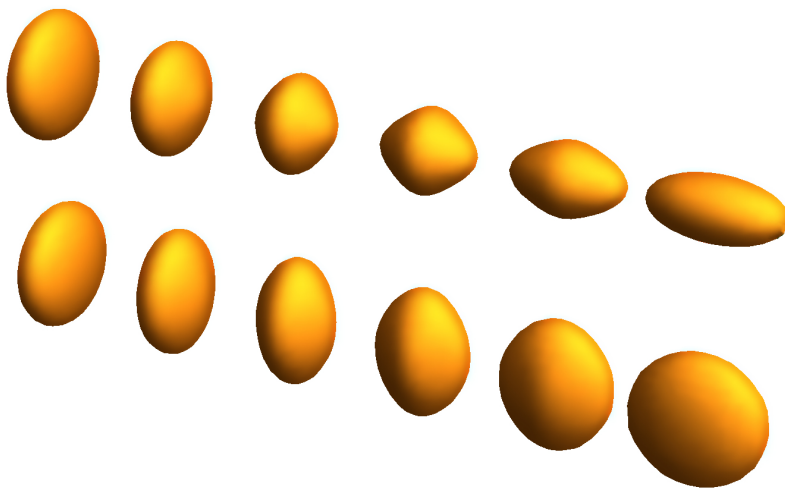
Figure 5.8: Three-dimensional direction-controlled interpolation of DTI-based ellipsoids. The figures on either end are genuine ellipsoids, the interpolated objects in-between are convex indicatrices. A preferred direction is needed in order to project each interpolated shape back to an ellipsoid. Top: Interpolation between an oblate and a prolate spheroid. In this configuration, the minor axis of the former is aligned with the major axis of the latter. Bottom: Interpolation between two oblate spheroids with a 90° relative turn around one of the major axes.

Extension to three dimensions is straightforward. There are many more scenarios for the relative configurations of two ellipsoids and many more possible preferred directions to be specified in order to single out an osculating indicatrix/figuratrix. Figure 5.8 shows two examples.

5.4   Discussion

In Section 5.1 we discussed interpolation in the context of the Riemannian and Finslerian frameworks, i.e., interpolation of quadratic forms (represented by ellipsoids), and interpolation of normed spaces (represented by convex bodies; indicatrices). Because of the relationships elucidated in Chapter 2, interpolation schemes proposed in the context of diffusion tensor imaging (DTI) can be applied straightforwardly to the Riemannian framework. Interpolation of DTI data has been extensively investigated, and we briefly discussed how the ideas behind geodesic interpolation in particular may be extended to Finsler manifolds. A more detailed treatment of interpolation of data in the Finslerian framework is future work.

To bridge the gap between the relatively well-understood Riemannian case, and the still uncharted Finslerian case, we propose a set of direction-controlled interpo-

lation schemes based on the Finslerian fundamental tensor, cf. Eq. (3.5). The fundamental tensor provides a mapping from Finsler functions to quadratic forms, by approximating the Finslerian norm around a given 'preferred' orientation. This enables the use of Riemannian interpolation techniques in the Finsler setting, and allows Finslerian interpolation techniques to be used in the Riemannian setting. The former of these has some significant disadvantages, with the most important being that this style of interpolation disregards important information in the Finsler function. The latter, on the other hand, can be potentially useful in applications where a preferred orientation is readily available, such as streamline tractography [26, 242].

Direction-controlled interpolation of Riemannian data entails an embedding of the data in a Finsler manifold. We can then apply more general interpolation schemes, and the interpolated Finsler functions can be projected back to the space of quadratic forms using the preferred orientation. In Section 5.3 we present some preliminary results, based on the multiplicative (geometric mean) interpolation of Finsler functions explained in Section 5.2. These results illustrate how the additional degrees of freedom provided by this embedding, embody a 'memory' for recollection of the orientation information available in the original tensor data, which would inevitably be destroyed in a parameter-free interpolation scheme.

The preferred orientation can be used to guide the interpolation, e.g. by including priors such as anatomical (atlas) data. Additionally, the direction-controlled scheme provides a way to incorporate a contextual bias—it assigns a relevance to a given ellipsoid not only on the basis of its spatial distance to a point, but also includes its approximate orientation. By going to higher order interpolation schemes, we also change the scale at which context is taken into account. Another way to look at this, is that direction-controlled interpolation allows for a trade-off between spatial and angular resolution. Taking into account a larger neighborhood—i.e., lowering the spatial resolution—allows an increased angular resolution. In future work we hope to frame this technique w.r.t. more rigorous solutions to include context, such as the higher order tensor inpainting method of Schultz [300] and the works on sub-Riemannian geometry and diffusion MRI by various authors [93, 111–114, 279, 280].

Chapter 5

Part III

Analysis

Chapter 6

Point

The Local Geometry

Based in part on:

T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. *The Finslerian character of the diffusion MRI signal*. In preparation, preprint available upon request.

This first chapter of the *Analysis* part of the thesis introduces a number of new local information descriptors, based on the geometrical viewpoint of Part I. These local descriptors provide voxel-wise characterizations of the observed diffusion process, and thus of the postulated geometry. Because the model in the Riemannian framework is conformal to the diffusion tensor imaging (DTI) model, for which there already exists a vast set of measures in the literature [5, 27, 273, 350, 372], we will focus here on the more general Finslerian framework proposed in Chapter 3.

The proposed descriptors are evaluated in the context of neuroimaging, where we are interested in the relation between the local geometry and the tissue structure. We will focus in particular on the local distribution of axon orientations, which has likely been the most sought-after tissue feature [179, 344, 346] in the short history of the diffusion MRI field. One expects this distribution to be inferable from the diffusion data, because it has been observed in so-called single fiber populations that diffusion is mainly diminished perpendicular to the dominant axon orientation [30].

The chapter is structured as follows. In Section 6.1 we discuss various scalar and orientation-dependent measures based on the Finslerian framework. We start in Section 6.1.1 with an orientation distribution function (ODF) that characterizes the orientation-dependent resistance encountered by diffusing spins, which we propose as a measure of the axonal orientation distribution. The first examples of scalar measures in Section 6.1.2.1 are generalizations of Riemannian measures, and the remaining scalar measures discussed in Sections 6.1.2.2 and 6.1.2.3 are intended to quantify the differences between the Riemannian and Finslerian viewpoints. Finally in Section 6.1.3 we briefly describe how features may be derived from the (orientation-dependent) fundamental tensor. Implementation details are provided in Section 6.2, and we present proof-of-concept experiments in Section 6.3. Section 6.4 concludes the chapter with a discussion on the potential applications of the proposed local descriptors. Preliminary results of the work presented in this chapter have been presented at the ISMRM [84, 87–89]

6.1 Theory

6.1.1 The barrier orientation distribution function

In the Riemannian and Finslerian frameworks introduced in Part I, the local geometry is typically visualized by means of the figuratrix, recall for example Fig. 3.9. In both geometries the figuratrix represents the local deformation of the space, i.e., the radius in a given direction shows the degree to which space is stretched in that direction. In this section we attempt to characterize the source of this anisotropy—the spatial correlations in interactions undergone by diffusing particles [248, 249].

(a) Model of linear barriers   **(b)** Model of planar barriers
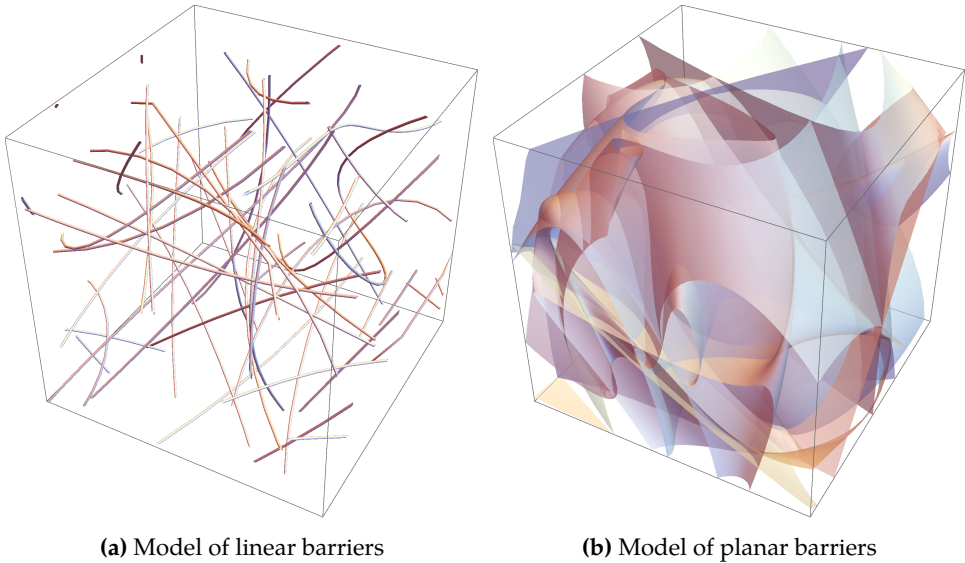
Figure 6.1: The structural models used in the derivation of the barrier orientation distribution functions. Particle dynamics within voxels (the black boxes in the figure) are assumed to be governed entirely by structures that function as barriers to diffusion (colored lines and surfaces), and whose local orientations can be modeled as if randomly drawn from a single distribution on the sphere. In the visualized examples this distribution is uniform.

6.1.1.1   Linear barriers

The primary vehicle we deploy for our correlation analysis is the barrier, which is a localized structure defined to have a linearly impeding effect on motion; a particle encountering twice the number of barriers will move half as far. In this section we assume that all barriers are fundamentally one-dimensional, and that the only relevant parameter at the diffusion length scale is the orientation of (the tangent to) the barrier (see e.g. Jensen et al. [179]). Considering the spatial homogeneity assumption (Assumption 1), we then submit that there exists a distribution Ψ_l on the sphere, such that $\Psi_l(\hat{r})$ represents the probability density that a one-dimensional barrier is (locally) oriented along $\hat{r} \in S^2$ in any given point in the voxel. See Fig. 6.1(a) for an illustration. A more formal definition of this probability distribution can be found in convex geometry literature, see e.g. the book by Stoyan [324]. We refer to Ψ_l as the linear barrier orientation distribution function (linear bODF), and its relation to the Finsler function F can be understood as follows.

The expected number of barriers with an orientation $\hat{r}'$ encountered along a unit displacement vector $\hat{r} \in S^2$ is given by $\sqrt{1 - |\hat{r} \cdot \hat{r}'|^2} \, \Psi_l(\hat{r})$, Fig. 6.2, where $\sqrt{1 - |\hat{r} \cdot \hat{r}'|^2}$ is the sine of the angle between $\hat{r}$ and $\hat{r}'$ and $\cdot$ denotes the Euclidean dot product. The
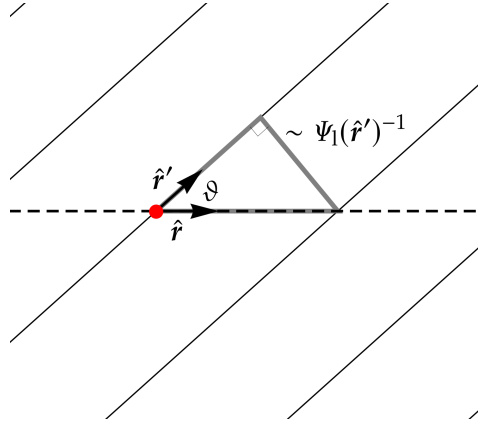
Figure 6.2: A schematic drawing illustrating the relation between the density of barriers with a given orientation $\hat{r}'$, and the largest distances particles can traverse along an orientation $\hat{r}$ in a barrier setting. The maximum distance particles can travel is inversely related to the average number of barriers per unit of (path) length, which in turn is inversely related to the average distance between barriers relative to $\hat{r}$. This distance is determined by $\Psi_l(\hat{r}')$ and the angle $\vartheta(\hat{r}, \hat{r}')$ between $\hat{r}$ and $\hat{r}'$. A formal derivation of this relationship is more nuanced (particularly in three or more dimensions) and is not presented in this manuscript.

fraction of the barriers that cross $\hat{r}$ is thus given by the sine transform

$$\mathcal{S}\left\{\Psi_l\right\}(\hat{r}) = \int_{S^2} \sqrt{1 - |\hat{r} \cdot \hat{r}'|^2}\, \Psi_l(\hat{r})\, d\sigma(\hat{r}'), \tag{6.1}$$

with $d\sigma$ the Lebesgue measure on the Euclidean unit sphere S^2. In light of the definition of a barrier, we can then state that $\mathcal{S}\left\{\Psi_l\right\}(\hat{r})/F(\hat{r})$ should be approximately constant, and it follows that Ψ_l is related to F as

$$\Psi_l(\hat{r}) \propto \mathcal{S}^{-1}\left\{F\right\}(\hat{r}), \tag{6.2}$$

where $\mathcal{S}^{-1}$ is the inverse sine transform. Ψ_l is then determined after normalization using

$$\int_{S^2} \Psi_l(\hat{r}')\, d\sigma(\hat{r}') = 1. \tag{6.3}$$

The linear barrier ODF of an example Finsler function is shown in Fig. 6.3.

6.1.1.2 Planar barriers

We could alternatively consider the orientation distribution of planar barriers Ψ_p, such that $\Psi_p(\hat{r})$ represents the probability that a two-dimensional barrier has a normal vector oriented (locally) along $\hat{r}$, see Fig. 6.1(b). An analogous derivation then gives

(a) $\partial\Omega$ **(b)** Ψ_1
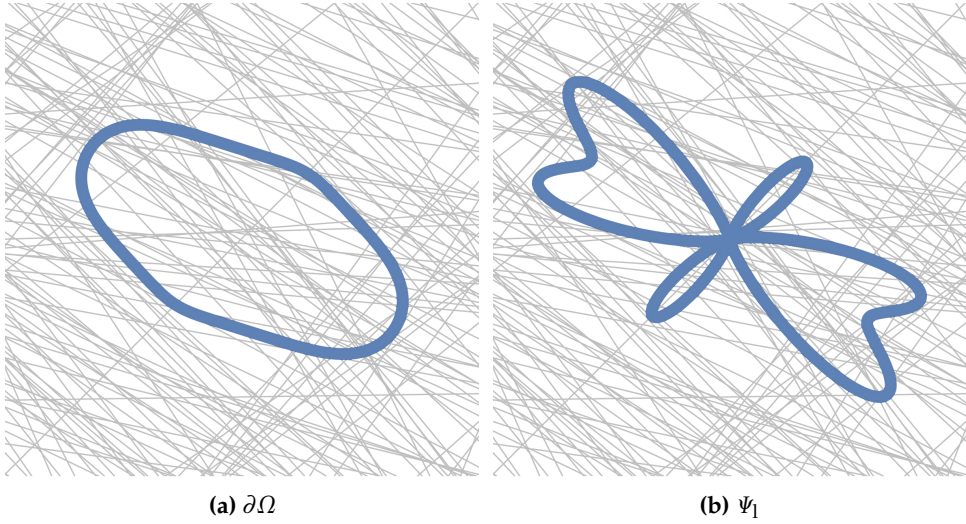
Figure 6.3: (a) An example boundary $\partial\Omega$ of a finite speed isotropic transport process, determined by the Finsler function F as described in Section 3.2.1.3. $\partial\Omega$ represents the largest displacements along different orientations in a locally Minkowskian manifold. (b) The inverse sine transform of the boundary in (a) produces the linear barrier ODF Ψ_1, cf. Eq. (6.2). The gray lines in the background are a rendering of linear barriers distributed according to Ψ_1. In two dimensions $C\{f\}(\hat{r}(\theta)) = S\{f\}\left(\hat{r}\left(\theta + \frac{\pi}{2}\right)\right)$, with θ the polar coordinate, so the planar barrier ODF corresponding to the boundary in (a) would simply be a rotated version of the shown linear barrier ODF.

the relation between F and the planar barrier ODF

$$\Psi_p(\hat{r}) \propto C^{-1}\{F\}(\hat{r}), \tag{6.4}$$

where C^{-1} is the inverse of the cosine transform C defined as

$$C\{\Psi_p\}(\hat{r}) = \int_{S^2} |\hat{r} \cdot \hat{r}'|\, \Psi_p(\hat{r}')\, d\sigma(\hat{r}'). \tag{6.5}$$

It can be shown that the planar barrier ODF is not fundamentally different from the corresponding linear bODF, in the sense that a set of linear barriers distributed uniformly along a great circle of the sphere is indistinguishable from an aligned planar barrier. In fact, the sine and cosine transforms of Eqs. (6.1) and (6.5) are related according to

$$\mathcal{R} \circ \mathcal{S}^{-1} \propto C^{-1}, \tag{6.6}$$

where $\circ$ denotes function composition and $\mathcal{R}$ denotes the Funk–Radon transform

$$\mathcal{R}\{f\}(\hat{r}) = \int_{S^2} \delta\left(|\hat{r} \cdot \hat{r}'|\right) f(\hat{r}')\, d\sigma(\hat{r}'), \tag{6.7}$$

with f a continuous function on the sphere and with $\delta(x) = 1$ if $x = 0$ and 0

otherwise. Eq. (6.6) follows when considering $C \circ \mathcal{R} \overset{?}{\propto} \mathcal{S}$. This simple relation means that it suffices to consider just the linear bODF $\Psi_l =: \Psi_b$ if we are interested in the local configuration of (either linear or planar) barriers. We will thus omit the linear/planar prefix from hereon.

Remark 8. *We conclude this discussion on linear and planar barriers by noting that there is a complementary model to describe anisotropy in a Finsler manifold. In this alternative case Ψ_l does not represent a distribution of barriers, but instead represents a distribution of 'oriented facilitators' whose density is linearly related to the largest displacements* along *the corresponding orientation. A similar derivation then leads to Ψ_l being related to F through the cosine transform that now links F with Ψ_p, and thus provides an alternative interpretation for Ψ_p. Analogously Ψ_l could alternatively be viewed as a distribution of planar 'facilitators'. In simple terms, the linear barrier ODF coincides with the planar facilitator ODF, and the planar barrier ODF coincides with the linear facilitator ODF. The facilitator ODF is known in microscopy as the 'Rose of Directions' [196].*

6.1.1.3 Axons as barriers

It is generally accepted that the diffusion anisotropy observed in the brain is to a large extent due to a dense packing of fibrous axonal membranes present in the tissue [30]. These axons have diameters of around $1\,\mu m$ [1], while the diffusion length scale in neuroimaging experiments is in the order of tens of microns. We additionally expect that the largest displacements in a voxel occur in the extra-axonal compartment, which would thus be represented by the Finsler function. Based on these points, we hypothesize that the axonal orientation density function can be well-approximated by the barrier orientation distribution function.

The barrier ODF is fundamentally different from commonly used existing approaches to compute axon orientations [2, 101, 148, 167, 195, 259, 261, 344, 347]. The closest alternative is probably the fiber ODF Ψ_f that can be obtained by means of spherical deconvolution [96, 146, 182, 302, 344], although the basic premise of spherical deconvolution-based approaches is almost polar opposite to ours—axons of different orientations are assumed to contribute to the signal independently. It is then assumed that the contribution of axons with a given orientation can be represented by a single building block—the single fiber response function U—so that the signal (for a unit wave vector $\hat{q}$ with a fixed b-value) can be written as a convolution:

$$S(\hat{q}) = \int_{S^2} U(\hat{q}, \hat{r}') \Psi_f(\hat{r}') \, d\sigma(\hat{r}'). \tag{6.8}$$

The response function U is typically estimated from the average signal in regions considered representative of uni-oriented fiber pathways [336, 344], and the fiber ODF can then be recovered from Eq. (6.8) using deconvolution [97, 342].

Due to the log transform in the relation between the signal and the Finsler function (Eq. (3.29)), fiber ODFs are generally incompatible with barrier ODFs. There does not exist a single fiber response function such that Ψ_l and Ψ_f always coincide.[1] Whether Ψ_l or Ψ_f is more accurate depends essentially on whether or not crossing axons can be modeled at the diffusion length scale as distinct compartments—whether the axons are interwoven or distinct.

6.1.2 Finslerian scalar measures

In this section we cover a number of scalar measures that can be derived from the Finsler function discussed in Chapter 3. These scalar measures are designed to capture basic information about the local geometry of the Finsler manifold.

6.1.2.1 Range measures

The reconstructed Finsler function $F : T_pM \to [0, \infty)$ follows from an estimate of the boundary $\partial\Omega$, which is an interesting quantity in itself. Basic questions addressed in DTI by looking at the tensor eigenvalues, such as where is the diffusion strongly direction-dependent (cf. Section 2.3.1), can be answered using $\partial\Omega$ as well. We consider the voxel-wise smallest and largest (estimated) maximum displacements $\Lambda_{\min}$ and $\Lambda_{\max}$, their difference Λ_{rel}, and the spherical mean Λ_{avg}. We define these scalar measures in terms of F (at an implicit point p) by Eqs. (3.22) and (3.42) as

$$\Lambda_{\min}(p) := \min_{F(r)=v\,\Delta} \|r\|, \tag{6.9}$$

$$\Lambda_{\max}(p) := \max_{F(r)=v\,\Delta} \|r\|, \tag{6.10}$$

$$\Lambda_{\mathrm{avg}}(p) := v\,\Delta \int_{\mathrm{S}^2} \frac{1}{F(\hat{r}')}\,\mathrm{d}\sigma(\hat{r}'), \tag{6.11}$$

$$\Lambda_{\mathrm{rel}}(p) := \frac{\Lambda_{\max}(p) - \Lambda_{\min}(p)}{\Lambda_{\max}(p)}, \tag{6.12}$$

where $\hat{r}$ is a normalized (dimensionless) vector in T_pM.

Just as in the Riemannian framework, Chapter 2, these measures can be interpreted independent of the diffusion as geometrical descriptors. This is made explicit after a division by $v\,\Delta$ (turning the Λ measures into dimensionless quantities) after which the norm F can be understood to simply represent the relative amount of scaling along a given direction. The $\Lambda_{\max}$ measure for example is proportional to the largest relative stretching of space at p, while Λ_{avg} gives the average amount of scaling.

---

[1]The barrier ODF is closer to a morphological deconvolution of the signal [47].

Remark 9. *Note that the Λ measures are quantitative; we get estimates of the maximum displacements in e.g.* mm*. The fixed unknown v in the expressions cancels with the fixed value set in the computation of F in Eq. (3.29).*

6.1.2.2 Non-Gaussianity

The Λ-measures have direct analogues in DTI by construction; the smallest and largest eigenvalue, the mean diffusivity (MD), and the fractional anisotropy (FA) convey similar information as the listed measures. The most obvious quantity that cannot be obtained with DTI is non-Gaussianity, which represents the portion of the signal not accurately described by the DTI tensor. As purely Gaussian diffusion implies that $d_l^m = 0$ for $l > 2$ in Eq. (3.43), i.e., that the indicatrix is ellipsoidal, we can define a non-Gaussianity measure similar to Özarslan et al. [259] as the dissimilarity

$$\mathrm{NG} := \sin \vartheta(\boldsymbol{u}, \boldsymbol{v}) = \sqrt{1 - \cos^2 \vartheta(\boldsymbol{u}, \boldsymbol{v})}, \tag{6.13}$$

where $\cos \vartheta := \cos \vartheta(\boldsymbol{u}, \boldsymbol{v})$ is the cosine similarity between the vectorizations $\boldsymbol{u} = (d_0^0, d_2^{-2}, \ldots, d_2^2)$ and $\boldsymbol{v} = (d_0^0, d_2^{-2}, \ldots)$ of the (truncated) spherical harmonic coefficients d_l^m. A non-Gaussianity $\mathrm{NG} = 0$ indicates Gaussian diffusion—a perfectly ellipsoidal indicatrix.

Remark 10. *Non-Gaussian behavior originates from a non-trivial interplay between local structural complexity, and the inevitable emergence of relevant higher order moments at large gradient strengths. It is thus a measure of 'signal complexity', not of 'structural complexity'.*

6.1.2.3 Non-Finslerianity

The non-Gaussianity measure of the previous section quantifies to what extent a Riemannian manifold is a suitable model for the measured signal, roughly by comparing high gradient strength measurements to their predicted values based on the Riemannian metric. We can similarly compare predictions based on the Finslerian structure derived from high gradient strength data, to observed behavior at low gradient strengths. This leads to a measure of non-Finslerianity as follows.

In Section 3.2.1.2 we explained how the Finsler function F provides a prediction for the central limit behavior that governs the low gradient strength segment of the signal through Eq. (3.19), under the condition that the measured diffusion is reasonably close to an isotropic transport process. Given an estimate F_κ of F, we can use Eq. (3.21) to define an estimated low gradient strength diffusion tensor

$$D_\kappa^{ij} := 3\,\alpha \int_{I_xM} y^i y^j\, \mu_x(\mathrm{d}\boldsymbol{y}), \tag{6.14}$$

where μ_x is the measure defined by F_κ and α is the diffusion coefficient (cf. Chapter 1).

If F_κ tends to F and Assumption 1 holds, then D_κ converges to the usual diffusion tensor D (Eq. (2.14)) that can be recovered from low gradient strength diffusion-weighted images ($\kappa \to 0$).

Remark 11. *Note that Eq. (6.14) (implicitly) depends on v and α, which are both unknown. Practically, this means that we can only estimate D_κ up to a constant. In the preliminary experiments presented in this chapter, we will define a global scaling between the diffusion tensors and their Finslerian estimates, such that the median mean diffusivities (over all voxels) of the two tensor fields are the same.*

The quantity of interest is now the difference between the tensor D, and the estimate D_κ recovered from the complete set of measurements. Here we will quantify this difference by means of the log-Euclidean distance [13]:

$$d(D, D_\kappa) := \sqrt{\mathrm{tr}\left(\left\{\log D - \log D_\kappa\right\}^2\right)}, \tag{6.15}$$

where log denotes the matrix logarithm and $\{\cdot\}^2$ the matrix square. Based on the log-Euclidean distance between D and D_κ we define a non-Finslerianity measure as

$$\frac{d(D, D_\kappa)}{\sqrt{\mathrm{tr}\left(\left\{\log D\right\}^2\right)}}. \tag{6.16}$$

In Section 3.5.1 we already identified two factors that can cause a high non-Finslerianity value—either the measured spin stochastics cannot be accurately modeled as a Finslerian isotropic transport process, or there are multiple partitions that contribute to the signal independently, i.e., Assumption 1 does not hold. Remark 11 adds a third one: if v or α is not constant over M, then the globally fixed scaling may introduce a position-dependent bias in the non-Finslerianity.

6.1.3 Riemannian measure glyphs

The last class of features we introduce in this chapter are the orientation-dependent DTI measures, based on the fundamental (dual) tensor that was already used in Chapter 5, see Fig. 5.2. The fundamental tensor associates a two-tensor to each orientation, from which we can compute scalar-valued descriptors like the mean diffusivity and the fractional anisotropy, cf. Section 2.3.1. Recall that the fundamental tensor describes the shape of the indicatrix/figuratrix near a given point, meaning that it describes the behavior of diffusing particles that move approximately in a given orientation.

Orientation-dependent DTI measures can be represented as glyphs by defining the radius of the glyph for a given orientation to be proportional to the measure of the associated tensor. Alternatively, we can compute the average value of a Riemannian

measure glyph f weighted by the bODF, defining a barrier-weighted DTI measure through the integral

$$\int_{S^2} f(\hat{r}')\Psi_1(\hat{r}')\,d\sigma(\hat{r}'),\tag{6.17}$$

which represents the expected value of the DTI measure along fiber pathways. We will generally look at measures based on the fundamental dual tensors, as these are natively closest to the diffusion tensors in DTI.

6.2 Methods

In this section we provide implementation details for the discussed local descriptors. In all instances it is assumed that we have expressions for H_κ and F_κ in terms of the spherical harmonic coefficients $c_l^m(\kappa)$ and $d_l^m(\kappa)$ respectively, as explained in Section 3.3. Spherical descriptors $f : M \times S^2 \to \mathbb{R}$ will be represented by glyphs $\{f(x, \hat{r})\,\hat{r} \,|\, \hat{r} \in S^2\}$ as explained in Section 2.3.2.

6.2.1 The barrier orientation distribution function

The barrier ODF requires the computation of inverse sine transform. For convex bodies expressed in terms of spherical harmonics, this transform can be computed through an application of the Funk–Hecke theorem [298, Appendix]. Using this theorem it follows that given F_κ in terms of $d_l^m(\kappa)$ as in Eq. (3.43), the coefficients $v_l^m(\kappa)$ in the spherical harmonic expansion of the barrier ODF are given by

$$v_l^m(\kappa) = -\frac{2}{\pi}\frac{\Gamma\left(\frac{l}{2}+1\right)\Gamma\left(\frac{l}{2}+2\right)}{\Gamma\left(\frac{l}{2}-\frac{1}{2}\right)\Gamma\left(\frac{l}{2}+\frac{1}{2}\right)}d_l^m(\kappa),\tag{6.18}$$

with Γ the gamma function. Though we do not consider this here, a unique boundary $\partial\Omega$ can be derived from a barrier (or any other type of) ODF by using the spherical harmonic factors to compute a sine transform.

Remark 12. *The coefficients $w_l^m(\kappa)$ of the planar barrier ODF are given by*

$$w_l^m(\kappa) = \pi^{-\frac{3}{2}}\Gamma\left(\frac{3}{2}-\frac{l}{2}\right)\Gamma\left(\frac{l}{2}+2\right)d_l^m(\kappa),\tag{6.19}$$

and the Funk–Radon transform can be similarly expressed as

$$u_l^m(\kappa) = \frac{2\pi^{\frac{3}{2}}}{\Gamma\left(\frac{1}{2}-\frac{l}{2}\right)\Gamma\left(\frac{l}{2}+1\right)}d_l^m(\kappa).\tag{6.20}$$

Together with Eq. (6.18), these allow us to easily verify the relation in Eq. (6.6).

In the experiments involving the barrier ODF we are particularly interested in the local maxima (or 'peaks') of the ODF—the orientations that we expect to correspond to dominant axon orientations. We identify local maxima using a simple gradient ascent method, initialized from several hundred randomly generated orientations for each voxel.

6.2.2 Finslerian scalar measures

The optimizations in the definitions of $\Lambda_{\max}$ and $\Lambda_{\min}$ (Eqs. (6.10) and (6.9)) are approximated by evaluating $\|\xi(q)\|$ for a large number of q, using the same sampling scheme as in Section 3.3.3. $\|\xi(q)\|$ is computed using the expressions in Appendix B. Λ_{avg} follows straightforwardly by orthonormality of the spherical harmonics, giving

$$\Lambda_{\text{avg}} = 2\sqrt{\pi}\,\frac{v\,\Delta}{d_0^0(\kappa)} = 2\sqrt{\pi}\,v\,\Delta\,c_0^0(\kappa). \tag{6.21}$$

The non-Gaussianity measure NG can also be expressed in terms of the spherical harmonic coefficients:

$$\cos\vartheta = \sqrt{\frac{\sum_{l=0}^{2}\sum_{m=-l}^{l}\left(d_l^m(\kappa)\right)^2}{\sum_{l=0}^{\infty}\sum_{m=-l}^{l}\left(d_l^m(\kappa)\right)^2}}. \tag{6.22}$$

The integration in Eq. (6.14) required for the non-Finslerianity measure, is computed numerically using the analytical expression for the determinant det g given in Appendix B.

6.2.3 Riemannian measure glyphs

The fundamental tensor is computed using the analytical expressions provided in Appendix B, again for the same icosahedral sampling used in Section 3.3.3. We then compute for each of these tensors the fractional anisotropy and mean diffusivity [372], and approximate the corresponding Riemannian measure glyphs by fitting spherical harmonics to the resulting values. The barrier-weighted scalar measures are computed directly from the spherical harmonic coefficients of the barrier ODF and the Riemannian measure glyphs, using orthogonality of the spherical harmonics, Eq. (B.2).

6.2.4   Data

Experiments are based on the Human Connectome Project (HCP) data described in Section 2.3.3, and on the artificial data described in Section 3.3.4. We also define two additional artificial ensemble average propagators (EAPs) based on the second order tensors $B := 10^{-3}\,\mathrm{diag}(3, 2, 1)$ and $C := 10^{-3}\,\mathrm{diag}(3, 3, 1)$, and define a less anisotropic version of $A(\vartheta)$ (recall Eq. (3.49)) as

$$\tilde{A}(\vartheta) := \frac{1}{2} \times 10^{-3} \begin{pmatrix} 3 + \cos 2\vartheta & -\sin 2\vartheta & 0 \\ -\sin 2\vartheta & 3 - \cos 2\vartheta & 0 \\ 0 & 0 & 2 \end{pmatrix}. \tag{6.23}$$

The spherical harmonic coefficients c_l^m and d_l^m of the (dual) Finsler functions are reconstructed as described in Section 3.3.3. The DTI results are computed as explained in Section 2.3. We also compare the bODFs to fiber ODFs, computed using constrained spherical deconvolution (CSD) [342] based on the $b = 3000\,\mathrm{s/mm^2}$ shell. CSD results are computed with MRtrix 3 (`mrtrix.org`), using default settings.

6.3   Results

6.3.1   The barrier orientation distribution function

6.3.1.1   Artificial data

We show the bODF obtained from the different DTI data sets A, B, and C in Fig. 6.4. The fractional anisotropy decreases from left to right, reflected by the decreased probability associated with the main orientation in the bODF. The rightmost column shows a circularly symmetric DTI ellipsoid, which translates as expected to a circularly symmetric bODF.

In Fig. 6.5, the bODFs corresponding to the profiles in Fig. 3.8 again show the expected behavior. The first row shows a profile of 'crossing barriers', based on the artificial diffusion tensors up to fourth order, where the angular resolution (the sharpness of the profile around the main orientation) increases significantly with κ. The importance of high quality, high b-value data, is apparent from these results, as this allows for larger κ and thus more accurate estimates F_κ of F. A similar trend can be observed for the multi-tensor data (Fig. 6.5, bottom row). Here the bODF is computed for the average of two identical tensors rotated to have a crossing angle of 45 degrees. For low κ only a single main orientation can be discerned from the bODF, while for higher κ the distinction becomes clearer. Note however, that because multi-tensor data cannot in general be interpreted in terms of largest displacements (recall Section 3.2.2), the main orientations in the profile do not necessarily correspond with
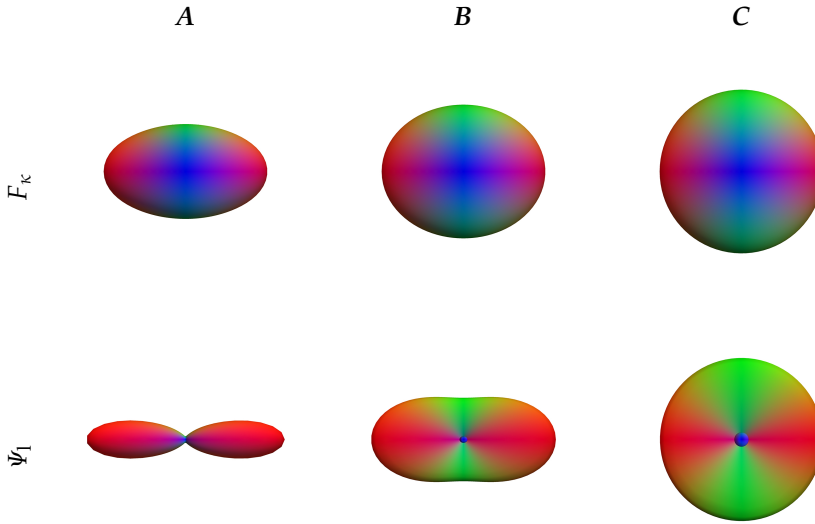
Figure 6.4: Level sets of F_κ and the corresponding bODFs Ψ_1 for different types of diffusion tensor data (*A*, *B*, and *C*, see Sections 3.3.4 and 6.2.4). In the DTI case, the bODF is oriented along the direction of the main eigenvector of the diffusion tensor.

the orientations of interest given by the main eigenvectors of the constituent tensors. We will refer to the resulting deviations from the expected main orientations as 'multi-compartment effects'.

To quantify this deviation, we compare the main orientations of the bODF with the angle between the orientations of the main eigenvectors of different multi-tensor data sets. For $\bar{\kappa} = 100$, the results of these comparisons are plotted in Fig. 6.6. We find that for angles smaller than approximately 25° the bODF has only a single main orientation, but when the bODF has two discernable peaks the total deviation from the expected orientations is less than 5°. This difference increases if κ is taken too large, although the smaller separation angles can then be resolved better (Fig. 6.7). In Figs. 6.6 and 6.7, where the underlying tensors are identical, the difference is distributed equally between the two main orientations. If the underlying tensors differ the total error remains comparable, but the orientation of the dominant tensor can be obtained with far greater accuracy (Figs. 6.8 and 6.9).

6.3.1.2 Real data

In Fig. 6.10(a) and (c) we show the barrier ODFs in the marked region of the HCP data, cf. Fig. 2.2(a). Fig. 6.10(e) contains state-of-the-art fiber orientation density functions (fODF) generated with CSD [342] for comparison purposes. Figs. 6.10(b), (d), and (f) show peaks derived from these ODFs.
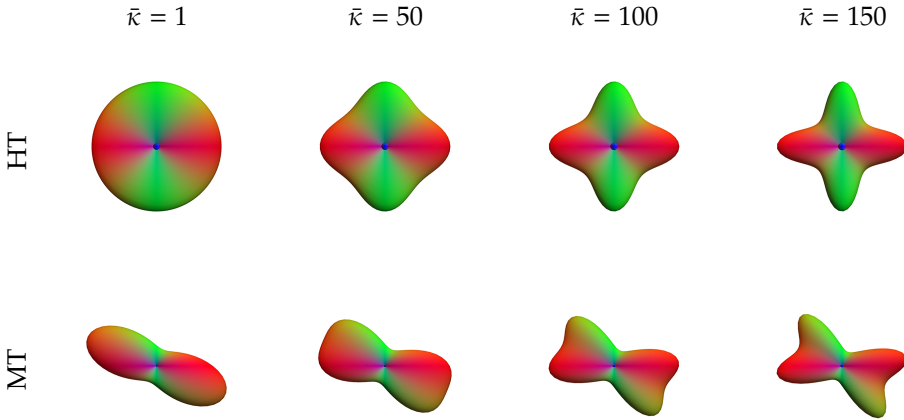
Figure 6.5: bODFs for different values of κ, based on the artificial higher order data (HT, top row) and multi-tensor data (MT, bottom row, $\vartheta = \pi/4$ in $A(\vartheta)$) described in Section 3.3.4. The displayed bODFs correspond to the level sets shown in Fig. 3.8. The dominant barrier directions become more apparent with increasing κ.

The barrier ODF seems to represent known anatomy quite well; we can clearly recognize the coherent parallel bundles in the corpus callosum, transitioning into the centrum semiovale where the callosal pathways intersect with the corticospinal tract, cf. Fig. 2.2(b). There is also a qualitative correspondence with the results obtained with CSD, which is generally considered to work well in this specific region. We find again that the parameter κ can be used to control the angular resolution of the bODF.

The barrier ODF can also be used to visualize the effect of the constrained reconstruction of cumulants discussed in Section 3.3.3. We can already observe—by comparing the level sets and barrier ODFs in Figs. 3.8 and 6.5—that the barrier ODF is a very sensitive measure; differences in the indicatrices which at first seem almost negligible (Fig. 3.8), have a highly significant effect on the corresponding barrier ODF profiles (Fig. 6.5). Based on this one can already imagine that non-convexity artifacts may result in a noticeable deterioration of the barrier ODF, and this is indeed clearly visible if we look at bODFs computed from unconstrained Finsler functions in the HCP data, see Fig. 6.11.

As explained in Chapter 3, we can have minor non-convexity artifacts even after using a constrained model reconstruction. This is due to the unconstrained spherical harmonic reconstruction that is used in the (dual) Finsler function representation. Oscillatory artifacts in level sets of the Finsler function, as highlighted in Fig. 3.11, produce low magnitude high frequency noise in the barrier ODF, see Fig. 6.12.
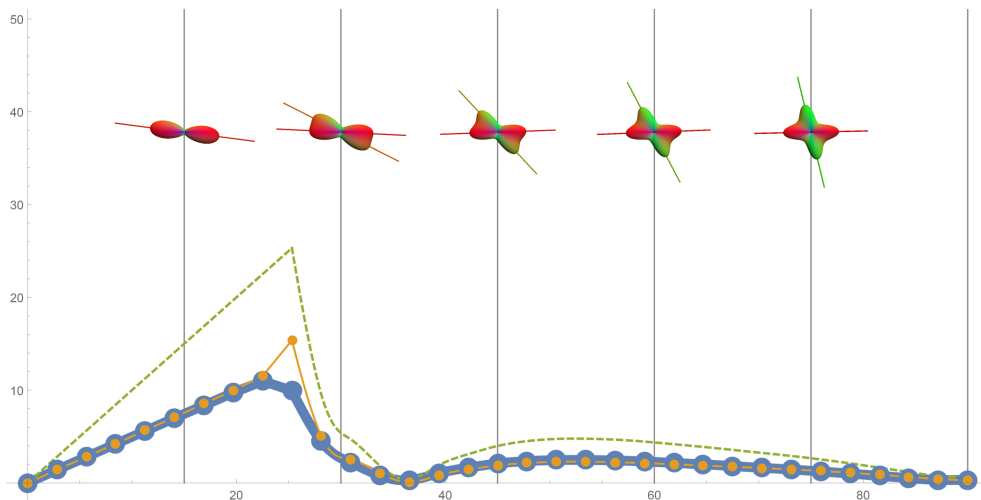
Figure 6.6: The effect of multi-compartment effects on the dominant orientations in the bODF. The horizontal axis shows the angle between the main eigenvectors of the compartmental diffusion tensors in degrees, defined by the angle ϑ in $A(\vartheta)$ (Eq. (3.49)). The orange and blue dots/lines show the angle between the measured orientations and the ground truth orientations (vertical axis, in degrees), and the dashed green line shows the total angular differences. For the considered example, with $\bar{\kappa} = 100$, the separation angle has to be greater than $25°$ for there to be two distinguishable dominant orientations.
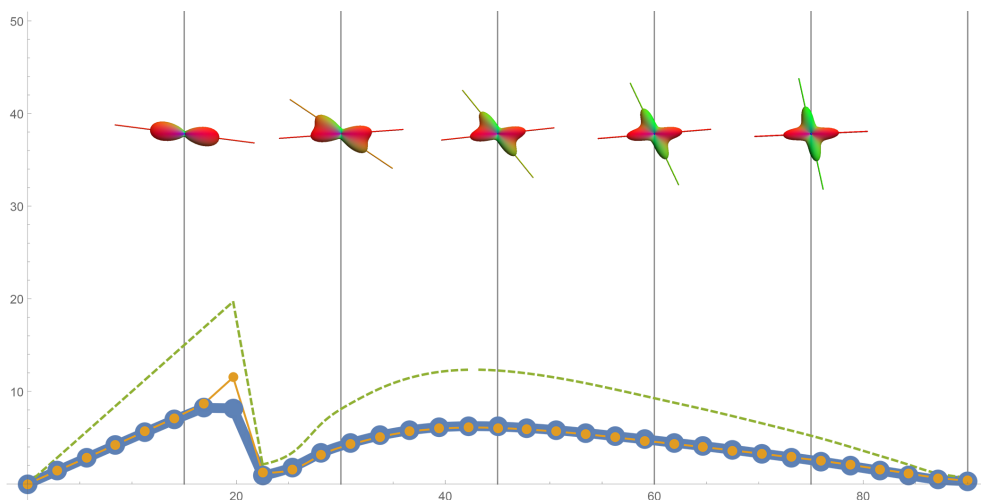
Figure 6.7: The effect of multi-compartment effects on the dominant orientations in the bODF. The orange and blue dots/lines show the angle between the measured orientations and the ground truth orientations (vertical axis), for increasing separation angles (horizontal axis, defined by the angle ϑ in $A(\vartheta)$, Eq. (3.49)). The dashed green line shows the total angular differences. For the considered example, with $\bar{\kappa} = 150$, the separation angle has to be greater than $20°$ for there to be two distinguishable dominant orientations.
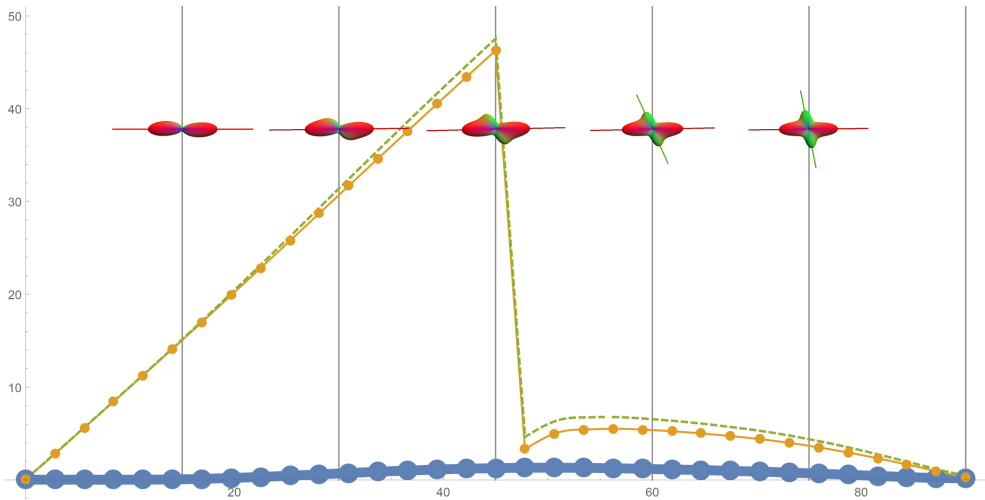
Figure 6.8: The effect of multi-compartment effects on the dominant orientations in the bODF. The orange and blue dots/lines show the angle between the measured orientations and the ground truth orientations (vertical axis), for increasing separation angles (horizontal axis, defined by the angle ϑ in $\tilde{A}(\vartheta)$, Eq. (6.23)). The dashed green line shows the total angular differences. For the considered example, with $\bar{\kappa} = 100$, the separation angle has to be greater than $45°$ for there to be two distinguishable dominant orientations. The difference w.r.t. Fig. 6.6 is due to the more isotropic tensors $\tilde{A}$ used to generate the multi-tensor data.
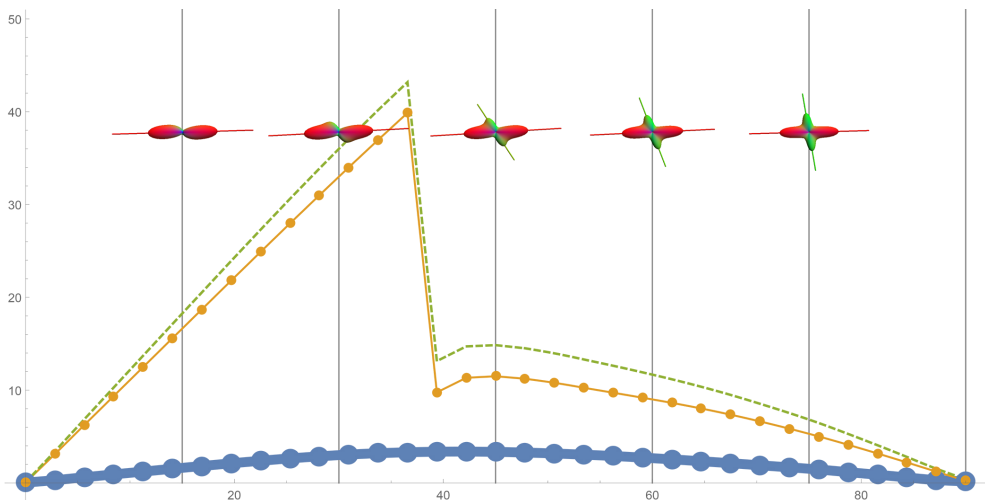
Figure 6.9: The effect of multi-compartment effects on the dominant orientations in the bODF. The orange and blue dots/lines show the angle between the measured orientations and the ground truth orientations (vertical axis), for increasing separation angles (horizontal axis, defined by the angle ϑ in $\tilde{A}(\vartheta)$, Eq. (6.23)). The dashed green line shows the total angular differences. For the considered example, with $\bar{\kappa} = 150$, the separation angle has to be greater than $37°$ for there to be two distinguishable dominant orientations.

(a) bODF with $\bar{\kappa} = 1$

(b) Peaks corresponding to (a)

(c) bODF with $\bar{\kappa} = 100$

(d) Peaks corresponding to (c)

(e) fODF obtained with CSD

(f) Peaks corresponding to (e)
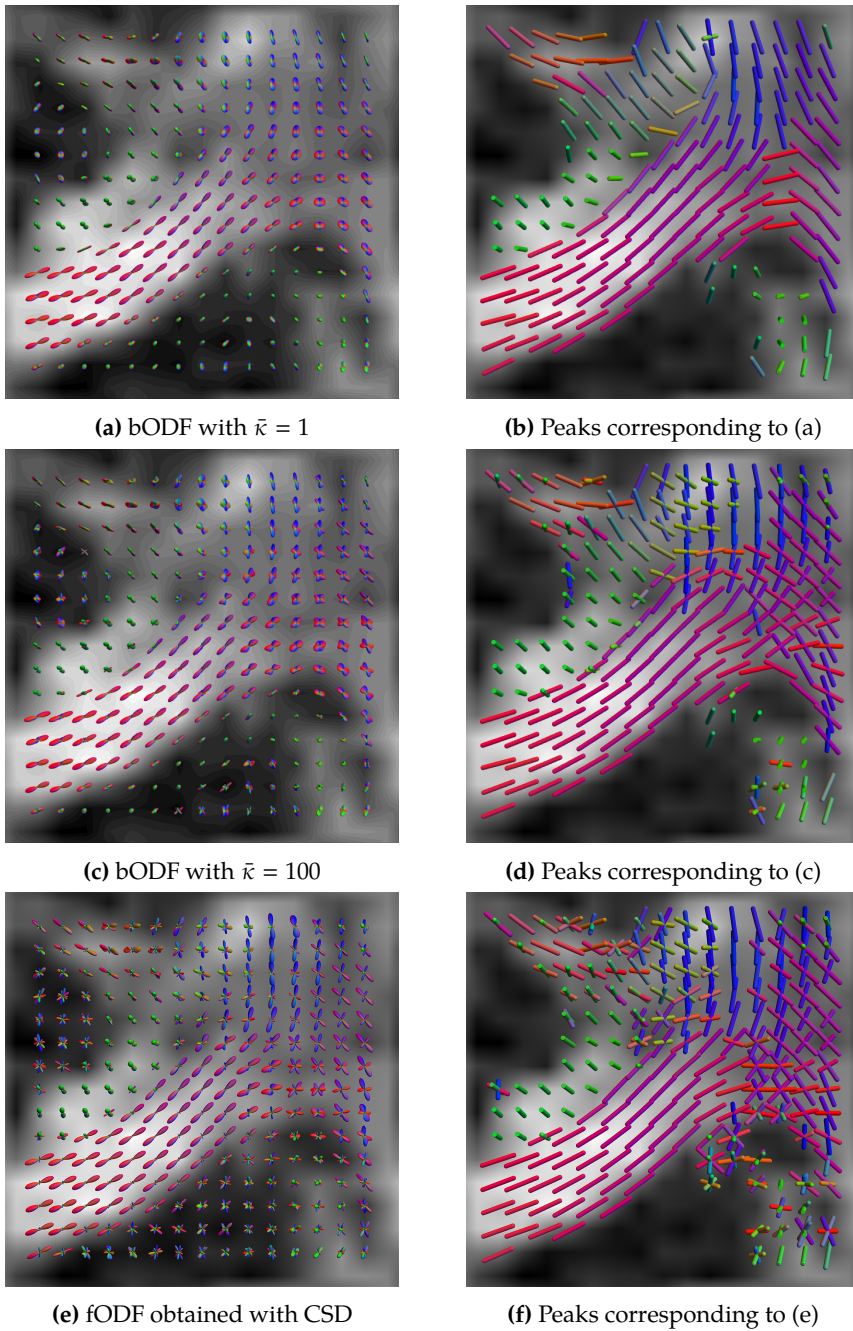
Figure 6.10: A comparison between the barrier ODF (top row, $\bar{\kappa} = 1$ and slice average $b_{\max}$ is $44\,\mathrm{s/mm^2}$, and middle row, $\bar{\kappa} = 100$ and average $b_{\max}$ is $3356\,\mathrm{s/mm^2}$) and the fiber ODF produced by CSD (bottom row) of a coronal slice of the HCP data set. The bODF figures illustrate how κ can be used to tune the angular resolution of the ODFs.

(a) bODF with $\bar{\kappa} = 100$ **(b)** Peaks corresponding to (a)
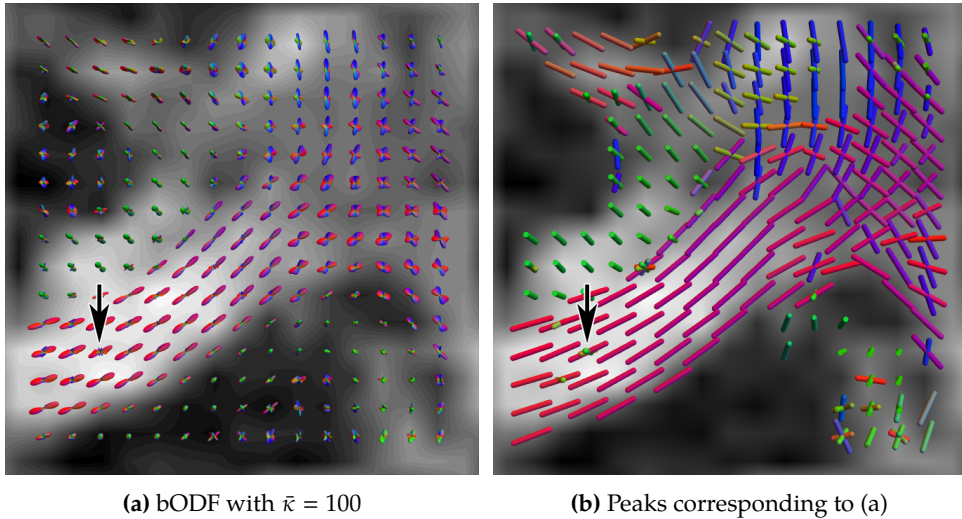
Figure 6.11: Part of a coronal slice of a HCP data set to be compared with Fig. 6.10 (middle row), showing barrier ODFs and the corresponding peaks obtained without regularization. All experiments in the main text are regularized, which remedies occurrences of the indicated artifacts in the barrier ODF caused by regions of mild non-convexity in the Finsler function, introduced in the spherical harmonic expansion (Eq. (3.43)).
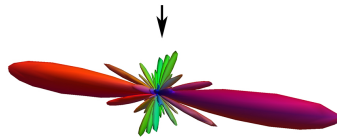
Figure 6.12: The barrier ODF corresponding to the level set of H_κ shown in Fig. 3.11. The oscillations indicated in the level set of the dual Finsler function produce the indicated high frequency noise in the barrier ODF.

6.3.2 Finslerian scalar measures

In Fig. 6.13 we show the Λ-measures for the F_κ (top row, $\bar{\kappa} = 100$) reconstructed from the coronal slice of a HCP data set (Fig. 2.2), together with the analogous measures from the DTI reconstruction based on the $b = 1000\,\text{s/mm}^2$ shell (bottom row). The contrast of the Finslerian measures is arguably better, but the intensity patterns between the corresponding measures are similar.

The non-Gaussianity NG introduced in Section 6.1.2.2 is shown in Fig. 6.14. This map shows high intensity in e.g. the corpus callosum, which can also be observed in radial and mean kurtosis maps [9, 178]. Note that non-Gaussianity does not say anything about the validity of DTI as a tissue model—it only provides information about DTI as a model of the signal. Non-Gaussianity is for example much lower in the centrum semiovale, which contains crossing fibers, than around the uni-oriented fibers of the corpus callosum.

As explained in Section 6.1.2.3, we can compute an estimate of the DTI tensor given a Finsler function F using Eq. (6.14). The comparison between the DTI tensor D derived from low b-value data, and the estimated diffusion tensor D_κ derived from high b-value data, is made in Fig. 6.15. Differences between the two are concentrated in gray matter, as can be seen in Fig. 6.16.

6.3.3 Riemannian measure glyphs

6.3.3.1 Artificial data

In Figs. 6.17 and 6.18 we show examples of the orientation-dependent fundamental tensor in relation to the barrier ODF for the artificial data used in Figs. 6.4 (left column) and 6.5 (top row). For the single-fiber (purely Riemannian) Finsler function shown in Fig. 6.17 the fundamental tensor is independent of orientation, and we obtain the same ellipsoidal level set visualization for each of the evaluated orientations. Naturally the resulting measure glyphs are spherical. For the crossing shown in Fig. 6.18, based on the artificial fourth order cumulants (Section 3.3.4), we find mildly prolate ellipsoids along the main orientations of the bODF, while perpendicular to these the ellipsoid is more oblate. The corresponding MD and FA glyphs are mildly anisotropic, and there appears to be a negative correlation between the bODF and the FA, and a positive correlation between the bODF and the MD.

6.3.3.2 Real data

In Fig. 6.19 and 6.20 we show the MD and FA measure glyphs for representative voxels in the corpus callosum and the centrum semiovale. These show quite conclusively that the previously purported relation between the MD and bODF is not as simple as suggested. The negative correlation with FA does seem to hold true, as confirmed by

(a) Λ_{min} **(b)** Λ_{max} **(c)** Λ_{rel} **(d)** Λ_{avg}

(e) λ_{min} **(f)** λ_{max} **(g)** FA **(h)** MD
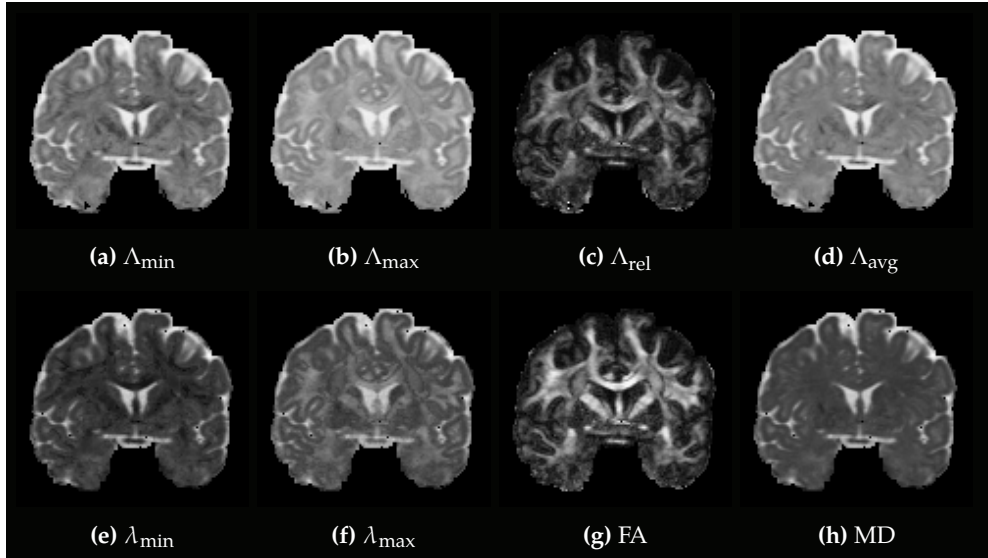
Figure 6.13: The top row shows the estimated smallest maximum displacement Λ_{min} (a), the estimated largest maximum displacement Λ_{max} (b), their relative difference Λ_{rel} (c), and the average estimated maximum displacement Λ_{avg} (d), with $\bar{\kappa} = 100$. These measures are the Finslerian analogues of the smallest (largest) eigenvector λ_{min} (λ_{maxf}), (e) and (f), fractional anisotropy FA (g), and mean diffusivity MD (h). The colors are scaled between the image minimum value (black) and image maximum value (white).
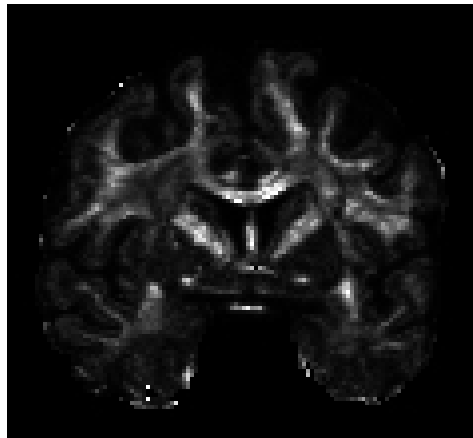
Figure 6.14: The non-Gaussianity NG computed for a single coronal slice of a HCP data set, cf. Eq. (6.13). Black indicates NG = 0, white indicates NG $\geq$ 0.1. The high intensity observed in e.g. the corpus callosum can also be seen in radial and mean kurtosis maps [9, 178].

(a) D **(b)** D_κ
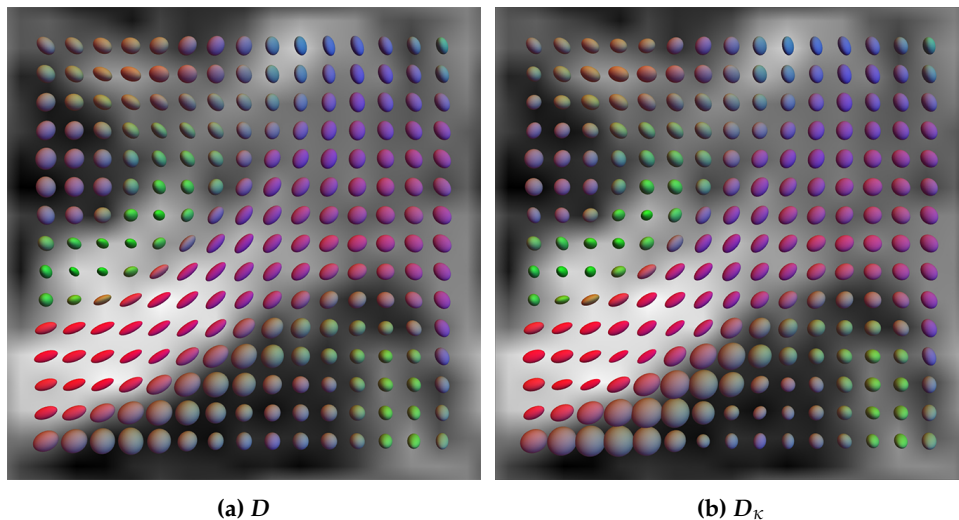
Figure 6.15: DTI ellipsoids derived from low b-value data (a), and ellipsoids derived from the estimated Finsler function F_κ using Eq. (6.14) (b).

(a) Non-Finslerianity   **(b)** (a), with highlighted white matter
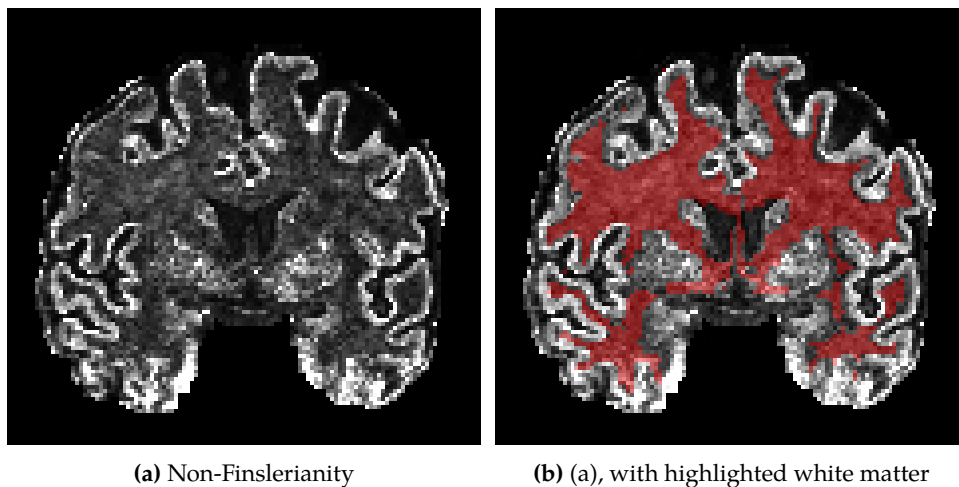
Figure 6.16: The non-Finslerianity measure (a) ($\bar{\kappa} = 100$), defined in Eq. (6.16) as the ratio between the log-Euclidean distance $d(D, D_\kappa)$ (Fig. 6.15) and the Frobenius norm of D. Black indicates that this ratio is zero, white indicates a ratio of one or greater. (b) is identical, but with a highlighted segmentation of the white matter to illustrate that high non-Finslerianity values are mainly found outside of white matter. Note that non-Finslerianity is also low in the ventricles, where diffusion is essentially Gaussian, and we thus expect no significant differences between low and high b-value data.

(a) Level set F_κ **(b)** bODF

(c)   **(d)**   **(e)**

(f) MD glyph   **(g)** FA glyph
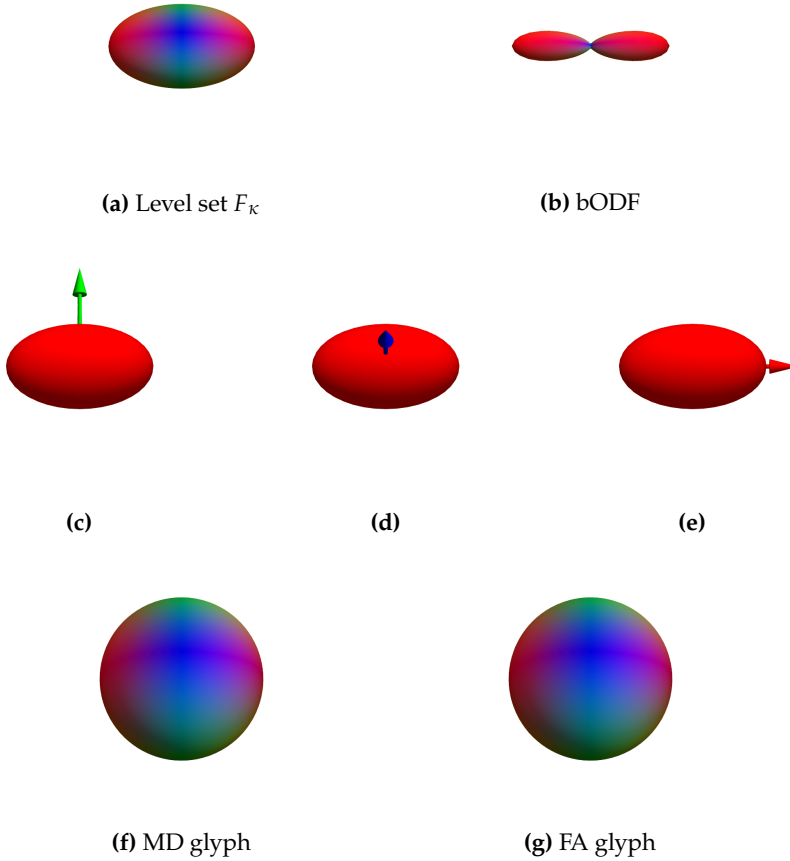
Figure 6.17: An example of Riemannian measure glyphs, computed for F_κ based on the second order data set defined by A with $\bar{\kappa} = 1$. (a) and (b) show the level set of F_κ and the corresponding bODF, (c), (d), and (e) show the ellipsoids derived from the fundamental tensor for the indicated orientations, and (f) and (g) show the mean diffusivity (MD) and fractional anisotropy (FA) measure glyphs.

(a) Level set F_κ

(b) bODF

(c)

(d)

(e)

(f) MD glyph

(g) FA glyph
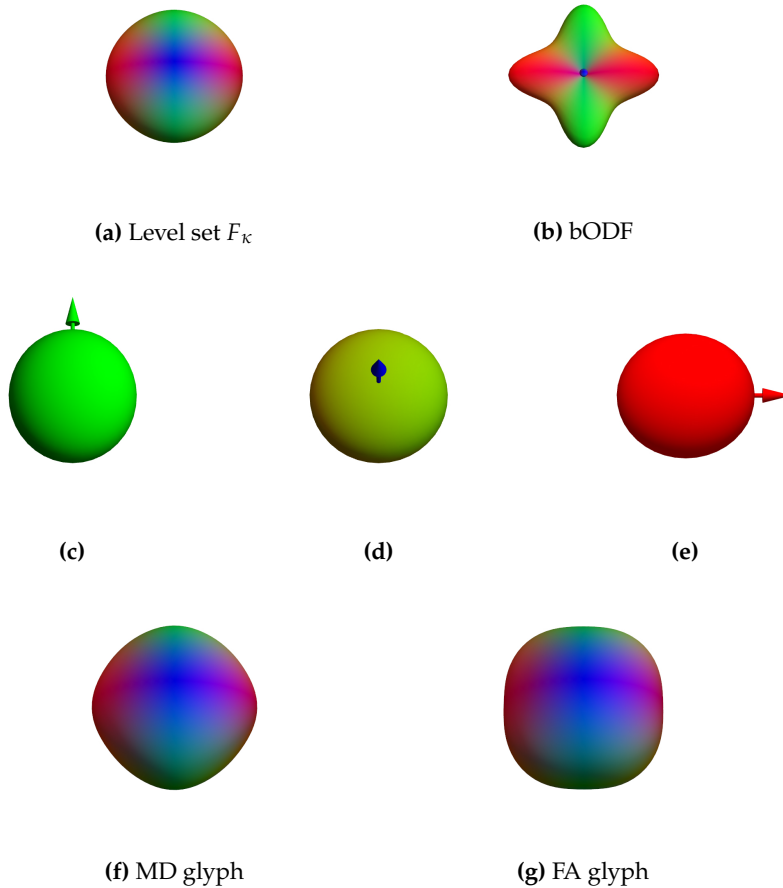
Figure 6.18: An example of Riemannian measure glyphs, computed for F_κ based on the second order data set defined by A with $\bar{\kappa} = 100$. (a) and (b) show the level set of F_κ and the corresponding bODF, (c), (d), and (e) show the ellipsoids derived from the fundamental tensor for the indicated orientations, and (f) and (g) show the mean diffusivity (MD) and fractional anisotropy (FA) measure glyphs.

(a) Level set F_κ **(b)** bODF   **(c)** MD glyph   **(d)** FA glyph
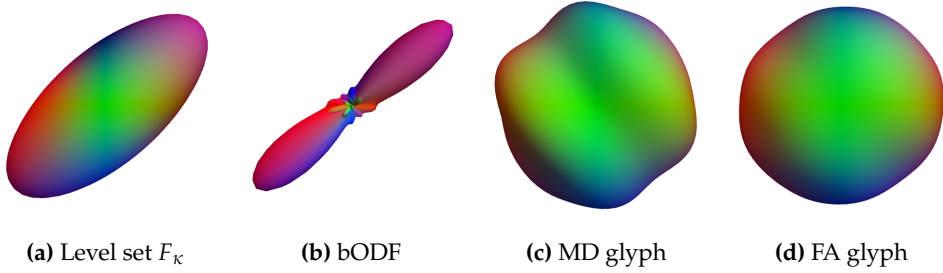
Figure 6.19: The level set of the estimated Finsler function ($\bar{\kappa} = 100$) and the barrier ODF for a single voxel in the corpus callosum of the HCP data set, together with the mean diffusivity (MD) and fractional anisotropy (FA) measure glyphs. For purely Gaussian diffusion, the measure glyphs are perfectly ellipsoidal.

(a) Level set F_κ **(b)** bODF   **(c)** MD glyph   **(d)** FA glyph
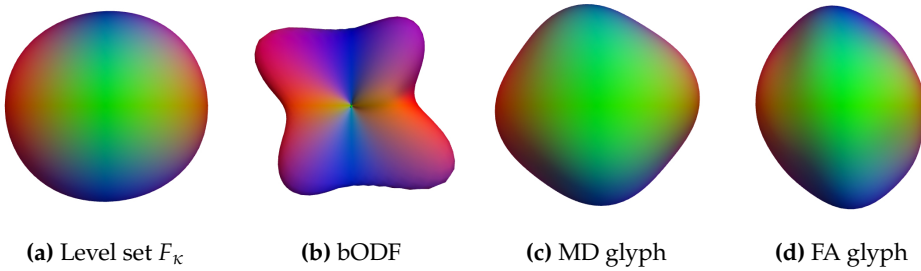
Figure 6.20: The level set of the estimated Finsler function ($\bar{\kappa} = 100$) and the barrier ODF for a single voxel in the centrum semiovale of the HCP data set, together with the mean diffusivity (MD) and fractional anisotropy (FA) measure glyphs.

preliminary experiments in real data. We recognize a clear drop in the FA of fundamental tensors evaluated at orientations perpendicular to expected fiber pathways orientations, see Fig. 6.21. Still, these differences are very subtle, and the barrier-weighted Riemannian measures do not differ significantly from the DTI-derived FA and MD (Fig. 6.22).

6.4   Discussion

As tissue structure is only indirectly inferrable from the observed diffusion, we have to be particularly careful when relating the two. For example, it is not always clear whether particular tissue properties can be reliably deduced from the data. But within the set of properties that actually have a measurable and resolvable effect on the diffusion-weighted signal, we are further limited by our models. A model generated from a very restrictive set of assumptions may be biased by effects that are not explicitly taken into account in the assumptions [175]. Because the geometries considered in this thesis are limited to position- and orientation-dependent features—more complex features will require more complex geometries—our results may be skewed
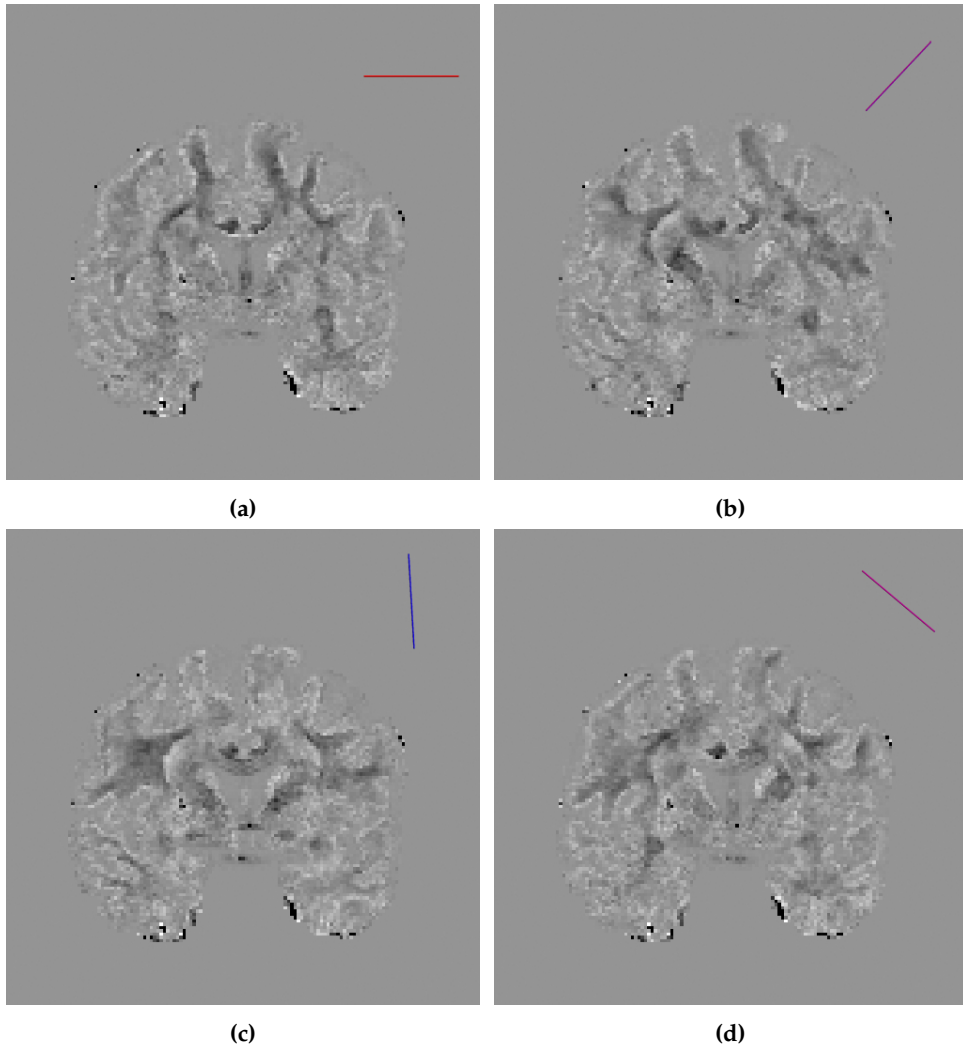
Figure 6.21: The orientation-dependent FA in a coronal slice of a HCP data set (Fig 2.2). The figures show the difference between the standard FA computed from DTI tensors, and the FA of the fundamental tensor associated with the orientations indicated at the top right in each image. Dark voxels correspond to a relative drop in FA, and bright voxels imply a relative increase. By looking at e.g. the cingulum bundles (pointing out of the slice) and the large u-fibers connecting the top of the left and right hemispheres through the corpus callosum, we may discern a negative correlation between the orientation-dependent FA and the axonal ODF.

(a) Barrier-weighted FA   **(b)** Barrier-weighted MD
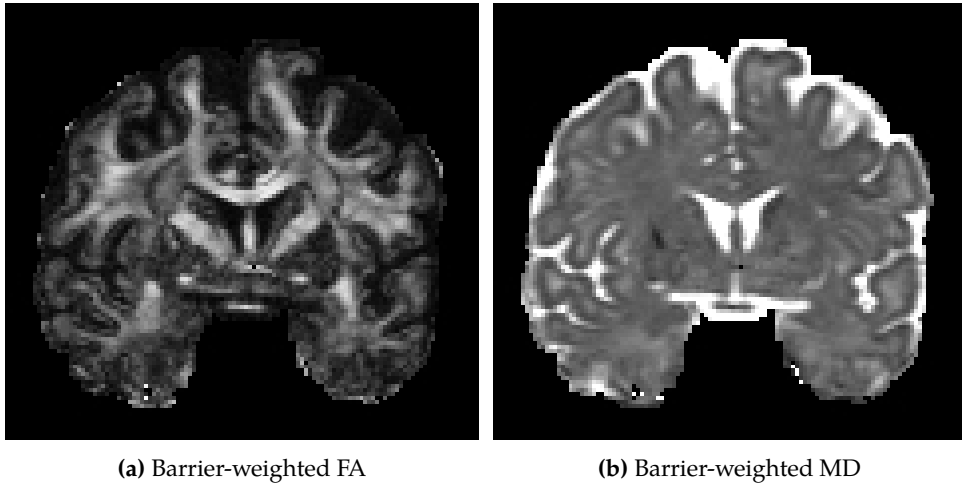
Figure 6.22: The barrier-weighted fractional anisotropy (FA) and mean diffusivity (MD). For the relatively small κ used in this experiment, these are strongly correlated with the standard DTI-based measures.

if for example there is a significant curvature-component discernable in the average spin's motion, e.g. a tendency to diffuse along trajectories with a specific curvature. Keeping this in mind, we briefly review some of the prospects and limitations of the proposed descriptors.

6.4.1 The barrier orientation distribution function

The barrier orientation distribution function (barrier ODF or bODF) introduced in Section 6.1.1 is a clear example of a new approach to an old problem, inspired by the introduction of the geometrical framework. Finding the distribution of fibrous structures in biological tissue seems like a highly specialized problem, but when framed as a geometrical problem we quickly recognize an overlap with problems in stochastic and convex geometry [145, 190, 196, 298, 320, 324]. The same problem has been considered in the context of microscopy in e.g. metallurgy [196], revealing the synergetic potential of the geometrical description.

 The barrier ODF represents the correlation of impedances in the underlying transport process along a given orientation. We have considered both one- and two-dimensional barriers, and have shown that these do not produce fundamentally different orientation distribution functions. Though their dimensionality does not manifest in the bODF, it should be noted that the correlation dimension is (theoretically) observable in diffusion-weighted MRI. As described in the work of Novikov et al. [248], this property becomes apparent in the long diffusion time limit of D (whose time dependence we do not consider in this thesis). Like all descriptors proposed in this chapter, the barrier ODF Ψ_1 can be interpreted in terms of the local Finsler

geometry. The relevant relation in this instance is Eq. (6.2), from which it follows that space is compressed proportional to the barrier density encountered by a typical particle moving along a given orientation.

The barrier ODF is defined as the inverse sine transform of the Finsler function, which can be implemented similarly to existing transforms in diffusion MRI such as those in Q-ball and fiber ball imaging [179, 347]. Alternative implementations are possible [196], but are not considered in this work. The sine transform can be viewed as a mapping from (sufficiently smooth centrally symmetric) convex bodies to point-wise non-negative functions on the sphere, and as such it plays a major role in e.g. convex and stochastic geometry [298, 324]. In diffusion MRI this relation allows us to map arbitrary ODFs to convex sets, in principle enabling the definition of a Finsler manifold from any model that provides a field of ODFs. We can then apply any geometrical approach—like the PDE-based processing of Pennec et al. [270] or geodesic interpolation [28, 128, 270]—to these ODF fields.

Initial experiments comparing the barrier ODF to the widely accepted fiber orientation density function (fODF) obtained with constrained spherical deconvolution (CSD) [342], Fig. 6.6, suggest that for crossings of 45° or more (the approximate resolution limit of CSD) the two ODFs conform with an error of less than 5°. For fiber pathways that cross at more acute angles the differences are more pronounced. Note that these differences cannot be interpreted as errors of either the barrier or fiber ODF—depending on which set of assumptions is more appropriate for a given application, either CSD results or bODF results will be more accurate. The reconstruction of fiber pathways based on the barrier ODF can be done using the same techniques as for the fiber ODF [33, 99, 287].

In Fig. 6.10 we show the bODF and fODF side-by-side in real data, where similar global trends can be recognized. This figure also illustrates again the effect of κ on the angular resolution. For small κ the barrier ODF is essentially limited to a single peak, while higher values allow the resolution of more complex pathway configurations (see also Fig. 6.6 and Figs. 6.7, 6.8, and 6.9). An extensive evaluation of the barrier ODF is a subject of future work.

6.4.2 Finslerian scalar measures

The scalar measures we described in Section 6.1.2 are intended to give a quick understanding of what the diffusion in a voxel looks like, and to what extent it can be modeled as either a Riemannian or a Finsler manifold. First, we defined a number of scalar measures in analogy to existing DTI and non-Gaussianity measures. These perform as expected, as seen in Figs. 6.13 and 6.14, but have to be interpreted somewhat differently. The merit of these maps compared to (or in conjunction with) existing ones, e.g. as biomarkers, remains to be investigated. Note that there exist alternative (geometrical) non-Gaussianity measures that we have not considered here. In earlier

work, Florack et al. [132] proposed for example a non-Gaussianity measure based on the Cartan one form C_i:

$$C_i = \sum_{j,k} \frac{1}{4} g^{jk} \frac{\partial^3 F^2}{\partial y^i \partial y^j \partial y^k}. \tag{6.24}$$

A Cartan scalar field can be then defined for example as

$$\sum_{i,j} \int_{I_x M} g^{ij}(\boldsymbol{x}, \boldsymbol{y}) C_i(\boldsymbol{x}, \boldsymbol{y}) C_j(\boldsymbol{x}, \boldsymbol{y}) \, \mu_x(\mathrm{d}\boldsymbol{y}), \tag{6.25}$$

which is identically zero if and only if F describes a Riemannian manifold.

We have further defined a new scalar measure called non-Finslerianity, also briefly discussed in Section 3.5.1. Barring partial volume effects and spatial variability in the medium characteristics, this measure is designed to show the extent to which a measured stochastic process is Finslerian in nature. Where this measure is low, we can be reasonably confident that the Finslerian framework adequately represents the information in the signal. If the measure is high, and partial volume effects are negligible, this likewise implies that the Finslerian framework does not capture the full information, meaning that a more complex framework would be necessary. We discussed such extensions of the Finslerian case in Section 3.5.3. Our initial results, Fig. 6.16, suggest that both partial volume effects (as defined in Assumption 1) and non-Finslerian characteristics are small in large parts of the white matter.

So far we have neglected partial volume effects, but if we assume instead that the measured processes are truly Finslerian, then the non-Finslerianity can actually be interpreted as a (lower bound) measure of these partial volume effects. In that case we could for instance add extra compartments to the model, with weights depending on the non-Finslerianity measure, or, preferably, move towards higher resolution data. In reality the true situation is likely somewhere in the middle, with the non-Finslerianity measure clouded by both effects to some extent.

6.4.3 Riemannian measure glyphs

The last of the descriptors proposed in this chapter is the Riemannian measure glyph, which is defined by one of various scalar measures computed for the orientation-dependent fundamental dual tensor associated with the Finsler function. This tensor describes differential structure of the figuratrix around a given orientation, see Fig. 5.2, and thus describes the behavior of particles that approximately move along that orientation.

Figs. 6.19, 6.20, and 6.21 show an apparent decrease in the orientation-dependent fractional anisotropy (FA) perpendicular to the expected fiber orientation—-similar to the relation between the barrier ODF and the indicatrix. The important conclusion

(a)   **(b)**
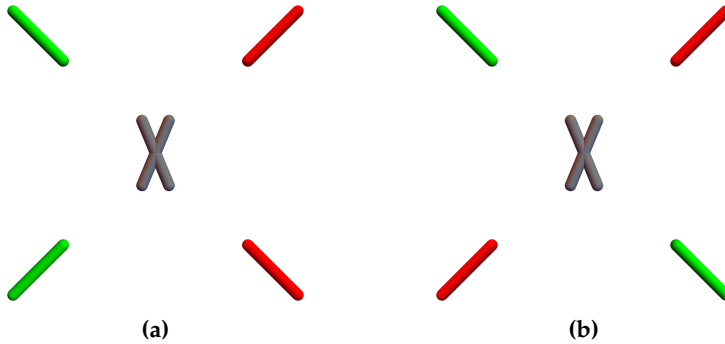
Figure 6.23: The tubes represent local orientations colored according to the underlying fiber pathways these orientations represent. The gray crossing in the center represents both the green and the red pathways, and the figures show a bending configuration (a) and a crossing configuration (b). A large segment of current tractography algorithms are based on the local peaks of ODFs, which do not provide sufficient information to resolve situations such as the one depicted here. Deterministic algorithms will at most resolve one of these configurations accurately, while probabilistic algorithms can at best produce both configurations as options. Orientation-dependent measures may help in identifying the correct paths, by providing additional information for each orientation that can be used to group orientations that likely represent the same pathways, inspired by various microstructure informed tractography approaches that have recently been proposed [22, 78, 79, 151].

that we can draw from this, is that the *local* shape of the diffusion profile appears to provide information about the underlying fiber pathways. Though observable effects at clinical b-values are subtle, this suggests that we may be able to extract fiber-specific (tractometric [34, 76, 82, 172, 173]) features from the data, that could for example be used to help correctly resolve ambiguous pathway configurations, cf. Fig. 6.23. This is particularly relevant in light of recent work [104, 225] that reports an alarmingly high percentage of false positive tracts in all existing tractography algorithms, which may be largely due to this kind of ambiguities.

We have not focused in this chapter on which tensor features specifically could be relevant in different contexts. The introduction of these measures should probably be application-driven, e.g. radial diffusivity (cf. Section 2.3.1) may capture some dispersion index, that could be useful as a biomarker in neurodegenerative diseases [72, 382]. We also note that the main eigenvector of the fundamental tensor is in general not aligned with the orientation supplied to the fundamental tensor, and could provide further information about the local structure. The main eigenvector has for example been used by Astola et al. [17] in a streamline-based tractography algorithm.

Keeping in mind the notion that the Riemannian measure glyphs are expected to be characteristic of particular fiber pathways, we could expect that the barrier-weighted measures might for example diminish 'artificial' drops in FA around complex fiber configurations. At the current κ however, this effect is minor—a close comparison between the barrier-weighted and the DTI-based measures only reveals a slightly improved contrast in the former, visible e.g. near the cortical gray matter.

Chapter 7

Curve

Inference of Spatial Relations

Based in part on:

A. Fuster, T. Dela Haije, A. Tristán-Vega, B. Plantinga, C.-F. Westin, and L. Florack. "Adjugate diffusion tensors for geodesic tractography in white matter". In: *Journal of Mathematical Imaging and Vision* 54.1 (2016), pp. 1–14. ISSN: 0924-9907, 1573-7683. DOI: `10.1007/s10851-015-0586-8`.

T. C. J. Dela Haije, P. Savadjiev, A. Fuster, A. Tristán-Vega, R. Schultz, R. Verma, L. M. J. Florack, and C.-F. Westin. *Structural connectivity analysis using Finsler geometry*. In preparation, preprint available upon request.

In this chapter we discuss how Riemannian and Finsler geometry, and specifically the related geodesic tractography, can be levied to analyze structural connections between different brain regions. We thus move from the purely local information discussed in the previous chapter, to information found by combining data acquired at different points in the brain. Though our focus is entirely on neuroimaging applications, we note that geodesics are of fundamental importance in geometrical frameworks.

In geodesic tractography white matter is represented by a manifold, as covered extensively in Part I, with the added conjecture that neural fibers coincide with certain geodesic curves. In this way the problem of tractography[26, 33, 68, 69, 99, 180, 242, 287, 299, 310], the general term for reconstructing fiber pathways from diffusion MRI data, becomes one of finding geodesics. This is attractive from a practical point of view, as it obviates the need for ad hoc stopping and bending criteria necessary in traditional fiber tracking algorithms. It also tends to be more robust to noise compared to other methods, and finally, it has the conceptual advantage that Riemannian and Finsler geometry provide well-understood and powerful theoretical machinery facilitating algorithmics.

On the other hand, the relationship between geodesics and neural tracts is not as clear as it is in alternative approaches. Geodesics can be computed between any two points in a (compact) manifold, while we know for a fact that not all regions in the brain are interconnected. This means that we will need to know in advance which points are connected, and/or determine some additional criteria that can be used to select only those geodesics that represent fiber pathways. Furthermore, there are arguments against the assumption that neural fibers should *in general* coincide with geodesics, such as the observation that standard Riemannian geodesics tend to move away from the centerline in bending fiber bundles. Recent work by Hao et al. [163] shows that this issue can be addressed by locally scaling the standard Riemannian metric, and along the same lines we review in this chapter the impact of using a scaled geometry (Chapter 2) for Riemannian tractography.

Finsler geometry was invoked in tractography works by Pichon et al. [271] and Melonakos et al. [227, 228], to take advantage of the additional information present in high gradient strength diffusion-weighted data that is not taken into account in Riemannian geometry. The authors computed shortest geodesic tracts based on different ad hoc relations with the diffusion MRI signal, which have since been used in other works [38]. A different definition of the Finsler function was employed by Astola et al. [17], who illustrate some more technical applications of Finsler geometry, while the work of Sepasian et al. [307] extends shortest geodesic tractography to allow multiple geodesic connections between points. Finally, the first geodesic connectivity analyses were performed by de Boer et al. [38], which also marks the first time Finsler geometry was used for a group analysis of diffusion MRI data. In geodesic connectivity, shape and diffusivity properties of geodesics are combined to get an

idea of the connection strength between regions. These results can be interesting, but a significant downside is the lack of publicly available tools that can perform these types of analyses. For this reason, along with the technical content of even introductory works on Finsler geometry, Finsler-based methods have been inaccessible to the majority of diffusion MRI researchers. In this chapter we aim to provide an informal introduction to the core concepts involved with geodesic tractography and connectivity computations, based on Antonio Tristán-Vega's Riemannian tractography module (RTM, `nitrc.org/projects/riemantract`) and the recently released Finsler tractography module (FTM, available at `github.com/tomdelahaije/fcm`, adapted from `nitrc.org/projects/finslertract`).

In Section 7.1 we describe the ideas behind geodesic tractography, before describing the used tractography and connectivity algorithms in Section 7.2. We look at Riemannian geodesics in a tumorous brain using different metrics, and as a practical illustration of Finsler-based analyses we compute the connectivity for a number of well-known major fiber bundles in a high resolution public data set. We also discuss how the connectivity algorithm might be used in group studies. As a proof of concept, we use our method to study group differences in autism spectrum disorder data with network-based analysis techniques. The results of these experiments are presented in Section 7.3. Finally we discuss some strengths and shortcomings of the approach in Section 7.4. Parts of the results in this chapter were presented at the CDMRI workshop [140, 142].

We rely on different implementations for practical computations in Riemannian and Finsler geometry, as the Riemannian-constrained versions of algorithms are typically more efficient than their Finslerian counterparts. The principles behind these algorithms are however the same, and so we will explain the ideas behind geodesic tractography and connectivity analyses only for the Finsler case.

7.1 Theory

7.1.1 The Finsler function as a cost function

Finsler geometry provides a means to measure distances on Finsler manifolds, recall Section 3.1. The distance between two points in a Finsler manifold is defined similar to the standard Euclidean distance, namely in terms of the length of the shortest curve connecting the two points. The length $\mathcal{L}_F(C)$ of a curve $C : [0, L] \to M$ is still a sum over infinitesimal line elements dC, but the associated length of each line element is now weighted with the Finsler function F depending on both its position and its orientation

$$\mathcal{L}_F(C) := \int_0^L F\left(C(t), \frac{\dot{C}(t)}{\|\dot{C}(t)\|}\right) \|\dot{C}(t)\| \, dt, \tag{7.1}$$

where $\dot{C}(t) = \frac{dC}{dt}(t)$. As the Finsler function F is (positively) homogeneous of degree one in its second argument (recall Section 3.1.1), it follows that Eq. (7.1) reduces for the isotropic Finsler function $F(C(t), \dot{C}(t)) = \|\dot{C}(t)\|$ to the Euclidean length of the curve,

$$\mathcal{L}_E(C) := \int_0^L \|\dot{C}(t)\| \, dt. \tag{7.2}$$

Homogeneity further implies that Eq. (7.1) is equivalent to Eq. (3.9), and gives

$$F\left(C(t), \frac{dC}{dt}(t)\right) \, dt = F(C(t), dC(t)), \tag{7.3}$$

which clarifies the interpretation of F as a function acting locally on an infinitesimal dC. At the same time we note that the homogeneity of F (Section 3.1.1) guarantees that the length $F(C(t), dC(t))$ associated to dC is determined solely by its orientation—not by its magnitude—and that $\mathcal{L}_F(C)$ is independent of the (proper) parametrization of C. As a technical aside we observe that homogeneity also implies that F is strictly speaking defined only when $\|\dot{C}(t)\| \neq 0$ for all t, and for this reason we have to assume a parametrization of C that avoids this issue. To simplify the following discussion we assume without loss of generality that $\|\dot{C}(t)\| = 1$, i.e., we assume that $\dot{C}(t)$ is an element of the sphere $S^2 \subset T_{C(t)}M$.

As discussed in Chapter 3, the Finslerian framework for diffusion MRI models the brain as a Finsler manifold, by deriving a Finsler norm from diffusion-weighted data. Different choices for the Finsler function have been proposed in the context of geodesic tractography based on the well-known correlation between the local amount of diffusion in a certain direction, and the large-scale structural orientation of white matter [30, 242]. By deriving (or, in this chapter, assuming) a correspondence between the Finslerian length of a curve and the amount of diffusion along a curve, we can leverage a rich set of Finsler geometrical tools for the analysis of diffusion data.

Bearing this in mind, the Finsler norm is generally defined such that some measure of diffusivity (e.g. a diffusion orientation distribution function [347], or dODF) at a given point and in a certain direction, is inversely related to the associated length. In other words, we have that a large diffusivity at a point $x \in M$ along a vector $y \in T_x M$, corresponds to a small norm $F(x, y)$. This leads to the useful alternative viewpoint of the Finsler norm as a kind of cost function. If we consider a displacement in direction y as a parameter that can be controlled, then F can be interpreted as associating a high cost to movement in a direction with low diffusivity, and a low cost to movement in directions of high diffusivity.

7.1.2 Geodesics

A prime example of analysis tools made available by the geometrical framework are the geodesics introduced in Sections 2.1.3 and 3.1.3. In diffusion MRI, geodesics are typically regarded as curves along which one encounters, in some sense, optimal diffusivity. More specifically, geodesics connecting two given points are those curves for which the length $\mathcal{L}_F$ is (locally) minimal, and can thus be viewed as the Finslerian analogue of 'straight lines'. The existence of a geodesic between any two points in the Finsler manifold is guaranteed [21, Chapter 6.6], which means that we can find optimal connections between any two points or regions of interest. For now we will assume that the shortest geodesic between two points, called the minimal geodesic, is uniquely defined.

In practice we determine minimal geodesics using a fast-sweeping algorithm [228] based on the principle of optimality, which states that given a unique minimal geodesic $C : [0, L] \rightarrow M$, the geodesic segment between the points $C(a)$ and $C(b)$, $a, b \in [0, L]$, is necessarily identical to the geodesic between these two points. If we write $\mathcal{L}_F^*(x)$ for the shortest (geodesic) distance from a point $x \in M$ to a seed region $B \subset M$ relative to F, then

$$\mathcal{L}_F^*(x) := \inf_C \left\{ \mathcal{L}_F(C) \mid C(0) = x, C(L) \in B \right\}$$

$$= \inf_C \left\{ \mathcal{L}_F^* \left(C(T) \right) + \int_0^T F\left(C(t), \dot{C}(t) \right) \, \mathrm{d}t \,\middle|\, C(0) = x \right\} \tag{7.4}$$

for all $T \in (0, L)$, cf. 7.1. From the Taylor expansion

$$\mathcal{L}_F^* \left(C(T) \right) = \mathcal{L}_F^*(x) + y_x \left(\mathcal{L}_F^* \right) T + \mathcal{O} \left(T^2 \right), \tag{7.5}$$

where y_x denotes the directional derivative along $y \in S^2$ at $x \in M$ (recall Section 1.3) and where we use the shorthand notation $x = C(0) \in M$ and $y = \dot{C}(0) \in S^2$, we find the Hamilton–Jacobi–Belmann equation in the limit $T \rightarrow 0$:

$$\inf_{y \in S^2} y_x \left(\mathcal{L}_F^* \right) + F(x, y) = 0. \tag{7.6}$$

Together with the initial condition

$$\mathcal{L}_F^*(x) = 0 \text{ for } x \in B, \tag{7.7}$$

repeated application of Eq. (7.6) allows us to compute the complete $\mathcal{L}_F^*$ map for all $x \in M$. Riemannian geodesics are computed in the exact same way, by simply defining an appropriate Riemannian Finsler function (see Section 3.1.1). Additional details on the algorithm follow in Section 7.2.1.4.
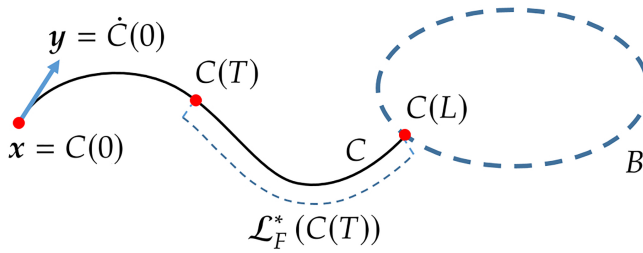
Figure 7.1: The principle of optimality states that segments of the minimal geodesic between two points are themselves geodesic. The black curve represents the optimal curve C (minimal geodesic) connecting the point x to the seed region B, i.e., the curve C that minimizes the Finslerian length functional $\mathcal{L}_F$. The distance $\mathcal{L}_F^*(x)$ from x to B is defined as the length of the optimal curve C that connects the two. If $\mathcal{L}_F^*(C(T))$ is known for $C(T)$ near x, the principle of optimality allows us to compute $\mathcal{L}_F^*(x)$ by solving the Hamilton–Jacobi–Bellman equation (Eq. (7.6)). As $\mathcal{L}_F^*(x) = 0$ for all $x \in B$, repeated application of Eq. (7.6) allows us to compute $\mathcal{L}_F^*$ for all $x \in M$.

This method is relatively fast, but as explained before it is limited in that it only finds the shortest geodesic out of the possibly many geodesics connecting two given points [305]. We will assume for now that the relevant information is captured by the minimal geodesic. More information on fast-sweeping algorithms can be found in the references [191, 192, 386].

Returning to the context of diffusion MRI, the optimal diffusivity along geodesics is made more precise by the inverse relation discussed in Section 7.1.1. With the cost function interpretation of F, we note that geodesics correspond to curves along which the accrued cost is minimal. Because of the inverse relation between the cost and the local diffusivity, geodesics thus minimize movement in directions of low diffusivity. Additionally, we see that the Finslerian length $\mathcal{L}_F$ of a curve approximately corresponds to the average reciprocal diffusivity along the curve.

7.2 Methods

Geodesics have been used as a tool in tractography, based on the hypothesis that (some small subset of) geodesics between two points coincide with the physical connections between them. This approach has certain practically advantageous features; it relies on the full diffusion information available (unlike deterministic streamline tractography for instance), and has essentially no parameters that have to be tuned. The `riemanntract` and `finslertract` packages available on `nitrc.org` can be used to perform geodesic tractography with the different metrics.

Geodesic connectivity analysis is based on the hypothesis that curves of optimal diffusivity between two points contain information on the likelihood of a structural connection. This does not necessarily imply that such curves trace existing physical

connections. In this section we discuss the options for the Riemannian metrics compared in Section 7.3, and the options for the Finsler function in the FTM as well as two available geodesic-based path and connectivity measures that have been studied in a number of recent works.

7.2.1 Implementation

7.2.1.1 The Riemannian metric

The standard choice $g = D^{-1}$ for the Riemannian metric [86, 141, 218, 254, 270], Eq. (2.15), directly relates the metric to the diffusivity along a given orientation. Geodesics obtained through this metric thus represent the most likely paths along which a spin undergoing Brownian motion moves between two points in the manifold. As an alternative we consider the scaled metric discussed in some detail in Chapter 2,

$$\tilde{g} := \det(D) \, D^{-1}. \tag{7.8}$$

We will refer to the scaled metric as the adjugate, since $\det(D) \, D^{-1} = \operatorname{adj} D$, and refer to the standard metric as the inverse. The adjugate metric was originally proposed in works by Fuster et al. [141, 142], and has since been considered in a number of works [299, 306].

7.2.1.2 The Finsler function

The choice for the Finsler function F in terms of the diffusion signal is the primary degree of freedom in the Finslerian framework, and determines to a large extent how geometrical features such as geodesics can be interpreted. The most common choice for F found in the literature was proposed by Melonakos et al. [227] and is given by

$$F_{\mathrm{old}}(x, y) := \left(\frac{S(x, y)}{\Psi_{\mathrm{d}}(x, y)} \right)^3, \tag{7.9}$$

with $y \in S^2$ and S the diffusion MRI signal acquired on a fixed b-value shell, and where Ψ_{d} is the diffusion ODF defined in terms of the Funk–Radon transform (e.g., [347], recall also Eq. (6.7)). The power 3 is used as a type of sharpening. The choice $F = F_{\mathrm{old}}$ has been shown to produce reasonable tractography results e.g. near the cingulum bundle [228], and it has been used in the literature by de Boer et al. [38]. However, it lacks a clear relation with existing metrics in the Riemannian framework—we expect the Finsler function in case of Gaussian diffusion to reduce to a Riemannian norm compatible with the structures derived in Chapter 2.

As an alternative we postulate a new choice for F:

$$F_{\text{new}}(x, y) := \frac{\text{MD}(x)}{\Psi_{\text{sa}}(x, y)}, \tag{7.10}$$

where MD is a generalization of the mean diffusivity used in diffusion tensor imaging (DTI, Chapter 2) [27] defined as the S^2-average apparent diffusion coefficient [263], and Ψ_{sa} is the solid angle dODF [2]. For purely Gaussian diffusion, F_{new} corresponds to (the cost function of) a sharpened version of the Riemannian metric given by the adjugate diffusion tensor [141, 142], cf. Chapter 2 and Section 7.2.1.1.

The relation between F_{new} and the adjugate $\tilde{g}$, Eq. (7.8), can be recognized in the exact expression for the normalized ODF [210]

$$\Psi_{\text{d}}(x, y) = \frac{\text{MD}(x)}{\sqrt{y^T \cdot \det D(x) \, D^{-1}(x) \cdot y}}, \tag{7.11}$$

where MD is the mean diffusivity, $\det D$ is the determinant of the diffusion tensor, $y \in S^2$ is a direction unit vector, and D^{-1} is the inverse of the diffusion tensor, recall Chapter 2. Next we note that the denominator in this equation is the norm of direction vector y in the Riemannian space equipped with the metric $\tilde{g}$:

$$\|y\|_{\tilde{g}} = \sqrt{y^T \cdot \det(D) \, D^{-1} \cdot y} = \sqrt{\sum_{i,j} \tilde{g}_{ij} y^i y^j}. \tag{7.12}$$

Combining Eqs. (7.11) and (7.12) we find

$$\|y\|_{\tilde{g}} = \frac{\text{MD}(x)}{\Psi_{\text{d}}(x, y)}, \tag{7.13}$$

and replacing the Ψ_{d} with the sharper Ψ_{sa} produces the proposed F_{new} metric (Eq. (7.10)).

Remark 13. *The work presented in this chapter precedes the results from Chapter 3 where we derived a fundamental relation between diffusion MRI data and Finsler geometry. In the remainder of this chapter we will focus on the definitions employed in the literature as described in this section. Note that the ad hoc Finsler functions of this chapter do not necessarily satisfy the defining properties of a Finsler function as given in Section 3.1.1, and an additional justification of their validity as input to any geodesic tractography algorithm is required, c.f. Melonakos et al. [227, App. A].*

7.2.1.3 Path measures

Geodesic-based connectivity analysis combines a variety of curve shape measures with measures derived from the diffusion signal along the curve into a single path measure. This shape measure could be as simple as the Euclidean length, while more advanced shape measures such as local curvature and torsion are possible but used less frequently. The diffusion signal is encoded in the Finslerian length, representative of the total diffusivity, or the Finslerian speed, representative of the local diffusivity, or in a set of statistical measures such as the quantiles, mean, and standard deviation of the Finsler function evaluated along the geodesic. The heuristic definition of connectivity in terms of path measures is discussed in Section 7.2.2.3.

The basic path measure used in the Finsler tractography module is defined as

$$C_{\mathrm{avg}}(C) := \frac{\mathcal{L}_F(C)}{\mathcal{L}_E(C)}, \tag{7.14}$$

which is the Finslerian generalization of the most commonly used measure in Riemannian geodesic connectivity analysis [18, 219]. This measure can be interpreted as the average cost incurred along the geodesic, which is expected to be low for curves between two well-connected regions.

The second available path measure is defined as the largest local cost along a geodesic, given by

$$C_{\max}(C) := \max_t F\left(C(t), \frac{\dot{C}(t)}{\|\dot{C}(t)\|}\right). \tag{7.15}$$

The $C_{\max}$ path measure was originally proposed in the Riemannian setting by Pechaud et al. [269]. Although this measure might be expected to be very sensitive to noise, it should be noted that it is based on the same (intrinsically smooth) geodesics as the C_{avg} measure, and is thus as stable as C_{avg}. The $C_{\max}$ measure highlights geodesics that have continuously strong diffusivity along their paths, in contrast to the C_{avg} measure for which a locally weak diffusivity might be offset by very strong diffusivities further along the geodesic. Again, a low value of the path measure implies a high connectivity.

7.2.1.4 Slicer modules

In the pseudo-code below (Alg. 1) we present a modification of Melonakos' fast-sweeping algorithm [227], with the additional steps needed to compute the proposed path measures (Eqs. (7.14) and (7.15)). In analogy with the notation for the distance map $\mathcal{L}_F^*$, which gives the shortest distance between each point and the seed region B, we define C_{avg}^* and $C_{\max}^*$ as the path measure maps that give for each point the path measure associated to the geodesic from that point to the considered seed region.

Similarly, we define $\mathcal{L}_E^*$ as the shortest Euclidean distance between each point and the seed region, and y^* as the tangent to the geodesic at each point. The algorithm is based on the observation that for a solution of Eq. (7.6), the Finslerian distance between x and the nearby point[1] $\exp(x, y)$ for $\|y\|$ small is simply given by $F(x, y)$. In the FTM interpolation is done linearly [227], but alternative methods are possible.

Algorithm 1: The fast-sweeping algorithm used to compute the maps $\mathcal{L}_F^*$, C_{avg}^*, C_{max}^*, and y^*, adapted from Melonakos et al. [227]. The implementation was based on the `finslertract` project of Antonio Tristán-Vega.

Data: A seed region B and a Finsler function F;
Result: The distance map $\mathcal{L}_F^*$, tangent map y^*, and path measure maps C_{avg}^* and C_{max}^*;
Initialize $\mathcal{L}_F^*(x \notin B) \leftarrow \infty$, $\mathcal{L}_F^*(x \in B) \leftarrow 0$, $\mathcal{L}_E^*(x) \leftarrow 0$;
repeat
    foreach *position x* **do**
        $(y^*)'(x) \leftarrow \mathrm{argmin}_{y \in \mathbb{S}^2} \mathcal{L}_F^*(x + y) + F(x, y)$;
        $(\mathcal{L}_F^*)'(x) \leftarrow \mathcal{L}_F^*(x + (y^*)'(x)) + F(x, (y^*)'(x))$;
        if $(\mathcal{L}_F^*)'(x) < \mathcal{L}_F^*(x)$ **then**
            $y^*(x) \leftarrow (y^*)'(x)$;
            $\mathcal{L}_F^*(x) \leftarrow (\mathcal{L}_F^*)'(x)$;
            $\mathcal{L}_E^*(x) \leftarrow \mathcal{L}_E^*(x + y^*(x)) + 1$;
            $C_{\mathrm{avg}}^*(x) \leftarrow \mathcal{L}_F^*(x) / \mathcal{L}_E^*(x)$;
            if $F(x, y^*(x)) > C_{max}^*(x + y^*(x))$ **then**
                $C_{\mathrm{max}}^*(x) \leftarrow F(x, y^*(x))$;
            else
                $C_{\mathrm{max}}^*(x) \leftarrow C_{\mathrm{max}}^*(x + y^*(x))$;
            end
        end
    end
until *convergence of $\mathcal{L}_F^*$*;

The FTM typically returns two scalar maps—a distance map $\mathcal{L}_F^*$ and the corresponding path measure map—based on an input diffusion-weighted data set that provides F, an optional mask image, and a label map containing labeled seed regions. The choice between F_{old} and F_{new} can be supplied as well, along with a choice of the path measures C_{avg} and C_{max} and further resolution and convergence parameters (e.g. the maximum order used in the spherical harmonic representation of F). The distance map provides at each position the shortest Finslerian distance to the seed region in accordance with Eq. (7.4), and the path measure map gives the path measure associated with the shortest geodesic connecting each point to the seed region. Geodesics can be computed if the internally generated tangent vector map is returned as well. The FTM is available at `github.com/tomdelahaije/fcm`, and can be used as a command-line tool.

The Riemannian tractography module (`nitrc.org/projects/riemantract`, by Antonio Tristán-Vega) computes Riemannian geodesics based on the adjugate metric, using an implementation similar to the one discussed for the FTM. Geodesic

---

[1]Recall that the exponential map $\exp(x, y)$ represents displacements in a Finsler manifold along the geodesic that is tangent to y in the point x, cf. Eq. (3.10).

tracking based on the normal metric was done with in-house software. Wavefront propagation methods like fast-sweeping have been used for tractography in various related works [17, 29, 51, 103, 169, 171, 266, 281, 305, 343]. Research into the optimal implementation of Finsler geodesic tracking is ongoing [231].

7.2.2 Experimental design

We only look at Riemannian geodesic tractography, as the advantages of geodesic tractography weigh heaviest in the low b-value regime of the Riemannian framework, and focus in particular on the difference between the standard choice for the metric, $g = D^{-1}$, and the scaled metric tensor $\tilde{g} = \det(D)\, D^{-1}$ proposed in Chapter 2. This choice is compared on synthetic data and on a patient data set. We further demonstrate the performance of the proposed Finsler-based connectivity analysis in two different settings: (1) a qualitative validation based on data from the WU-Minn Human Connectome Project (HCP), and (2) a quantitative network analysis of connectivity in autism spectrum disorder.

7.2.2.1 Geodesic tractography

Synthetic data   The synthetic data set is generated by placing tensors $\frac{1}{2}A(\vartheta)$ (Section 3.3.4) along curves to form fiber pathways, where ϑ is chosen such that the principal eigenvector is parallel to the nearest part of the centerline. Each voxel from which the distance to the centerline is smaller than 1.5 voxels is considered to be part of the fiber. The centerline is constructed by joining a half circle of radius 5 voxels, a horizontal straight line of length 5, a quarter circle of radius 8 and finally a straight vertical line of length 5. Any voxels outside a range of 1.5 voxels from the centerline are replaced by isotropic tensors with a mean diffusivity of $4.5 \times 10^{-3}\,\mathrm{mm^2/s}$. The ratio between the mean diffusivity of the isotropic tensors and the largest eigenvalues of the pathway tensors (the axial diffusivity) is based on average values at the interface between the corticospinal tract (CST) and the cerebrospinal fluid (CSF) as found in the literature, see Table 7.1. The data set and the geodesics are visualized together using vIST/e (`bmia.bmt.tue.nl/software/viste`).

Real data   In addition to the synthetic data we will evaluate the Riemannian tractography performance on a patient data set. The data set has 64 gradient directions and a b-value of 3000 s/mm^2, with 128×128×60 voxels and a voxel size of 1.75×1.75×2 mm^3, and was acquired from a patient with a tumor located next to the ventricles. We have segmented the cerebrospinal fluid inside the ventricles together with the tumor for our visualizations. Seeds are placed in the cerebral peduncles, with a number of target regions in the motor cortex and near the cingulum. Tracts are visualized with 3D Slicer [272].

Table 7.1: Mean and axial diffusivity values in cerebrospinal fluid (CSF) and white matter (WM) in the corticospinal tract (CST), expressed in units of 10^{-3} mm^2/s. Literature references are indicated in the table. WM-CST diffusivity value in [229] corresponds to the posterior limb of the internal capsule, which contains a.o. corticospinal fibers.

Regions	MD	AD
CSF	3.2 [7, 213]	–
WM-CST	–	1.0 [229]
		1.1 [36]
		1.2 [188]

Table 7.2: Parameter values used for deterministic and probabilistic CSD-based fiber tracking as implemented in the MRtrix package.

Parameters	Deterministic CSD	Probabilistic CSD
Step size	0.5 mm	0.2 mm
Minimum radius of curvature	0.5 mm	2.0 mm
fODF amplitude cutoff	0.1	0.15
fODF amplitude cutoff for initiation	0	0

For comparison we also include constrained spherical deconvolution-based [342] deterministic and probabilistic fiber tracking using MRtrix [341] (CSD, `mrtrix.org`). See Section 6.1.1.3 for a brief explanation on CSD. The deterministic tractography algorithm follows the peaks of the fiber orientation distribution (fODF), and in the probabilistic case tractography is based on orientations sampled from the fODF at each step. The algorithm parameters are optimized for this particular data set by visual inspection, and can be found in Table 7.2. The cerebral peduncles are again used as seed regions, and only fibers reaching the target regions in the motor cortex are selected.

7.2.2.2 Qualitative validation of geodesic-based connectivity

The qualitative Finsler connectivity validation experiments presented in Section 7.3.2 are based on the pre-processed data of a single healthy subject in the WU-Minn Human Connectome Project [351], released as part of the HCP 500 Subject Release. The data was acquired on a modified 3T Siemens scanner, with 1.25 mm isotropic voxels. The diffusion-weighted images are acquired on three shells ($b = \{1000, 2000, 3000\}$ s/mm^2) with 90 uniformly distributed gradient directions each, together with 18 baseline images. We express F_{old} and F_{new} in terms of spherical harmonics (maximum order 6), based on the $b = 3000$ s/mm^2 shell. T1 data with an isotropic voxel size of 0.7 mm was also available. Further details on the acquisition protocol can be found on the HCP web site and in the references [8, 152, 351].

All maps based on the HCP data are seeded from manually selected voxels within the white matter, one voxel per bundle, based on the DTI white matter atlas by Catani

and Thiebaut de Schotten [59]. Seeds are placed in four well-known major white matter bundles: the cingulum, the arcuate fasciculus, the corticospinal tract, and the splenium of the corpus callosum.

7.2.2.3   Network analysis in autism spectrum disorder

In recent years, the view that the functional and structural systems of the brain can be modeled as complex networks has motivated a large amount of research on the application of graph theoretical concepts to brain network analysis [44, 291]. The standard graph network model of the brain consists of a set of nodes, which represent a partitioning of the cortex and other gray matter structures. These nodes are connected via a set of edges, or links, that represent structural and/or functional connections between gray matter partition units. Such a graph model of the brain's network organization can be constructed from a variety of imaging modalities such as structural MRI, diffusion MRI, functional MRI, or EEG/MEG. In this setting, a characterization of the organization of the different computational nodes and the functional interaction between them is achieved via graph theoretical analysis [44, 291].

In order to provide an illustration of how the Finsler connectivity framework could be applied to population studies in a clinical setting, we present here a proof of concept network-based analysis of our Finsler connectivity approach, applied to autism spectrum disorder (ASD) data. To do so, we use a paradigm that is commonly used in network-based studies of brain connectivity. Specifically, we build a graph model of the brain, where graph nodes represent gray matter regions as defined through a FreeSurfer parcellation based on the Desikan–Killiany atlas [105]. The edge weights in this model represent the Finsler connectivity between pairs of FreeSurfer-defined gray matter regions. Once this network model is constructed for each subject, we compute its local efficiency measure. The local efficiency measure [291] is one of the many standard graph-theoretical measures that are commonly computed in network analysis studies, in order to quantitatively summarize the network structure. It is possible to compute other measures as well, but since this experiment is meant only as an illustration, we focus on a single measure that has been previously implicated in ASD [221, 292].

Subjects and data acquisition   Diffusion and structural MRI data were acquired from 69 typically developing male controls (TDC, age range: 8.0–14.4years, mean: 10.7, standard deviation: 1.8) and 46 age-matched male autism spectrum disorder patients (ASD, age range: 8.1–14.1years, mean: 10.7, standard deviation: 2.0). A t-test for difference in age between the two groups resulted in a p-value of 0.98.

All imaging was performed using a Siemens 3T Verio scanner with a 32 chan-nel head coil. Structural images were acquired on all subjects using an MP-RAGE

imaging sequence (TR = 19 s, TE = 2.54 ms, TI = 0.9 s, 0.8 mm in-plane resolution, 0.9 mm slice thickness). In addition, a single shell high b-value acquisition was performed using a monopolar Stejskal–Tanner diffusion-weighted spin echo, echo-planar imaging sequence with the following parameters: TR = 14.8 s, TE = 110 ms, b = 3000 s/mm^2, 2 mm isotropic resolution, with 64 gradient directions and with two baseline images. The diffusion-weighted images of each subject were filtered using a joint linear minimum mean squared error filter to suppress Rician noise [345]. Eddy current correction was performed using registration of each volume to one of the baseline images. The same data has been used for example in a study by Caruyer and Verma [56].

FreeSurfer parcellation   FreeSurfer is a freely available software toolbox that reconstructs mesh-based models of the cortical surface [80, 127], and provides a parcellation of the cortex into neuro-anatomical areas using both the geometrical model of the cortical surface and neuro-anatomical convention [126].

For the present set of experiments, we computed a FreeSurfer parcellation for each subject based on the Desikan–Killiany atlas [105], which resulted in a total of 86 cortical and sub-cortical gray matter regions. FreeSurfer parcellations are initially defined in each subject's T1 standard space. In order to allow for these gray matter regions of interest to be used as seed regions for Finsler connectivity analysis, we registered the FreeSurfer parcellation to the diffusion MRI space using non-linear registration.

Connectivity matrix construction   The path measures discussed in Section 7.2.1.3 are length-based, i.e., low values indicate a short connection distance, which is equivalent to a high connection strength. From these path measures we can heuristically derive a new measure that signifies a general notion of connectivity, the values of which are bounded to lie between 0 and 1 such that small values indicate low connection strength.

For each subject in our study, we constructed an 86×86 connectivity matrix C, whose ij-th entry C_{ij} represents the connectivity between the i-th and j-th FreeSurfer regions. Each entry C_{ij} is computed as follows. First, seeding in the i-th region, we run our Finsler connectivity algorithm to produce a path measure map for the entire brain. Then, the values for the path measure in the spatial extent of the j-th region are averaged for each $j \neq i$, which results in the average path measure values $m_{i \to j}$. This process is repeated for all seed regions i. Finally, we construct the connectivity matrix C:

$$C_{ij} = \begin{cases} \exp\left(-\nu \frac{m_{i \to j} + m_{j \to i}}{2}\right) & \text{if } i \neq j \\ 1 & \text{if } i = j \end{cases}. \tag{7.16}$$

Note that C is a symmetric matrix by construction, so each pair of regions has a uniquely defined connectivity value. The scaling parameter ν is set to 0.1 for the presented experiments, producing a reasonable spread of connectivity values over the range $[0, 1]$.

As an example, Fig. 7.2 shows a coronal slice through a path measure map computed on one of the TDC subjects. This map was seeded in the left caudal middle frontal gyrus region, as defined by a FreeSurfer segmentation using the Desikan–Killiany atlas as described above. The seed voxels that intersect this particular coronal slice are shown in white color outlined with a black boundary. The path measure map itself is shown in a modified temperature color map, such that bright red colors indicate a low path measure (high connectivity) and a dark blue color indicates a high path measure value. Thus, in order to compute $m_{i \to j}$, where i indicates the caudal middle frontal region, and j indicates any of the other FreeSurfer-defined regions, we average the values of this path measure map over the voxels that comprise FreeSurfer region j. Intuitively, regions that are well-connected to region i will result in low values for $m_{i \to j}$, which in turn will lead to higher connectivity values C_{ij} according to Eq. (7.16).

In addition to illustrating the concept of a path measure map and the corresponding computation of $m_{i \to j}$, Fig. 7.2 also provides a comparison between the F_{new}- and the F_{old}-based metrics. With the caudal middle frontal gyrus as a seed region, the newly proposed F_{new}-based metric results in the recovery of the known transcallosal connectivity to the opposite hemisphere (Fig. 7.2, left). In contrast, the F_{old}-based metric produces a map (Fig. 7.2, right) that has drawbacks similar to known artifacts observed with tractography on single-tensor DTI data. In particular, the connectivity does not reach the contralateral cortical regions, as it appears to stop in the well-known region of three-way crossings between the corpus callosum, the corticospinal tract and the superior longitudinal fasciculus. On the other hand, connectivity appears to 'leak' into the corticospinal tract of both hemispheres, as well as into other white matter tracts that do not have direct anatomical connectivity with the seed region. Furthermore, the F_{new}-based metric correctly identifies CSF areas, such as the ventricles, and assigns a high path measure (low connectivity). In contrast, the F_{old}-based metric does not detect CSF and as a result propagates connectivity through CSF regions, which is clearly anatomically incorrect. The differences between the two metrics is addressed in more detail in Section 7.3.2.

Computation of local efficiency and statistical analysis Once the connectivity matrix C is computed for each subject, we compute its local efficiency measures using the Brain Connectivity Toolbox [291]. The local efficiency of a network is a quantity computed at each node of the network, such that it quantifies the network's resistance to failure at the local scale. In other words, it quantifies the importance of

F_{new} F_{old}
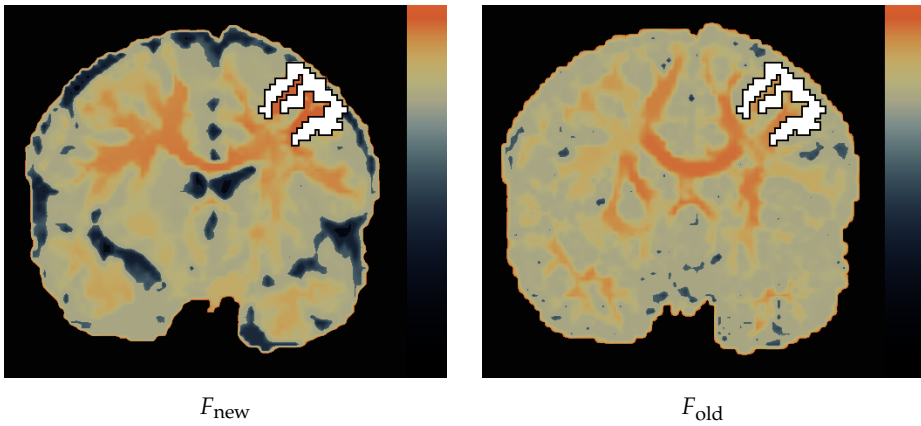
Figure 7.2: A comparison of C_{max} path measure maps, using the F_{new}- and F_{old}-based metrics on data from one of the TDC subjects. These maps are seeded in the left caudal middle frontal gyrus region, as defined by FreeSurfer. The seed voxels that intersect this particular coronal slice are shown in white. Bright red voxels are strongly connected to the seed region according to the used path measures, while dark voxels are weakly connected.

a graph node by measuring how well information is exchanged by the immediate neighbors of the node when it is removed. Thus, in the present experiment, this computation results in a 1×86 vector, such that its i-th element corresponds to the local efficiency measure of the i-th FreeSurfer region.

In the present experiment, we are interested in performing a statistical test for a group difference between the TDC and ASD groups of subjects based on their local efficiency vectors. To reduce the number of multiple comparisons, we do not test each region individually but perform a one-way multivariate analysis of variance (MANOVA) to compare the mean vectors for the two groups. This is a statistical test for the null hypothesis that the mean local efficiency vectors of the TDC and ASD groups are the same. If we can reject this null hypothesis, we conclude the two groups differ in terms of their local efficiency measure, although the test does not identify specific regions that may be responsible for this difference.

7.3 Results

7.3.1 Geodesics in the scaled manifold

7.3.1.1 Results on synthetic data

Figs. 7.3, 7.4, and 7.5 show geodesic tractography results for the synthetic data corrupted with increasing amounts of Rician noise, parameterized by the standard deviation σ of the underlying normal distribution. Each of the figures compares the

scaled metric proposed in Chapter 2 with the standard metric $g = D^{-1}$, and with sharpened versions of both using power sharping similar to Eq. (4.23):

$$(D, s) \mapsto (\det D)^{\frac{1-s}{3}} D^s, \tag{7.17}$$

where s is the sharpening power. Seeds and targets are placed at the inflection point and at the boundaries of the fiber pathways.

For this particular data set, we find that for both the short and for the long fibers, geodesics obtained with the inverse diffusion tensor take a shortcut through the isotropic background, failing to describe the fibers in all situations (Figs. 7.3(a), 7.4(a) and 7.5(a)). With a sharpening factor $s = 2$ this remains the case, except for the longer fiber in the specific case of $\sigma = 0.15$, Fig. 7.4(b). It is not until we use a sharpening factor of $s = 4$ that geodesics computed for the inverse sharpened tensor nicely follow the expected trajectories, although a slight degradation is observed for $\sigma = 0.3$ (Fig. 7.5(c)). This effect for sharpened metrics had been shown in the work of Sepasian et al. [305] for a sharpening power $s = 2$, and by Hao et al. [163] for $s = 3$ (but with the background masked).

On the other hand, geodesics from adjugate tensors, with and without sharpening, follow the synthetic fibers rather well in all scenarios and without taking shortcuts through the isotropic background. Without sharpening the geodesics do still tend to move away from the center line a bit, but the effect is much less severe than in the inverse tensor case. A slight degradation is observed for the $s = 4$ adjugate sharpened tensor when there is significant noise ($\sigma = 0.3$, Fig. 7.5(c)), which also happened for the standard (sharpened) metric. Comparing geodesics from the unsharpened adjugate tensor (Figs. 7.3(a), 7.4(a), and 7.5(a)) and the $s = 4$ inverse or adjugate sharpened tensor (Figs. 7.3(c), 7.4(c), and 7.5(c)) we observe that the sharpened ones follow the fibers more closely in the noiseless case. However, in the noisy case these degrade by taking a shortcut of about one voxel while the adjugate unsharpened ones remain almost unchanged. This is explained by an increased sensitivity to noise after sharpening, a shortcoming that was also pointed out in the work by Hao et al. [163].

7.3.1.2 Results on real data

In Figs. 7.6, 7.7 and 7.8 we show candidate fibers reaching the trunk and foot motor area of the cortex (upward bundle) and the lip area (bundle bending to the left), which ought to correspond to the corticospinal and corticobulbar tracts. Results above the ventricles are seeded consistent with the left and right cingulum. In Fig. 7.6 we show tractography results for metrics given by the inverse and adjugate diffusion tensor, and the outcome for inverse sharpened diffusion tensors is given in Fig. 7.7. Results obtained with our approach, Fig. 7.6(b), seem to better resemble the anatomy of the stipulated white matter bundles. Additionally, the curvature of the candidate fibers

(a) $s = 1$ **(b)** $s = 2$ **(c)** $s = 4$
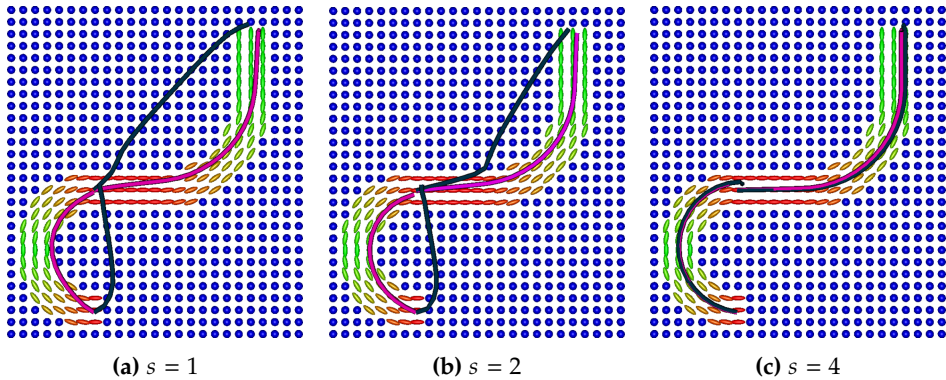
Figure 7.3: Results on synthetic data, for metrics given by the inverse diffusion tensor (black) and the adjugate diffusion tensor (magenta) with sharpening powers $s = 1$ (a), $s = 2$ (b), and $s = 4$ (c). The ellipsoids are color-coded based on the direction of the diffusion tensor main eigenvector. Geodesics obtained with the standard inverse diffusion tensor metric fail to describe fibers except for $s = 4$, while the scaled metric based on the adjugate tensor produces geodesics that follow the fibers well, with sharpening further improving results.

(a) $s = 1$ **(b)** $s = 2$ **(c)** $s = 4$
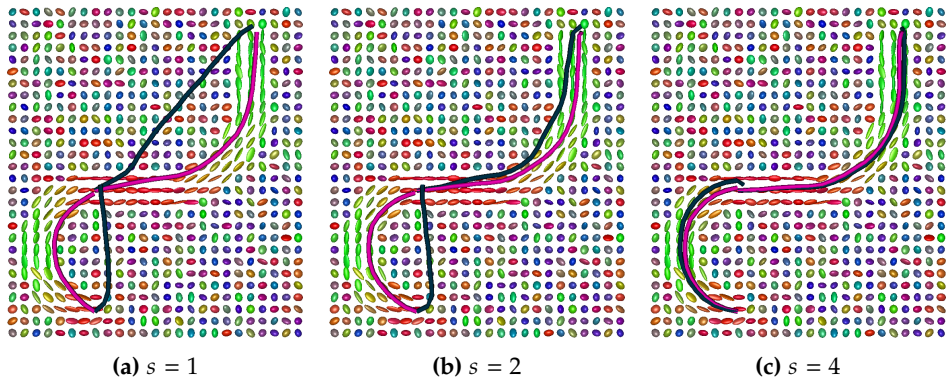
Figure 7.4: Results on synthetic data with a Rician noise level of $\sigma = 0.15$ (standard deviation of the underlying normal distribution), for metrics given by the inverse diffusion tensor (black) and the adjugate diffusion tensor (magenta) with sharpening powers $s = 1$ (a), $s = 2$ (b), and $s = 4$ (c). Color-coding as in Fig. 7.3. Again, geodesics obtained with the standard metric fail to describe the fibers except for $s = 4$, although the longer fiber $s = 2$ tracking does improve w.r.t. the noiseless case. Adjugate-based geodesics follow the expected trajectories well, and sharpening improves results further.

(a) $s = 1$ **(b)** $s = 2$ **(c)** $s = 4$
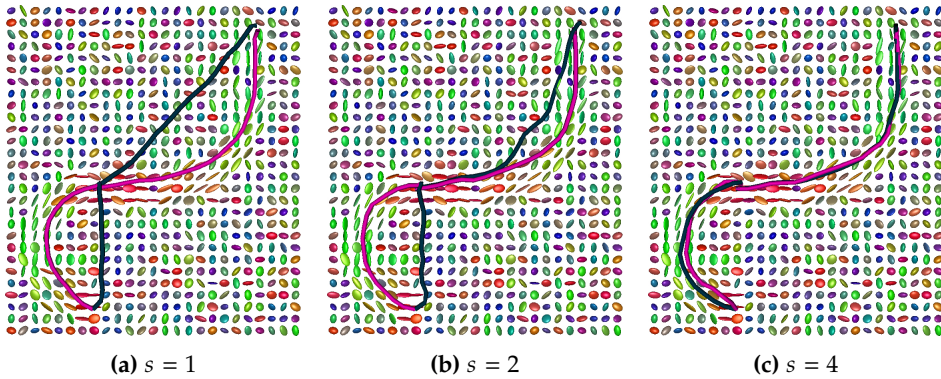
Figure 7.5: Results on synthetic data with Rician noise of $\sigma = 0.3$ (standard deviation of the underlying Gaussian distribution), for metrics given by the inverse diffusion tensor (black) and the adjugate diffusion tensor (magenta) with sharpening powers $s = 1$ (a), $s = 2$ (b), and $s = 4$ (c). Color-coding as in Fig. 7.3. Inverse-based geodesics fail to describe fibers except for $s = 4$, while adjugate-based geodesics follow the fibers well. For both $s = 4$ sharpened metrics tractography worsens noticeably compared to Fig. 7.4.

is smoother and the bundles are more coherent. A particularly interesting result is the fact that our candidate fibers circumvent the ventricles, known to be void of fibers, while the majority of geodesics obtained with the (sharpened) inverse tensor go through them. Note that for sharpened inverse tensors, Fig. 7.7, fewer bundles cross the CSF than in the original inverse diffusion tensor case, Fig. 7.6(a), though the problem is not completely overcome. The results obtained with sharpening power $s = 2$, Fig. 7.7(a), are consistent with our synthetic data experiments, but for $s = 4$ results are worse than expected, Fig. 7.7(b), likely due to more significant noise in the real data.

In Fig. 7.6 we also see that tracts obtained with the newly proposed adjugate do not go through the tumor. This is consistent with our findings concerning the CSF, as the observed diffusion inside the tumor is quite comparable. Our results may reflect real fibers being pushed aside by a tumor, or white matter inside the tumor being severely degraded. However we have no means to establish if one of these scenarios is true, or whether we have simply missed some tracts that were actually present in the tumor, and therefore we cannot draw any decisive conclusions about the validity of our results in this case.

In Fig. 7.8 we show results for a metric given by the sharpened adjugate diffusion tensor. Tractography improves again drastically compared to the sharpened inverse results shown in Fig. 7.7, with none of the tracts cross isotropic diffusion regions, but there is little gained over the ($s = 1$) adjugate diffusion tensor results shown in Fig. 7.6(b).

In Fig. 7.9 we show CSD-based deterministic and probabilistic tractography results for the corticospinal tract, together with those obtained from geodesic tractography

(a)   (b)
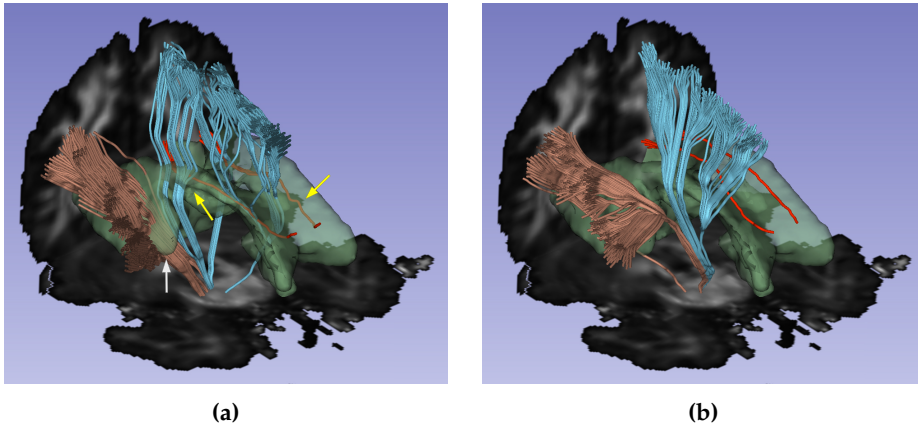
Figure 7.6: Candidate fibers for the corticobulbar and corticospinal tracts (brown and blue, respectively), and the cingulum (red), in an anterior view. A tumor is located next to the ventricles on the left-hand side. Candidate fibers are obtained through geodesic tractography as explained in Section 7.2.2.1, using the inverse diffusion tensor as a metric (a) and using the adjugate diffusion tensor (b). Candidate fibers going through the ventricles or the tumor are indicated with yellow and white arrows respectively. Bundles obtained with the scaled metric proposed in Chapter 2, shown in (b), avoid both the CSF in the ventricles and the tumor.

based on the adjugate diffusion tensor. In all cases fibers circumvent the ventricles, but the CSD fiber pathways barely reach the anterior part of the (right) motor cortex. This could be improved by tuning the parameters (e.g. by taking step size 0.5 mm, minimum radius of curvature ≥ 5 mm, fODF amplitude cutoff ≤ 0.01 mm), but this results in unrealistically straight fibers that tend to move through the ventricles and jump from one hemisphere to the other.

7.3.2 Qualitative validation of geodesic-based connectivity

All connectivity and path measure maps are shown in radiological convention. Path measure maps are shown using a temperature color map, where dark blue indicates a low connectivity/high path measure, and bright red indicates high connectivity/low path measure. Note that the results for the different path measures $C_{\max}$ and C_{avg} cannot be compared directly, and that the connectivity highlights regions that are connected to a seed point, which should be a superset of the bundle in which the seed is placed. The color maps are shown on the right of each figure.

7.3.2.1 Cingulum

The first major bundle we consider is the cingulum, which consists of a set of fibers that project from the cingulate gyrus to the entorhinal cortex. In the work of Melonakos et al. [228] geodesic tractography was successfully used to trace this bundle, so we can expect our geodesic-based connectivity analysis to produce reasonable
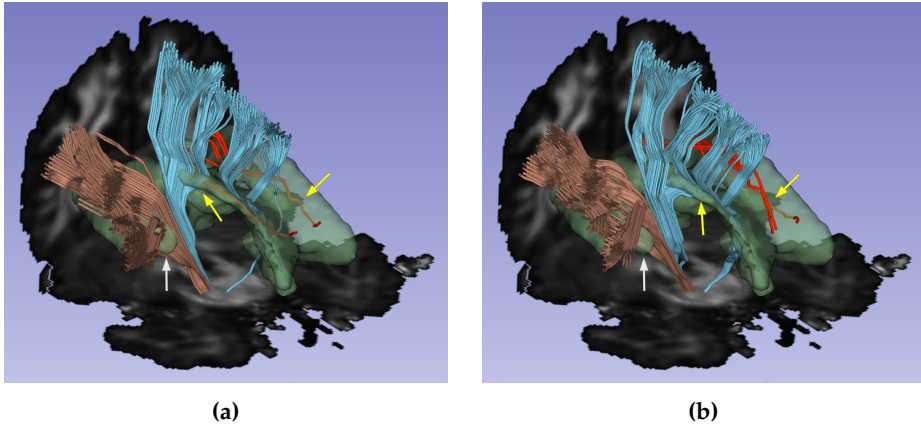
Figure 7.7: As in Fig. 7.6, but now showing results for metric given by (a) inverse sharpened diffusion tensor $(\det \boldsymbol{D})^{1/3}\boldsymbol{D}^{-2}$ and (b) inverse sharpened diffusion tensor $(\det \boldsymbol{D})\boldsymbol{D}^{-4}$. Note that results from sharpened tensors improve compared to those without sharpening in Fig. 7.6(a) (i.e., less tracts cross isotropic diffusion regions), but the problem is not completely overcome as in our approach.
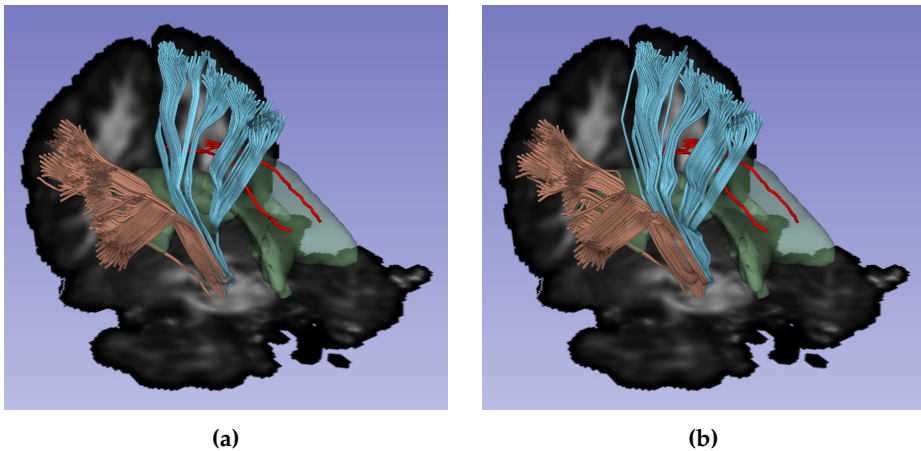
Figure 7.8: As in Fig. 7.7, but now showing results for metric given by (a) adjugate sharpened diffusion tensor $(\det \boldsymbol{D})^{4/3}\boldsymbol{D}^{-2}$ and (b) adjugate sharpened diffusion tensor $(\det \boldsymbol{D})^{2}\boldsymbol{D}^{-4}$. Note that results from adjugate sharpened tensors improve drastically compared to those from inverse sharpening in Fig. 7.7 (i.e., none of the tracts cross isotropic diffusion regions). The outcome is very similar to that from the adjugate diffusion tensor, Fig. 7.6(b).

(a) **(b)**
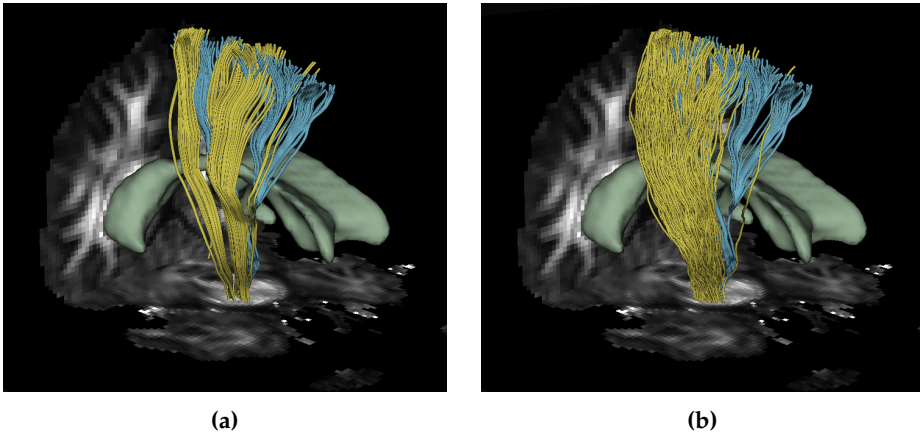
Figure 7.9: CSD-based tractography results for the corticospinal tract (yellow), together with those obtained by geodesic tractography from the adjugate diffusion tensor (blue). (a) deterministic CSD (b) probabilistic CSD. In all cases fibers circumvent the ventricles. CSD fibers do not reach the anterior part of the (right) motor cortex.

results for seeds placed in this bundle. The distance map from which geodesics can be computed is shown in Fig. 7.10 for both $F = F_{\text{old}}$ and the newly proposed F_{new}-based metric. The Finslerian distance increases as expected from the seed outward, and the boundaries between white matter, gray matter, and cerebral spinal fluid can be identified very roughly. It is however difficult to judge the relative merit of the different metrics from these maps alone. Because all the connectivity measures under investigation in this section are derived from the same $\mathcal{L}_F^*$ distance maps, and because the distance maps themselves provide little information, we omit the distance maps henceforth.

The connectivity maps shown in Figs. 7.11 and 7.12 provide more information. First, a high-level comparison between Figs. 7.11 and 7.12 reveals that the C_{max} connectivity measure (Eq. (7.15)) provides much more anatomical detail than the C_{avg} connectivity measure (Eq. (7.14)). In Fig. 7.11 we see furthermore that the standard C_{avg} path measure map seeded in the cingulum bundle leaks into the corpus callosum, and to large sections of the posterior part of the brain. This effect is much less in the C_{max} map shown in Fig. 7.12, which also correctly assigns low connectivity values to CSF voxels. Overall, the C_{max} map is much more specific to the known anatomy of the cingulum bundle. This allows us to perform a more detailed comparison between the $F = F_{\text{old}}$ metric and our newly proposed $F = F_{\text{new}}$ metric. We see that the former does result in leakage from the cingulum into the posterior parts of the corpus callosum (the splenium), and from there into large sections of posterior white matter. The choice $F = F_{\text{new}}$ results in a path measure map that closely follows the known anatomy of the cingulum bundle, without leaking into the corpus callosum.

$$F_{\text{new}} \qquad\qquad F_{\text{old}}$$
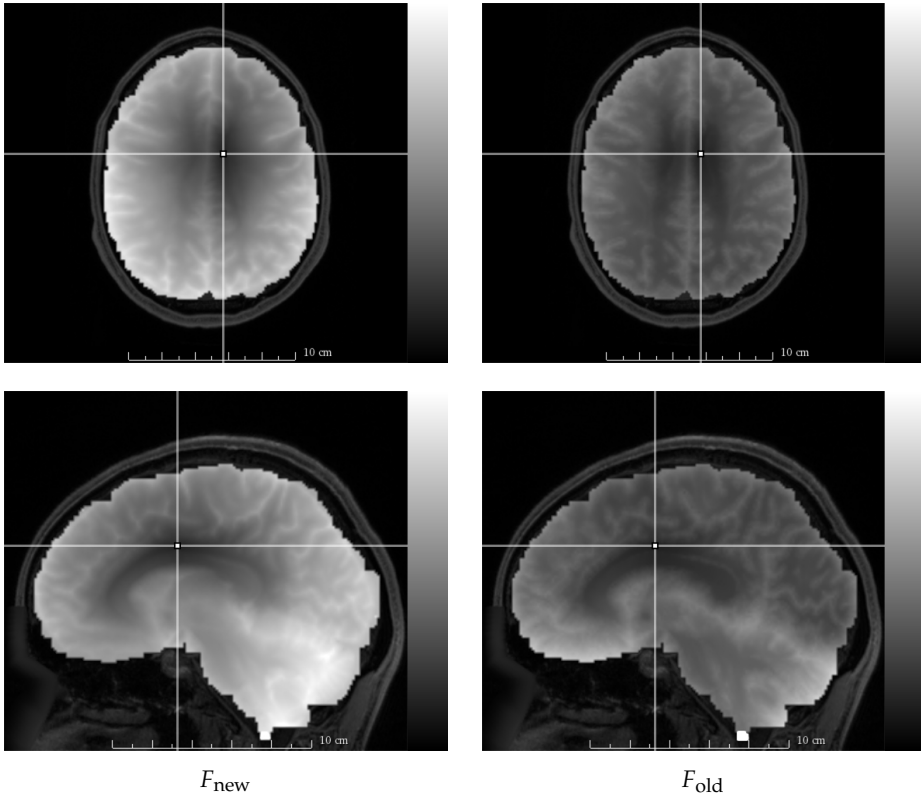
Figure 7.10: Axial (top row) and sagittal (bottom row) slices of distance maps $\mathcal{L}_F^*$, Eq. (7.4), for $F = F_{\text{new}}$ (left column, Eq. (7.10)) and $F = F_{\text{old}}$ (right column, Eq. (7.9)) seeded at a single voxel in the cingulum (annotated point) of the HCP data set.

Finally we note that all maps suffer to a certain extent from errors in the mask, which can cause erroneous and problematic high path measure values near the boundaries of the mask. These false positives are clearly visible at the edges in Fig. 7.11 (left) and near the brainstem in Fig. 7.12 (bottom). With the C_{avg} measure these errors can propagate throughout the brain, which is particularly grievous in combination with the new F_{new}-based metric. With the C_{max} path measure this propagation is however completely suppressed, which means the errors remain localized.

7.3.2.2 Arcuate fasciculus

Maps of the two path measures, C_{avg} and C_{max}, are shown for the arcuate fasciculus in Figs. 7.13 and 7.14. The arcuate fasciculus is a functionally important bundle, involved in aspects of language processing. It connects frontal cortical areas with the superior temporal gyrus. It also includes connections to the inferior parietal lobe. The maps for this bundle very clearly highlight a major issue with the widely used C_{avg} measure:

F_{new} F_{old}
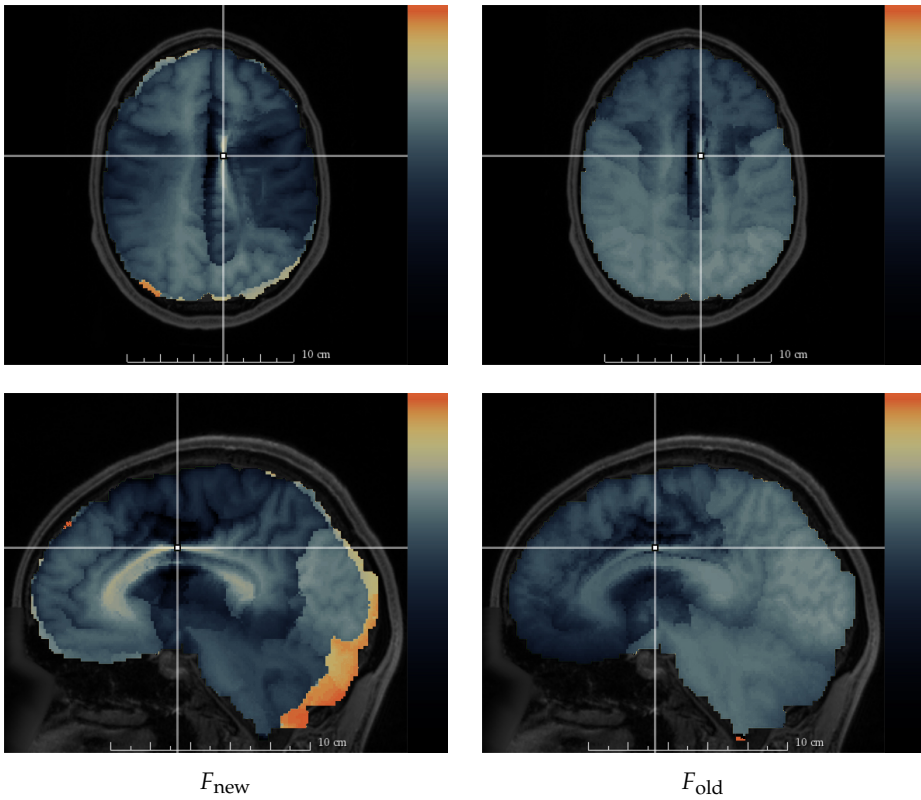
Figure 7.11: Axial (top row) and sagittal (bottom row) slices of maps based on the C_{avg} path measure (Eq. (7.14)), derived from the data shown in Fig. 7.10. The left column shows the results for the newly proposed $F = F_{\text{new}}$ metric (Eq. (7.10)), and the right column shows results for the $F = F_{\text{old}}$ metric (Eq. (7.9)). The choice $F = F_{\text{new}}$ suffers from artificially high values at the edges, due to errors in the mask.

despite the seed being placed in the right hemisphere, the left hemisphere shows an overall *stronger* connection to the seed than the voxels in the right hemisphere. This problem is to a large extent resolved with the introduction of the C_{max} measure. For the remaining bundles we show only results with the C_{max} measure, though it should be noted that the C_{avg} measure can produce good results *locally* as shown in the top-left map in Fig. 7.13.

Again, a comparison between the C_{max}-based maps with $F = F_{\text{old}}$ and $F = F_{\text{new}}$ reveals a superior performance of the new F_{new}-based metric (barring the artifacts due to masking errors). With the F_{old}-based metric, there is again leakage into the corpus callosum and then into the opposite hemisphere, which is not observed with the F_{new}-based metric. As it is difficult to judge the reconstruction of the arcuate fasciculus' connectivity neighborhood from 2D slices, we show a 3D reconstruction of a thresholded path measure map obtained with $F = F_{\text{new}}$ in Fig. 7.15, which shows

F_{new} F_{old}
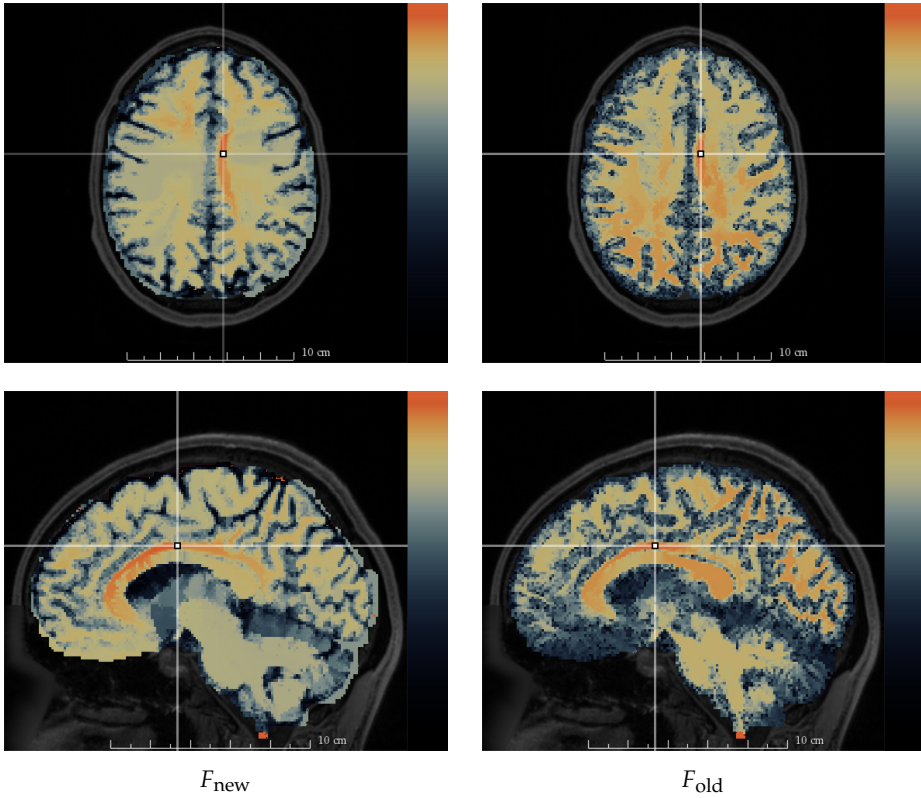
Figure 7.12: Axial (top row) and sagittal (bottom row) slices of maps based on the $C_{\max}$ path measure (Eq. (7.15)), derived from the data shown in Fig. 7.10. The left column shows the results for the newly proposed metric $F = F_{\text{new}}$ (Eq. (7.10)), and the right column shows results for the $F = F_{\text{old}}$ metric (Eq. (7.9)). Note that the F_{old} metric leads to much more significant 'leakage' into the corpus callosum. The $C_{\max}$ path measure shown in this figure is largely unaffected by masking errors, cf. Fig. 7.11.

a fairly complete reconstruction of the arcuate fasciculus and its connected regions, with minimal leakage into other bundles. This reconstruction was impossible with $F = F_{\text{old}}$.

7.3.2.3 Cortico-spinal tract

The next tract we examine is the corticospinal tract (CST), which is a major fiber tract that conducts sensorimotor signals between the cortex and the spinal cord. $C_{\max}$ path measure maps obtained with a single voxel seed in the CST are provided in Fig. 7.16. These maps show some of the advantages and disadvantages of the different metrics. The traditionally used F_{old} metric appears to be better equipped to avoid leakage into the opposite hemisphere, at the level of the brain stem, and specifically at the pons, where pontine crossing fibers may cause the connectivity to cross the midline into the other hemisphere [357]. To see this, compare the images in the bottom row of

F_{new} F_{old}
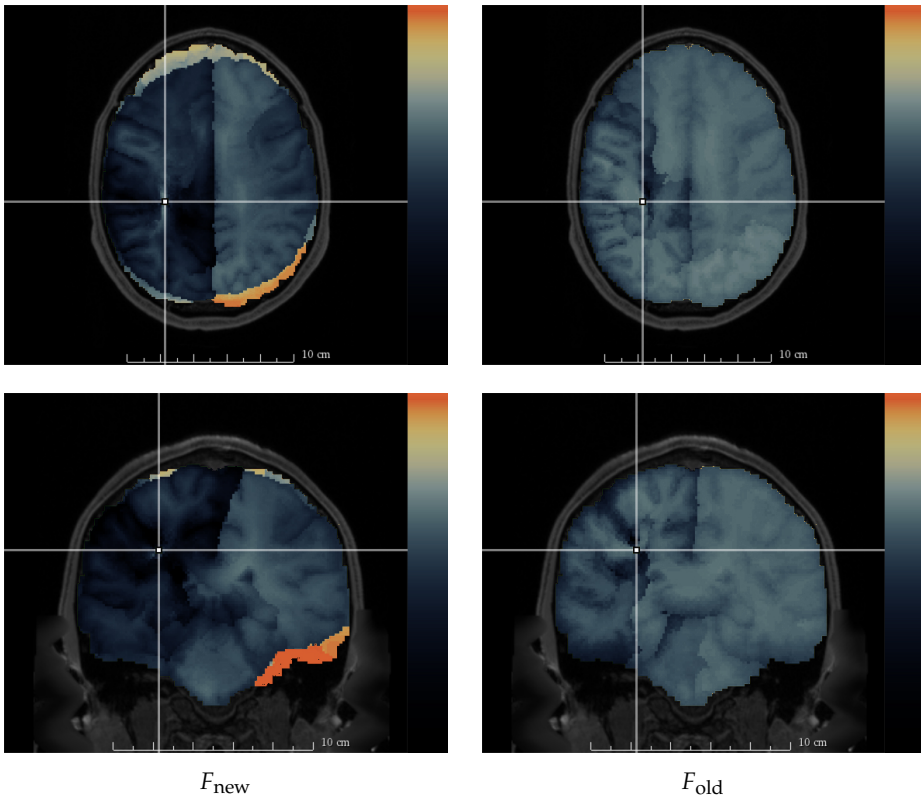
Figure 7.13: Axial (top row) and coronal (bottom row) slices of C_{avg}-based maps (Eq. (7.14)) for $F = F_{\text{new}}$ (left column, Eq. (7.10)) and $F = F_{\text{old}}$ (right column, Eq. (7.9)) seeded in the arcuate fasciculus (annotated point) of the HCP data set. Connections over the corpus callosum typically have large average diffusivities, which results in the undesired high values in the unseeded hemisphere.

Fig. 7.16. This result is anatomically incorrect and is a well-known issue with many tractography algorithms [357]. While the F_{new} metric is more sensitive to this issue, the regions of high connectivity extend much further in the superior direction towards the cortex, better reproducing the known fiber fanning in this region. Furthermore, F_{new} based connectivity of the CST shows only a minor leak into the corpus callosum, while the F_{old} metric produces a much more extensive leak into the posterior part of the corpus callosum and into the opposite hemisphere.

7.3.2.4   Corpus callosum

The last bundle we consider is the corpus callosum, a massive white matter highway that connects the two hemispheres. The C_{max}-based path measure maps for the F_{old} and F_{new} metrics were seeded in the splenium (posterior part) of the corpus callosum, as shown in Fig. 7.17. With the F_{new} metric, we observe that the path measure map

F_{new} F_{old}
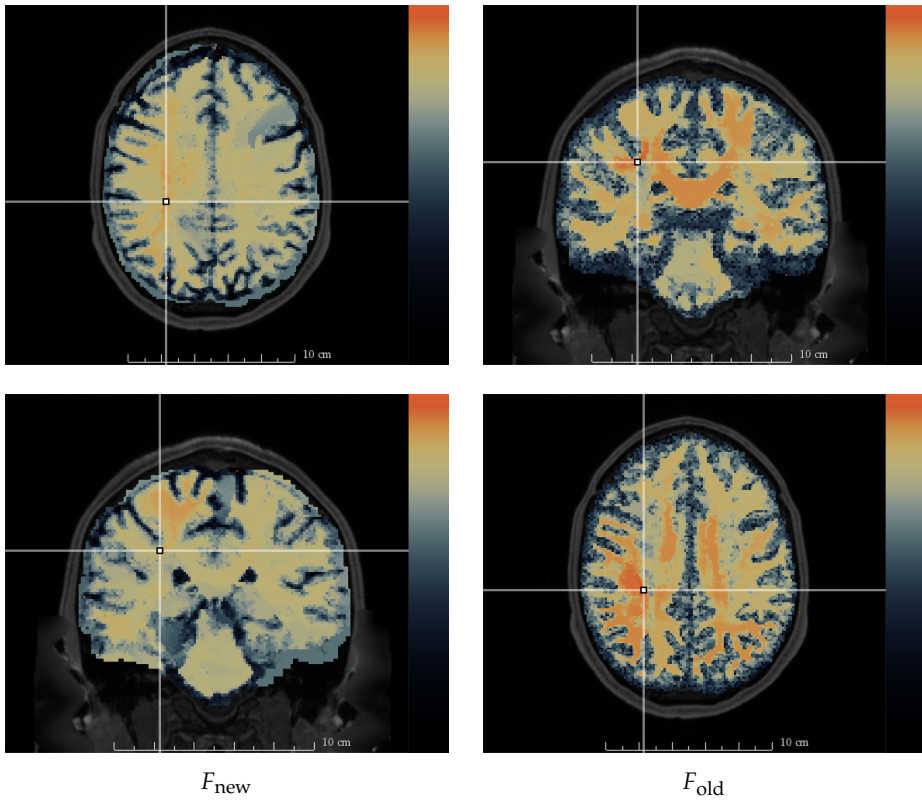
Figure 7.14: C_{max}-based maps (Eq. (7.15)) for $F = F_{\text{new}}$ (left column, Eq. (7.10)) and $F = F_{\text{old}}$ (right column, Eq. (7.9)) corresponding to Fig. 7.13. The C_{max} path measure does not have a bias for cross-hemispheric connections.

(a)   **(b)**
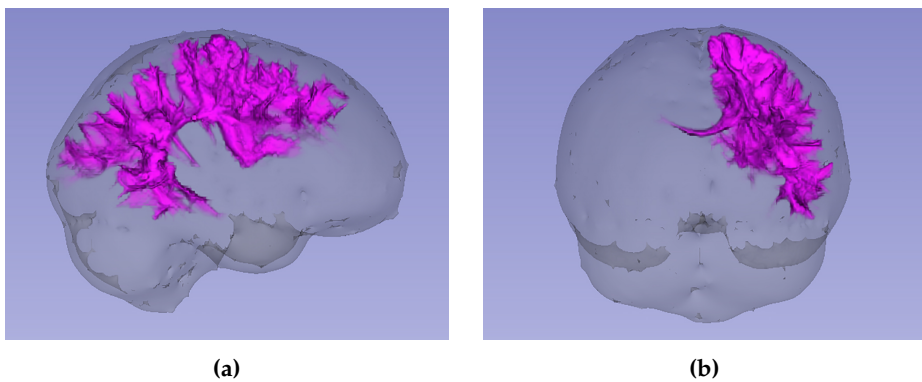
Figure 7.15: Two views of a rendering of the regions connected to the arcuate fasciculus obtained by thresholding the C_{max}-based map (Eq. (7.15)) (with $F = F_{\text{new}}$) shown in Fig. 7.14.

F_{new} F_{old}
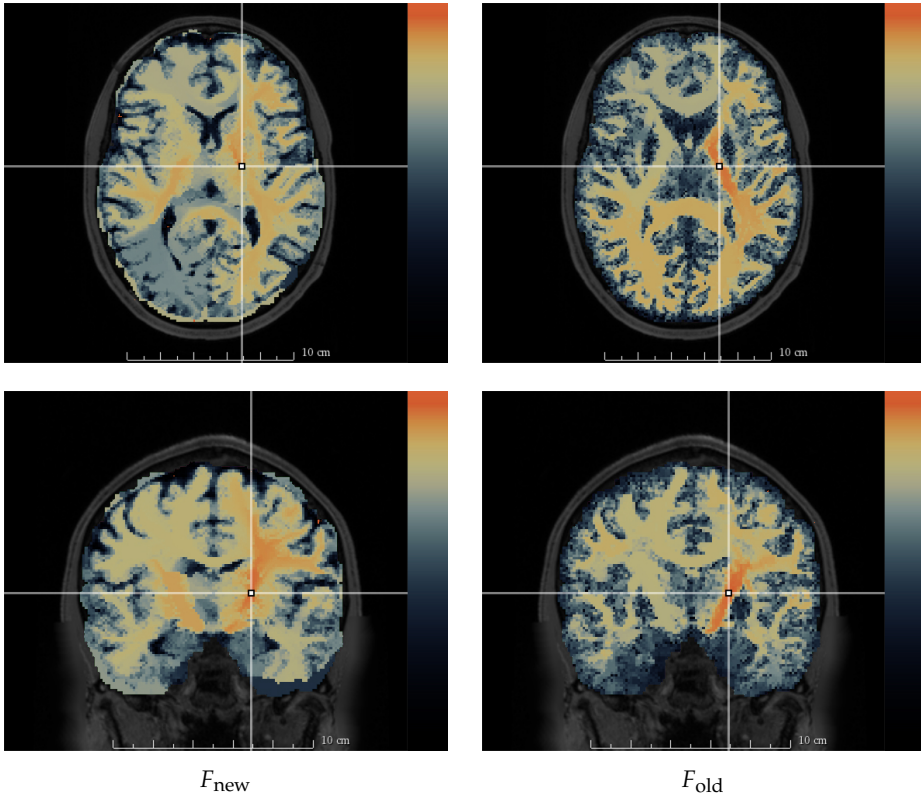
Figure 7.16: Axial (top row) and coronal (bottom row) slices of C_{max}-based maps (Eq. (7.15)) for $F = F_{\text{new}}$ (left column, Eq. (7.10)) and $F = F_{\text{old}}$ (right column, Eq. (7.9)) seeded in the corticospinal tract (annotated point) of the HCP data set.

remains concentrated in the posterior parts of the corpus callosum, while with the F_{old} metric it progresses much further in the anterior direction, which should not be happening with a seed located posteriorly.

7.3.3   Network analysis in autism spectrum disorder

We performed the network analysis study with the newly proposed F_{new}-based Finsler metric, using three different maximum orders for the spherical harmonic representation of the ODF. For comparison we also repeated the experiment, for the same three different spherical harmonic orders, using the F_{old}-based metric proposed by Melonakos et al. [227], Eq. (7.9). The resulting p-values are presented in Table 7.3. All experiments used the C_{max}-based connectivity measure.
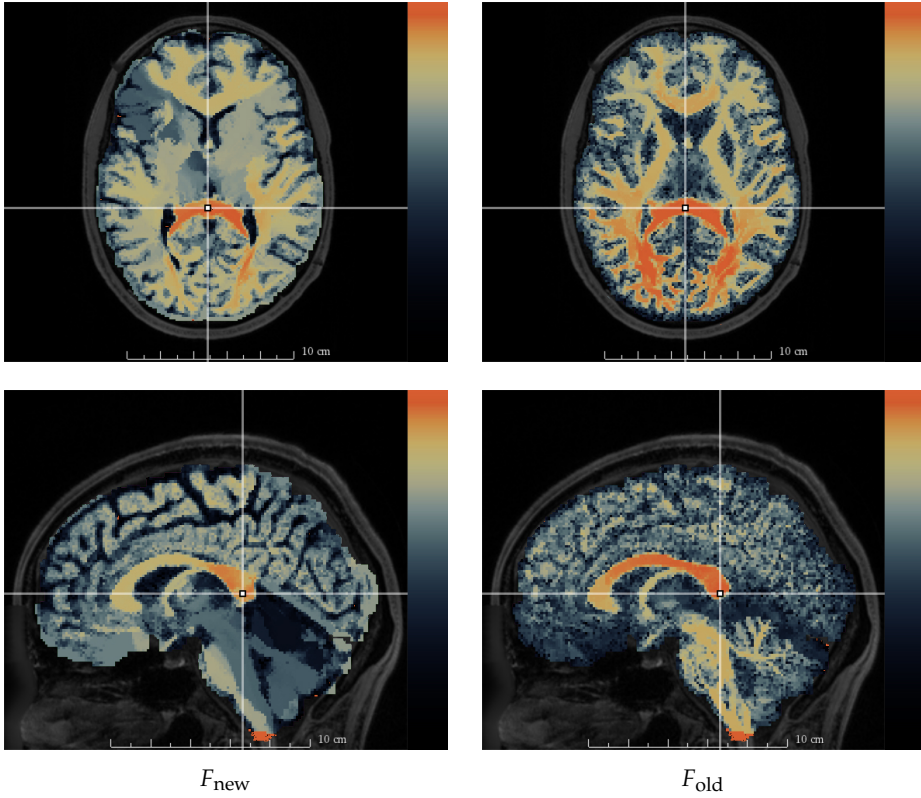
Figure 7.17: Axial (top row) and sagittal (bottom row) slices of C_{max}-based maps (Eq. (7.15)) for $F = F_{new}$ (left column, Eq. (7.10)) and $F = F_{old}$ (right column, Eq. (7.9)) seeded in the splenium of the corpus callosum (annotated point) of the HCP data set. The F_{new}-based metric follows anatomy more closely than the F_{old}-based one, exemplified here by the high values of the latter observed in the frontal part of the corpus callosum.

Order	F_{new}	F_{old}
2	**0.0035**	**0.0061**
4	**1.7×10^{-4}**	0.0087
6	0.035	0.030

Table 7.3: p-values for the MANOVA comparison between the local efficiency vectors of the TDC and ASD groups. Column heading F_{new} indicates connectivity computed using $F = F_{new}$ metric, whereas F_{old} denotes the metric proposed by Melonakos et al. [227]. Connectivity matrices were computed using a spherical harmonic representation of the diffusion ODF with three different orders: 2, 4 and 6. All connectivity measures were computed with the C_{max} path measure. Since a total of six tests were performed, the Bonferroni-corrected threshold for significance was $0.05/6 \approx 0.0083$. p-values below this threshold (shown in **bold**) indicate rejection of the null hypothesis that there is no difference between the two groups.

7.4   Discussion

In this chapter we have looked at geodesic tractography based on Riemannian geometry, where we have illustrated potential benefits that come with moving to a scaled manifold as explained in Chapter 2. We have also described a new, publicly available 3D Slicer module that can be used to perform Finsler geodesic tractography and connectivity mapping, using a newly proposed (ad hoc) Finsler function inspired by the scaled Riemannian metric. The current release is unfortunately not directly compatible with the Finslerian framework that we discussed in Chapter 3, which is one of the main topics of interest for future work. Finally we looked at the different path measures C_{avg} and C_{max} that have been proposed to compute geodesic connectivity maps, showing that there are clear advantages, both theoretical and practical, to using the latter one. In this last section we will discuss these results in some detail.

7.4.1   Geodesic tractography

We computed geodesics on both synthetic and patient data, wherewith we compare the adjugate-based metric $\tilde{g}$ (Eq. (7.8)) and the standard inverse-based metric (Eq. (2.15)). We also compare the results in patient data with probabilistic and deterministic streamline tractography, using the industry-standard constrained spherical deconvolution (CSD [342]) model. Note that we do not mask regions of isotropic diffusivity in our processing pipeline.

Results on the synthetic data, Figs. 7.3, 7.4, and 7.5, show that geodesics computed with the standard inverse diffusion tensor fail to describe the fibers even for noise-free data without significant sharpening ($s = 4$, Eq. (7.17)). Geodesics obtained with the adjugate-based metric, either with or without sharpening, follow the expected trajectories rather well in all considered noise scenarios and without taking shortcuts through the isotropic background. Both metrics become increasingly sensitive to noise when a high sharpening power is used, which has been noted in other works as well [163], while for low noise levels sharpening improves tractography substantially. In our experiments, the adjugate-based metric appears less sensitive to noise than the standard metric, in particular when the applied sharpening is limited.

In real brain data, Figs. 7.6, 7.7, and 7.8, geodesic tracts obtained with the adjugate avoid isotropic diffusion regions such as ventricles, while the standard metric produces tracts that almost always move through these isotropic regions. This remains the case even with significant sharpening, which provides little benefit for the selected seed regions. As the ventricles are entirely devoid of fibers, we find that the adjugate-based metric leads to better results than the standard metric in this scenario, consistent with the conclusions from the synthetic data experiments. The positive performance of the adjugate on real diffusion data was also noted in the
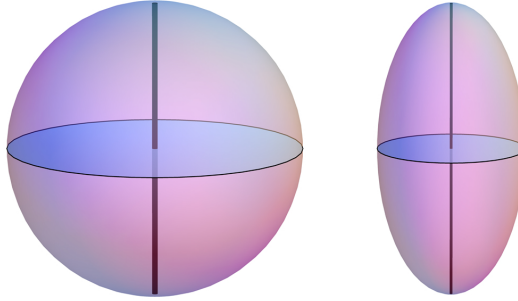
Figure 7.18: Graphical sketch of the quadratic forms corresponding to a typical diffusion tensor D in an isotropic region (left) and in a vertically oriented fiber bundle (right); the vertical axis corresponds to $\lambda = \lambda_1$.

recently published comparison by Schober et al. [299]. Finally, we obtain comparable results for the corticospinal tract from scaled Riemannian geodesics and CSD tractography, cf. Fig. 7.9. There are minor differences between the two, the most notable being that CSD tracts barely reach the more anterior part of the motor cortex, but without ground truth data we cannot draw any decisive conclusions from this. A comparison to tractography results using the Finslerian framework proposed in Chapter 3 is future work.

7.4.2   The scaled Riemannian manifold

The rather different behaviour of the two considered Riemannian metrics can be intuitively explained by the following argument, where we omit the effect of sharpening for simplicity. Consider two neighbouring voxels with a typical diffusion tensor $D = \text{diag}(\lambda, \lambda, \lambda)$ in an isotropic region and $D = \text{diag}(\lambda_1, \lambda_2, \lambda_3)$, with $\lambda_1 > \lambda_2 = \lambda_3$, in a vertically oriented fiber bundle (see Fig. 7.18). Consider first the case $\lambda = \lambda_1$. With the standard metric, the Riemannian cost (Eq. (2.15)) of travelling along an infinitesimal, vertically oriented line element, is proportional to $1/\lambda_1$. As such, the standard metric obviously assigns the same cost in both the isotropic and anisotropic voxels. Using the scaled metric however, the cost of moving along those same line elements is proportional to the area of the orthogonal cross-section (indicated by the shaded equatorial planes) i.e., to the product of λ_2 and λ_3. This leads to a relatively smaller cost in the anisotropic voxel, and we find that the metric favors tracking through anisotropic regions over isotropic regions.

This argument holds as well when the isotropic regions have greater diffusivities than those present in nearby anisotropic voxels ($\lambda > \lambda_1$), such as in the presented synthetic experiments and (commonly) in real diffusion data. In fact, the classical metric is only able to avoid isotropic regions when $\lambda < \lambda_1$, while the scaled metric

favors anisotropic regions up to the limit $\lambda \leq \lambda_2, \lambda_3$. In this scenario the area of the orthogonal cross section in the isotropic case becomes equal to or smaller than the anisotropic one, and so does the Riemannian cost. However, such scenarios seemingly take place in real data only in regions with complex fiber configurations, which we of course do not want to avoid.

In terms of practicality, adjugate-based tractography has essentially no free parameters (such as sharpening powers or those related to the fiber orientation distribution in CSD) and hence works out-of-the-box. Though such free parameters do offer some flexibility when extracting specific bundles, they typically have to be chosen in an ad hoc and application-dependent way, and a globally satisfactory setting might not exist. Additionally, the scaled metric obviates to a large extent the need to mask cerebrospinal fluid, as geodesics tend to avoid such region by construction. Note that these advantages have no bearing on the respective validity of these metrics as models for large-scale diffusion, as discussed Section 2.2.3.

After its proposal [141], the adjugate metric was evaluated in the context of geodesic tractography of subcortical u-fibers in the related work by Sepasian et al. [306], showing an improved reconstruction of fiber pathways that are challenging for traditional diffusion tensor tractography methods [58]. In future work it would be interesting to see how the adjugate metric fares against deconvolution sharpened metrics [103] and to the different scaling approach suggested by Hao et al. [163].

7.4.3 Geodesic connectivity analysis

We have also described the basic ideas for geodesic connectivity analyses, and introduced the open-source Finsler tractography module (FTM) for Finsler geodesic tractography and connectivity studies. The FTM is based on the `finslertract` project of Antonio Tristán-Vega, and includes changes that reflect recent advances in Riemannian geodesic tractography [141, 269]. The connectivity analysis capacities of the module were evaluated on Human Connectome Project (HCP) data, with results depending heavily on the choice of Finsler function and path/connectivity measure.

7.4.3.1 Influence of the Finsler function

We considered two different Finsler metrics, one based on a newly proposed Finsler function F_{new} derived from the work in Chapter 2 [141] in the Riemannian setting, and the F_{old}-based metric originally proposed by Melonakos et al. [227], which is the one typically used in literature. While a visual assessment of the distance maps obtained with each approach is not informative (see e.g. Fig. 7.10), the corresponding path measure maps highlight some interesting differences. In particular, the maps obtained using the Finsler function $F = F_{\text{new}}$ are more faithful to the known anatomy of tracts, as illustrated by the examples in Section 7.3.2. Although both the F_{new}-

and F_{old}-based maps suffer some 'leakage' problems, i.e., high connectivity values spreading to nearby but unrelated tracts, the F_{new}-based metric is much more robust to this issue compared to the F_{old}-based metric. This is especially clear near the corpus callosum, as can be seen in the cingulum results shown in Fig. 7.12. Thus, in addition to its more rigorous theoretical foundation, the F_{new} Finsler function typically results in connectivity maps which are anatomically more reliable.

In contrast to the F_{old}-based Finsler metric, the new metric is designed to correspond to a theoretically well-founded Riemannian metric [141], cf. Section 7.2.1.2. However, the fact that the generalized Finslerian metrics used here are defined in an ad hoc manner remains a significant issue with the interpretation of Finsler geometrical analyses. Given the improved results obtained with the relatively simple F_{new} Finsler function, it will be worthwhile to investigate the application of the fundamental work presented in Chapter 3 to the current analysis pipeline in future work.

7.4.3.2   Influence of the path measure

In Section 7.2.1.3 we explained how the C_{max} measure used in the Riemannian setting by e.g. Pechaud et al. [269] can be applied in the Finslerian setting. This measure is based on the 'weakest link' of the geodesic. That is, geodesics between strongly connected points should have a continuously low cost along the entire tract, while even small regions of high cost along the geodesic are taken to significantly decrease the likelihood of the points being structurally connected. Compared to the more common C_{avg} measure, which considers the average of some diffusivity measure along a geodesic, the C_{max} path measure has a number of important advantages, both theoretical and practical.

Primarily, the C_{max} measure associates a relatively low connectivity to geodesics taking shortcuts, a notorious problem of geodesic tractography [163, 305]. In the same vein, provided a sufficiently fine spatial resolution, low connectivity is associated with geodesics that jump from one fiber bundle to another across a small region with high cost. Because of this, in practice, regions of high connectivity tend to concentrate much more on the seeded fiber bundles, which again significantly reduces leakage artifacts. One can appreciate this especially in the difference between the C_{avg} and C_{max} path measure maps seeded in the arcuate fasciculus (Figs. 7.13 and 7.14), where the consistently low cost in the corpus callosum results in an above average C_{avg}-based connectivity for all geodesics that cross the corpus callosum to the other hemisphere. This in fact highlights another theoretical advantage; the C_{max} measure has the very natural property of being monotonic, i.e., it cannot increase in connectivity with distance along a path. Combined, these properties result in C_{max}-based connectivity maps that are generally much closer to anatomy than maps produced using the C_{avg} measure.

However, the C_{avg} measure can lead to greater contrast than C_{max} at a *local* level, as can be seen in the left columns (F_{new}-based metric) of the arcuate fasciculus results, Figs. 7.13 and 7.14. C_{max}-based maps typically have low homogeneous connectivity throughout the white matter, which essentially results in a default situation in which everything is (at least weakly) connected to everything else, while C_{avg}-based maps can efficiently extract the local, more direct structural connections.

The large differences between the two path measures highlight an issue with the currently employed definitions. The different types of information captured by these measures make it clear that other path measures, or more likely combinations of various measures, can and should be developed to obtain a more complete characterization of the structural connectivity captured by the geodesics. Because the validation of connectivity measures is very challenging, a possible next step could be a structured inclusion of a complete set of descriptive measures, e.g. shape measures of increasing complexity, subject to natural constraints like scale and orientation invariance. The introduction of anatomical priors, which has become more common in recent works [315], could also be used to further improve geodesic-based connectivity analysis.

7.4.4 Group differences in autism spectrum disorder

Finally, we have studied group differences in a graph-theoretical analysis of brain networks in autism spectrum disorder (ASD), where we found significant differences in the local network efficiency between the ASD group and the normal developing controls. Results seemed slightly more robust with the new metric, while both metrics suffered from overfitting when the spherical harmonics order was taken greater than four, Table 7.3. Differences in local network efficiency have been widely reported in previous works, and our results corroborate previous findings of abnormalities [221, 292]. Note that this analysis was not intended to be exhaustive, but rather to illustrate the application of our new connectivity framework in a standard brain network analysis setting. In future work, we will conduct a more thorough analysis of brain networks based on Finsler connectivity.

7.4.5 Concluding remarks

We have discussed several modifications to the standard geodesic tractography and connectivity pipelines for diffusion MRI, motivated by theoretical arguments presented in Part I. First, we showed that using a scaled Riemannian metric has practical benefits in geodesic tractography for fiber pathways that lie near regions of high diffusivity. Next, we concluded that replacing the standard choices in the connectivity pipeline with the weakest link path measure C_{max} (Eq. (7.15)) and the newly proposed Finsler function F_{new} (Eq. (7.10)), resolves a number of clearly identifiable issues.

These changes guarantee monotonicity of the connectivity along a path, effectively eliminate artificially high connectivity along paths that cross the corpus callosum, and significantly reduce leakage of high connectivity to nearby pathways. Further improvement, in both tractography and connectivity experiments, is likely to result from the canonical definition of the Finsler metric which we established in Chapter 3 but did not use here. The modified algorithms are publicly available as 3D Slicer modules (`github.com/tomdelahaije/fcm`, `nitrc.org/projects/riemantract`).

The main problem addressed only superficially in this work is the validation of geodesic tractography and the associated connectivity analyses. As with available alternatives, the absence of ground truth data makes it very challenging to validate and compare results. Recent developments have focused on the generation of simulated data with a known ground truth [245], and a lot of promising work is being done to improve histological validation [304].

Chapter 8

Surface

Quantifying Sheet Structure

Based in part on:

C. M. Tax, T. Dela Haije, A. Fuster, C.-F. Westin, M. A. Viergever, L. Florack, and A. Leemans. "Sheet Probability Index (SPI): Characterizing the geometrical organization of the white matter with diffusion MRI". In: *NeuroImage* 142 (Nov. 2016), pp. 260–279. ISSN: 10538119. DOI: `10.1016/j.neuroimage.2016.07.042`. Joint first authors Tax and Dela Haije.

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 literature [57, 364, 365], and motivated the work we present in this final chapter. Wedeen et al. [365] analyzed adjacency and crossings between cerebral fiber pathways using diffusion MRI and found that the pathways form a three-dimensional grid structure. More specifically, the authors used diffusion spectrum imaging (DSI) [50, 367], which infers information on the ensemble average propagator (EAP) 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 two-dimensional grid- or sheet-like structures (similar to the "warp and weft of a fabric") that are layered in space "like pages of a book" [364, 365]. This sheet structure was consistently recognized across species and scales, and throughout the white matter.

Catani et al. [57] suggested that the observed grid pattern is most likely an artifact, attributed to the limitations of DSI used in Wedeen et al. [365]. The authors showed that diffusion orientation distribution functions (dODFs) as derived from the DSI EAPs 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 the work of Wedeen et al. [365] and inadvertently makes "the grid structure of interwoven sheets a very likely configuration" [57]. By using another diffusion MRI technique called spherical deconvolution [96, 97, 342], 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). Additionally, Catani et al. noted that the presented results [365] are mainly qualitative, and that diffusion MRI-based pathways cannot be equated to true axons.

Wedeen et al. [364] 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" [364].

This debate has gained a lot of attention from the diffusion MRI 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 impor-

tant role in axonal path-finding during embryogenesis by guiding growing fibers, thereby greatly reducing the complexity of such processes [365]. 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 [15, 17, 98, 124, 216, 285, 295, 338, 382].

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 chapter, we first formalize the terminology to clarify the definition of sheet structure as proposed by Wedeen et al. [365] and the condition for its existence; this is done in Section 8.1. 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 Sections 8.2.1 and 8.2.2. An intuitive way to visualize the SPI throughout the brain is described in Section 8.2.3. We evaluate this method with simulations and real diffusion MRI data sets as described in Sections 8.2.4 and 8.3. Finally, we use the proposed method to investigate and discuss some of the claims made in Wedeen et al. [364, 365] and Catani et al. [57] in Sections 8.3 and 8.4. Note that even though we use diffusion MRI 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 [20]. Preliminary results of this work have been presented at the BASP workshop and at the ISMRM [328–331].

Remark 14. *We move away from the geometrical frameworks and their applications, which stood central in the majority of previous chapters. The presented work on the sheet structure hypothesis only requires a reliable way to reconstruct fiber pathways, and we focus here on a number of well-established techniques that have played a role in the debate thus far.*

8.1 Theory

Several definitions and interpretations of sheets in the brain exist in the literature [199, 301, 355, 380, 385]. Some major white matter tracts resemble a thin sheet-like structure by themselves: 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 [316] or a surface [380]. In contrast, sheets formed by crossing or intertwining fibers as proposed by Wedeen et al. [365] (hereafter shortened to sheets) are composed of two families of tracts that cross each other on the same surface in certain regions along their trajectories. As a result, large bundles of tracts can in theory still form

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

We first present the theoretical background required to understand some of the key mathematical concepts in relation to the investigation of the brain's sheet structure. In Section 8.1.1 we present a definition of sheet structure in terms of integral curves of vector fields, and in Section 8.1.2 we discuss the relevant measure used to assess the presence of sheet structure, namely the Lie bracket. Section 8.1.3 covers the relation between the Lie bracket of two vector fields and flows along them, which we use later on in this chapter to compute the Lie bracket from diffusion MRI data. Finally we discuss in Section 8.1.4 the necessary condition on the Lie bracket for sheet structure to exist.

8.1.1 Definition of sheet structure

Consider a set of unit vectors at each position of the brain $M \subset \mathbb{R}^3$, which can for example be obtained from the dODF, fODF, or barrier ODF (Section 6.1.1) using any diffusion MRI reconstruction technique. In streamline tractography, discussed briefly in earlier chapters, one integrates a three-dimensional (Lipschitz continuous) unit vector field v defined on a subset $N_v \subset M$, generating streamlines (or integral curves) $\Phi_t^v(p)$[1] passing at time $t = 0$ through the initial position $p = \Phi_0^v(p)$, such that the local tangents to Φ_t^v are given by the vector field v. In other words, for every $p \in N_v$, and every t such that $\Phi_t^v(p) \in N_v$, the following equality holds (Fig. 8.1(a)):

$$\frac{\mathrm{d}}{\mathrm{d}t} \Phi_t^v(p) = v_{\Phi_t^v(p)}. \tag{8.1}$$

Here, $v_p \in \mathbb{R}^3$ denotes the vector at position p (recall Section 1.3), and $\Phi_{\beta t}^v = \Phi_t^{\beta v}$ for $\beta > 0$. The set of integral curves for a given vector field v and all $p \in N_v$ defines a one-dimensional foliation of N_v. Cf. Spivak [319] for further details.

Sheets are defined by extending this idea of integrability 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 whose tangent plane at p is parallel to the plane spanned by v_p and w_p for all $p \in S$, see Fig. 8.1(b). If v and w form sheets throughout a set $N_S \subset N_v \cap N_w$, the integral surfaces form the leafs of a two-dimensional foliation of N_S. Three vector fields u, v, and w can similarly form a grid structure, i.e., a three-dimensional foliation of a common set. Grid structure exists when each pair of vector fields forms a sheet.

Points on a single sheet structure formed by v and w can be reached from any other point on the sheet through a flow along v, followed by a flow along w. Grid structure similarly ensures that any two points on the grid are connected through

---

[1]Here we adopt the flow operator notation Φ^v to represent the streamlines.
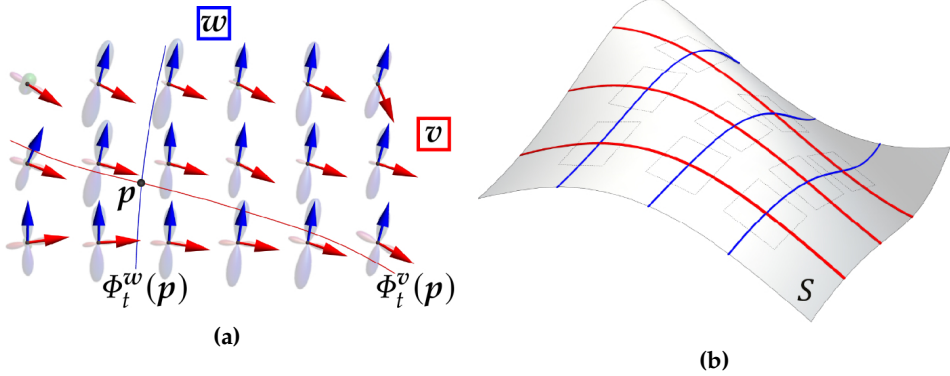
Figure 8.1: (a) An example set of fiber ODFs (semi-transparent) along with their peak directions (arrows), which form the vector fields v (red) and w (blue). By integrating these vector fields, one can reconstruct at each position p the integral curves $\Phi_t^v(p)$ (red curve) and $\Phi_t^w(p)$ (blue curve). (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).

at most three flows along the vector fields. Grid structure on the whole of M would thus imply that any two points in the brain can be connected by means of such simple flows: the integral curves along u, v, and w are coordinate lines on M.

Integration to a sheet or grid structure clearly requires intricate long-range correlations in the vector fields, and continuity of the vector fields is not sufficient for sheets to exist. The exact requirement for the existence of sheet structure is formalized in the Frobenius theorem, which imposes a condition involving the Lie bracket $[v, w]$ of the vector fields v and w.

8.1.2 The Lie bracket

The Lie bracket of two (non-zero) vector fields v and w on M is defined by its action on smooth functions f:

$$[v, w]_p(f) = v_p w(f) - w_p v(f). \tag{8.2}$$

The Lie bracket is a linear operator that satisfies Leibniz' product rule, and so $[v, w]_p$ can be regarded as a directional derivative, i.e., $[v, w]_p$ is an element of the tangent 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. (8.2) trivially gives

$$[v, w]_p = v_p w - w_p v, \tag{8.3}$$

which is commonly known as the commutator of the vector fields.

8.1.3 Flows and the Lie bracket

Intuitively, the Lie bracket at a given location p can be understood as the deviation from p when trying to move around in a small loop along the integral curves of v and w. To see this, let v and w be non-collinear vector fields $M \to TM$, and define an effective time-dependent Lipschitz continuous vector field x given by

$$x(t) = a(t)v + b(t)w, \tag{8.4}$$

with a and b the loop characteristics; almost everywhere differentiable scalar functions that are nowhere identically zero simultaneously, and with antiderivatives A and B that satisfy for some $T > 0$

$$A(T) = B(T) = 0, \tag{8.5}$$

Next we define the integral curve α_p of x as the solution to

$$\begin{cases} \frac{\partial \alpha_p}{\partial t}(\beta, t) = \sqrt{\beta}\, x(t)\, \alpha_p(\beta, t) \\ \alpha_p(\beta, 0) = p \end{cases}. \tag{8.6}$$

where $\beta \geq 0$ is a parameter used to scale the path loops. A smaller β effectively decreases the norm of $x(t)$, resulting in shorter paths with the same loop characteristics specified by the functions a and b. α_p can be expanded [205] to obtain

$$\alpha_p(\beta, t) = p + \sqrt{\beta} \int_0^t x_p(t')\, dt' + \beta \int_0^t x_p(t') \int_{t'}^T x(t'')\, dt''\, dt' + O\left(\beta^{3/2}\right), \tag{8.7}$$

which after substitution of Eqs. (8.4) and (8.5) reduces at $t = T$ to

$$\alpha_p(\beta) := \alpha_p(\beta, T) = p - \beta \left(\int_0^T a(t')B(t')\, dt'\, v_p w + \int_0^T A(t')b(t')\, dt'\, w_p v \right) + O\left(\beta^{3/2}\right). \tag{8.8}$$

Integration by parts finally gives

$$\lim_{\beta \to 0} \frac{\alpha_p(\beta) - p}{\beta} = \frac{d\alpha_p}{d\beta}(0) = c_{a,b}(T)\,[v, w]_p, \tag{8.9}$$

where the proportionality constant $c_{a,b}$ is given by

$$c_{a,b}(T) := \int_0^T A(t')b(t')\, dt'. \tag{8.10}$$

Consider for example a loop consisting of four 'legs' (i.e., a quadrilateral), that

starts at point p by following the integral curve $\Phi_t^v(p)$ along v for some distance $\sqrt{\beta}$. From the end point $\Phi^v_{\sqrt{\beta}}(p)$ of the first leg, continue along the vector field w for the same distance $\sqrt{\beta}$. From the resulting end point of this second leg $\Phi^w_{\sqrt{\beta}} \circ \Phi^v_{\sqrt{\beta}}(p)$, where $\circ$ indicates concatenation, move 'backwards' along the integral curves of $-v$ and $-w$ in order, again with the same time $t = \sqrt{\beta}$. The final end point is then $\left(\Phi^{-w}_{\sqrt{\beta}} \circ \Phi^{-v}_{\sqrt{\beta}} \circ \Phi^{w}_{\sqrt{\beta}} \circ \Phi^{v}_{\sqrt{\beta}}\right)(p)$, see Fig. 8.2(a). This loop is the integral curve of the β-scaled effective vector field x with

$$
a(t) = \begin{cases} 1 & \text{for } 0 \leq t < 1 \\ 0 & \text{for } 1 \leq t < 2 \\ -1 & \text{for } 2 \leq t < 3 \\ 0 & \text{for } 3 \leq t \leq 4 \end{cases}
\tag{8.11}
$$

and

$$
b(t) = \begin{cases} 0 & \text{for } 0 \leq t < 1 \\ 1 & \text{for } 1 \leq t < 2 \\ 0 & \text{for } 2 \leq t < 3 \\ -1 & \text{for } 3 \leq t \leq 4 \end{cases} .
\tag{8.12}
$$

We then find that the curve α_p formed by appropriate scaling through β is related to the Lie bracket as

$$
\frac{\partial \alpha_p}{\partial \beta}(0) = [v, w]_p.
\tag{8.13}
$$

The asymptotic behavior for small β defines the closure

$$
R_p(\beta) := \alpha_p(\beta) - p \sim \beta[v, w]_p \qquad (\beta \to 0),
\tag{8.14}
$$

that we rely on in our implementation of the Lie bracket on discrete data.

8.1.4 The Frobenius theorem

The requirements for sheet structure to exist are given by the Frobenius theorem, which can be understood intuitively by following the line of thinking in the previous section. The Lie bracket was introduced as the local deviation from p after small flows over v and w, and if these vector fields are tangent to a sheet in a neighborhood of p then the end points α_p of these flows must lie on the sheet structure as well (Fig. 8.2(c)). In the limit of smaller and smaller loop sizes this leads to the Frobenius theorem [207, 319], which states that two vector fields 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
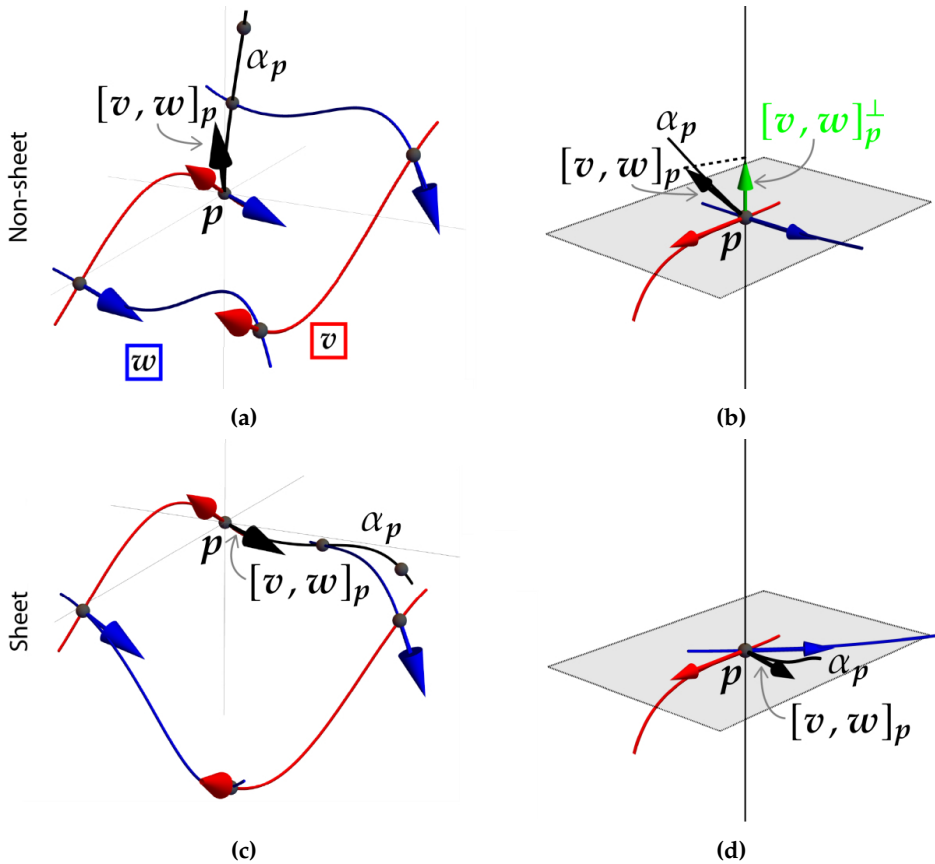
Figure 8.2: (a) A loop composed of integral curves of v (red arrows) and w (blue arrows) with p as the starting position. α_p is the curve formed by the end points of all loops (by varying the scaling parameter β) starting at the point p. $[v, w]_p$ is the Lie bracket at point p, and is related to the difference vector $R_p(\beta) := \alpha_p(\beta) - p$ according to Eq. (8.14). In this scenario, the vector fields v and w cannot be integrated to form a sheet structure. (b) The Lie bracket depicted in (a) does not lie in the plane spanned by v_p and w_p (gray). Hence the normal component of the Lie bracket $[v, w]_p^{\perp}$ defined in Eq. (8.15) is non-zero (green arrow), and the vector fields cannot be integrated to form a sheet structure. (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 $[v, w]_p$ depicted in (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.

spanned by v_p and w_p, cf. Fig. 8.2(d). As a comparison, Fig. 8.2(b) shows the case (corresponding to Fig. 8.2(a)) 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 [365]. We do this by computing the projection of the vector $[v, w]_p$ onto the outer product $v_p \times w_p$ (normalized):

$$[v, w]_p^{\perp} := [v, w]_p \cdot \frac{v_p \times w_p}{\|v_p \times w_p\|}. \tag{8.15}$$

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 between v and w.

8.2   Methods

We can examine the existence of sheet structure by (1) estimating the integral curves and loops in Eq. (8.9), (2) estimating the Lie bracket based on Eq. (8.14), and (3) extracting the normal component of the estimated Lie bracket (Eq. (8.15)) 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 8.1 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 an ODF). These issues complicate the actual computation of a Lie bracket.

In Section 8.2.1 we describe the approach to compute the discrete Lie bracket, inspired by the qualitative reconstruction of path neighborhoods in Wedeen et al. [365]. Our method can deal with noisy vector fields and addresses the problem of sorting vectors in a neighborhood of a point p. In Section 8.2.2 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 8.2.3 we define the sheet tensor, which allows us to visualize the Lie bracket for every pair of vector fields. Finally, in Section 8.2.4 we describe the simulated and acquired MRI data used for the experiments.

8.2.1   The discrete Lie bracket

We propose here to calculate a discrete Lie bracket by approximating the integral curves in Eq. (8.9) with tractography [328, 329, 368], and by computing a large number of loops with configurations as in Fig. 8.1(a). The tractography process and the

inclusion of multiple loop configurations handle noise in the Lie bracket computation implicitly (but partially).

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 8.2.1.1 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 8.2.1.2 we outline the frame tractography that performs clustering during tract propagation. Finally, the estimation of the Lie bracket is discussed in Section 8.2.1.3.

8.2.1.1 Clustering of frames

Consider an ordered set of n vectors $[y_1, \ldots, y_n]$ at a position q, and m vectors $\{z_1, \ldots, z_m\}$ in some point p near q[1]. The ordered set $[y_1, \ldots, y_n]$ serves as a reference frame, i.e., we assume that n vector fields $u, v, \ldots$ are present in the local neighborhood of q that satisfy $u_q = y_1, v_q = y_2, \ldots$. The aim of the clustering algorithm is to find a permutation of the frame (an ordered set $[z_{P_1}, \ldots, z_{P_n}, 0, \ldots]$ that corresponds to the reference frame, so that we can take $u_p = z_{P_1}, v_p = z_{P_2}, \ldots$ for some permutation P of $[1, \ldots, n]$ (see Fig. 8.3 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 that represents the total element-wise similarity between the frames $[y_1, \ldots, y_n]$ and $[z_{P_1}, \ldots, z_{P_n}]$. In the algorithm below we will use the 'total cosine similarity', which is defined as the sum of the cosines of the angles between corresponding vectors. The global steps in the clustering algorithm can be found in Alg. 2

8.2.1.2 Frame tractography

Given a step size Δh and a distance h, we can approximate the flow along a unit vector field $x \in \{u, v, \ldots\}$ with streamline tractography [242]. The algorithm explained in Alg. 3 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.

---

[1]If $m < n$, we append $n - m$ zero vectors to the list $\{z_1, \ldots, z_m\}$, so in the following we can take $m \geq n$.
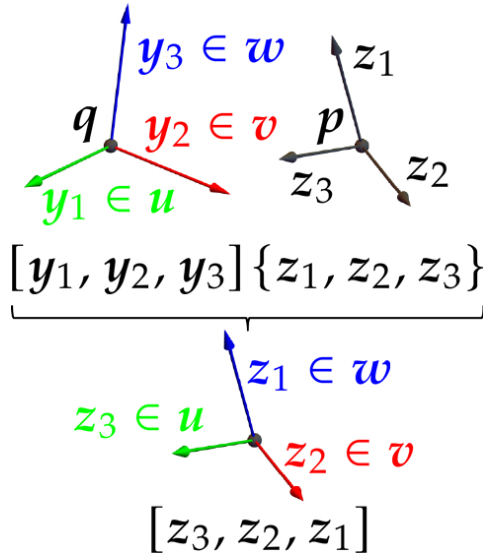
Figure 8.3: Example of the clustering of frames. We have an ordered set of vectors $[y_1, y_2, y_3]$ at a position q, and 3 vectors $\{z_1, z_2, z_3\}$ in some point p near q. We assume that they are assigned to the vector fields u, v, and w as follows: $u_q = y_1$, $v_q = y_2$, and $w_q = y_3$. Frame clustering yields the ordered set $[z_3, z_2, z_1]$.

Algorithm 2: Frame clustering algorithm

Data: An ordered frame $[y_1, \ldots, y_n]$ and a set of vectors $\{z_1, \ldots, z_m\}$;
Parameters: Angle threshold t (default $\cos 35°$);
Result: The ordered frame $[\tilde{z}_1, \ldots, \tilde{z}_n]$;
```
/* Find the n-permutation Π ∈ Aut({1,...,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]) for which the total similarity is
   maximized, and recall that yᵢ and zⱼ are unit or zero vectors                              */
```
$\Pi \leftarrow \mathrm{argmax}_{P \in \mathrm{Aut}(\{1,\ldots,m\})} \sum_{i=1}^{n} |y_i \cdot z_{P_i}|$;
$\tilde{z} \leftarrow [z_{\Pi_1}, \ldots, z_{\Pi_n}]$ **foreach** $i \in \{1, \ldots, n\}$ **do**
```
    /* Apply an angle threshold t                                                            */
```
    if $|y_i \cdot \tilde{z}_i| > t$ **then**
       | $\tilde{z}_i \leftarrow 0$;
```
    /* Align the directions of the vectors in the frames                                     */
```
    else if $y_i \cdot \tilde{z}_i < 0$ **then**
       | $\tilde{z}_i \leftarrow -\tilde{z}_i$;
    end
end

Algorithm 3: Frame tracking algorithm

Data: Vector fields $\{z_1, z_2, \dots\}$ defined in a neighborhood of an initial point q_0;
Parameters: The step size Δh and the distance h, which give $L \leftarrow \lfloor h/\Delta h \rfloor$;

Result: The end point q_L of the approximate flow $\widehat{\Phi}_h^x(q_0)$;

```
/* Initialize counter                                              */
```
$l \leftarrow 0$;
```
/* Assume the frame at q_0 is ordered                              */
```
$[u_{q_0}, v_{q_0}, \dots] \leftarrow [z_1, z_2, \dots]_{q_0}$;
```
/* Until the maximum number of steps is reached                    */
```
while $l \leq L$ **do**
    ```/* Take a step of size Δh in the direction of x_{q_l}   */```
    $q_{l+1} \leftarrow q_l + \Delta h \, x_{q_l}$;
    ```/* Cluster the frame at q_{l+1} using the frame at q_l   */```
    $[u_{q_{l+1}}, v_{q_{l+1}}, \dots] \leftarrow \mathrm{cluster}\left([u_{q_l}, v_{q_l}, \dots], \{z_1, z_2, \dots\}_{q_{l+1}}\right)$ (Alg. 2);
    ```/* Return an error if clustering is unsuccessful   */```
    if $v_{q_{l+1}} = 0$ **then**
        | Error;
    end
    ```/* Increment counter   */```
    $l \leftarrow l + 1$;
end

8.2.1.3 The closure and the Lie bracket

To calculate the discrete Lie bracket we reconstruct approximate flow loops through repeated application of Alg. 3, 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. (8.9). 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. 8.4):

$$\widehat{R}_1 := \left(\widehat{\Phi}_{h_2}^{-w} \circ \widehat{\Phi}_{h_1}^{-v} \circ \widehat{\Phi}_{h_2}^{w} \circ \widehat{\Phi}_{h_1}^{v}\right)(p) - p, \tag{8.16}$$

$$\widehat{R}_2 := p - \left(\widehat{\Phi}_{h_1}^{-v} \circ \widehat{\Phi}_{h_2}^{-w} \circ \widehat{\Phi}_{h_1}^{v} \circ \widehat{\Phi}_{h_2}^{w}\right)(p), \tag{8.17}$$

$$\widehat{R}_3 := \left(\widehat{\Phi}_{h_2}^{w} \circ \widehat{\Phi}_{h_1}^{v}\right)(p) - \left(\widehat{\Phi}_{h_1}^{v} \circ \widehat{\Phi}_{h_2}^{w}\right)(p). \tag{8.18}$$

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 the algorithm). Note that we thus sample all four 'quadrants' surrounding point p.

When all difference vectors are estimated, we can compute an estimate of the Lie bracket $\widehat{[v, w]}_p$ with a simple linear least squares fit corresponding to Eq. (8.14), and the normal component of the estimated Lie bracket follows from Eq. (8.15). In practice, we only compute $\widehat{[v, w]}_p$ when the number of successfully estimated loops and corresponding difference vectors exceeds a minimum threshold.
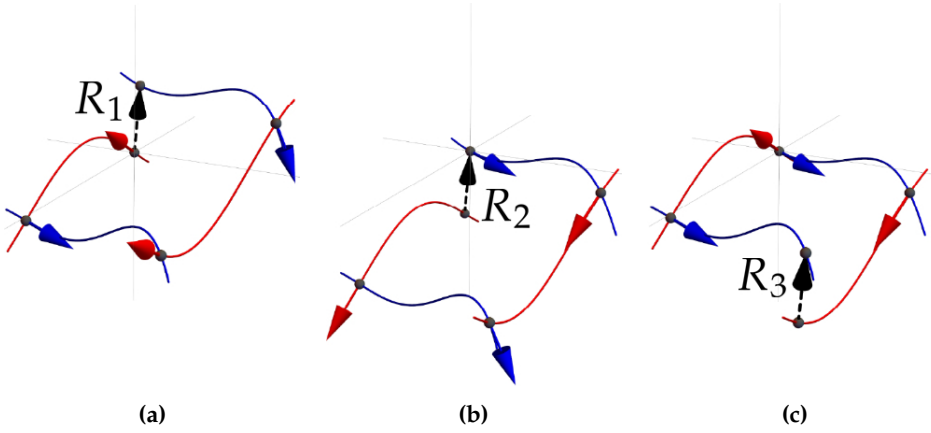
Figure 8.4: Loops that lead to $\widehat{R}_1 := \left(\widehat{\Phi}_{h_2}^{-w} \circ \widehat{\Phi}_{h_1}^{-v} \circ \widehat{\Phi}_{h_2}^{w} \circ \widehat{\Phi}_{h_1}^{v}\right)(\boldsymbol{p}) - \boldsymbol{p}$, $\widehat{R}_2 := \boldsymbol{p} - \left(\widehat{\Phi}_{h_1}^{-v} \circ \widehat{\Phi}_{h_2}^{-w} \circ \widehat{\Phi}_{h_1}^{v} \circ \widehat{\Phi}_{h_2}^{w}\right)(\boldsymbol{p})$ and $\widehat{R}_3 := \left(\widehat{\Phi}_{h_2}^{w} \circ \widehat{\Phi}_{h_1}^{v}\right)(\boldsymbol{p}) - \left(\widehat{\Phi}_{h_1}^{v} \circ \widehat{\Phi}_{h_2}^{w}\right)(\boldsymbol{p})$.

8.2.2 Sheet probability index

The algorithm described in Section 8.2.1 allows us to compute estimates of the Lie bracket normal component $\widehat{[\cdot,\cdot]}^{\perp}$ for every combination of vectors at every position in the brain, and according to the Frobenius theorem (Section 8.1.4) 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 to achieve the desired contrast) for the estimated distribution $\mathcal{N}(\mu, \sigma^2)$. 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 diffusion MRI sets, so we consider residual bootstraps as an alternative (see Section 8.2.4.2).

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 ($[\cdot,\cdot]_{\boldsymbol{p}}^{\perp}$ is likely close to 0), while a low value indicates that there are significant deviations from sheet structure ($[\cdot,\cdot]_{\boldsymbol{p}}^{\perp}$ likely differs significantly from 0 relative to the noise level).

8.2.3 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 $\binom{n}{2}$ SPIs. 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

$$S_p = \frac{P_\lambda}{\lambda_{\max}} \left(v_p \otimes v_p + w_p \otimes w_p \right).$$

(8.19)

Here $\lambda_{\max}$ denotes the largest eigenvalue of the tensor $\left(v_p \otimes v_p + w_p \otimes w_p \right)$ and $\otimes$ denotes the 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 (LR) direction gives a red tensor, normal in inferior-superior (IS) direction gives a blue tensor, and normal in anterior-posterior (AP) direction gives a green tensor, cf. Section 2.3.2) [383]. Furthermore, the size of the ellipsoid is determined by the SPI, where a larger SPI gives larger ellipsoids, and the division by $\lambda_{\max}$ fixes the largest semi-axis of the ellipsoid for a given SPI. The shape represents the angle between v and w. Fig. 8.5 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.

8.2.4 Data

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

8.2.4.1 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. 8.6(a)):

$$u = \left(-\sin\phi_1, \cos\phi_1\cos\theta_2, \cos\phi_1\sin\theta_2 \right),$$

(8.20)

$$v = \left(\cos\phi_2\cos\theta_1, -\sin\phi_2, \cos\phi_2\sin\theta_1 \right),$$

(8.21)

$$w = \left(\cos\phi_2\cos\theta_2, -\sin\phi_2, -\cos\phi_2\sin\theta_2 \right).$$

(8.22)
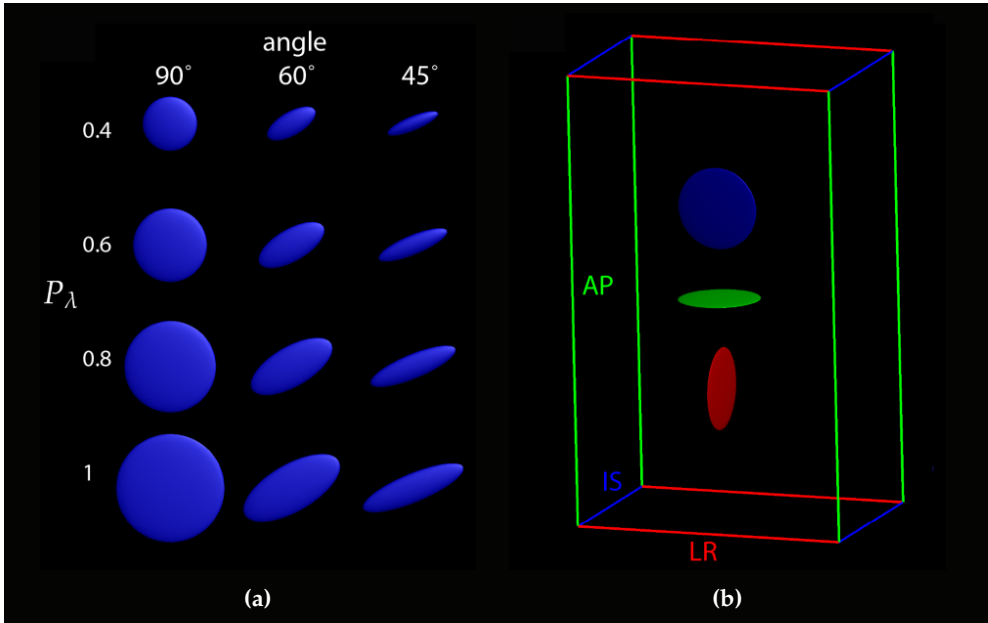
Figure 8.5: (a) Sheet tensors with a normal direction pointing towards the reader, for different angles between u_p and v_p and different sheet probabilities. Here, v is always oriented in left-right direction. (b) Sheet tensors with different orientations are colored according to their third eigenvector. AP is the anterior-posterior direction, IS is inferior-superior, and LR is left-right.

Here $\theta_i = \arctan \frac{x^i}{\sqrt{\rho^2-(x^1)^2-(x^2)^2}}$ and $\phi_i = \arccos \frac{x^i}{\rho}$ (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 the Lie bracket normal component $[u, v]_p^\perp = 0$ for all $p \in \{x \in \mathbb{R}^3 \mid (x^1)^2 + (x^2)^2 < \rho^2, x^3 = z\}$. u and w generally do not form a sheet, and the normal component of the Lie bracket $[u, w]_p$ at $p \in \{x \in \mathbb{R}^3 \mid (x^1)^2 + (x^2)^2 < \rho^2, x^1 \neq 0, x^2 \neq 0\}$ is given by

$$[u, w]_p^\perp = \frac{6x^1x^2\left(-\rho^2 + (x^1)^2 + (x^2)^2\right)}{\sqrt{\left(\rho^2 - (x^1)^2\right)\left(\rho^2 - (x^2)^2\right)\left(\rho^6 - 8\rho^2(x^1)^2(x^2)^2 + 4(x^1)^2(x^2)^2\left((x^1)^2 + (x^2)^2\right)\right)}}.$$

$$(8.23)$$

A plot of $[u, w]_p^\perp$ as a function of x^1 and ρ is shown in Fig. 8.6(b), where we take $x^2 = -x^1$ so that $[u, w]_p^\perp$ is generally greater 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
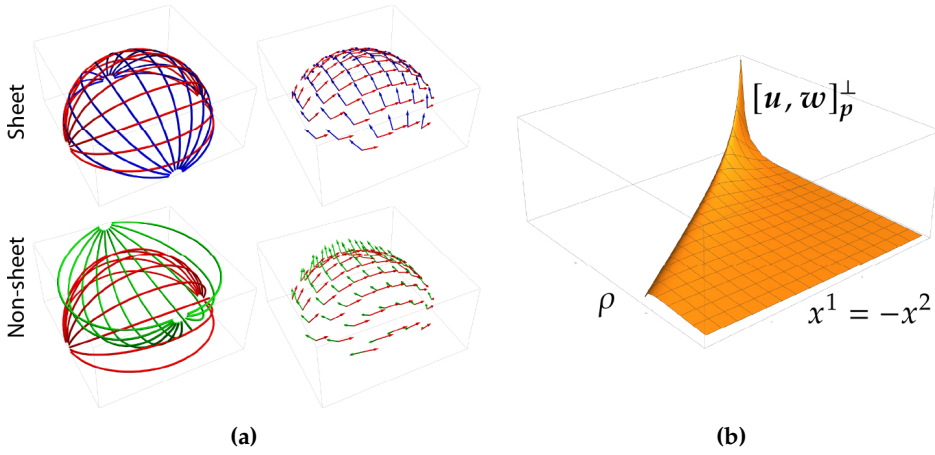
Figure 8.6: (a) Vector fields u (red), v (blue), and w (green), where u and v form a sheet and u and w do not. The left column shows a subset of integral curves, and vectors sampled on the upper hemisphere are shown on the right. This pattern of vector fields is repeated in the vertical direction. (b) Plot of $[u, w]_p^\perp$ as a function of ρ and x^1, with $x^2 = -x^1$.

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 [63, 226] with probability density function

$$f_W\left(\pm\tilde{v}_q; v|_q, k\right) = M\left(\frac{1}{2}, \frac{3}{2}, k\right)^{-1} e^{k\left(\tilde{v}_q \cdot v_q\right)^2}. \tag{8.24}$$

Here, $M\left(\frac{1}{2}, \frac{3}{2}, \cdot\right)^{-1}$ is the Kummer function [226], $\tilde{v}_q$ 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 smaller perturbations).

8.2.4.2 Diffusion MRI data

Simulations Diffusion MRI signals were simulated using a Zeppelin-Stick-Dot model [119] 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 data sets: single shell data sets with 90 directions and $b = 3000\,\text{s/mm}^2$ suitable for spherical deconvolution, and Cartesian sampled data sets with 514 directions (maximum b-value of $10000\,\text{s/mm}^2$) and one $b = 0\,\text{s/mm}^2$ point suitable for DSI (protocol corresponds to the MGH HCP DSI data, see Section 8.2.4.2).

Real data We use different diffusion MRI 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 [152, 318, 351]; (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 [136]; (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 [308]; and (4) a Cartesian sampled data set (514 directions) with b-values up to $10000 \, \text{s/mm}^2$ of the MGH-USC HCP with an isotropic voxel size of 2 mm (`humanconnectomeproject.org/data/inventory`).

Processing Data sampled on a single shell was processed using constrained spherical deconvolution (CSD, spherical harmonics up to order 8, see Section 6.1.1.3) [342] in ExploreDTI [215]. The response function for the simulated data was generated from the Zeppelin-Stick-Dot model, and the response function for real data was computed using recursive calibration [336]. Peaks were extracted using a Newton optimization algorithm [180] 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 diffusion MRI data from a single set of noisy measurements [181]. The peaks extracted from the different bootstrap realizations were clustered using the method described in Section 8.2.1.1, taking the peaks extracted from the original data as reference frames.

Cartesian sampled data was analyzed using the DSI model [367], which was reconstructed with Diffusion ToolKit using default settings [359]. 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.

8.3 Results

The results for analytical vector field simulations are presented in Section 8.3.1, for diffusion MRI simulations in Section 8.3.2, and for real diffusion MRI data in Section 8.3.3.

8.3.1 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
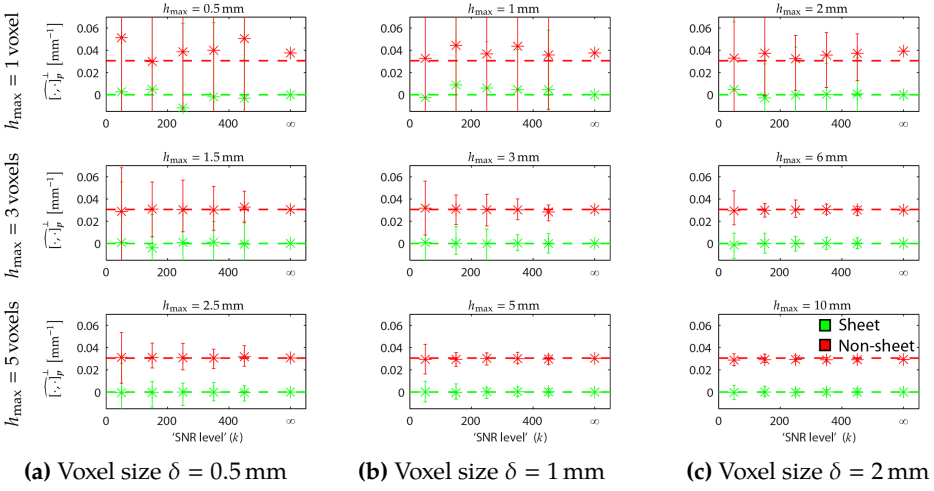
Figure 8.7: $\widehat{[\cdot,\cdot]}_p^{\perp}$ for different voxel sizes $\delta = \{0.5, 1, 2\}$ mm ((a), (b), (c)) and different settings for $h_{\max} = \{1, 3, 5\}$ voxels (the corresponding $h_{\max}$ in mm is noted above each plot). Each plot shows the mean and range of the estimates in the case of sheet (green, $[u, v]_p^{\perp} = 0$ indicated by the dashed line) and non-sheet (red, $[u, w]_p^{\perp} = 0.031$) for different SNR levels (i.e., the concentration parameter k, higher k means a smaller perturbation of the vectors). We used 50 noise iterations, $\kappa = 1/\rho = 1/26\,\mathrm{mm}^{-1}$, and $\boldsymbol{p} = (10, -10, 0)$.

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.

8.3.1.1 The influence of spatial resolution and noise

Fig. 8.7 shows results for different voxel sizes $\delta = \{0.5, 1, 2\}$ mm (a-c), different settings for $h_{\max}$ (rows) and different SNR levels k (the concentration parameter in Eq. (8.24), higher k indicates a smaller 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 $\kappa = 1/\rho = 1/26\,\mathrm{mm}^{-1}$ and estimate the normal component of the Lie bracket at $\boldsymbol{p} = (10, -10, 0)$ 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 (green, $[u, v]_p^{\perp} = 0\,\mathrm{mm}^{-1}$) and non-sheet (red, $[u, w]_p^{\perp} = 0.031\,\mathrm{mm}^{-1}$). 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 $h_{\max} = 3$ voxels compared to $h_{\max} = 1$ voxel, but remains similar when further increasing to $h_{\max} = 5$ voxels. We can see that $h_{\max} = 1$ voxel is generally too low to obtain a reasonable accuracy and precision, and to distinguish
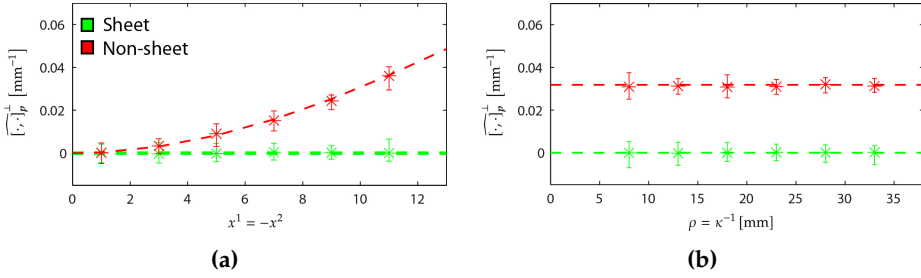
Figure 8.8: (a) Mean and range of $\widehat{[\cdot,\cdot]}_p^\perp$ for different points $p = (x^1, -x^1, 0)$ to vary the Lie bracket normal component magnitude $[u, w]_p^\perp$ ($\kappa = 1/26\,\text{mm}^{-1}$). (b) Mean and range of $\widehat{[\cdot,\cdot]}_p^\perp$ for different curvatures $\kappa = 1/\{8, 13, 18, 23, 28, 33\}$. Different curvatures were achieved by changing ρ in Eq. (8.20), and $[u, w]_p^\perp$ was kept constant using Eq. (8.23) 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 $k = 350$, $\delta = 1\,\text{mm}$, and $h_{\max} = 5$ voxels.

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, $\delta = 1\,\text{mm}$ and $h_{\max} = 5$ voxels, $\delta = 0.5\,\text{mm}$).

8.3.1.2 The influence of the Lie bracket normal component

Fig. 8.8(a) 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 $k = 350$, $\delta = 1\,\text{mm}$, and $h_{\max} = 5$ voxels. Here and in further analyses, we apply clustering of the vector fields as described in Section 8.2.1.1 (using the known vector fields as prior information gave similar results, not shown here).

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 $[u, w]_p^\perp \gtrsim 0.015\,\text{mm}^{-1}$.

8.3.1.3 The influence of curvature

Fig. 8.8(b) 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. (8.23)). $[u, v]_p^\perp$ (sheet) is evaluated at the same points for reference. We set $k = 350$, $\delta = 1\,\text{mm}$, and $h_{\max} = 5$ 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.

8.3.2 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 diffusion MRI images instead of perturbing vectors), the influence of the interpolation technique (nearest neighbor vs. fODF interpolation), and the influence of diffusion MRI 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.

8.3.2.1 The influence of noise and interpolation

Here, we extract fODFs and peak directions using CSD from the single shell simulated data. Fig. 8.9(a) 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 throughout the chapter, but include Fig. 8.10 for comparison which shows the same results with linear interpolation on the fODF spherical harmonic coefficients.

Similar to the vector field simulations in Fig. 8.7, 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. 8.9(a) at $h_{max} = 3$ voxels. This does not occur when using fODF interpolation (Fig. 8.10), 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.

8.3.2.2 The influence of diffusion MRI technique: CSD vs. DSI

In Fig. 8.9(b), 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
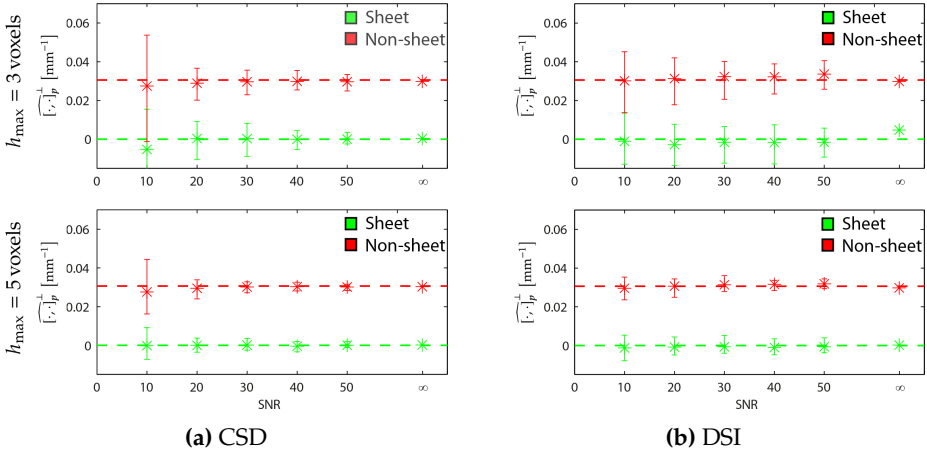
Figure 8.9: Mean and range of $\widehat{[\cdot,\cdot]}_{\boldsymbol{p}}^{\perp}$ for different settings of $h_{\max}$ (rows) and different SNR. We set the curvature $\kappa = 1/\rho = 1/26\,\mathrm{mm}^{-1}$, voxel size $\delta = 1\,\mathrm{mm}$ (giving $h_{\max} = \{3,5\}$ voxels $= \{3,5\}\,\mathrm{mm}$), and evaluate $\widehat{[\cdot,\cdot]}_{\boldsymbol{p}}^{\perp}$ at $\boldsymbol{p} = (10,-10,0)$. 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.
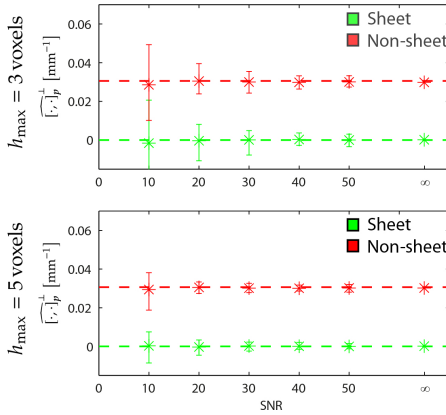
Figure 8.10: Mean and range of $(\widehat{[\cdot,\cdot]}_{\boldsymbol{p}}^{\perp}$ for different settings of $h_{\max}$ (rows) and different SNR. We set the curvature $\kappa = 1/\rho = 1/26\,\mathrm{mm}^{-1}$, voxel size $\delta = 1\,\mathrm{mm}$, and evaluate $\widehat{[\cdot,\cdot]}_{\boldsymbol{p}}^{\perp}$ at $\boldsymbol{p} = (10,-10,0)$. Here we used trilinear interpolation of the fODF spherical harmonic coefficients and extracted the peaks during tractography. This is in contrast to Fig. 8.9, where peaks were pre-extracted and nearest neighbor interpolation was used. Dashed lines indicate the true Lie bracket normal component.
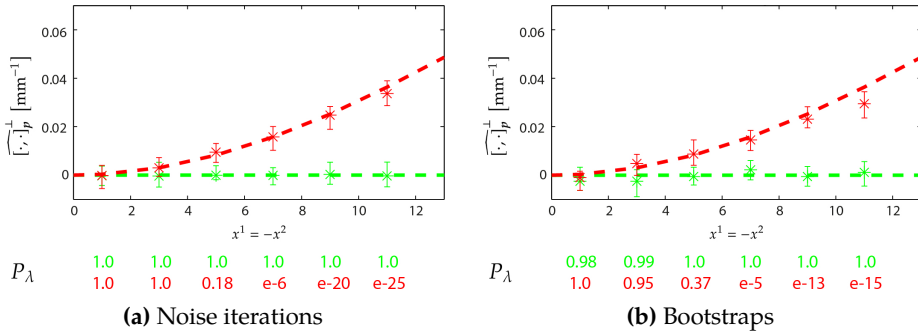
Figure 8.11: Mean and range of $\widehat{[\cdot,\cdot]}_p^{\perp}$ and the SPI P_{λ} for different points $p = (x^1, -x^1, 0)$ to vary the Lie bracket normal component magnitude $[u, w]_p^{\perp}$. $[u, v]_p^{\perp}$ is evaluated at the same points for reference. We set the SNR to 20, $\delta = 1\,\text{mm}$, $h_{\max} = 5\,\text{voxels}$, and $\kappa = 1/26\,\text{mm}^{-1}$.

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 SNR greater than 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 data sets have over five times fewer measurements (90 vs. 514) and a lower diffusion weighting (maximum $b = 3000\,\text{s/mm}^2$ compared to $b = 10000\,\text{s/mm}^2$).

8.3.2.3 Sheet probability index

Here, we used the peak directions resulting from CSD on the single shell simulated data. Fig. 8.11 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 the SNR to 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.

8.3.3 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 8.3.3.1 we explain the interpretation of these maps and their relation to the Lie bracket by means of an example data set, and in Section 8.3.3.2 we investigate inter-subject variability. In the remaining Sections 8.3.3.3, 8.3.3.4, and 8.3.3.5, we consider the influence spatial resolution, diffusion weighting, and diffusion MRI technique, respectively.

8.3.3.1 Sheet probability index

Fig. 8.12(a) (left) shows Lie bracket normal component estimates of the two largest fODF peaks in a single slice of a WU-Minn HCP data set. A fractional anisotropy (FA) color map of the same slice is shown for reference in Fig. 8.12(b). 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 $\widehat{[\cdot,\cdot]}_p^\perp \approx 0$. 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 red and green arrow look spatially continuous, whereas the area indicated by the blue 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. 8.12(a) on the right: high SPI (red arrow, $\widehat{[\cdot,\cdot]}_p^\perp$ concentrated around zero), 'medium' SPI (blue arrow, $\widehat{[\cdot,\cdot]}_p^\perp$ spread), and low SPI (green arrow, $\widehat{[\cdot,\cdot]}_p^\perp$ concentrated away from zero). The histograms illustrate that the normality assumption used to calculate the SPI is reasonable. Figs. 8.12(c) and 8.12(d) show SPI maps for the largest fODF peaks.

Fig. 8.13(a) 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 []). Fig. 8.14 shows similar results for $h_{max} = 3$ voxels, 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. 8.13(b) and 8.13(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 $\left(\widehat{\Phi}_{h_2}^w \circ \widehat{\Phi}_{h_1}^v\right)(p)$ and $\left(\widehat{\Phi}_{h_1}^v \circ \widehat{\Phi}_{h_2}^w\right)(p)$, to be specific). Fig. 8.13(b) shows a sheet formed by the corpus callosum (CC) and the corticospinal tract (CST) in the left hemisphere (see also []). In addition, the white arrow highlights an area in which
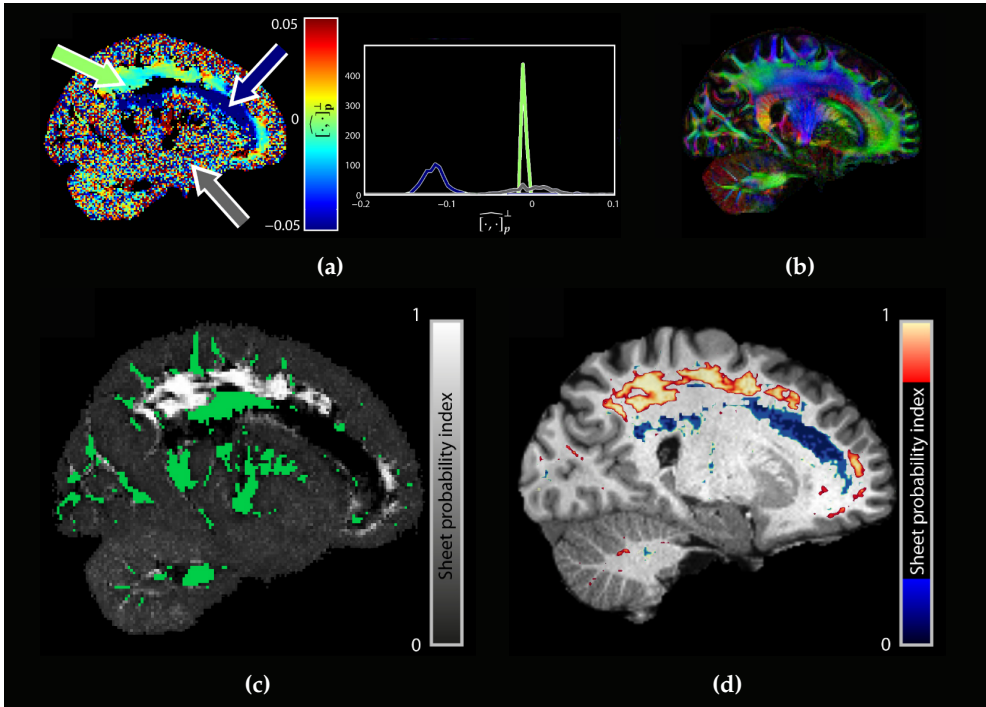
Figure 8.12: (a) A single bootstrap of the Lie bracket normal component (two largest fODF peaks) in a single slice, together with histograms of the normal component at the indicated locations in (high- (green arrow), medium- (grey arrow), and low-sheet probability (blue arrow) area). (b) A direction 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- (red, $P_\lambda > 0.5$) and low-sheet probability areas (blue, $P_\lambda < 0.1$) shown as an overlay on an anatomical scan.

crossing sheets are found. Fig. 8.13(c) 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. 8.13(a)) are shown in Fig. 8.13(d). 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. 8.13(a).

Fig. 8.15 shows a quick comparison between crossing angles and the corresponding SPI in one of the HCP data sets analyzed using CSD. In the resolution range of CSD (angles $45° - 90°$) there is no obvious bias towards $90°$ crossings. Finally,
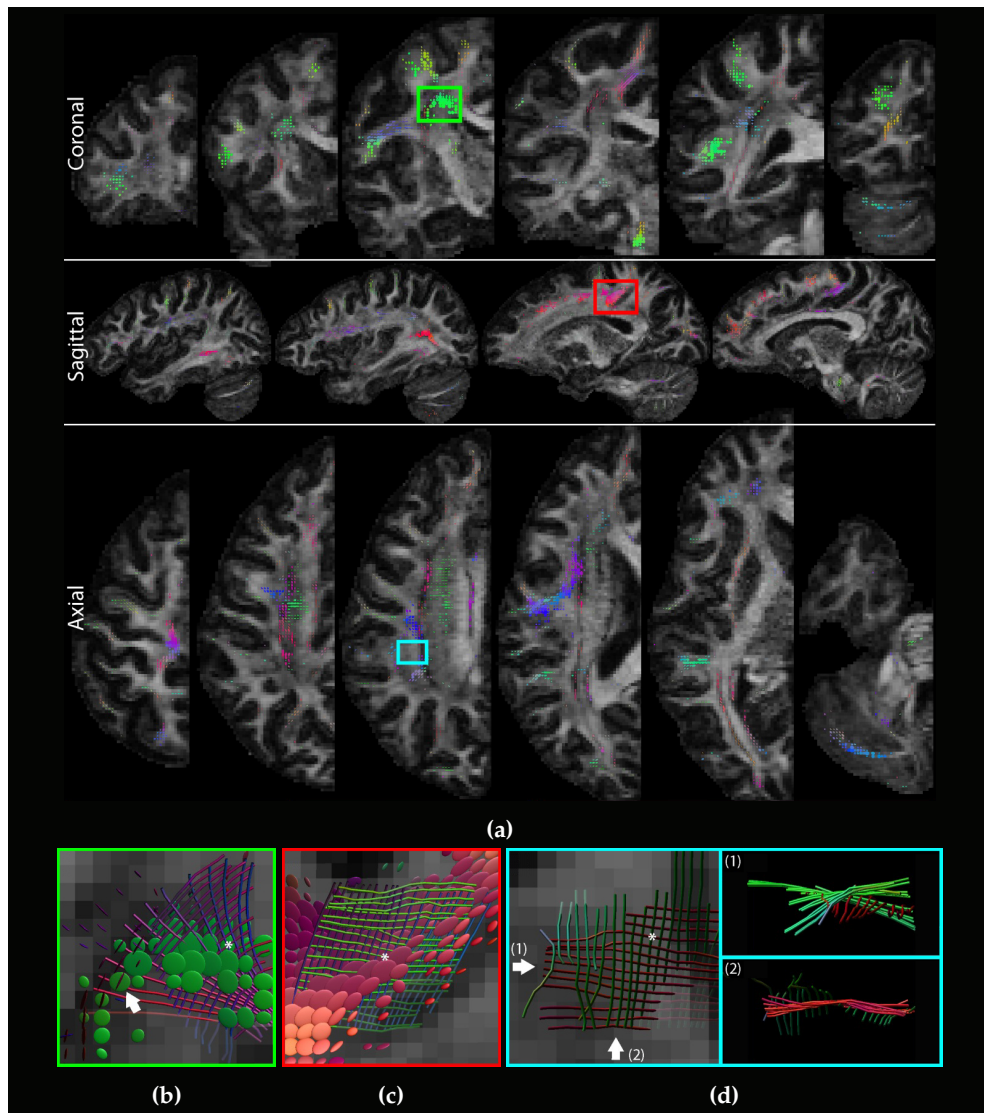
Figure 8.13: (a) Sheet tensors (λ = 0.008) on different coronal (top), sagittal (middle), and axial (bottom) slices. Ellipsoids with $P_\lambda < 0.2$ are not shown for clarity, and the sheet tensors are colored according to Fig. 8.5(b). (b) High-SPI area with streamlines (paths $\left(\widehat{\Phi}^w_{h_2} \circ \widehat{\Phi}^v_{h_1}\right)(\boldsymbol{p})$ and $\left(\widehat{\Phi}^v_{h_1} \circ \widehat{\Phi}^w_{h_2}\right)(\boldsymbol{p})$ 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).
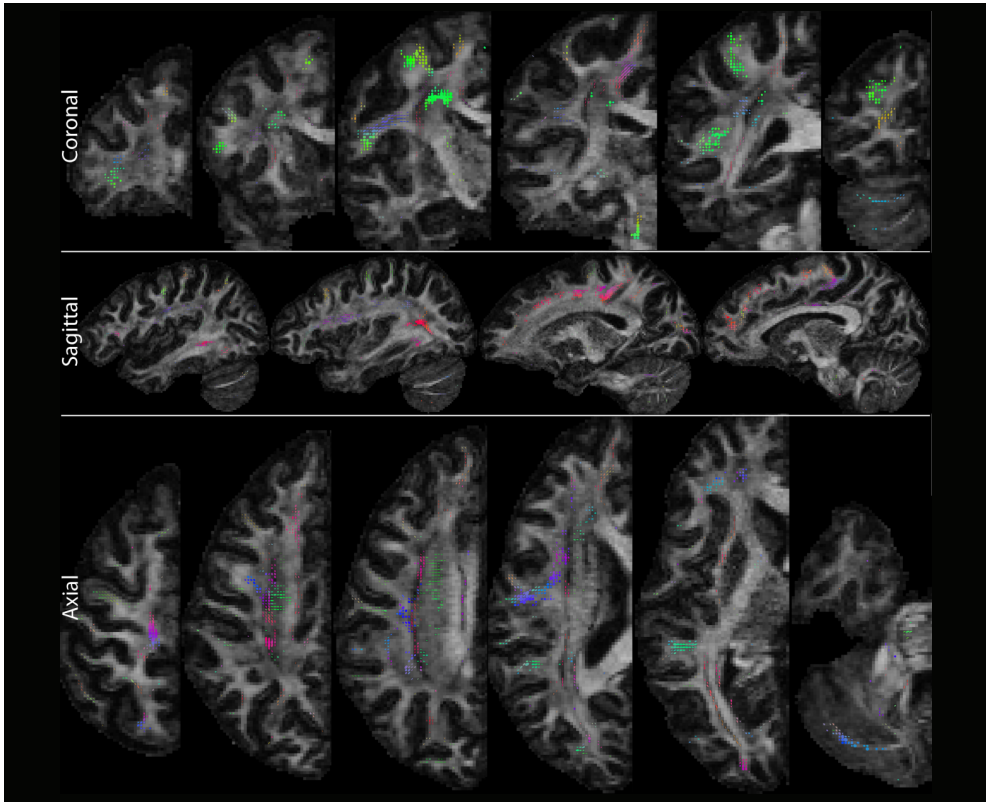
Figure 8.14: Sheet tensors (satisfying $P_\lambda > 0.2$, $\lambda = 0.008$) on different coronal (top), sagittal (middle), and axial (bottom) slices, computed with $h_{\max} = 3$ voxels.

Fig. 8.16 shows a comparison between the SPI and three diffusion tensor imaging (DTI) shape indices (recall Section 2.3.1), which have been used in previous work to detect sheet-like regions in DTI data [355].

8.3.3.2 Inter-subject variability

Fig. 8.17 shows results for 3 HCP subjects (first three rows, $h_{\max} = 5$ voxels = 6.25 mm) and the MASSIVE data set (last row, $h_{\max} = 2.5$ voxels = 6.25 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. In Fig. 8.18 we show a quantitative comparison between the HCP subjects after registration of the corresponding FA images, which illustrates a moderate overlap between the maximum SPI maps after binarization.
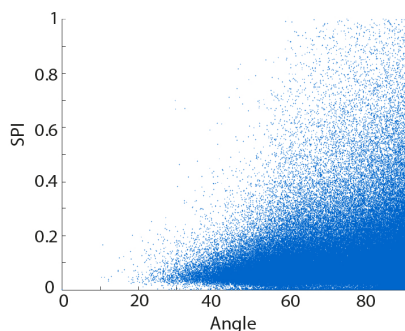
Figure 8.15: SPI plotted against angle. High SPI can be encountered for angles ranging from $40° - 90°$, which is in the range of the angular resolution of the used reconstruction technique. There are only slightly more voxels with high SPI for angles closer to $90°$.

| (a) | (b) | (c) |
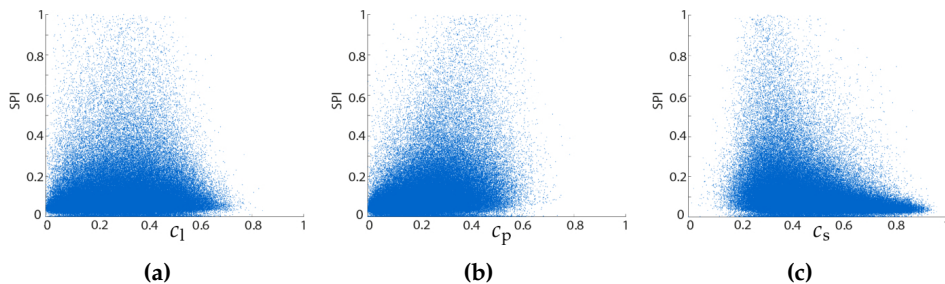
Figure 8.16: Scatter plots of the SPI and three DTI measures, cf. Section 2.3.1. (a) Linear coefficient c_l. (b) Planar coefficient c_p. (c) Spherical coefficient c_s.

8.3.3.3 The influence of spatial resolution

Fig. 8.17 compares the MASSIVE data set (2.5 mm isotropic voxels) with the HCP data sets (1.25 mm isotropic voxels), where we kept the maximum distance h_{max} constant at 6.25 mm. In the MASSIVE data set the same high-SPI areas can be recognized as in the HCP data.

Fig. 8.19 shows results for the same subject as in Fig. 8.13 (voxel size 1.25 mm, $h_{max} = 5$ voxels $= 6.25$ mm) and Fig. 8.14 (voxel size 1.25 mm, $h_{max} = 3$ voxels $= 3.75$ mm), but now we down-sampled the data spatially (voxel size 2.5 mm, $h_{max} = 2.5$ voxels $= 6.25$ mm). The same sheet structures can still be recognized, but some finer-scale structures get lost.

In Fig. 8.20 we illustrate the effect of changes in the angle threshold t, Alg. 2, which determines the maximum curvature allowed at a specific spatial scale.
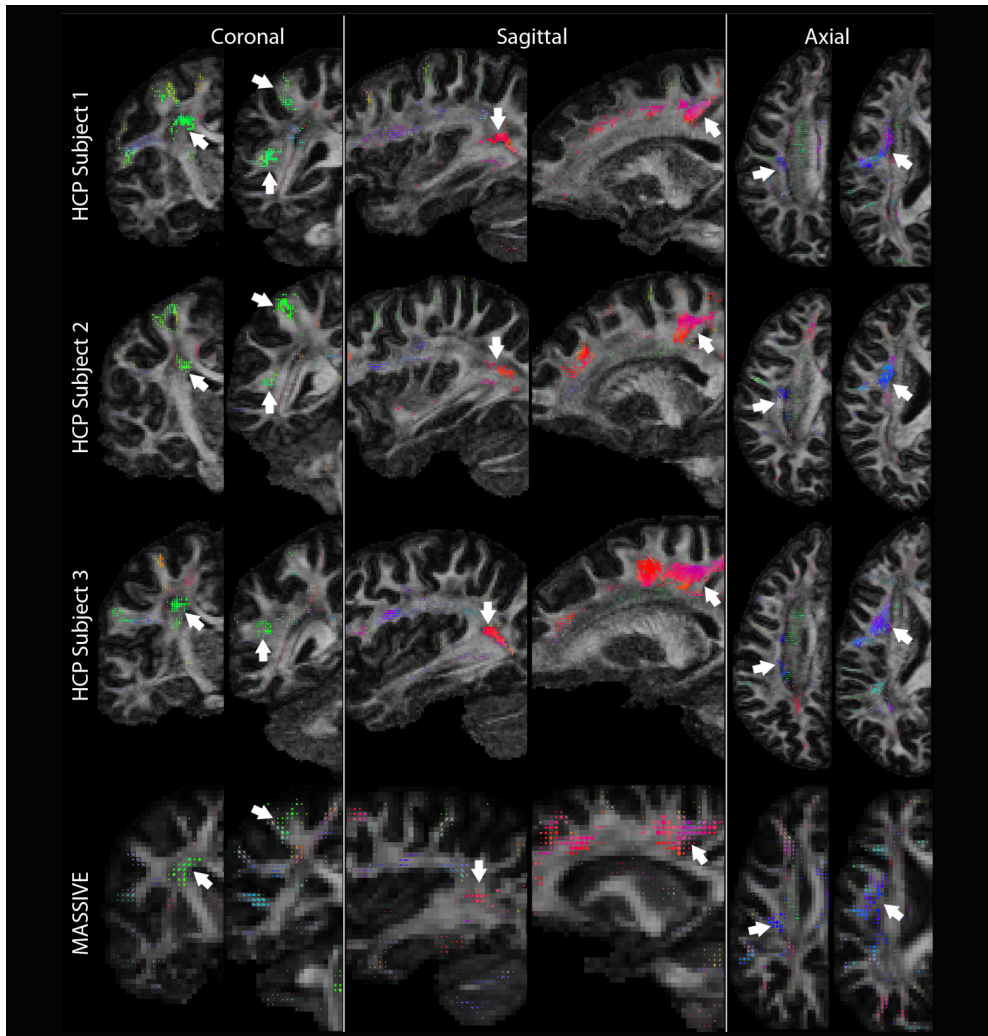
Figure 8.17: Visual comparison of sheet structures between subjects and spatial scales (tensors with $P_\lambda < 0.2$ are not shown for clarity, colors according to Fig. 8.5(b), and we set $\lambda = 0.008$). Examples of visually similar sheet structures are indicated by the arrows.

(a) FA-based registration pipeline

(b) FA after registration

(c) SPI after registration

(d) Thresholded SPI
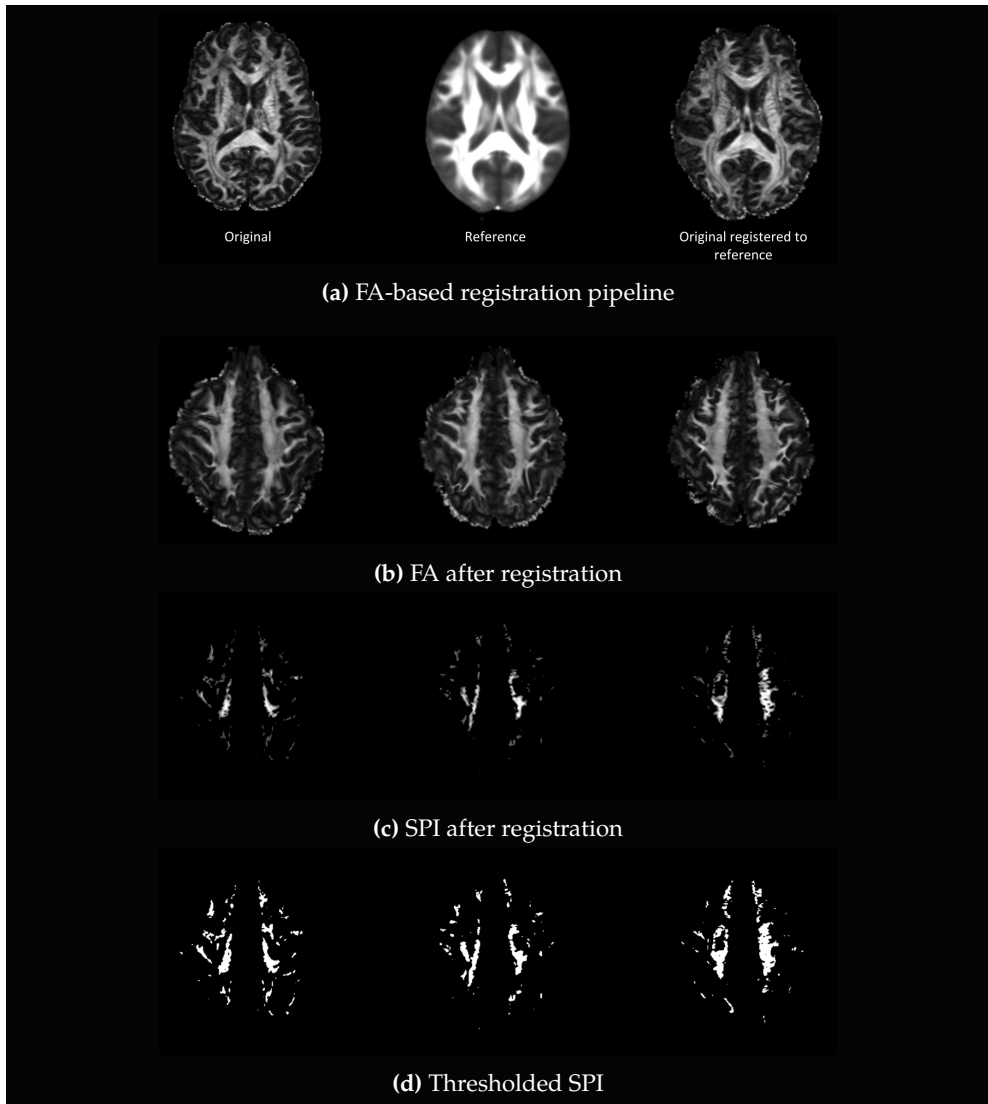
Figure 8.18: Preliminary results of a quantitative comparison between three HCP subjects. (a) An FA map for one subject before (left) and after (right) registration to the FMRIB58 template (middle) using Elastix [201] with the Oxford-optimized settings described in [155]. (b) Registered FA maps for the three HCP data sets. (c) The voxel-wise maximum-SPI maps corresponding to (a). 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.
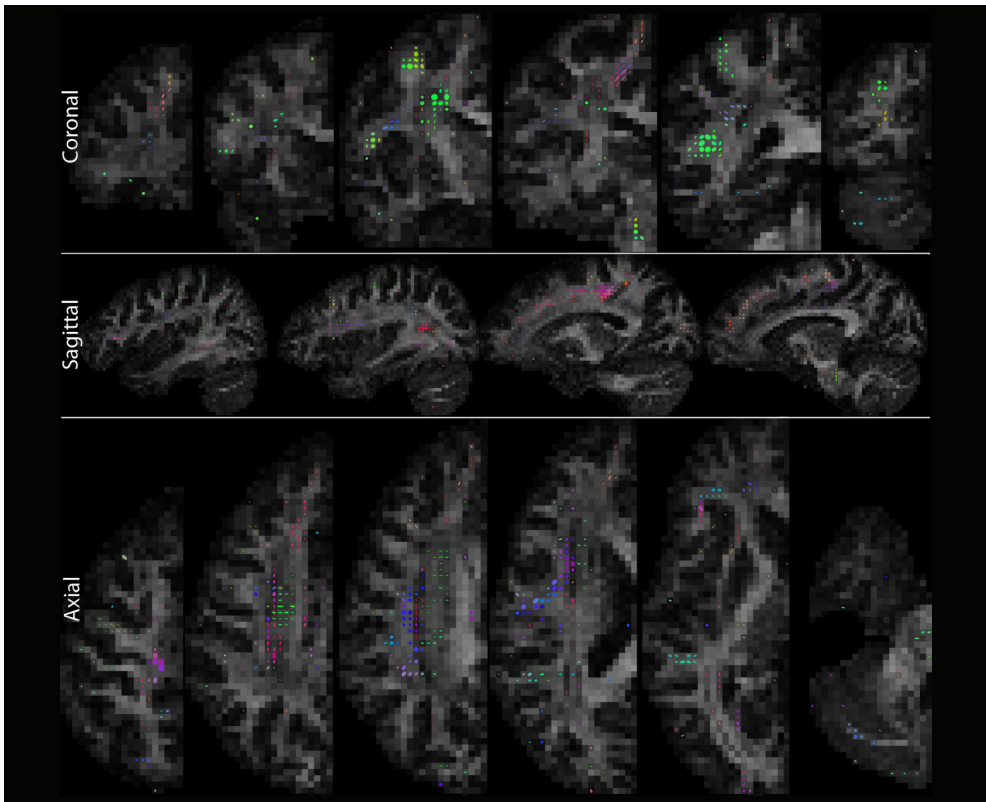
Figure 8.19: Sheet tensors (satisfying $P_\lambda > 0.2$, $\lambda = 0.008$) on different coronal (top), sagittal (middle), and axial (bottom) slices, computed with $h_{\max} = 2.5$ voxels on a sub-sampled data set.

(a) $t = \cos 15°$ **(b)** $t = \cos 25°$ **(c)** $t = \cos 35°$ **(d)** $t = \cos 45°$
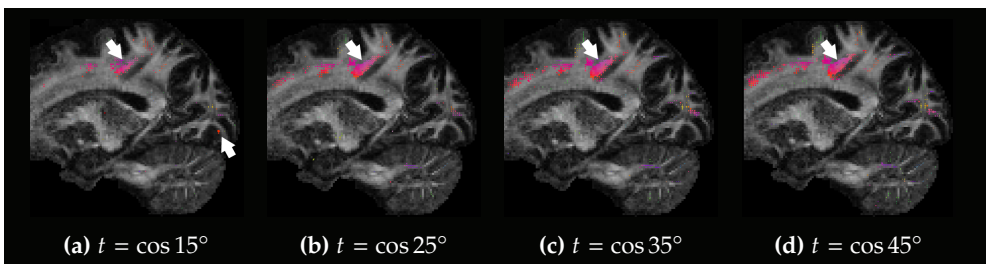
Figure 8.20: SPI for different angle thresholds. Setting the angle threshold too low (15°) can result in failure to detect curved sheets or spurious high SPI values because of early termination of paths (white arrows).
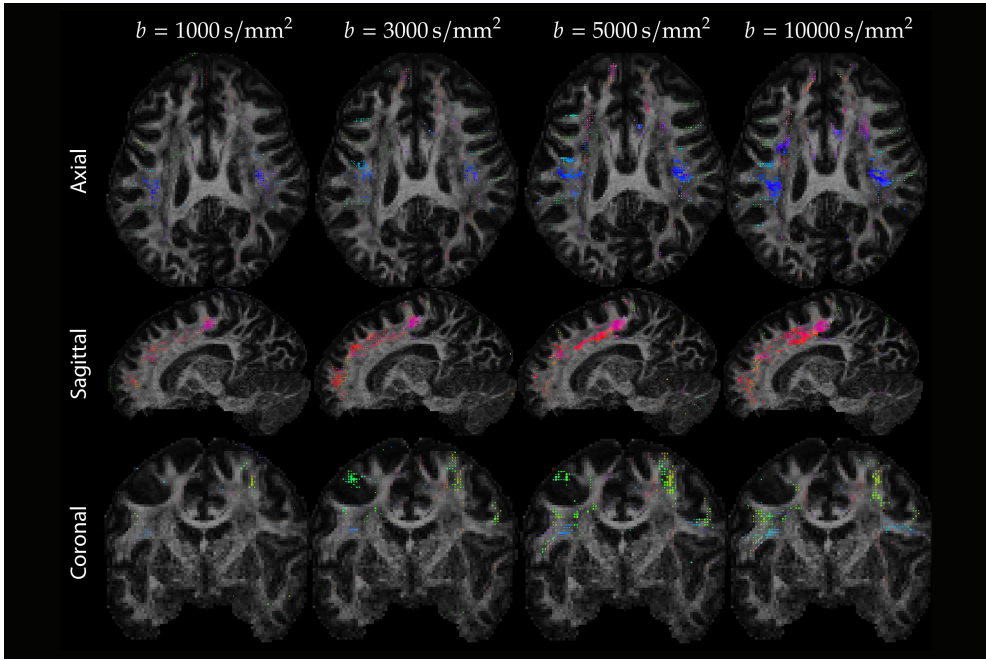
Figure 8.21: The detection of sheet structure with different diffusion weightings in the MGH HCP data set, overlaid on the FA of the $b = 1000\,\text{s/mm}^2$ shell (tensors with $P_\lambda < 0.2$ are not shown for clarity, colors according to Fig. 8.5(b), and we set $\lambda = 0.008$).

8.3.3.4 The influence of diffusion weighting

Fig. 8.21 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 $3000\,\text{s/mm}^2$.

8.3.3.5 The influence of diffusion MRI technique: CSD vs. DSI

Fig. 8.22(a) shows results of a single Lie bracket computation in tissue on the MGH DSI data set ($h_{\max} = 3\,\text{voxels} = 6\,\text{mm}$). Here, we visualize a tensor if $\left|\widehat{[\cdot, \cdot]}_p^\perp\right| \leq 0.008$ (all tensors have the same size), and we color the voxel red if the minimum $\left|\widehat{[\cdot, \cdot]}_p^\perp\right|$ in that voxel is larger than 0.025 (which would indicate that there is likely no sheet structure locally, see e.g. Fig. 8.11). The arrows indicate high-SPI areas that could also be identified in the previous experiments. Many high $\left|\widehat{[\cdot, \cdot]}_p^\perp\right|$ estimates are found in
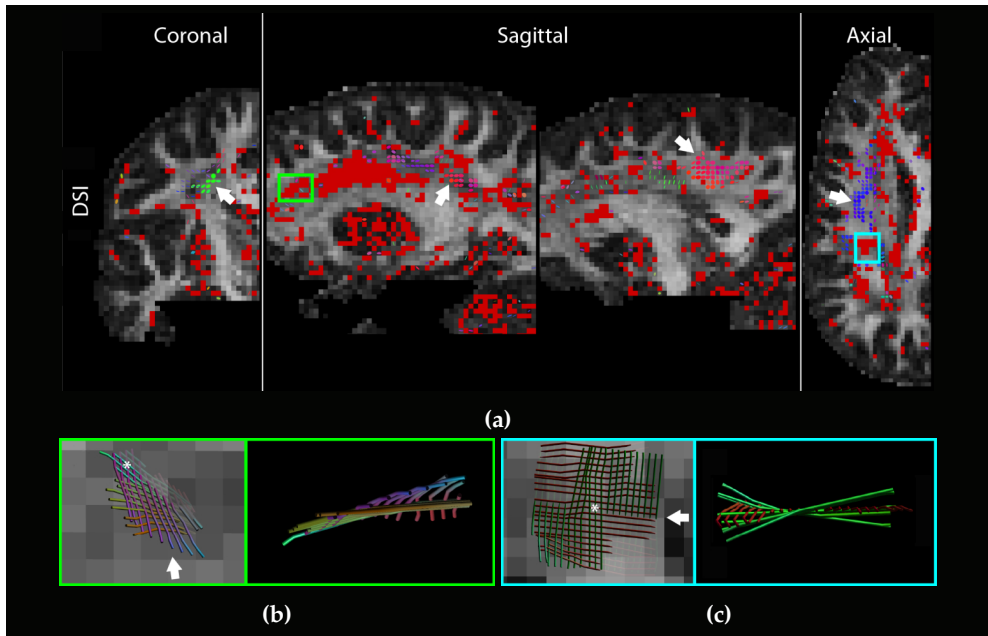
Figure 8.22: (a) A map of $\left|\widehat{[\cdot,\cdot]}_p^{\perp}\right|$ for the MGH DSI data set ($h_{\max}$ = 3voxels). A tensor is visualized if $\left|\widehat{[\cdot,\cdot]}_p^{\perp}\right| \leq 0.008$ with colors according to Fig. 8.5(b) (tensors have the same size, arrows indicate example areas), and we color the voxel red if the minimum $\left|\widehat{[\cdot,\cdot]}_p^{\perp}\right|$ in that voxel is larger than 0.025. Rectangles show spatially continuous example areas of high minimal $\left|\widehat{[\cdot,\cdot]}_p^{\perp}\right|$, 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).

the gray 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. 8.22(b) and 8.22(c).

8.4   Discussion

The hypothesis that brain pathways cross nearly orthogonally forming two-dimensional sheet-like structures is an active topic of debate [57, 364, 365]. 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. 8.15 for a plot of the angle against the sheet prob-

ability index (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, diffusion MRI models, and other acquisition parameters.

8.4.1 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. 8.7 and 8.8). 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. 8.9 and 8.11). We found that the performance was dependent on factors such as noise, voxel size, curvature, true Lie bracket normal component magnitude, and the chosen diffusion MRI 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 [185], also a concern raised by Catani et al. [57].

Based on our results from real diffusion MRI data (Figs. 8.11, 8.12, 8.13, 8.17, 8.21, and 8.22), 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 [365] (Figs. 8.13(b) and 8.13(c)). 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 (Figs. 8.13(d), 8.22(b) and 8.22(c)). This discrepancy clearly shows the added value of quantitative analysis in the investigation of such structures in the brain. Whereas Wedeen et al. [364, 365] 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. 8.11, 8.12, 8.13, 8.17, 8.21, and 8.22). 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. 8.12(c)) 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. 8.17), indicating that our framework provides reproducible results. In agreement with Wedeen et al. [364, 365], we found high SPI values in crossing regions of the corpus callosum with the cingulum (see []) and SLF 1-3 (e.g. Fig. 8.13(c)), 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. 8.13(b)). 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 [365] with non-orthogonal angles [57]. 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. [364] 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. 8.13(b)). 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 [364], 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.

8.4.2 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. 8.7 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 used to compute $\widehat{[v, w]}_p$, cf. Section 8.2.1.3, which 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 8.2.1.3). 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. 8.7), 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" [365], we can thus conclude that it is only possible to reliably detect sheet structures larger than the voxel scale.

Fig. 8.7 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. 8.7) 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. 8.8(b)).

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 [61]. 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. 8.20). 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.

8.4.3 The impact of the diffusion MRI technique

The impact of the diffusion MRI method—the use of diffusion spectrum imaging (DSI) versus e.g. constrained spherical deconvolution (CSD)—has also been a big part of the debate [57, 364]. 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" [57], 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 and thus potentially reducing the accuracy and precision (a similar effect as shown in Fig. 8.7, where a lower h_{max} in voxels equals a lower amount of reconstructed paths). We visually confirmed that in the DSI experiment (Fig. 8.22) 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. 8.17), resulting in a lower amount of voxels where the Lie bracket could be computed. This is in agreement with Catani et al. [57], 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 diffusion MRI data (based on the best scoring model of experiments in Ferizi et al. [119], which also included high b-values). The results shown in Fig. 8.9 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. [364] that "DSI should present the lower risk of bias".

A comparison between the CSD results of the MASSIVE data with voxel size $\delta = 2.5$ mm (Fig. 8.17, bottom row) and the MGH DSI data with $\delta = 2$ mm (Fig. 8.22) reveals similar large-scale high-SPI areas. These sheet areas could also be observed in the WU-Minn Human Connectome Project (HCP) subjects (Fig. 8.17, 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. 8.22).

A higher diffusion weighting generally causes an increase in the number of voxels with a significantly high SPI value (Fig. 8.21). 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" [364]. Increased diffusion weighting, however, also comes at the cost of a lower SNR. Generally, a b-value of $3000\,\text{s/mm}^2$ 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 diffusion MRI data, one requires a robust and reliable means to determine these peak directions. Our method is general and not limited to a particular diffusion MRI technique or acquisition scheme, and we therefore believe that its performance can be improved with ongoing technical developments in the field.

8.4.4   Further methodological considerations

The least squares reconstruction of $\widehat{[v,w]}_p$ (Section 8.2.1.3) is sensitive to errors due to the approximation of Eq. (8.14), and to 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 fiber ODF interpolation (strictly speaking the actual diffusion measurements would have to be interpolated), which is more precise (Fig. 8.9(a) and Fig. 8.10) 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. [78, 286]). 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 (Section 8.2.1.3)

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 $h_1\ h_2$ 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 diffusion MRI data sets is mostly not feasible. In Fig. 8.11 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. 8.22). 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. [355, 385], which compute stream surfaces at points where the DTI tensor has a high planarity coefficient c_p [372], see Fig. 8.16. 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).

In this manuscript we view the Lie bracket in terms of the integral curves of vector fields (Eq. (8.14)), which is equivalent to a combination of differential operators on the vector fields known as the commutator (see Eq. (8.3) 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. (8.3) 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) [384]. 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. 8.7). 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. Preliminary results of these efforts have been presented at the ISMRM [333, 334].

8.4.5 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" [365], 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 [268]. The hypothesis that sheet structures have a close connection to development, axonal path finding, and the chemotactic gradients of early embryogenesis [365] 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 fractional anisotropy (FA) images of HCP subjects and calculated the overlap of thresholded maximum SPI maps indicated only a moderate overlap (Fig. 8.18). 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 [253]. Instead, point correspondences should be found 'in the space of sheet structures'. Ongoing developments in registration of tensor fields (e.g. [381] which could potentially be applied to sheet tensors), ODFs and multi-fascicle models (e.g. [284, 326]), and tractography data (e.g. [147, 253]) could contribute towards this end.

Although we find that the diffusion MRI data investigated in this work supports the existence of- sheet-like structures at certain locations in the brain, it should be noted that the diffusion MRI 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 [70] and polarized light imaging [20] 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.

Chapter 8

Conclusion

Chapter 9

Discussion

Neurogeometry and Diffusion MRI

The three parts of this thesis—*Modeling*, *Processing*, and *Analysis*—cover a small subset of topics where techniques and ideas related to geometry are of use to diffusion MRI. We have attempted throughout to root these techniques in a solid theoretical foundation, either by formalizing existing approaches, or by proposing new techniques from first principles. In this concluding chapter we review the key contributions presented in this work, and place them in the context of current challenges and existing solutions.

In Part I we considered Riemannian and Finsler manifolds as geometrical modeling frameworks for the analysis of diffusion-weighted MRI data. The fundamental relevance of these geometries in the context of diffusion modeling is discussed from a mathematical point of view in Section 9.1, and their practical applicability to diffusion MRI, specifically to neuroimaging, is discussed in Section 9.2. Potential applications of the large gradient strength approximation employed in the derivations of Part I are discussed in Section 9.3. In Section 9.4 we discuss the usage of the proposed geometrical features in the analysis of diffusion-weighted data, and in Section 9.5 we briefly look at generalizations of the Finslerian framework. In Section 9.6 we consider finally the large-scale geometrical organization of the brain, as discussed in Chapter 8.

9.1   Relevance of geometry in diffusion

Given a stochastic process with well-defined moments, there is a sensible way in which one can define various associated geometries. In the case of Riemannian geometry this has been known for some time (see e.g. the book by Ikeda et al. [168]), where the second order (central) moments of a process, i.e., the mean squared displacements, determine a Riemannian metric tensor. When a process is governed entirely by its second order moments, i.e., if it is normally distributed with zero mean, then this relation is invertible—the Riemannian metric contains all statistical information about the stochastic process. Generically, distances and other geometrical concepts that can be derived from the metric are directly related to particle dynamics, e.g. shortest paths in the Riemannian manifold represent the most likely route along which particles in Brownian motion move between points. We speak of a geometrical framework: one can apply techniques from the associated geometrical theory to process and analyze the given statistical data.

The Riemannian geometrical framework can be extended to the case of Finsler geometry using the reasoning presented in Chapter 3, where we establish a fundamental relation between the cumulant generating function of a (compactly supported) stochastic process, and the geometrical structure of an associated Finsler manifold. Specifically, we find that the level sets of the cumulant generating function can be interpreted as Finslerian spheres, which uniquely determine the local geometry of the manifold. These level sets asymptotically approach the (convex hull of the) largest

displacements in the process, and where Riemannian geometry describes the *most likely* displacements in a process, Finsler geometry thus describes the *largest*—and in some sense *least likely*—displacements.

By viewing diffusion as a stochastic process—i.e., as a sequence of random displacements that are governed by (approximately) the same probability distributions—the fundamental relevance of these two frameworks is revealed. The Riemannian framework captures the central limit [66] characteristics of the diffusion process: as the number of collisions between particles tends to infinity, the probability distribution of the process will asymptotically approach a Gaussian distribution, and the bulk motion of particles can then be described completely by the mean squared displacements. The Finslerian framework, on the other hand, is related to rare events in the diffusion process, which are typically considered a topic of the theory of large deviations [353]. Finsler geometry captures statistical properties of a diffusion process that are not accurately described by its asymptotic distribution. The precise relation between large deviations theory and Finsler geometry is a subject of ongoing research.

9.2 Geometrical frameworks in diffusion MRI

In Part I we discussed how geometrical frameworks can be applied in diffusion MRI. In diffusion MRI we measure an ensemble average propagator (EAP), which represents the probability of a particle in an ensemble undergoing a given translation over a fixed diffusion time. We can measure an EAP in ensembles centered around different points in an imaged sample, and by combining the associated metrics at all points we can define a manifold that represents statistical properties of the ensembles in the sample. Subsequent processing of the data can then be done on this manifold, where the geometrical framework naturally handles spatial correlations in the data that are otherwise difficult to manage.

In most diffusion MRI experiments, anisotropy in an observed stochastic process is not an *intrinsic* property of the observed particles. Instead, the displacement probability distribution typically reflects the anisotropy of a static ambient structure that is probed by an isotropically diffusing medium. If the diffusion characteristics of the medium are known, geometrical frameworks naturally allow us to factor these out. The geometry then solely represents the structure embedded in the medium, which can be investigated independently from the medium used to probe it. Additional processing may be useful to filter out the effect of other irrelevant interactions, as described in Chapter 4. It is widely accepted that the structure probed by diffusion MRI in neuroimaging is predominantly made up of axonal membranes [30].

It is useful to think of a reconstructed geometry as the 'simplest prototypical sample' that explains the observations. The complexity in the stochastic processes under-

lying the measurements is encoded in (typically minor) deformations of a Euclidean base space. This interpretation, as well as the intuitive relations between metric and particle motion described in Section 9.1, are provably valid when the displacement probability distributions of different particles in the observed ensemble are similar. In diffusion MRI terminology, this means that the presented frameworks can be applied straightforwardly if the metric can be approximately inferred from the contributions of a single diffusion compartment. This assumption is not universally valid in practice, and in situations where there is more than one relevant diffusion compartment the interpretation of the geometrical frameworks may have to be adapted. We do note that this assumption is only necessary for the described *interpretations*—as long as the moments/cumulants of the EAP are well-defined, geometrical frameworks can be used to process and analyze the data.

The human brain is generally accepted to contain two (essentially non-exchanging) compartments that have an observable impact on the diffusion-weighted signal in typical acquisitions: the intra-axonal compartment, and the extra-axonal compartment [14, 30, 206, 246, 250]. In the neuroimaging experiments presented in Chapter 6 we assume that the only relevant compartment is the extra-axonal compartment, and that the intra-axonal compartment can be safely neglected. This is motivated by the derivations in Chapter 3 that show how large gradient strength acquisitions are biased towards large displacements, which can be expected to occur in the relatively unrestricted extra-axonal compartment [206, 250]. The promising results obtained with the barrier orientation distribution function (barrier ODF or bODF) in Section 6.3.1 that rely on this assumption are also encouraging in this regard. However, although there may be strong arguments supporting the notion that the Finsler function should in theory represent the extra-axonal compartment, there are some practical considerations to take into account. After all, the asymptotic relation between the Finsler function and the diffusion MRI signal in Eq. (3.29) formally only applies in the unattainable large gradient strength limit $G \to \infty$, and our interpretation of the results in Chapters 3 and 6 is based on the intuition that the more practical estimates provided by Eq. (3.36) can be used to approximate the true Finsler function. In future work we intend to generalize Eq. (3.29) from an asymptotic *relation* to an asymptotic *expansion* [108, 278], i.e., a series expansion that describes a function as its argument tends towards infinity. If this is possible, we can clarify exactly how our assumption of infinite gradient strengths impacts our analysis, and perhaps even derive more accurate models that resolve the dissonance between the practical cumulant expansion (valid near the origin) and our theoretical relations (valid away from the origin).

9.3 The strong gradient extremal path approximation

A recent overview paper by Novikov et al. [250] lists three regimes to categorize diffusion MRI models based on the diffusion time Δ. In the limit $\Delta \to 0$ very few spins have had an opportunity to interact with their surroundings, and only the surface-to-volume ratio between barriers and medium is required to fully describe the average spin dynamics [236]. The second regime is the limit $\Delta \to \infty$, where the central limit theorem applies to all diffusion compartments separately, recall also Assumption 1. The propagator thus becomes a weighted average over the per-compartment Gaussian (attractor) distributions, which forms the basis of a large number of current models [14, 32, 119, 124, 250, 265, 297, 317, 321, 376, 382]. The Riemannian framework (Chapter 2) fits in this regime, but in general geometrical frameworks such as proposed and discussed in this thesis belong to the last regime, that of intermediate diffusion times. In this regime spins have time to interact with their immediate surroundings, but coarse-graining is not yet complete, i.e., the central limit theorem only gives a first order approximation for the observables in a compartment. Theoretical developments in the intermediate Δ regime are often complex, involving e.g. "qualitatively distinct behavior along [123, 248] and transverse [45, 123] to the neurites in the brain" [250].

In Chapter 3, where we derive the Finslerian framework, we approached this challenge by focusing on the large gradient strength limit $G \to \infty$. This again greatly simplifies the diffusion MRI signal equations, and the application of the $G \to \infty$ approximation extends beyond the narrow pulse approximation that we assumed throughout this work. The easiest way to see how, is by considering the path integral formulation of diffusion-weighted MRI [234]:

$$S[G] = \int_{C \in \tilde{\Omega}} e^{i\gamma \int_0^T \langle G(t), C(t) \rangle \, dt} \, d\rho(C), \tag{9.1}$$

where the argument $G : \mathbb{R} \to \left(\mathbb{R}^3 \right)^*$ of the signal is a function of time with $G(t) = 0$ outside the range $[0, T]$. The path integral sums over the set $\tilde{\Omega}$ of all possible spin paths $C : [0, T] \to \mathbb{R}^3$, $d\rho(C)$ is the associated path probability measure, and G is a general gradient sequence that scales linearly with a gradient strength parameter G and which satisfies

$$\int_0^T G(t) \, dt = 0. \tag{9.2}$$

Considering the same limit as in Section 3.2.2 then gives

$$\log S[-iG] \sim \gamma \sup_{C \in \tilde{\Omega}} \int_0^T \langle G(t), C(t) \rangle \, dt \quad (G \to \infty). \tag{9.3}$$

(a) **(b)**
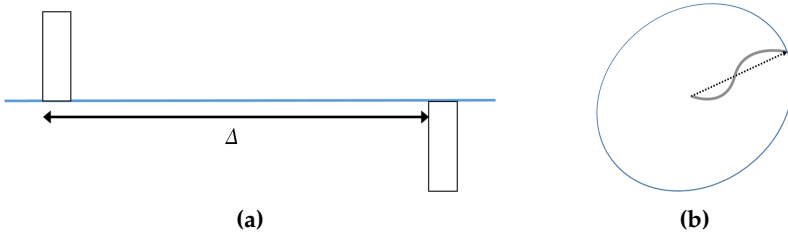
Figure 9.1: (a) The gradient sequence for the single diffusion encoding PGSE sequence, where we assume that the pulse width δ is negligible. We show effective gradients, including effects of any refocusing pulse(s). (b) An extremal path (gray) that maximizes $\int_0^T \langle G(t), C(t) \rangle \, dt$ for the gradient sequence shown in (a), the boundary containing all spins after diffusing for a time Δ (blue), and a vector representation of the gradient co-vector (dashed black arrow). In the large gradient strength approximation, only spins that follow paths close to the extremal path(s) are relevant to the diffusion MRI signal formation. In this case the extremal paths are only determined by the effective displacement of spins, which has to be maximal relative to the gradient orientation.

Eq. (9.3) shows that mainly the extremal paths (relative to the functional $\langle G, C \rangle$) are relevant in the formation of the diffusion signal S at large gradient strengths. This could for instance be useful in the analysis of generalized diffusion encoding (GDE) [110, 313, 371, 379] sequences like double or multi diffusion encoding (DDE or MDE) sequences [64, 73, 235, 255, 256, 309] or q-space trajectory imaging sequences [370], whose interpretation is the subject of ongoing research efforts [267].

Let us consider some basic examples. In the narrow pulse approximation of the standard PGSE sequence (Fig. 9.1(a), also called a single diffusion encoding sequence or SDE) we have a fixed gradient covector G with norm $\|G\| = G$ and

$$G(t) = G\big(\delta(t) - \delta(t - \Delta)\big), \tag{9.4}$$

with $\delta(x) = 1$ if $x = 0$ and 0 otherwise, and Eq. (9.3) clearly reduces to Eq. (3.27). In the path integral formalism, this relation states that the cumulant generating function is proportional to the largest displacement along the gradient orientation, i.e., the signal is dependent on the paths whose 'projection' on the gradient orientation is maximal. In Chapter 3 we assumed that the stochastic behavior of spins could be modeled as a Finslerian (locally Minkowskian) isotropic transport process, so that these extremal paths would be straight lines, but in the general case these paths could be arbitrarily curved, Fig. 9.1(b).

Take now the following simple variation of Eq. (9.4) shown in Fig. 9.2(a):

$$G(t) = \begin{cases} G\big(\delta(t) - \frac{1}{\Delta}\big) & \text{if } 0 \le t < \Delta \\ 0 & \text{else} \end{cases}. \tag{9.5}$$

We will refer to this for now as an asymmetric diffusion encoding sequence (ADE).
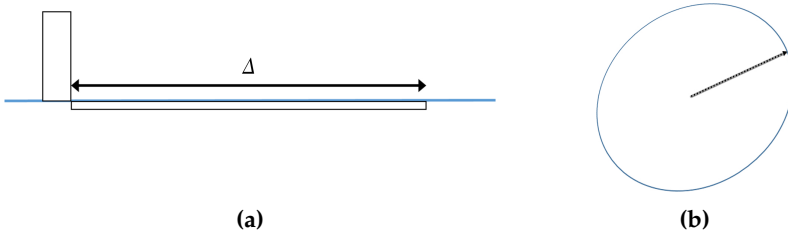
Figure 9.2: (a) The gradient sequence for an ADE sequence, where we assume that the width δ of the initial pulse is negligible. We show effective gradients, including effects of any refocusing pulse(s). (b) An extremal path (gray) that maximizes $\int_0^T \langle G(t), C(t) \rangle \, dt$ for the gradient sequence shown in (a), the boundary containing all spins after diffusing for a time Δ (blue), and a vector representation of the gradient co-vector (dashed black arrow). In the large gradient strength approximation, only spins that follow paths close to the extremal path(s) are relevant to the diffusion MRI signal formation. In this case the extremal paths are restricted to straight lines with an orientation determined by the applied gradient.

In this case, the cumulant generating function reduces asymptotically to

$$\log S[-iG] \sim \gamma \sup_{C \in \tilde{\Omega}} \int_0^T \langle G(t), C(t) - C(0) \rangle \, dt \qquad (G \to \infty). \tag{9.6}$$

In the PGSE sequence gradients were turned off for the duration of the displacement, and as a result the actual paths resulting in the large relevant displacements was unimportant. But with ADE the gradients are always on, and the extremal paths are then straight lines, see Fig. 9.2(b). A slightly different but similar conclusion could be drawn for e.g. oscillating gradient spin echo (OGSE) sequences [296]. Because of their targeted sensitivity to spins of interest, it may well be that ADE (and OGSE) sequences are practically more suited for Finsler-based modeling. In the path integral formalism, a comparison between the results of ADE and PGSE could be a measure of the diffusion dispersion around the gradient orientation. ADE sequences could potentially be adapted to sensitize the scanner to arbitrary spin paths.

Finally take the standard DDE sequence, a repeated SDE sequence separated by a time constant τ, parameterized by two diffusion times Δ_1, Δ_2, and two gradient vectors G_1, G_2 scaled by a global gradient strength parameter G:

$$G(t) = GG_1\big(\delta(t) - \delta(t - \Delta_1)\big) + GG_2\big(\delta(t - \Delta_1 - \tau) - \delta(t - \Delta_1 - \tau - \Delta_2)\big), \tag{9.7}$$

see Fig. 9.3(a). If we assume a small τ, we find that extremal paths for this sequence are a concatenation of the corresponding extremal paths in the single diffusion encoding case, cf. Figs. 9.1(b) and 9.3(b). We could thus interpret these measurements as being sensitive to a specific (unsigned) curvature, which could have potential applications in neuroimaging as a method to distinguish between kissing and crossing fiber pathways.
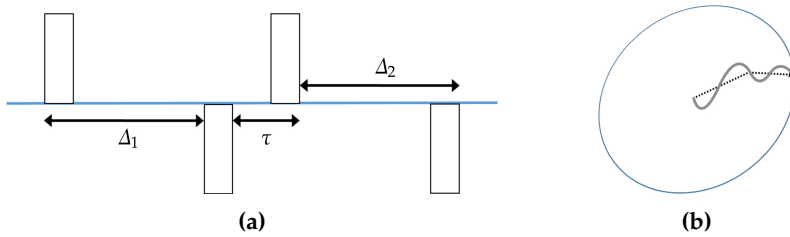
Figure 9.3: (a) The gradient sequence for a basic DDE sequence, where we assume that the pulse widths are negligible. We show effective gradients, including effects of any refocusing pulse(s). (b) An extremal path (gray) that maximizes $\int_0^T \langle G(t), C(t) \rangle \, dt$ for the gradient sequence shown in (a), the boundary containing all spins after diffusing for a time $\Delta_1 + \tau + \Delta_2$ (blue), and a vector representation of the gradient co-vector (dashed black arrow). In the large gradient strength approximation, only spins that follow paths close to the extremal path(s) are relevant to the diffusion MRI signal formation. If τ is small, the extremal paths are determined by the extremal paths of the corresponding SDE sequences.

9.4 Applications of geometrical frameworks

The benefit of this new geometrical perspective comes both from an abstraction that helps to distill the essence of a given problem, and from the immediate access to an expanded body of existing solutions and methods for these abstract problems. In neuroimaging applications one wishes for example to understand what information about the brain is captured by the signal, and how to perform a meaningful analysis of the data. Regarding the brain as a specific type of manifold as discussed in Part I, can help in solving practical problems such as interpolation and tissue characterization.

One of the important tools provided by geometrical frameworks are geodesics, which function as a natural means to investigate spatial correlations in the data. The interpretation of geodesics depends both on the geometry and on the application. In the case of the Riemannian framework for diffusion MRI, geodesics represent the trajectories along which water molecules are most likely to move. As such, they are also related to the diffusion-governed redistribution of cells and nutrients, which has been used for example by Mosayebi et al. [243] to predict the bulk growth of gliomas in the human brain. With the formal definition of the Finslerian framework, Chapter 3, we can now also interpret Finslerian geodesics in terms of the observed stochastics— they represent the fastest route connecting two points in the manifold, as traversed by particles in Brownian motion. We can easily envision these geodesics playing an even more important role in tumor growth predictions. By the same arguments used in the Riemannian case we can argue that Finslerian geodesics would describe the peripheral spread of an invading tumor. This could be relevant e.g. for radiation therapy planning, where an accurate delineation of the tumor extent is essential.

Diffusion MRI literature related to geodesics [17, 29, 51, 103, 128, 141, 163, 169, 171, 218, 227, 228, 254, 266, 269, 271, 281, 305–307, 343] has so far been predominantly concerned with the study of structural (neuronal) connections in the brain, i.e., ge-

odesic tractography and/or connectivity. With few exceptions [171, 218, 254] these works rely on different metric definitions (in terms of the diffusion-weighted data) from those used in modeling, because the standard 'modeling' metric (Chapter 2) can be shown to produce inaccurate results in various realistic scenarios [163, 171]. In Chapter 7 considered at a typical case where Riemannian geodesic tractography has problems, and showed that the scaled Riemannian metric [141] proposed in Chapter 2 leads to improved tractography by discouraging passage through isotropic diffusion regions. We further show in the Finsler case that we obtain more plausible connectivity maps by using new connectivity measures based on recent developments in Riemannian connectivity analysis [269].

However, despite the advancements achieved by using modified metrics, there remain fundamental unanswered questions about the tenets of geodesic tractography. To our knowledge, there is no solid case for the hypothesis that axonal fiber pathways correspond to geodesics of some 'holy grail' geometry, nor are there verifiable conditions or assumptions that would validate this hypothesis. This makes it very difficult to either provide arguments in favor of geodesic tractography, or to categorically reject it. Given recent results that underscore the inefficacy of state-of-the-art tractography algorithms [104, 225], and because there are no practical means to validate tractography output, the problem should be approached much more methodically than is currently the case. Though likely not in the form of geodesic tractography, we do think that the abstraction provided by the concept of a geometrical framework may be of use in this regard.

Aside from geodesics there are a number of other potentially relevant features that can be derived from the presented frameworks. Firstly there is the barrier ODF, which we proposed in Chapter 6 as an alternative axonal orientation distribution function. (The bODF can be seen to some extent as separate from the Finslerian framework, although the two are intimately related.) Like the fiber ODF, the barrier ODF may be of use in a variety of (non-geodesic-based) tractography algorithms. Then there are various global features, which (as described in Section 2.5.3.2) can generally be expressed in terms of the differential structure of the Riemannian or Finsler manifold. One example that we mentioned previously would be the geodesic deviation (or Jacobi field [21]), which gives a measure of the degree of divergence in nearby geodesics. This measure has been used in DTI [18], for example as a stopping criterion for tractography [306], and could be useful to identify branching points of fiber pathways.

9.5 Extensions of the Finslerian framework

In Section 1.2.1 we introduced geometrical modeling in relation to constrained deformations of Euclidean space. From this perspective, Riemannian geometry can be

interpreted as a scaling deformation of the local space (parameterized by the scaling factors along three orthogonal directions), and Finsler geometry as an elastic scaling along arbitrary orientations. Increasingly complex deformations lead to more complex (and more powerful) geometries, and following the successful generalization from Riemannian geometry to Finsler geometry, one might wonder which other geometries might be useful in describing stochastic processes.

One logical next step in the development of geometrical frameworks would be a generalization to Kawaguchi spaces, which allow e.g. *bending* of the local Euclidean space. Lengths of curves in a Kawaguchi space depend not only on the curve's position and orientation, but also on higher order derivatives of the curve at each point. The length functional in a Kawaguchi space is thus of the form $F\big(C(t), \dot{C}(t), \ddot{C}(t), \dots\big)$, where dots denote derivatives of a suitably differentiable curve $C : [0, L] \to M$, and where F satisfies a complex combination of conditions to ensure that the length integral

$$\mathcal{L}_F(C) := \int_0^L F\big(C(t), \dot{C}(t), \ddot{C}(t), \dots\big) \, \mathrm{d}t \tag{9.8}$$

is invariant under changes of the parameter t [193]. Restricting F to position and orientation dependencies reduces it to a Finsler function. Because higher order derivatives at a point provide increasingly accurate information about a neighborhood, Kawaguchi spaces that depend on increasingly higher order derivatives form a natural hierarchy of features to describe the local structure.

While the benefits of Kawaguchi-based frameworks may be significant, there are serious challenges both from a theoretical and from a practical point of view. Even if there is e.g. a measurable 'curvature component' detectable in a diffusion process, it will probably be very difficult to extract the intricate associated geometric structures, and in the context of diffusion MRI possibly involve more complicated pulse sequences [255, 267, 370].

9.6   The geometrical organization of the brain

The origin of the human brain's extraordinary capabilities is one of the major unresolved questions of today. The recently proposed sheet structure hypothesis [365] promises a possible starting point for an explanation in terms of the brain's formation at early stages of development. The study of Wedeen et al. [365], based on an extensive investigation into the structure of the brain using diffusion MRI, identified sheet-like structures of interwoven fiber pathways indicating an inherently grid-like organization of neuronal connections. The axes of this grid appeared to correspond to the principal axes of development, suggesting that the three chemotactic gradients of embryogenesis are the primary forces guiding early axonal growth. The validity of this hypothesis would have wide-ranging implications far beyond the field of

diffusion MRI.

The extensive presence of sheet structure in the brain has been debated since its proposal [57, 95, 327, 331, 363–365], in part due to a lack of quantitative characterization. In this thesis and related works [92, 328–335] we have formalized the necessary and sufficient condition for sheet structure to exist, leading to the proposal of the sheet probability index (SPI) which indicates the extent to which the data locally supports sheet structure at a given spatial scale. We performed an extensive validation of the SPI on artificial vector fields, which showed that we can quantitatively distinguish sheet from non-sheet structure, with spatial resolution and SNR being important factors that influence the accuracy and precision. Diffusion MRI experiments on real data reveal a high SPI at various locations in the brain at the investigated scale, but low SPI areas were also identified. Several high SPI areas could consistently be recognized across subjects, scanners, and spatial resolutions, independent of the employed diffusion model (i.e., CSD vs. DSI). Since tractography pathways do not represent true axons, validation of sheet structure with other technologies is necessary, and we note that the proposed method can be extended to quantify sheet structure in other types of directional data as well.

We already discussed a number of potential improvements of the currently employed SPI methodology in Section 8.4. Here we finally treat a possible reformulation of the sheet structure hypothesis based on the Funk–Radon transform, which we briefly invoked when discussing the barrier ODF in Chapter 6. Given an ODF that represents the orientations of the axons in a voxel, we find that the Funk–Radon transform associates a measure of perpendicular axon density to each orientation. As a result, the Funk–Radon transform of the ODF (e.g. the planar barrier ODF if one considers the bODF) could be interpreted as a sheet orientation distribution function. The existence of sheet structure is then guaranteed not by conditions on the Lie bracket (recall Chapter 8) but e.g. by continuity of the sheet ODF peaks, which is a much simpler property to verify. This would additionally allow us to perform 'sheetography' using standard ODF-based tractography algorithms.

9.7 Conclusion

Diffusion MRI is used to investigate spatial correlations in complex diffusion processes and (implicitly) in the underlying structure, and differential geometry naturally provides the tools to analyze different aspects of the imaging data. With more and more local properties being included in modern diffusion models, structured reasoning in terms of geometry can be useful to identify relevant parameters and to guide the development of novel models, acquisitions, and analysis techniques. Recent evidence suggesting an innately sheet-like geometrical organization of fiber pathways in the human brain, exemplifies the importance of geometry in neuroimaging.

Appendix A

Implementation

Classes` and Diffusica`

This appendix contains information on the two Wolfram Language packages written to support the research presented in this thesis. Most of the code used to generate the result shown in this manuscript is bundled in the Diffusica` package, which is in essence a general coding framework with guidelines on where and how to include medical image processing code. Diffusica` is built on top of the Classes` package, which provides the class functionality used to represent models, data, etc. The Classes` package has also been used as the basis for e.g. the LieAnalysis` package (`lieanalysis.nl`). Both Classes` and Diffusica` are open source under the Apache 2.0 license. The sections in this appendix are included almost verbatim as tutorials in the documentation of these packages.

A.1   The Classes` package

The Classes` package provides a basic implementation of hierarchical classes, including property inheritance and invariant declaration. In this section we will discuss how to set up a package with Classes`, and how to work within a Classes`-based package.

A.1.1   Classes setup

The Classes` package provides tools to easily create packages that work with objects or instances. These instances are of the form *class* [ < |*elements*| >], where the head *class* defines the class to which the instance belongs, and the *elements* in the `Association` object represent the data in the instance. The classes in the Classes` package are

designed to resemble as closely as possible existing constructs like **FittedModel** and **CompiledFunction**. The standard definition of a class is done through three functions (listed below). **DeclareClass** designates specific symbols as a class symbol that can function as a head of a class instance. **DeclareDefaults** is used to define the default elements in an instance that provide values in case none are provided explicitly. **DeclareInvariant** defines a pure function called the invariant, that defines when a class instance is valid. In the following sections we will explain how these functions can be used to create a basic classes setup, and how custom properties and functions can be defined for these classes.

DeclareClass	designate a symbol as a class
DeclareDefaults	sets the defaults for a class symbol
DeclareInvariant	sets the invariant for a class symbol

Functions used to setup classes.

A.1.1.1   Basic setup

Start with loading the Classes` package.

Load the Classes` context

```
Needs["Classes`"]
```

Top-level classes   In this section we will walk through the standard setup of classes, with the first step the declaration of a top-level class using **DeclareClass**.

Declare the top-level class **TopClass**

```
DeclareClass[TopClass]
```

Once a symbol is declared a class, **ClassQ** will return **True**.

Check if **TopClass** is a declared class

```
ClassQ[TopClass]
```

True

From a declared class symbol *class*, class instances can be generated by evaluating *class*[].

Create a simple instance of the **TopClass** class

```
tc = TopClass[]
```

TopClass [⊞ ◁ Base: TopClass
Valid: True
Attributes: {}]
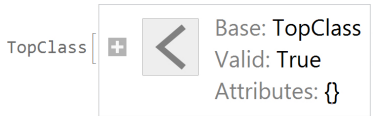
The declared class only has the bare-bone definitions that make up a class, and these are typically expanded firstly by declaring defaults for the elements that are to be included in the class instances. Classes are based internally on associations, and key-value pairs used in classes have the same makeup as their counterparts in **Association** objects.

Provide the defaults for **TopClass**

```
DeclareDefaults[TopClass,
    <|
        "FirstElement" -> 1,
        "DelayedElement" :> 2,
        Plot[x, {x,-1,1}] -> 3
    |>
]
```

For a given instance, values can be assigned to elements by a user as will be discussed in Section A.1.2. It is generally necessary to constrain the allowed values; for a class representing a person it makes sense to allow strings in elements that represent a first or last name, but not for elements that represent age or height. These constraints can be enforced with invariants, which are pure functions that have to return **True** when applied to an **Association** of the instance elements. The tag **ValidQ::elfail** is reserved for messages regarding invariants. It is advised to define failure messages as specifically as possible, to help users with debugging.

Appendix A

```
DeclareInvariant[TopClass,
    {
        IntegerQ[#[["FirstElement"]]] /.
        False :> (
            Message[
                ValidQ::elfail,
                "FirstElement",
                "TopClass",
                "be an integer"
            ];
            False
        ) &
    }
]
```

Once default elements are defined for a class, they take effect even for previously defined objects.

The defaults for **TopClass** are immediately accessible for the previously defined symbol

```
Retrieve[tc, "FirstElement"]
```

1

Invariants can be manually checked using the function **ValidQ**.

Verify that the default elements of **TopClass** satisfy the invariant

```
ValidQ[tc]
```

True

If an instance does not satisfy the invariant, **ValidQ** returns **False** and a warning is printed

```
tcInvalid = Affix[tc, "FirstElement" -> Pi];
ValidQ[tcInvalid]
```

⋯ Affix: The affix operation produced an invalid instance of class TopClass.

⋯ ValidQ: The expression fails to satisfy the supplied invariant.

False

The `ValidQ::elfail` messages are typically suppressed, but can be printed using the option `"Verbose"`.

The `"Verbose"` option of ValidQ provides information to narrow down why an instance is invalid

```
ValidQ[tcInvalid, "Verbose" -> True]
```

⋯ ValidQ: The element FirstElement fails to satisfy the invariant of the class TopClass; its value should be an integer.

⋯ ValidQ: The expression fails to satisfy the supplied invariant.

False

Child classes   A child class can be defined that inherits from a previously defined class.

Declare the class `ChildClass` as a child of `TopClass`

```
DeclareClass[ChildClass, TopClass]
```

Create a simple instance of the `ChildClass` class

```
cc = ChildClass[]
```

ChildClass[➕ ◁ Base: TopClass
Valid: True
Attributes: {}]

Child classes inherit formatting rules, class and instance properties, default elements, and invariants.

A child class inherits defaults from its parent

```
cc[["FirstElement"]]
```

1

The inherited defaults can be redefined by evaluating `DeclareDefaults` with the child class as an argument. `DeclareDefaults` can also be used to specify additional defaults.

Appendix A

```
DeclareDefaults[ChildClass,
    <|
        "FirstElement" -> 10,
        "SecondElement" -> Pi
    |>
];

cc[["FirstElement"]]
```

10

DeclareInvariant can also be used for child classes, though invariants of the parent cannot be removed locally; child classes by construction have to satisfy the invariants of its parents. Instead, **DeclareInvariant** can be used to further narrow down the range of values allowed in class elements. Because **DeclareDefaults** and **DeclareInvariant** check if existing defaults and invariants are satisfied, one usually declares the defaults before defining an invariant.

Declare the invariant for **ChildClass**, which poses additional constraints on top of the invariant already specified for **TopClass**

```
DeclareInvariant[ChildClass, {NumericQ[#[["SecondElement"]]] &}];
```

The child class can be verified to be a valid class with **TopClass** as its parent

```
ValidQ[cc, TopClass]
```

True

```
ValidQ[tc, ChildClass]
```

False

A.1.1.2 Instantiation

Now that we have our classes defined, we need an easy way to create instances of these classes. The default way to do this is by applying the class symbol to an association or a list of rules. This works relatively well, but it is not always convenient to specify different elements whose values are always defined in the same way. In this scenario,

we can overload the class symbol to translate a simpler input into an instance. These generation methods are implemented at the same level as class properties that we will discuss in the next section, so some care should be taken that there is no overlap with e.g. the pre-defined class properties.

Overload a class symbol to simplify instance generation

```
TopClass[int_Integer] := TopClass["FirstElement" -> int]
```

```
TopClass[12]
```


A.1.1.3   Information extraction

The rough setup is now complete, we have a top-level class and a child class that inherits from it, and we have functionality to easily generate class instances. Next we add functionality to these classes in the form of properties and functions. We distinguish between two type of properties: class properties and instance properties.

Note that properties and functions that create new class instance should never create invalid instances.

Class properties   Class properties are defined independent of the elements that make up a class instance, and are stored as `DownValues` of the class symbol.

Define the class property `"ClassString"`, which returns a string with the class name

```
TopClass["ClassString"] = "TopClass";
```

These properties can be defined with the same flexibility as other Wolfram Language definitions, though they are typically used for very class-specific and fixed definitions.

Like other properties and rules, class properties are passed down from the parent class to the child class, so we can request class properties of the child that are only defined specifically for the parent. One simple example of how this helps in practice is with the pre-defined `"Base"` property, which is by default defined only for base classes (classes without parents).

Appendix A

Class properties are inherited by child classes

`ChildClass["ClassString"]`

"TopClass"

Class properties can be overwritten on a per-child basis. Children of a class with an altered class property will inherit from the first parent in the hierarchy which has a definition for the property, and if no definition can be found a message is generated.

Unknown properties remain unevaluated and generate a message

`ChildClass["UnknownString"]`

... **ChildClass**: No definition/value was found for the property UnknownString.

ChildClass["UnknownString"]

Instance properties Completely analogous to class properties, we can define instance properties for class instances, which are easily identified by their heads.

Define the instance property `"FirstElementPlusTwo"`, which adds two to the element `"FirstElement"`

`inst_TopClass["FirstElementPlusTwo"] := inst[["FirstElement"]] + 2;`

`tc["FirstElementPlusTwo"]`

3

Instance properties are inherited just like class properties, and can be redefined for specific purposes just the same. Inheritance is essentially implemented with a catch-all definition for class symbols and instances, which forwards queries to the parent if no definition is found for the class itself.

The instance property `"FirstElementPlusTwo"` is inherited by the children of `TopClass`

`cc["FirstElementPlusTwo"]`

12

Instance properties do not necessarily need to be strings, though for many ap-

plications the majority of properties likely are. The standard usage of for example `CompiledFunction` objects would be an example of (what we loosely refer to as) a property that is not a string, but instead a pattern.

Functions   Properties provide a convenient way to directly query a class instance, but are too constricted for general functionality. As an alternative, functions can be defined for class instances in the usual way.

Define a function that returns the class defaults for the class to which an instance belongs

```
getClassDefaults[inst_?ValidQ] := inst[[0]]["ClassDefaults"]
```

```
getClassDefaults[cc]
```

<| "FirstElement" -> 10, "SecondElement" -> Pi |>

More specific patterns can be used to narrow down function definitions to specific classes, or to specific branches in a hierarchy.

Functions can similarly be defined to act on specific class symbols. There are no strict guidelines for which functionality should be implemented as a property, and which as a function, and the most natural division in practice typically depends on the situation.

A.1.1.4   Formatting

By default, class instances are formatted using the function defined by the `"Format"` property, which is inherited by the class' children. To setup inheritable formatting rules, the user would thus have to overwrite this property. Formatting rules specified through other means are not typically inherited.

A.1.1.5   Classes in package development

When deployed in a package environment, it is convenient to place all class definitions in their own contexts. The class definitions introduced before would be included in a package structure is shown in Code A.1 and A.2.

Functions that apply to specific classes can be placed within the class' context, but they could also be placed in their own context. The rule of thumb when developing packages on top of the Classes` package, is that functions should be as non-intrusive as possible. Core functionality should be implemented as functions that act on standard Wolfram Language expressions, and class functions should only be shells that extract the necessary information from a class instance, feed those to the core functions, and restructure the output into another class instance when needed.

Appendix A

```
(*Create a separate context for every class*)
BeginPackage["myClasses`TopClass`", {"Classes`"}];

(*Declares a top level class TopClass*)
DeclareClass[TopClass];

(*Provide the defaults of TopClass*)
DeclareDefaults[TopClass,
    <|"FirstElement" -> 1, "DelayedElement" :> 2|>
];

(*Declare the invariant of TopClass*)
DeclareInvariant[TopClass,
    {
        IntegerQ[#[["FirstElement"]]] /.
        False :> (
            Message[
                ValidQ::elfail,
                "FirstElement",
                "TopClass",
                "be an integer"
            ];
            False
        ) &
    }
];

(*Define properties for TopClass in the `Private` context to prevent
cluttering of the public myClasses` context*)
Begin["`Private`"];

(*Generative methods*)
TopClass[int_Integer] := TopClass["FirstElement" -> int]

(*Class properties*)
TopClass["ClassString"] = "TopClass";

(*Instance properties*)
inst_TopClass["FirstElementPlusTwo"] := inst[["FirstElement"]] + 2;

(*End `Private` context*)
End[];

(*End package context*)
EndPackage[];
```

Code A.1

236

Structure of a Classes`-based package (cont.)

```
(*Place the child class in a subcontext of the top-level class*)
BeginPackage["myClasses`TopClass`ChildClass`",
    {
        "Classes`",
        "myClasses`TopClass`"
    }
];

(*Declare a second class that inherits properties, invariant, and
defaults from TopClass*)
DeclareClass[ChildClass, TopClass];

(*Introduce one new element ("SecondElement"), and change the
default value of another ("FirstElement")*)
DeclareDefaults[ChildClass,
    <|
        "FirstElement" -> 10,
        "SecondElement" -> Pi
    |>
];

(*The declared invariants are added on top of the invariant of the
parent class*)
DeclareInvariant[ChildClass, {NumericQ[#[["SecondElement"]]] &}];

(*Begin the `Private` context*)
Begin["`Private`"];

(*Again we can add custom properties, and overwrite others*)
inst_ChildClass["SecondElementDivTwo"] := inst[["SecondElement"]]/2;

(*End the `Private` context*)
End[];

(*End the package*)
EndPackage[];
```

Code A.2

237

A.1.2   Working with classes

In this section we walk through the basics of working with packages built on top of Classes`, which would be the typical situation for an end-user.

There are two fundamental objects in the Classes` package, class symbols and class instances. Class symbols are simply symbols that have been designated to be classes with **DeclareClass**, and to which all functionality of a specific class is attached. Class instances are the actual class objects that a user works with, which represent mutable data that are naturally grouped together. Among other things, this section covers how to modify and extract data stored in class instances. The basic functions used to work with class symbols and instances are listed below.

ClassQ	check if a symbol is a declared class
ValidQ	validate a class instance
Affix	add or change an element value in an instance
AffixTo	add or change an element value in a symbol representing an instance
ClassSet	operator form of **AffixTo** that works similar to **Set**
ClassSetDelayed	operator form of **AffixTo** that works similar to **SetDelayed**
Retrieve	retrieve values of elements in an instance

Basic functions that work with class symbols or instances.

A.1.2.1   Classes basics

Start with loading the Classes` package, and setting up a very simple class structure that will be used throughout this section.

Load the Classes` context and create some classes.

```
Needs["Classes`"];
DeclareClass[TopClass];
DeclareDefaults[TopClass, <|"FirstElement" -> 1|>];
DeclareClass[ChildClass, TopClass];
inst_TopClass["FirstElementSquared"] := Normal[inst]["FirstElement"]^2
```

Class symbols   The first basic object in a Classes`-based package are the class symbols. These are symbols that have some basic functionality that simplifies their usage as class symbols, which are typically created using the function `DeclareClass`. The function `ClassQ` can be used to check whether a certain symbol is a class symbol.

Check if a symbol is a class

`ClassQ[TopClass]`

True

All functionality for class symbols and class instances is stored as `DownValues`, `UpValues`, etc. for the class symbol. These definitions are typically hidden from the user with the attribute `ReadProtected`, but unless the symbol is locked users can access this information by removing the `ReadProtected` attribute.

Class instances   Class instances are the objects that can be created as belonging to a specific class. These instances work essentially like `Association` objects, with key-value pairs in which information is stored. Instances are most easily created by mapping the class symbol on an `Association` object.

Create a class instance

`tc = TopClass[<|"FirstElement" -> 10|>]`

TopClass [⊞ ◀ Base: TopClass / Valid: True / Attributes: {}]

Class instances can be reduced to an `Association` using `Normal`.

Convert a class instance to an `Association`

`Normal[tc]`

<| "FirstElement" -> 10 |>

Values stored in class instances can be accessed with the function `Retrieve`, or, alternatively, using the overloaded Wolfram Language function `Part`.

Appendix A

```
Retrieve[tc, "FirstElement"]
```

10

```
tc[["FirstElement"]]
```

10

If necessary the overloaded **Part** functionality can be disabled using **Block**.

Extracting information using **Part**, with and without overloading

```
FullForm[tc]
```

TopClass[Association[Rule["FirstElement", 10]]]

```
tc[[1]]
```

10

```
Block[{Classes`Private`$PartOverload = False}, tc[[1]]]
```

<| "FirstElement" -> 10 |>

Values can be modified and added using **Affix** and **AffixTo**, which work on class instances similar to how **Append** and **AppendTo** work on **Association** objects.

Modify and add values in a class instance

Affix[tc, "FirstElement" -> 12]

TopClass ⎡ ⊞ ◁ Base: TopClass
Valid: True
Attributes: {} ⎤
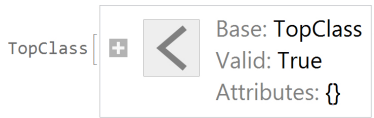
Normal[tc]

<| "FirstElement" -> 10|>

AffixTo[tc, "SecondElement" -> 12]

TopClass ⎡ ⊞ ◁ Base: TopClass
Valid: True
Attributes: {} ⎤

Normal[tc]

<| "FirstElement" -> 10, "SecondElement" -> 12 |>

For now it is not possible to use **Set** and **SetDelayed** to modify class instances, but **ClassSet** and **ClassSetDelayed** function in much the same way. **ClassSet** and **ClassSetDelayed** have the infix forms $\overset{<}{=}$ and $\overset{<}{:=}$.

Modify and add values in a class instance using **ClassSet**

tc[["FirstElement"]] $\overset{<}{=}$ 1

1

Normal[tc]

<| "FirstElement" -> 1, "SecondElement" -> 12 |>

ClassSet can also be used to convert an instance of one class to an instance of another class. This is done by changing its head, which can be accessed with **Part**.

Convert a class instance

$\texttt{tc[[0]]} \overset{\leq}{=} \texttt{ChildClass}$

ChildClass

tc

ChildClass [⊞ < Base: TopClass
Valid: True
Attributes: {}]
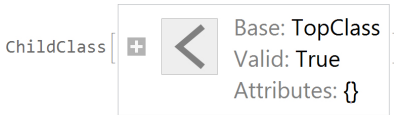
Values stored in class instances are generally constrained by the class invariant, which is a pure function that gives **True** if the **Association** object derived from a class instance represents a valid structure. This invariant can be defined using the function **DeclareInvariant**. Invalid instances should never occur unless users incorrectly modify an instance themselves. Package functions should never generate invalid instances. The function **ValidQ** can be used to validate instances, and the option **"Verbose"** can be used to help figure out why an instance is invalid.

Validate an instance using **ValidQ**

ValidQ[tc]

True

Affix, **AffixTo**, **ClassSet**, and **ClassSetDelayed**, are all low-level functions, which can in principle generate invalid instances. This is the case, because functions are allowed (though discouraged) to use these low-level functions to generate invalid instances as intermediate results. **Affix** and other functions will generate error messages if an invalid instance is created.

Note that there is no real encapsulation in Classes`. All elements in an instance can always be modified.

A.1.2.2 Properties

Properties are the basic means for the user to interact with class symbols and instances. They provide a natural way to quickly extract information from an instance in a form that is useful for a user. The second way of interacting with class symbols and instances is through functions, but the details of those functions are specific to the Classes`-based packages and should be covered in the corresponding documentation.

All properties of a parent class are defined (by inheritance) for a child class, but not all properties for a child class necessarily exist for every parent. In rare cases, properties that exist for a parent class can be redefined, re-implemented, or even re-purposed for child classes.

Class properties   Class properties are `DownValues` of a class symbol, and provide a simple way to store class specific information. There are some default class properties defined at declaration.

Class properties are data and methods stored as `DownValues` of a class symbol

`TopClass["Class"]`

TopClass

`TopClass["Instance"]`

`TopClass["Defaults"]`

<| "FirstElement" -> 1|>

String properties, which typically make up the majority of class properties, can be listed with the `"Properties"` property. Non-string class properties are generally methods used in class instance generation, which should be covered in the package's documentation.

List all string properties

`TopClass["Properties"]`

{"Base", "Class", "ClassDefaults", "ClassInvariant", "Defaults", "Format", "Instance", "Invariant", "ParentDefaults", "ParentInvariant", "Parents", "Properties"}

Instance properties   Instance properties are implemented as `SubValues` of a class symbol, so that they are accessed in the same way as class properties.

Instance properties provide access to information derived from the data stored in a class instance

```
tc["FirstElementSquared"]
```

1

String properties, which typically make up the majority of class properties, can be listed with the `"Properties"` property. The few non-string properties should be clearly described in the package documentation, and are typically almost self-evident, as in e.g. `CompiledFunction`.

List all string properties of an instance

```
tc["Properties"]
```

{"FirstElementSquared", "Properties"}

Properties are generally implemented with the assumption that an instance is valid, and errors can occur if this is not the case. There is no strict requirement for properties to ensure validity of an input instance, though this may be enforced differently in different packages.

A.1.2.3 Customization

The more advanced use of the Classes` package, even for some end-users, is the custom addition of instance properties, classes, and functions. One other typical example would be a user-side copy of a class, to which they can add their own definitions and modifications without having to modify any package functionality.

A child of a class, without any changes to the defaults and the invariant, can function as a copy of the class

```
DeclareClass[myTopClass, TopClass]
```

A.2   The Diffusica` package

Diffusica` is a modular package built on top of the Classes` package, intended for medical image processing and focused at the moment mainly on the analysis of diffusion-weighted MRI data. The tutorial presented in this section does not cover most of the more advanced functionality included in Diffusica`, of which a lot is still under development and thus not fully documented. Some of the functions in the

package are not yet implemented, and are referenced here to guide the development of these functions when there comes a need to include them.

A.2.1 Package design

In this tutorial we cover the design of the Diffusica`  package. We discuss the use of classes to represent data, how algorithms can be designed to process these classes and their instances, and how the Diffusica` package is structured to facilitate quick prototyping in the development of such algorithms.

Included functionality is still limited at the current stage, as the package is intended to be used for collaborative development.

Load the Diffusica` context

```
Needs["Diffusica`"]
```

A.2.1.1 The Object class

The `Object` class is used to associate information to positions in a common space $\{z, y, x, \dots\}$, and is the core structure used to represent data in Diffusica`. Data inside an `Object` instance is stored internally in an object space $\{z', y', x', \dots\}$.

Different types of data sets impose different constraints on the data, and so these different types can be implemented as child classes of `Object`. Basic algorithms acting on these data sets can then be implemented as instance properties, while more complex algorithms can be implemented as class functions. The `Object` class is stored in the Diffusica`Classes` context, and child classes are included in the appropriate subcontexts, e.g. in Diffusica`Classes`Object`.

The `Object` class has the following default elements:

Common space and object space The `"Transformation"` element contains a function used to translate points $\{z', y', x', \dots\}$ in the object space to a common space. This transformation function is typically affine, but in future releases will also allow arbitrary non-rigid transformations. The object space always has `"DimensionlessUnit"` units to facilitate storage in `PackedArray` objects, though the common space can have arbitrary units as specified through the `"Transformation"` element. Child classes of `Object` typically constrain the dimensions and units of the common space.

Localized data The `"Data"` and `"Positions"` elements provide a flexible way to store localized information. The `"Data"` element is a `PackedArray` object of arbitrary dimensions. The `"Positions"` element is either an integer denoting the dimensionality of the `"Data"` array, or a `PackedArray` object of compatible dimensions that provides the

Appendix A

`"Background"`	`Automatic`	The background value used when there is no defined data
`"Data"`	`0`	Packed array or number containing the data stored in the object
`"History"`	`{}`	A record of functions applied to the object
`"Labels"`	`{}`	A list of rules that associates labels to specific values in `"Data"`
`"Positions"`	`0`	Packed array of positions in common space associated to `"Data"`, or an integer indicating the dimensionality of the data
`"Source"`	`Null`	The source from which the instance is derived
`"Transformation"`	`Automatic`	The mapping from object space to common space
`"InverseTransformation"`	`Automatic`	The mapping from common space to object space

Elements in the **Object** class.

(index-wise) corresponding positions explicitly. In the latter case, the last dimension of the `"Positions"` array determines the dimensionality of the stored data. If the `"Data"` array has fewer dimensions than indicated by the `"Positions"` element, only the first indices are used to determine a value. With the default unspecified `"Data"`, `"Positions"` can be used to store purely positional information such as points and curves. Using **Explicate**, both the `"Data"` and the `"Positions"` elements can be given explicit values, such that for every index in `"Data"` there is a corresponding element in `"Positions"` that provides the corresponding position.

Extracting values   The data associated to a point $\{z, y, x, \dots\}$ in the common space can be extracted using *instance*$[z, y, x, \dots]$, which by default uses a nearest neighbor algorithm in the common space. A default value, determined by the `"Background"` element, is returned if there is no data associated with a given position. If `"Background"` is **Automatic**, a zero array of the proper dimensions is used. Child classes of **Object** may use different interpolation schemes by default.

Labels   The `"Labels"` element is a list of rules that transforms values in the `"Data"` element to labels. These labels can be used for example to identify different data sets, or to provide localized visualization options.

Meta data   The `"History"` and `"Source"` elements are used to store processing meta data of an object. The `"Source"` element provides the expression or file from which the instance is derived, and `"History"` contains a list of all functions that have been applied to the instance. This information should in principle be sufficient to recreate an object instance.

Object processing   There is a large set of basic `Object` processing functionality included in the Diffusica`Processing` context. These functions are based on the Wolfram Language's handling of `Image` and `Image3D` expressions, and many functions (e.g. `ImageAdd`) have `Object` counterparts (e.g. `DataAdd`). These functions are all written specifically to take into account the common space. For example, `DataAdd` will perform addition in the common space, not in the object space.

A.2.1.2   Functions

Diffusica` follows the design principles of the Wolfram language and the Classes` package—basic information about an object is implemented through class and instance properties, while more general and more complicated functionality is accessible through high level functions. These functions are typically wrappers that provide access to class-specific implementations.

Diffusica` has the following symbols reserved for these wrapper functions. Not all of the listed functions are implemented in the current version of the package, and new functions are likely to be added in future releases.

`Construction`	Generate data from a model
`Enhancement`	Enhance data
`Reconstruction`	Fit a model to a data set
`Registration`	Compute a registration to an atlas or a data set
`Simulation`	Generate simulated data
`Tractography`	Compute tractograms
`Visualize`	Visualize class instances

Symbols reserved for wrapper functions in the Diffusica` package.

Appendix A

A.2.1.3 Formats

Diffusica` is designed to include **Import**/**Export** converters to support specific data formats. Because of the current emphasis on diffusion MRI, a converter for the NIfTI file format is included for example. The NIfTI converter can be downloaded separately from `github.com/tomdelahaije`.

Example usage of the NIfTI converter.

```
{meta, data} = Import["C:\\path\\avg152T1_RL_nifti.nii.gz", "NIfTI"][[;; , 2]];
Dataset[meta]
```

| "AuxiliaryFileName" → "none" |
| "BitDepth" → 8 |
| "ByteOrdering" → 1 |
| "DataOffset" → 352 |
| "DataScaling" → Identity |
| "DataType" → "Byte" |
| "DefaultCoordinates" → "MNI152NormalizedCoordinates" |
| "Description" → "FSL3.2beta" |
| "Dimensions" → {91, 109, 91} |
| "ExtendedHeader" → None |
| "ExtensionOffset" → 352 |
| "ExtensionParsing" → False |
| "Extensions" → {} |
| "HeaderLength" → 348 |
| "IntensityRange" → {0., 255.} |
| "Intent" → Undefined |
| "IntentName" → Undefined |
| "NIfTIUnused" → {"", "", 0, 0, "r", 0, 0} |
| "NIfTIVersion" → 1 |
| "NumberOfDimensions" → 3 |

K <   showing 1–20 of 31   > >|

SlicePlot is a Diffusica` built-in function used to visualize arbitrary planes in three-dimensional data sets.

SlicePlot[data]

248
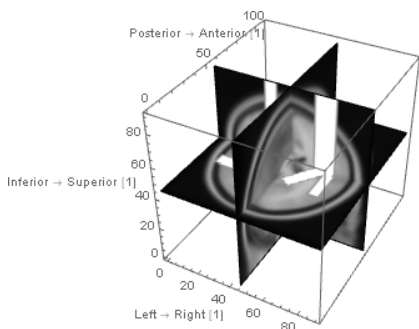
Converters are stored in the Diffusica`Convert` context, but are always registered to the Wolfram Language **Import**/**Export** functions. If applicable, as is the case for the NIfTI format, files can often also be converted directly to an `Object` instance using instance generation.

Appendix A

Appendix B

Expressions

Spherical Harmonics Expressions and Identities

Throughout this thesis we have employed spherical harmonics Y_l^m to describe functions on the sphere. We work with the convention

$$Y_l^m(\theta, \phi) := \sqrt{\frac{2l+1}{4\pi}} \sqrt{\frac{(l-m)!}{(l+m)!}} e^{im\phi} P_l^m(\cos \theta), \qquad (B.1)$$

with $l = 0, 1, \dots$ and $m = -l, -l+1, \dots, l$, and with P_l^m the associated Legendre polynomial. The spherical harmonics are orthonormal, i.e., we have

$$\int_{S^2} \overline{Y_{l_1}^{m_1}(\hat{r})} Y_{l_2}^{m_2}(\hat{r}) \, d\sigma(\hat{r}) = \delta_{l_1, l_2} \delta_{m_1, m_2}, \qquad (B.2)$$

where the bar indicates complex conjugation, $d\sigma$ is the Lebesgue measure on the sphere S^2, and δ is the Kronecker delta with $\delta_{i,j} = 1$ if $i = j$ and 0 otherwise. We further rely on

$$Y_l^m = Y_{-(l+1)}^m \qquad (B.3)$$

to extend the definition of spherical harmonics to negative integers l. Finally, we typically express real-valued function on the sphere with the modified [101] spherical harmonic basis $\tilde{Y}_l^m$ defined as

$$\tilde{Y}_l^m(\theta, \phi) := \begin{cases} \sqrt{2} \, \mathrm{Re} \left\{ Y_l^m(\theta, \phi) \right\} & m < 0 \\ Y_l^0(\theta, \phi) & m = 0 \\ \sqrt{2} \, \mathrm{Im} \left\{ Y_l^m(\theta, \phi) \right\} & m > 0 \end{cases}. \qquad (B.4)$$

Appendix B

The fundamental dual tensor defined in Eq. (3.40) can be expressed in terms of the spherical harmonic expansion coefficients c' in $H^2(r, \theta, \phi) = r^2 \sum_{l=0}^{\infty} \sum_{m=-l}^{l} c_l'^m Y_l^m(\theta, \phi)$, such that $h^{ij}(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} c_l'^m Q_l^m(\theta, \phi)$ for some alternative basis Q_l^m. The coefficients c' can be expressed in terms of c using the Clebsch–Gordan coefficients. Simply applying the (inverse) change of coordinates given in Eq. (3.32) to the function $r^2 Y'(\theta, \phi)$, with $Y'(\theta, \phi) := \sum_{l=0}^{\infty} \sum_{m=-l}^{l} c_l'^m Y_l^m(\theta, \phi)$, and computing the proper derivatives then gives the following expression for $h^{ij}(\theta, \phi)$:

$$
\begin{cases}
h^{11}(\theta, \phi) = a_1 \cos^2 \phi + a_2 \sin^2 \phi - a_3 \sin(2\phi) + a_0, \\
h^{12}(\theta, \phi) = a_3 \cos^2 \phi - a_3 \sin^2 \phi + \dfrac{a_1 - a_2}{2} \sin(2\phi), \\
h^{22}(\theta, \phi) = a_2 \cos^2 \phi + a_1 \sin^2 \phi + a_3 \sin(2\phi) + a_0, \\
h^{13}(\theta, \phi) = a_5 \cos \phi + a_6 \sin \phi, \\
h^{23}(\theta, \phi) = a_5 \sin \phi - a_6 \cos \phi, \\
h^{33}(\theta, \phi) = a_4 + a_0,
\end{cases}
\tag{B.5}
$$

where

$$
a_0 := a_0(\theta, \phi) = Y'(\theta, \phi), \tag{B.6}
$$

$$
a_1 := a_1(\theta, \phi) = \frac{1}{2} \left(Y'_{\theta\theta}(\theta, \phi) \cos^2 \theta + Y'_\theta(\theta, \phi) \sin 2\theta \right), \tag{B.7}
$$

$$
a_2 := a_2(\theta, \phi) = \frac{1}{2} \left(Y'_\theta(\theta, \phi) \cot \theta + Y'_{\phi\phi}(\theta, \phi) \csc^2 \theta \right), \tag{B.8}
$$

$$
a_3 := a_3(\theta, \phi) = \frac{1}{2} \left(Y'_{\theta\phi}(\theta, \phi) \cot \theta - Y'_\phi(\theta, \phi) \cos 2\theta \csc^2 \theta \right), \tag{B.9}
$$

$$
a_4 := a_4(\theta, \phi) = \frac{1}{2} \left(Y'_{\theta\theta}(\theta, \phi) \sin^2 \theta - Y'_\theta(\theta, \phi) \sin 2\theta \right), \tag{B.10}
$$

$$
a_5 := a_5(\theta, \phi) = \frac{1}{2} \left(Y'_\theta(\theta, \phi) \cos 2\theta - Y'_{\theta\theta}(\theta, \phi) \cos \theta \sin \theta \right), \tag{B.11}
$$

$$
a_6 := a_5(\theta, \phi) = \frac{1}{2} \left(Y'_{\theta\phi}(\theta, \phi) - 2 Y'_\phi(\theta, \phi) \cot \theta \right), \tag{B.12}
$$

and where the subscripts in Y' denote differentiation. The (artificial) singularities at $\theta = 0, \pi$ that occur in these expressions can be avoided by applying the identities given in the work of Eshagh [118], see also Janssen et al. [170]. Eq. (3.41) is most easily expressed in terms of the coefficients c, where we write $Y(\theta, \phi) := \sum_{l=0}^{\infty} \sum_{m=-l}^{l} c_l^m(\kappa) Y_l^m(\theta, \phi)$ to get

$$
\xi(r, \theta, \phi) = r Y(\theta, \phi) \begin{pmatrix} Y(\theta, \phi) \sin \theta \cos \phi + Y_\theta(\theta, \phi) \cos \theta \cos \phi - Y_\phi(\theta, \phi) \csc \theta \sin \phi \\ Y(\theta, \phi) \sin \theta \sin \phi + Y_\theta(\theta, \phi) \cos \theta \sin \phi + Y_\phi(\theta, \phi) \csc \theta \cos \phi \\ Y(\theta, \phi) \cos \theta \end{pmatrix}.
\tag{B.13}
$$

This gives

$$\left\|\xi(r, \theta, \phi)\right\| = rY(\theta, \phi)\sqrt{Y^2(\theta, \phi) + Y_\theta^2(\theta, \phi) + Y_\phi^2(\theta, \phi)\csc^2\theta}. \tag{B.14}$$

The determinant of the fundamental dual tensor h is finally given by

$$\det h(\theta, \phi) = b_0^4\left[(b_0 + b_1 + b_3)(b_0 + b_5) - (b_2 - b_4)^2\right], \tag{B.15}$$

with

$$b_0 := b_0(\theta, \phi) = Y(\theta, \phi), \tag{B.16}$$
$$b_1 := b_1(\theta, \phi) = Y_\theta(\theta, \phi)\cot\theta, \tag{B.17}$$
$$b_2 := b_2(\theta, \phi) = Y_\phi(\theta, \phi)\cot\theta\csc\theta, \tag{B.18}$$
$$b_3 := b_3(\theta, \phi) = Y_{\theta\theta}(\theta, \phi), \tag{B.19}$$
$$b_4 := b_4(\theta, \phi) = Y_{\theta\phi}(\theta, \phi)\csc\theta, \tag{B.20}$$
$$b_5 := b_5(\theta, \phi) = Y_{\phi\phi}(\theta, \phi)\csc^2\theta. \tag{B.21}$$

Note that these expressions can be trivially applied to the Finsler function expression in Eq. (3.43) and dual expressions.

Appendix B

Bibliography

[1] F. Aboitiz, A. B. Scheibel, R. S. Fisher, and E. Zaidel. "Fiber composition of the human corpus callosum". In: *Brain Research* 598.1-2 (Dec. 1992). PMID: 1486477, pp. 143–153. ISSN: 0006-8993.

[2] I. Aganj, C. Lenglet, and G. Sapiro. "ODF reconstruction in q-ball imaging with solid angle consideration". In: *Biomedical Imaging: From Nano to Macro, 2009. ISBI'09. IEEE International Symposium on*. 2009, pp. 1398–1401.

[3] S. Aja-Fernández, A. Tristán-Vega, and C. Alberola-López. "Noise estimation in single- and multiple-coil magnetic resonance data based on statistical models". In: *Magnetic Resonance Imaging* 27.10 (Dec. 2009), pp. 1397–1409. ISSN: 0730725X. DOI: `10.1016/j.mri.2009.05.025`.

[4] A. L. Alexander, K. M. Hasan, M. Lazar, J. S. Tsuruda, and D. L. Parker. "Analysis of partial volume effects in diffusion-tensor MRI". In: *Magnetic Resonance in Medicine* 45.5 (2001), pp. 770–780.

[5] A. L. Alexander, K. Hasan, G. Kindlmann, D. L. Parker, and J. S. Tsuruda. "A geometric analysis of diffusion tensor measurements of the human brain". In: *Magnetic Resonance in Medicine* 44.2 (Aug. 2000), pp. 283–291. ISSN: 0740-3194, 1522-2594. DOI: `10.1002/1522-2594(200008)44:2<283::AID-MRM16>3.0.CO;2-V`.

[6] A. L. Alexander, J. E. Lee, M. Lazar, and A. S. Field. "Diffusion tensor imaging of the brain". In: *Neurotherapeutics* 4.3 (July 2007), pp. 316–329. ISSN: 1933-7213, 1878-7479. DOI: `10.1016/j.nurt.2007.05.011`.

[7] D. Alexander, C. Pierpaoli, P. Basser, and J. Gee. "Spatial transformations of diffusion tensor magnetic resonance images". In: *IEEE Transactions on Medical Imaging* 20.11 (Nov. 2001), pp. 1131–1139. ISSN: 02780062. DOI: `10.1109/42.963816`.

[8] J. L. Andersson, S. Skare, and J. Ashburner. "How to correct susceptibility distortions in spin-echo echo-planar images: Application to diffusion tensor imaging". In: *NeuroImage* 20.2 (Oct. 2003), pp. 870–888. ISSN: 10538119. DOI: `10.1016/S1053-8119(03)00336-7`.

[9] M. Ankele and T. Schultz. "Quantifying microstructure in fiber crossings with diffusional kurtosis". In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2015, pp. 150–157.

[10] P. L. Antonelli and T. J. Zastawniak. *Fundamentals of Finslerian diffusion with applications*. Dordrecht: Springer Netherlands, 1999. ISBN: 978-94-010-6023-3 94-010-6023-1 978-94-011-4824-5 94-011-4824-4.

[11] M. ApS. *The MOSEK command line tool*. Denmark, May 2016.

[12] G. B. Arfken and H.-J. Weber. *Mathematical methods for physicists*. 6th ed. Boston: Elsevier, 2005. ISBN: 978-0-12-059876-2 978-0-12-088584-8.

[13] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. "Log-Euclidean metrics for fast and simple calculus on diffusion tensors". In: *Magnetic resonance in medicine* 56.2 (2006), pp. 411–421.

[14] Y. Assaf and P. J. Basser. "Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain". In: *NeuroImage* 27.1 (Aug. 2005), pp. 48–58. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2005.03.042.

[15] Y. Assaf, T. Blumenfeld-Katzir, Y. Yovel, and P. J. Basser. "Axcaliber: A method for measuring axon diameter distribution from diffusion MRI". In: *Magnetic Resonance in Medicine* 59.6 (June 2008), pp. 1347–1354. ISSN: 07403194, 15222594. DOI: 10.1002/mrm.21577.

[16] H.-E. Assemlal, D. Tschumperlé, L. Brun, and K. Siddiqi. "Recent advances in diffusion MRI modeling: Angular and radial reconstruction". In: *Medical image analysis* 15.4 (2011), pp. 369–396.

[17] L. J. Astola and L. M. J. Florack. "Finsler geometry on higher order tensor fields and applications to high angular resolution diffusion imaging". In: *International Journal of Computer Vision* 92.3 (2011), pp. 325–336.

[18] L. Astola, L. Florack, and B. ter Haar Romeny. "Measures for pathway analysis in brain white matter using diffusion tensor images". In: *Proceedings of the Biennial International Conference on Information Processing in Medical Imaging*. Springer, 2007, pp. 642–649.

[19] L. Astola, A. Fuster, and L. Florack. "A Riemannian scalar measure for diffusion tensor images". In: *Pattern Recognition* 44.9 (Sept. 2011), pp. 1885–1891. ISSN: 00313203. DOI: 10.1016/j.patcog.2010.09.009.

[20] H. Axer, M. Axer, T. Krings, and D. G. v. Keyserlingk. "Quantitative estimation of 3-D fiber course in gross histological sections of the human brain using polarized light". In: *Journal of neuroscience methods* 105.2 (2001), pp. 121–131.

[21] D. Bao, S.-S. Chern, and Z. Shen. *An introduction to Riemann-Finsler geometry*. New York [etc.]: Springer, 2000. ISBN: 0-387-98948-X 978-0-387-98948-8.

[22] M. Barakovic, D. Romascano, T. B. Dyrby, D. C. Alexander, M. Descoteaux, J.-P. Thiran, and A. Daducci. "Estimation of bundle-specific axon diameter distributions with microstructure informed tractography". In: (2016).

[23] A. Barmpoutis, B. Vemuri, T. Shepherd, and J. Forder. "Tensor splines for interpolation and approximation of dt-mri with applications to segmentation of isolated rat hippocampi". In: *IEEE Transactions on Medical Imaging* 26.11 (Nov. 2007), pp. 1537–1546. ISSN: 0278-0062. DOI: 10.1109/TMI.2007.903195.

[24] P. J. Basser, J. Mattiello, and D. Le Bihan. "Estimation of the effective self-diffusion tensor from the NMR spin echo". In: *Journal of Magnetic Resonance, Series B* 103.3 (Mar. 1994), pp. 247–254. ISSN: 10641866. DOI: 10.1006/jmrb.1994.1037.

[25] P. J. Basser, J. Mattiello, and D. Le Bihan. "MR diffusion tensor spectroscopy and imaging". In: *Biophysical journal* 66.1 (Jan. 1994). PMID: 8130344, pp. 259–267. ISSN: 0006-3495. DOI: 10.1016/S0006-3495(94)80775-1.

[26] P. J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. "In vivo fiber tractography using DT-MRI data". In: *Magnetic resonance in medicine* 44.4 (2000), pp. 625–632.

[27] P. J. Basser and C. Pierpaoli. "Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI". In: *Journal of magnetic resonance. Series B* 111.3 (1996), pp. 209–219.

[28] P. G. Batchelor, M. Moakher, D. Atkinson, F. Calamante, and A. Connelly. "A rigorous framework for diffusion tensor calculus". In: *Magnetic Resonance in Medicine* 53.1 (Jan. 2005), pp. 221–225. ISSN: 0740-3194, 1522-2594. DOI: 10.1002/mrm.20334.

[29] P. G. Batchelor, D. L. Hill, D. Atkinson, and F. Calamante. "Study of connectivity in the brain using the full diffusion tensor from MRI". In: *Information Processing in Medical Imaging*. Ed. by G. Goos, J. Hartmanis, J. van Leeuwen, M. F. Insana, and R. M. Leahy. Vol. 2082. Berlin, Heidelberg: Springer Berlin Heidelberg, 2001, pp. 121–133. ISBN: 978-3-540-42245-7 978-3-540-45729-9.

[30] C. Beaulieu. "The basis of anisotropic water diffusion in the nervous system - A technical review". In: *NMR in Biomedicine* 15.7-8 (Nov. 2002), pp. 435–455. ISSN: 0952-3480, 1099-1492. DOI: 10.1002/nbm.782.

[31] S. M. A. Becker, K. Tabelow, H. U. Voss, A. Anwander, R. M. Heidemann, and J. Polzehl. "Position-orientation adaptive smoothing of diffusion weighted magnetic resonance data (POAS)". In: *Medical Image Analysis* 16.6 (Aug. 2012), pp. 1142–1155. ISSN: 13618415. DOI: 10.1016/j.media.2012.05.007.

[32] T. E. J. Behrens, M. W. Woolrich, M. Jenkinson, H. Johansen-Berg, R. G. Nunes, S. Clare, P. M. Matthews, J. M. Brady, and S. M. Smith. "Characterization and propagation of uncertainty in diffusion-weighted MR imaging". In: *Magnetic Resonance in Medicine* 50.5 (Nov. 2003), pp. 1077–1088. ISSN: 0740-3194, 1522-2594. DOI: 10.1002/mrm.10609.

[33] T. Behrens, H. Johansen-Berg, S. Jbabdi, M. Rushworth, and M. Woolrich. "Probabilistic diffusion tractography with multiple fibre orientations: What can we gain?" In: *NeuroImage* 34.1 (Jan. 2007), pp. 144–155. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2006.09.018.

[34] S. Bells, M. Cercignani, S. Deoni, Y. Assaf, O. Pasternak, C. Evans, A. Leemans, and D. Jones. ""Tractometry" – Comprehensive multi-modal quantitative assessment of white matter along specific tracts". In: *Proceedings of the 19th Annual Meeting of the ISMRM*. Montréal, 2011, p. 0678.

[35] J. Bergh and J. Löfström. *Interpolation spaces: An introduction*. Ed. by S. S. Chern, J. L. Doob, J. Douglas, A. Grothendieck, E. Heinz, F. Hirzebruch, E. Hopf, S. Mac Lane, W. Magnus, M. M. Postnikov, F. K. Schmidt, W. Schmidt, D. S. Scott, K. Stein, J. Tits, B. L. van der Waerden, B. Eckmann, and J. K. Moser. Vol. 223. Grundlehren der mathematischen Wissenschaften. Berlin, Heidelberg: Springer Berlin Heidelberg, 1976. ISBN: 978-3-642-66453-3 978-3-642-66451-9.

[36] C. R. V. Blain, S. Brunton, V. C. Williams, A. Leemans, M. R. Turner, P. M. Andersen, M. Catani, B. R. Stanton, J. Ganesalingham, D. K. Jones, S. C. R. Williams, P. N. Leigh, and A. Simmons. "Differential corticospinal tract degeneration in homozygous 'D90A' SOD-1 ALS and sporadic ALS". In: *Journal of Neurology, Neurosurgery & Psychiatry* 82.8 (Apr. 2011), pp. 843–849. ISSN: 0022-3050. DOI: 10.1136/jnnp.2010.236018.

[37] F. Bloch. "Nuclear induction". In: *Physical Review* 70.7-8 (Oct. 1946), pp. 460–474. ISSN: 0031-899X, 1536-6065. DOI: 10.1103/PhysRev.70.460.

[38] R. de Boer, M. Schaap, F. van der Lijn, H. A. Vrooman, M. de Groot, A. van der Lugt, M. A. Ikram, M. W. Vernooij, M. M. Breteler, and W. J. Niessen. "Statistical analysis of minimum cost path based structural brain connectivity". In: *NeuroImage* 55.2 (Mar. 2011), pp. 557–565. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2010.12.012.

[39] B. Borchers. "CSDP, a C library for semidefinite programming". In: *Optimization Methods and Software* 11.1-4 (Jan. 1999), pp. 613–623. ISSN: 1055-6788, 1029-4937. DOI: 10.1080/10556789908805765.

[40] J.-P. Bouchaud and A. Georges. "Anomalous diffusion in disordered media: Statistical mechanisms, models and physical applications". In: *Physics Reports* 195.4-5 (Nov. 1990), pp. 127–293. ISSN: 03701573. DOI: 10.1016/0370-1573(90)90099-N.

[41] M. Bozzali, M. Franceschi, A. Falini, S. Pontesilli, M. Cercignani, G. Magnani, G. Scotti, G. Comi, and M. Filippi. "Quantification of tissue damage in AD using diffusion tensor and magnetization transfer MRI". In: *Neurology* 57.6 (Sept. 2001). PMID: 11571355, pp. 1135–1137. ISSN: 0028-3878.

[42] M. Bozzali, G. Giulietti, B. Basile, L. Serra, B. Spanò, R. Perri, F. Giubilei, C. Marra, C. Caltagirone, and M. Cercignani. "Damage to the cingulum contributes to Alzheimer's disease pathophysiology by deafferentation mechanism". In: *Human Brain Mapping* 33.6 (June 2012), pp. 1295–1308. ISSN: 10659471. DOI: 10.1002/hbm.21287.

[43] R. Brown. "A brief account of microscopical observations made in the months of June, July and August 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies". In: *The Philosophical Magazine and Annals of Philosophy*. 2nd ser. 4.21 (1828), pp. 161–173. ISSN: 1941-5850, 1941-5869. DOI: 10.1080/14786442808674769.

[44] E. Bullmore and O. Sporns. "Complex brain networks: Graph theoretical analysis of structural and functional systems". In: *Nature reviews. Neuroscience* 10.3 (Mar. 2009). PMID: 19190637, pp. 186–198. ISSN: 1471-0048. DOI: 10.1038/nrn2575.

[45] L. M. Burcaw, E. Fieremans, and D. S. Novikov. "Mesoscopic structure of neuronal tracts from time-dependent diffusion". In: *NeuroImage* 114 (July 2015), pp. 18–37. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2015.03.061.

[46] B. Burgeth, L. Pizarro, and S. Didas. "Edge-enhancing diffusion filtering for matrix fields". In: *New Developments in the Visualization and Processing of Tensor Fields*. Springer, 2012, pp. 51–67.

[47] B. Burgeth and J. Weickert. "An explanation for the logarithmic connection between linear and morphological system theory". In: *International Journal of Computer Vision* 64.2-3 (Sept. 2005), pp. 157–169. ISSN: 0920-5691, 1573-1405. DOI: 10.1007/s11263-005-1841-z.

[48] P. T. Callaghan, C. D. Eccles, and Y. Xia. "NMR microscopy of dynamic displacements: k-space and q-space imaging". In: *Journal of Physics E: Scientific Instruments* 21.8 (1988), p. 820.

[49] P. T. Callaghan. *Principles of nuclear magnetic resonance microscopy*. Oxford: Clarendon Press, 1991. ISBN: 0-19-853997-5 978-0-19-853997-1.

[50] P. T. Callaghan, D. MacGowan, K. J. Packer, and F. O. Zelaya. "High-resolution q-space imaging in porous structures". In: *Journal of Magnetic Resonance (1969)* 90.1 (1990), pp. 177–182.

[51] J. S. Campbell, K. Siddiqi, B. C. Vemuri, and G. B. Pike. "A geometric flow for white matter fibre tract reconstruction". In: *Proceedings of the 2002 IEEE International Symposium on Biomedical Imaging*. Washington, DC: IEEE, 2002, pp. 505–508.

[52] E. J. Canales-Rodríguez, Y. Iturria-Medina, Y. Alemán-Gómez, and L. Melie-García. "Deconvolution in diffusion spectrum imaging". In: *NeuroImage* 50.1 (Mar. 2010), pp. 136–149. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2009.11.066.

[53] H. Carr and E. Purcell. "Effects of diffusion on free precession in nuclear magnetic resonance experiments". In: *Physical Review* 94.3 (May 1954), pp. 630–638. ISSN: 0031-899X, 1536-6065. DOI: 10.1103/PhysRev.94.630.

[54] É. J. Cartan. *Geometry of Riemannian spaces*. Brookline, Mass.: Math Sci Press, 1983. ISBN: 0-915692-34-1 978-0-915692-34-7 0-317-18190-4 978-0-317-18190-6.

[55] É. J. Cartan. "Sur les espaces de Finsler". In: *Comptes Rendus de l'Académie des Sciences, Paris* 196 (1933), pp. 582–586. ISSN: 0001-4036.

[56] E. Caruyer and R. Verma. "On facilitating the use of HARDI in population studies by creating rotation-invariant markers". In: *Medical Image Analysis* 20.1 (Feb. 2015), pp. 87–96. ISSN: 13618415. DOI: 10.1016/j.media.2014.10.009.

[57] M. Catani, I. Bodi, and F. Dell'Acqua. "Comment on "The geometric structure of the brain fiber pathways"". In: *Science* 337.6102 (Sept. 2012), pp. 1605–1605. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1223425.

[58] M. Catani, F. Dell'Acqua, F. Vergani, F. Malik, H. Hodge, P. Roy, R. Valabregue, and M. Thiebaut de Schotten. "Short frontal lobe connections of the human brain". In: *Cortex* 48.2 (Feb. 2012), pp. 273–291. ISSN: 00109452. DOI: 10.1016/j.cortex.2011.12.001.

[59] M. Catani and M. Thiebaut de Schotten. "A diffusion tensor imaging tractography atlas for virtual in vivo dissections". In: *Cortex* 44.8 (Sept. 2008), pp. 1105–1132. ISSN: 00109452. DOI: 10.1016/j.cortex.2008.05.004.

[60] E. Caverzasi, M. L. Mandelli, S. J. DeArmond, C. P. Hess, P. Vitali, N. Papinutto, A. Oehler, B. L. Miller, I. V. Lobach, S. Bastianello, M. D. Geschwind, and R. G. Henry. "White matter involvement in sporadic Creutzfeldt-Jakob disease". In: *Brain* 137.12 (Dec. 2014), pp. 3339–3354. ISSN: 0006-8950, 1460-2156. DOI: 10.1093/brain/awu298.

[61] M. Chamberland, K. Whittingstall, D. Fortin, D. Mathieu, and M. Descoteaux. "Real-time multi-peak tractography for instantaneous connectivity display". In: *Frontiers in Neuroinformatics* 8 (May 2014). ISSN: 1662-5196. DOI: 10.3389/fninf.2014.00059.

[62] L.-C. Chang, D. K. Jones, and C. Pierpaoli. "RESTORE: Robust estimation of tensors by outlier rejection". In: *Magnetic Resonance in Medicine* 53.5 (May 2005), pp. 1088–1095. ISSN: 0740-3194, 1522-2594. DOI: 10.1002/mrm.20426.

[63] Y.-H. Chen, D. Wei, G. Newstadt, M. DeGraef, J. Simmons, and A. Hero. "Parameter estimation in spherical symmetry groups". In: *IEEE Signal Processing Letters* 22.8 (Aug. 2015), pp. 1152–1155. ISSN: 1070-9908, 1558-2361. DOI: 10.1109/LSP.2014.2387206.

[64] Y. Cheng and D. G. Cory. "Multiple scattering by NMR". In: *Journal of the American Chemical Society* 121.34 (Sept. 1999), pp. 7935–7936. ISSN: 0002-7863, 1520-5126. DOI: 10.1021/ja9843324.

[65] S.-S. Chern. "Finsler geometry is just Riemannian geometry without the quadratic restriction". In: *Notices of the American Mathematical Society* 43.9 (1996), pp. 959–963. ISSN: 0002-9920.

[66] G. S. Chirikjian. *Stochastic models, information theory, and Lie groups*. Classical results and geometric methods Gregory S. Chirikjian; Vol. 1. OCLC: 552383784. Boston, Mass.: Birkhäuser, 2009. ISBN: 978-0-8176-4803-9 978-0-8176-4802-2.

[67] M.-D. Choi, T. Y. Lam, and B. Reznick. "Sums of squares of real polynomials". In: *Proceedings of Symposia in Pure mathematics*. Vol. 58. American Mathematical Society, 1995, pp. 103–126.

[68] D. Christiaens, M. Reisert, T. Dhollander, S. Sunaert, P. Suetens, and F. Maes. "Global tractography of multi-shell diffusion-weighted imaging data using a multi-tissue model". In: *NeuroImage* 123 (Dec. 2015), pp. 89–101. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2015.08.008.

[69] H.-W. Chung, M.-C. Chou, and C.-Y. Chen. "Principles and limitations of computational algorithms in clinical diffusion tensor MR tractography". In: *American Journal of Neuroradiology* (Mar. 2010). ISSN: 0195-6108, 1936-959X. DOI: 10.3174/ajnr.A2041.

[70] K. Chung, J. Wallace, S.-Y. Kim, S. Kalyanasundaram, A. S. Andalman, T. J. Davidson, J. J. Mirzabekov, K. A. Zalocusky, J. Mattis, A. K. Denisin, S. Pak, H. Bernstein, C. Ramakrishnan, L. Grosenick, V. Gradinaru, and K. Deisseroth. "Structural and molecular interrogation of intact biological systems". In: *Nature* 497.7449 (Apr. 2013), pp. 332–337. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature12107.

[71] K. A. Clark, K. H. Nuechterlein, R. F. Asarnow, L. S. Hamilton, O. R. Phillips, N. S. Hageman, R. P. Woods, J. R. Alger, A. W. Toga, and K. L. Narr. "Mean diffusivity and fractional anisotropy as indicators of disease and genetic liability to schizophrenia". In: *Journal of Psychiatric Research* 45.7 (July 2011), pp. 980–988. ISSN: 00223956. DOI: 10.1016/j.jpsychires.2011.01.006.

[72] N. Colgan, B. Siow, J. O'Callaghan, I. Harrison, J. Wells, H. Holmes, O. Ismail, S. Richardson, D. Alexander, E. Collins, E. Fisher, R. Johnson, A. Schwarz, Z. Ahmed, M. O'Neill, T. Murray, H. Zhang, and M. Lythgoe. "Application of neurite orientation dispersion and density imaging (NODDI) to a tau pathology model of Alzheimer's disease". In: *NeuroImage* 125 (Jan. 2016), pp. 739–744. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2015.10.043.

[73] D. G. Cory, A. N. Garroway, and J. B. Miller. "Applications of spin transport as a probe of local geometry". In: *Polymer Preprints*. Vol. 31. 1. Boston: American Chemical Society, Apr. 1990, pp. 149–150.

[74] P. Coupé, J. V. Manjón, M. Chamberland, M. Descoteaux, and B. Hiba. "Collaborative patch-based super-resolution for diffusion-weighted images". In: *NeuroImage* 83 (Dec. 2013), pp. 245–261. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2013.06.030.

[75] P. Coupé, P. Yger, S. Prima, P. Hellier, C. Kervrann, and C. Barillot. "An optimized blockwise nonlocal means denoising filter for 3D magnetic resonance images". In: *Medical Imaging, IEEE Transactions on* 27.4 (2008), pp. 425–441.

[76] M. Cousineau, E. Garyfallidis, M.-A. Côté, P.-M. Jodoin, and M. Descoteaux. "Tract-profiling and bundle statistics: A test-retest validation study". In: *Proceedings of the 24th Annual Meeting of the ISMRM*. Singapore, 2016, p. 3436.

[77] D. Cyranoski. "Marmosets are stars of Japan's ambitious brain project". In: *Nature* 514.7521 (Oct. 2014), pp. 151–152. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/514151a.

[78] A. Daducci, A. Dal Palu, L. Alia, and J.-P. Thiran. "COMMIT: Convex optimization modeling for micro-structure informed tractography". In: *IEEE Transactions on Medical Imaging* (2014), pp. 1–1. ISSN: 0278-0062, 1558-254X. DOI: 10.1109/TMI.2014.2352414.

[79] A. Daducci, A. Dal Palú, M. Descoteaux, and J.-P. Thiran. "Microstructure informed tractography: Pitfalls and open challenges". In: *Frontiers in Neuroscience* 10 (June 2016). ISSN: 1662-453X. DOI: 10.3389/fnins.2016.00247.

[80] A. M. Dale, B. Fischl, and M. I. Sereno. "Cortical surface-based analysis". In: *NeuroImage* 9.2 (Feb. 1999), pp. 179–194. ISSN: 10538119. DOI: 10.1006/nimg.1998.0395.

[81] S. B. Damelin and A. J. Devaney. "Local Paley–Wiener theorems for functions analytic on unit spheres". In: *Inverse Problems* 23.2 (Apr. 2007), pp. 463–474. ISSN: 0266-5611, 1361-6420. DOI: 10.1088/0266-5611/23/2/001.

[82] S. De Santis, M. Drakesmith, S. Bells, Y. Assaf, and D. K. Jones. "Why diffusion tensor MRI does well only some of the time: Variance and covariance of white matter tissue microstructure attributes in the living human brain". In: *NeuroImage* 89 (Apr. 2014), pp. 35–44. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2013.12.003.

[83] T. C. J. Dela Haije. *Reconstruction of convex polynomial diffusion MRI models using semi-definite programming*. Oral. Ghent, 2015.

[84] T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. "A new fiber orientation distribution function". In: *Proceedings of the 25th Annual Meeting of the ISMRM*. Honolulu, HI, 2017, p. 3368.

[85] T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. *Convex polynomial diffusion MRI models and semi-definite programming*. Poster. Toronto, 2015.

[86] T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. "Finslerian diffusion and the Bloch-Torrey equation". In: *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*. Ed. by I. Hotz and T. Schultz. Cham: Springer International Publishing, 2015, pp. 21–35. ISBN: 978-3-319-15089-5 978-3-319-15090-1.

[87] T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. "Reconstruction of convex polynomial diffusion MRI models using semi-definite programming". In: *Abstract Book of the 7th Annual ISMRM Benelux Chapter Meeting*. Ghent, 2015.

[88] T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. "Reconstruction of convex polynomial diffusion MRI models using semi-definite programming". In: *Proceedings of the 23rd Annual Meeting of the ISMRM*. Toronto, 2015, p. 2821.

[89] T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. "The apparent range of spin movement in diffusion MRI data". In: *Proceedings of the 24th Annual Meeting of the ISMRM*. Singapore, 2016, p. 2084.

[90] T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. *The Finslerian character of the diffusion MRI signal*.

[91] T. C. J. Dela Haije, P. Savadjiev, A. Fuster, A. Tristán-Vega, R. Schultz, R. Verma, L. M. J. Florack, and C.-F. Westin. *Structural connectivity analysis using Finsler geometry*.

[92] T. C. J. Dela Haije, C. M. W. Tax, A. Fuster, C.-F. Westin, M. A. Viergever, L. M. J. Florack, and A. Leemans. "Does sheet happen? Mapping the brain's "sheet probability index" with diffusion MRI". In: *Abstract Book of the 8th Annual ISMRM Benelux Chapter Meeting*. Eindhoven, 2016, p–040.

[93] T. C. J. Dela Haije, R. Duits, and C. M. W. Tax. "Sharpening fibers in diffusion weighted MRI via erosion". In: *Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data*. Ed. by C.-F. Westin, A. Vilanova, and B. Burgeth. Mathematics and Visualization. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 97–126. ISBN: 978-3-642-54300-5 978-3-642-54301-2.

[94] T. Dela Haije, N. Sepasian, A. Fuster, and L. Florack. "Adaptive enhancement in diffusion MRI through propagator sharpening". In: *Computational Diffusion MRI*. Ed. by A. Fuster, A. Ghosh, E. Kaden, Y. Rathi, and M. Reisert. Cham: Springer International Publishing, 2016, pp. 131–143. ISBN: 978-3-319-28586-3 978-3-319-28588-7.

[95] F. Dell'Acqua. *Discussion & debate: "Sheet or no sheet?" - Against*. Singapore, May 2016.

[96] F. Dell'Acqua, G. Rizzo, P. Scifo, R. A. Clarke, G. Scotti, and F. Fazio. "A model-based deconvolution approach to solve fiber crossing in diffusion-weighted MR imaging". In: *IEEE Transactions on Biomedical Engineering* 54.3 (Mar. 2007), pp. 462–472. ISSN: 0018-9294. DOI: 10.1109/TBME.2006.888830.

[97]   F. Dell'Acqua, P. Scifo, G. Rizzo, M. Catani, A. Simmons, G. Scotti, and F. Fazio. "A modified damped Richardson–Lucy algorithm to reduce isotropic background effects in spherical deconvolution". In: *NeuroImage* 49.2 (Jan. 2010), pp. 1446–1458. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2009.09.033.

[98]   F. Dell'Acqua, A. Simmons, S. C. Williams, and M. Catani. "Can spherical deconvolution provide more information than fiber orientations? Hindrance modulated orientational anisotropy, a true-tract specific index to characterize white matter diffusion: Hindrance Modulated Orientational Anisotropy". In: *Human Brain Mapping* 34.10 (Oct. 2013), pp. 2464–2483. ISSN: 10659471. DOI: 10.1002/hbm.22080.

[99]   M. Descoteaux, R. Deriche, T. R. Knosche, and A. Anwander. "Deterministic and probabilistic tractography based on complex fibre orientation distributions". In: *IEEE Transactions on Medical Imaging* 28.2 (Feb. 2009), pp. 269–286. ISSN: 0278-0062, 1558-254X. DOI: 10.1109/TMI.2008.2004424.

[100]   M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche. "Apparent diffusion coefficients from high angular resolution diffusion imaging: Estimation and applications". In: *Magnetic Resonance in Medicine* 56.2 (2006), pp. 395–410.

[101]   M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche. "Regularized, fast, and robust analytical q-ball imaging". In: *Magnetic Resonance in Medicine* 58.3 (2007), pp. 497–510.

[102]   M. Descoteaux, R. Deriche, D. Le Bihan, J.-F. Mangin, and C. Poupon. "Multiple q-shell diffusion propagator imaging". In: *Medical Image Analysis* 15.4 (Aug. 2011), pp. 603–621. ISSN: 13618415. DOI: 10.1016/j.media.2010.07.001.

[103]   M. Descoteaux, C. Lenglet, and R. Deriche. "Diffusion tensor sharpening improves white matter tractography". In: *Proceedings of SPIE (Medical Imaging 2007: Image Processing)* 6512 (2007), 65121J.

[104]   M. Descoteaux, J. Sidhu, E. Garyfallidis, J.-C. Houde, P. F. Neher, B. Stieltjes, and K. H. Maier-Hein. "False positive bundles in tractography". In: *Proceedings of the 24th Annual Meeting of the ISMRM*. Singapore, 2016, p. 0790.

[105]   R. S. Desikan, F. Ségonne, B. Fischl, B. T. Quinn, B. C. Dickerson, D. Blacker, R. L. Buckner, A. M. Dale, R. P. Maguire, B. T. Hyman, M. S. Albert, and R. J. Killiany. "An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest". In: *NeuroImage* 31.3 (July 2006), pp. 968–980. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2006.01.021.

[106]   L. Devroye. *Non-uniform random variate generation*. New York: Springer, 1986. ISBN: 0-387-96305-7 3-540-96305-7 0-387-96305-7.

[107]   R. A. Dineen, C. M. Bradshaw, C. S. Constantinescu, and D. P. Auer. "Extra-hippocampal subcortical limbic involvement predicts episodic recall performance in multiple sclerosis". In: *PLoS ONE* 7.10 (Oct. 2012). Ed. by C. Beaulieu, e44942. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0044942.

[108]   R. B. Dingle. *Asymptotic expansions: Their derivation and interpretation*. London, New York: Academic Press, 1973. ISBN: 0-12-216550-0.

[109]   G. Douaud, S. Jbabdi, T. E. Behrens, R. A. Menke, A. Gass, A. U. Monsch, A. Rao, B. Whitcher, G. Kindlmann, P. M. Matthews, and S. Smith. "DTI measures in crossing-fibre areas: Increased diffusion anisotropy reveals early white matter alteration in MCI and mild Alzheimer's disease". In: *NeuroImage* 55.3 (Apr. 2011), pp. 880–890. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2010.12.008.

[110] I. Drobnjak and D. C. Alexander. "Optimising time-varying gradient orientation for microstructure sensitivity in diffusion-weighted MR". In: *Journal of Magnetic Resonance* 212.2 (Oct. 2011), pp. 344–354. ISSN: 10907807. DOI: 10.1016/j.jmr.2011.07.017.

[111] R. Duits, E. J. Creusen, A. Ghosh, and T. C. J. Dela Haije. *Diffusion, convection and erosion on $\mathbb{R}^3 \rtimes S^2$ and their application to the enhancement of crossing fibers*. Technical Report 18. Eindhoven: Eindhoven University of Technology, 2011.

[112] R. Duits, T. C. J. Dela Haije, E. J. Creusen, and A. Ghosh. "Morphological and linear scale spaces for fiber enhancement in DW-MRI". In: *Journal of Mathematical Imaging and Vision* (Nov. 2012). ISSN: 0924-9907, 1573-7683. DOI: 10.1007/s10851-012-0387-2.

[113] R. Duits, T. C. J. Dela Haije, A. Ghosh, E. J. Creusen, A. Vilanova, and B. M. ter Haar Romeny. "Fiber enhancement in diffusion-weighted MRI". In: *Scale Space and Variational Methods in Computer Vision*. Springer, 2012, pp. 1–13.

[114] R. Duits and E. Franken. "Left-invariant diffusions on $\mathbb{R}^3 \rtimes S^2$ and their application to crossing-preserving smoothing on HARDI-images". In: *CASA report, TU/e* 18 (2009), pp. 23–27.

[115] T. B. Dyrby, H. Lundell, M. W. Burke, N. L. Reislev, O. B. Paulson, M. Ptito, and H. R. Siebner. "Interpolation of diffusion weighted imaging datasets". In: *NeuroImage* 103 (Dec. 2014), pp. 202–213. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2014.09.005.

[116] C. Eichner, S. F. Cauley, J. Cohen-Adad, H. E. Möller, R. Turner, K. Setsompop, and L. L. Wald. "Real diffusion-weighted MRI enabling true signal averaging and increased diffusion contrast". In: *NeuroImage* 122 (Nov. 2015), pp. 373–384. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2015.07.074.

[117] A. Einstein. "Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen". In: *Annalen der Physik* 322.8 (1905), pp. 549–560. ISSN: 00033804, 15213889. DOI: 10.1002/andp.19053220806.

[118] M. Eshagh. "Alternative expressions for gravity gradients in local north-oriented frame and tensor spherical harmonics". In: *Acta Geophysica* 58.2 (Jan. 2010). ISSN: 1895-7455, 1895-6572. DOI: 10.2478/s11600-009-0048-z.

[119] U. Ferizi, T. Schneider, E. Panagiotaki, G. Nedjati-Gilani, H. Zhang, C. A. M. Wheeler-Kingshott, and D. C. Alexander. "A ranking of diffusion MRI compartment models with in vivo human brain data: Diffusion MRI Compartment Models". In: *Magnetic Resonance in Medicine* 72.6 (Dec. 2014), pp. 1785–1792. ISSN: 07403194. DOI: 10.1002/mrm.25080.

[120] A. Fick. "Über diffusion". In: *Annalen der Physik* 170.1 (1855), pp. 59–86.

[121] R. H. J. Fick, R. Duits, E. J. Creusen, T. C. J. Dela Haije, P. P. W. Ossenblok, B. M. ter Haar Romeny, and A. Vilanova. "Contextual enhancements on DW-MRI". In: *HARDI reconstruction challenge 2013 proceedings*. San Francisco, Apr. 2013, p. 24.

[122] R. H. Fick, D. Wassermann, E. Caruyer, and R. Deriche. "MAPL: Tissue microstructure estimation using Laplacian-regularized MAP-MRI and its application to HCP data". In: *NeuroImage* (Apr. 2016). ISSN: 10538119. DOI: 10.1016/j.neuroimage.2016.03.046.

[123] E. Fieremans, L. M. Burcaw, H.-H. Lee, G. Lemberskiy, J. Veraart, and D. S. Novikov. "In vivo observation and biophysical interpretation of time-dependent diffusion in human white matter". In: *NeuroImage* 129 (Apr. 2016), pp. 414–427. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2016.01.018.

[124] E. Fieremans, J. H. Jensen, and J. A. Helpern. "White matter characterization with diffusional kurtosis imaging". In: *NeuroImage* 58.1 (Sept. 2011), pp. 177–188. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2011.06.006.

[125] P. Finsler. "Über Kurven und Flächen in allgemeinen Räumen". PhD thesis. Zürich: Göttingen, 1918.

[126] B. Fischl. "Automatically parcellating the human cerebral cortex". In: *Cerebral Cortex* 14.1 (Jan. 2004), pp. 11–22. ISSN: 1460-2199. DOI: 10.1093/cercor/bhg087.

[127] B. Fischl, M. I. Sereno, and A. M. Dale. "Cortical surface-based analysis". In: *NeuroImage* 9.2 (Feb. 1999), pp. 195–207. ISSN: 10538119. DOI: 10.1006/nimg.1998.0396.

[128] P. T. Fletcher and S. Joshi. "Riemannian geometry for the statistical analysis of diffusion tensor data". In: *Signal Processing* 87.2 (2007), pp. 250–262.

[129] L. M. J. Florack and A. Fuster. "Riemann-Finsler geometry for diffusion weighted magnetic resonance imaging". In: *Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data*. Ed. by C.-F. Westin, A. Vilanova, and B. Burgeth. Mathematics and Visualization XV. Springer, 2014, pp. 189–208.

[130] L. Florack and L. Astola. "A multi-resolution framework for diffusion tensor images". In: *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW'08. IEEE Computer Society Conference on*. 2008, pp. 1–7.

[131] L. Florack, E. Balmashnova, L. Astola, and E. Brunenberg. "A new tensorial framework for single-shell high angular resolution diffusion imaging". In: *Journal of Mathematical Imaging and Vision* 38.3 (2010), pp. 171–181.

[132] L. Florack, T. Dela Haije, and A. Fuster. *Cartan scalars in Finsler-DTI for higher order local brain tissue characterization*.

[133] L. Florack, T. Dela Haije, and A. Fuster. "Direction-controlled DTI interpolation". In: *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*. Ed. by I. Hotz and T. Schultz. Cham: Springer International Publishing, 2015, pp. 149–162. ISBN: 978-3-319-15089-5 978-3-319-15090-1.

[134] L. R. Frank. "Anisotropy in high angular resolution diffusion-weighted MRI". In: *Magnetic Resonance in Medicine* 45.6 (June 2001), pp. 935–939. ISSN: 0740-3194, 1522-2594. DOI: 10.1002/mrm.1125.

[135] L. R. Frank. "Characterization of anisotropy in high angular resolution diffusion-weighted MRI". In: *Magnetic Resonance in Medicine* 47.6 (June 2002), pp. 1083–1099. ISSN: 0740-3194, 1522-2594. DOI: 10.1002/mrm.10156.

[136] M. Froeling, C. M. Tax, S. B. Vos, P. R. Luijten, and A. Leemans. ""MASSIVE" brain dataset: Multiple acquisitions for standardization of structural imaging validation and evaluation". In: *Magnetic Resonance in Medicine* (May 2016). ISSN: 07403194. DOI: 10.1002/mrm.26259.

[137] A. Fuster, L. Astola, and L. Florack. "A Riemannian scalar measure for diffusion tensor images". In: *Computer Analysis of Images and Patterns*. 2009, pp. 419–426.

[138] A. Fuster, T. C. J. Dela Haije, and L. M. J. Florack. "On the Riemannian rationale for diffusion tensor imaging". In: *Proceedings of the International Biomedical and Astronomical Signal Processing (BASP) Frontiers workshop*. Villars-sur-Ollon, Jan. 2013, p. 62.

[139] A. Fuster, A. Tristán-Vega, T. C. J. Dela Haije, C.-F. Westin, and L. M. J. Florack. *On the Riemannian rationale for DTI*. Poster. Villars-sur-Ollon, Jan. 2013.

[140] A. Fuster, A. Tristán-Vega, T. C. J. Dela Haije, C.-F. Westin, and L. M. J. Florack. "A novel Riemannian metric for geodesic tractography in DTI". In: *CDMRI 2013*. Nagoya, Sept. 2013, pp. 47–54.

[141] A. Fuster, T. Dela Haije, A. Tristán-Vega, B. Plantinga, C.-F. Westin, and L. Florack. "Adjugate diffusion tensors for geodesic tractography in white matter". In: *Journal of Mathematical Imaging and Vision* 54.1 (2016), pp. 1–14. ISSN: 0924-9907, 1573-7683. DOI: 10.1007/s10851-015-0586-8.

[142] A. Fuster, A. Tristán-Vega, T. Dela Haije, C.-F. Westin, and L. Florack. "A novel Riemannian metric for geodesic tractography in DTI". In: *Computational Diffusion MRI and Brain Connectivity*. Springer, 2014, pp. 97–104.

[143] J. K. Gahm and D. B. Ennis. "Dyadic tensor-based interpolation of tensor orientation: Application to cardiac DT-MRI". In: *Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges*. Ed. by D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, O. Camara, T. Mansi, M. Pop, K. Rhode, M. Sermesant, and A. Young. Vol. 8330. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 135–142. ISBN: 978-3-642-54267-1 978-3-642-54268-8.

[144] J. K. Gahm, N. Wisniewski, G. Kindlmann, G. L. Kung, W. S. Klug, A. Garfinkel, and D. B. Ennis. "Linear invariant tensor interpolation applied to cardiac diffusion tensor MRI". In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2012* 15.02 (2012), pp. 494–501.

[145] R. J. Gardner, M. Kiderlen, and P. Milanfar. "Convergence of algorithms for reconstructing convex bodies and directional measures". In: *The Annals of Statistics* 34.3 (June 2006), pp. 1331–1374. ISSN: 0090-5364. DOI: 10.1214/009053606000000335.

[146] E. Garyfallidis, S. St-Jean, M. Paquette, P. Coupé, and M. Descoteaux. "Constrained spherical deconvolution on signal and ODF values". In: *ISBI'13, HARDI Reconstruction Challenge* (2013).

[147] E. Garyfallidis, O. Ocegueda, D. Wassermann, and M. Descoteaux. "Robust and efficient linear registration of white-matter fascicles in the space of streamlines". In: *NeuroImage* 117 (Aug. 2015), pp. 124–140. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2015.05.016.

[148] A. Ghosh and R. Deriche. "Extracting geometrical features & peak fractional anisotropy from the ODF for white matter characterization". In: *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*. IEEE, 2011, pp. 266–271.

[149] A. Ghosh, R. Deriche, and M. Moakher. "Ternary quartic approach for positive 4th order diffusion tensors revisited". In: *Biomedical Imaging: From Nano to Macro, 2009. ISBI'09. IEEE International Symposium on*. 2009, pp. 618–621.

[150] A. Ghosh, T. Milne, and R. Deriche. "Constrained diffusion kurtosis imaging using ternary quartics & MLE". In: *Magnetic Resonance in Medicine* 71.4 (2014), pp. 1581–1591.

[151] G. Girard, R. Fick, M. Descoteaux, R. Deriche, and D. Wassermann. "AxTract: Microstructure-driven tractography based on the ensemble average propagator". In: *Information Processing in Medical Imaging*. 2015, pp. 675–686.

[152] M. F. Glasser, S. N. Sotiropoulos, J. A. Wilson, T. S. Coalson, B. Fischl, J. L. Andersson, J. Xu, S. Jbabdi, M. Webster, J. R. Polimeni, D. C. Van Essen, and M. Jenkinson. "The minimal preprocessing pipelines for the Human Connectome Project". In: *NeuroImage* 80 (Oct. 2013), pp. 105–124. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2013.04.127.

[153] A. Gorban. "Detailed balance in micro- and macrokinetics and micro-distinguishability of macro-processes". In: *Results in Physics* 4 (2014), pp. 142–147. ISSN: 22113797. DOI: 10.1016/j.rinp.2014.09.002.

[154] J. S. Gounarides, A. Chen, and M. J. Shapiro. "Nuclear magnetic resonance chromatography: Applications of pulse field gradient diffusion NMR to mixture analysis and ligand–receptor interactions". In: *Journal of Chromatography B: Biomedical Sciences and Applications* 725.1 (Apr. 1999), pp. 79–90. ISSN: 03784347. DOI: 10.1016/S0378-4347(98)00512-X.

[155] M. de Groot, M. W. Vernooij, S. Klein, M. A. Ikram, F. M. Vos, S. M. Smith, W. J. Niessen, and J. L. Andersson. "Improving alignment in tract-based spatial statistics: Evaluation and optimization of image registration". In: *NeuroImage* 76 (Aug. 2013), pp. 400–411. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2013.03.015.

[156] H. Gudbjartsson and S. Patz. "The Rician distribution of noisy MRI data". In: *Magnetic Resonance in Medicine* 34.6 (1995), pp. 910–914.

[157] E. M. Haacke. *Magnetic resonance imaging: physical principles and sequence design*. New York: Wiley, 1999. ISBN: 0-471-35128-8 978-0-471-35128-3.

[158] P. Hagmann, L. Jonasson, P. Maeder, J.-P. Thiran, V. J. Wedeen, and R. Meuli. "Understanding diffusion MR imaging techniques: From scalar diffusion-weighted imaging to diffusion tensor imaging and beyond". In: *Radiographics* 26.Supplement 1 (Oct. 2006), S205–S223. ISSN: 0271-5333, 1527-1323. DOI: 10.1148/rg.26si065510.

[159] E. L. Hahn. "Spin echoes". In: *Physical Review* 80.4 (1950), p. 580.

[160] M. G. Hall and T. R. Barrick. "From diffusion-weighted MRI to anomalous diffusion imaging". In: *Magnetic Resonance in Medicine* 59.3 (Mar. 2008), pp. 447–455. ISSN: 07403194, 15222594. DOI: 10.1002/mrm.21453.

[161] A. Hanyga and M. Seredyńska. "Anisotropy in high-resolution diffusion-weighted MRI and anomalous diffusion". In: *Journal of Magnetic Resonance* 220 (July 2012), pp. 85–93. ISSN: 10907807. DOI: 10.1016/j.jmr.2012.05.001.

[162] X. Hao, R. Whitaker, and P. Fletcher. "Adaptive Riemannian metrics for improved geodesic tracking of white matter". In: *Information Processing in Medical Imaging*. 2011, pp. 13–24.

[163] X. Hao, K. Zygmunt, R. T. Whitaker, and P. T. Fletcher. "Improved segmentation of white matter tracts with adaptive Riemannian metrics". In: *Medical Image Analysis* 18.1 (Jan. 2014), pp. 161–175. ISSN: 13618415. DOI: 10.1016/j.media.2013.10.007.

[164] S. Havlin and D. Ben-Avraham. "Diffusion in disordered media". In: *Advances in Physics* 51.1 (Jan. 2002), pp. 187–292. ISSN: 0001-8732, 1460-6976. DOI: 10.1080/00018730110116353.

[165] L. Hörmander. *Distribution theory and Fourier analysis*. 2nd ed. The analysis of linear partial differential operators 1. Berlin: Springer, 1990. ISBN: 3-540-52345-6 3-540-52345-6 0-387-52345-6 3-540-52343-X 0-387-52343-X.

[166] M. A. Horsfield and D. K. Jones. "Applications of diffusion-weighted and diffusion tensor MRI to white matter diseases - a review". In: *NMR in Biomedicine* 15.7-8 (Nov. 2002), pp. 570–577. ISSN: 0952-3480, 1099-1492. DOI: 10.1002/nbm.787.

[167] A. P. Hosseinbor, M. K. Chung, Y.-C. Wu, and A. L. Alexander. "Bessel Fourier Orientation Reconstruction (BFOR): An analytical diffusion propagator reconstruction for hybrid diffusion imaging and computation of q-space indices". In: *NeuroImage* 64 (Jan. 2013), pp. 650–670. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2012.08.072.

[168] N. Ikeda and S. Watanabe. *Stochastic differential equations and diffusion processes*. Amsterdam; New York; Tokyo; New York, NY: North-Holland Pub. Co. ; Kodansha ; Sole distributors for the U.S.A. and Canada, Elsevier North-Holland, 1981. ISBN: 0-444-86172-6 978-0-444-86172-6.

[169] M. Jackowski, C. Y. Kao, M. Qiu, R. T. Constable, and L. H. Staib. "White matter tractography by anisotropic wavefront evolution and diffusion tensor imaging". In: *Medical Image Analysis* 9.5 (Oct. 2005), pp. 427–440. ISSN: 13618415. DOI: 10.1016/j.media.2005.05.008.

[170] M. H. J. Janssen, T. C. J. Dela Haije, F. C. Martin, E. J. Bekkers, J. Olivan Bescos, and R. Duits. *The Hessian of axially symmetric functions on* SE(3) *and application in 3D image analysis*. 2017.

[171] S. Jbabdi, P. Bellec, R. Toro, J. Daunizeau, M. Pélégrini-Issac, and H. Benali. "Accurate anisotropic fast marching for diffusion-based geodesic tractography". In: *International Journal of Biomedical Imaging* 2008 (2008), pp. 1–12. ISSN: 1687-4188, 1687-4196. DOI: 10.1155/2008/320195.

[172] S. Jbabdi and H. Johansen-Berg. "Tractography: Where do we go from here?" In: *Brain Connectivity* 1.3 (2011), pp. 169–183.

[173] S. Jbabdi, T. E. Behrens, and S. M. Smith. "Crossing fibres in tract-based spatial statistics". In: *NeuroImage* 49.1 (Jan. 2010), pp. 249–256. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2009.08.039.

[174] S. St-Jean, P. Coupé, and M. Descoteaux. "Non local spatial and angular matching: Enabling higher spatial resolution diffusion MRI datasets through adaptive denoising". In: *Medical Image Analysis* 32 (Aug. 2016), pp. 115–130. ISSN: 13618415. DOI: 10.1016/j.media.2016.02.010.

[175] I. O. Jelescu, J. Veraart, E. Fieremans, and D. S. Novikov. "Degeneracy in model parameter estimation for multi-compartmental diffusion in neuronal tissue: Degeneracy in Model Parameter Estimation of Diffusion in Neural Tissue". In: *NMR in Biomedicine* 29.1 (Jan. 2016), pp. 33–47. ISSN: 09523480. DOI: 10.1002/nbm.3450.

[176] M. Jenkinson and S. Smith. "A global optimisation method for robust affine registration of brain images". In: *Medical image analysis* 5.2 (2001), pp. 143–156.

[177] J. H. Jensen and J. A. Helpern. "Effect of gradient pulse duration on MRI estimation of the diffusional kurtosis for a two-compartment exchange model". In: *Journal of Magnetic Resonance* 210.2 (June 2011), pp. 233–237. ISSN: 10907807. DOI: 10.1016/j.jmr.2011.03.012.

[178] J. H. Jensen, J. A. Helpern, A. Ramani, H. Lu, and K. Kaczynski. "Diffusional kurtosis imaging: The quantification of non-gaussian water diffusion by means of magnetic resonance imaging". In: *Magnetic Resonance in Medicine* 53.6 (June 2005), pp. 1432–1440. ISSN: 0740-3194, 1522-2594. DOI: 10.1002/mrm.20508.

[179] J. H. Jensen, G. Russell Glenn, and J. A. Helpern. "Fiber ball imaging". In: *NeuroImage* 124 (Jan. 2016), pp. 824–833. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2015.09.049.

[180] B. Jeurissen, A. Leemans, D. K. Jones, J.-D. Tournier, and J. Sijbers. "Probabilistic fiber tracking using the residual bootstrap with constrained spherical deconvolution". In: *Human Brain Mapping* 32.3 (Mar. 2011), pp. 461–479. ISSN: 10659471. DOI: 10.1002/hbm.21032.

[181] B. Jeurissen, A. Leemans, J.-D. Tournier, D. K. Jones, and J. Sijbers. "Investigating the prevalence of complex fiber configurations in white matter tissue with diffusion magnetic resonance imaging: Prevalence of multifiber voxels in WM". In: *Human Brain Mapping* 34.11 (Nov. 2013), pp. 2747–2766. ISSN: 10659471. DOI: `10.1002/hbm.22099`.

[182] B. Jeurissen, J.-D. Tournier, T. Dhollander, A. Connelly, and J. Sijbers. "Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data". In: *NeuroImage* 103 (Dec. 2014), pp. 411–426. ISSN: 10538119. DOI: `10.1016/j.neuroimage.2014.07.061`.

[183] H. Johansen-Berg and T. E. J. Behrens. *Diffusion MRI: From quantitative measurement to in-vivo neuroanatomy*. Amsterdam: Elsevier Science, 2014. ISBN: 978-0-12-405509-4.

[184] D. K. Jones, M. A. Horsfield, and A. Simmons. "Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging". In: *Magnetic resonance in medicine* 42.3 (Sept. 1999). PMID: 10467296, pp. 515–525. ISSN: 0740-3194.

[185] D. K. Jones. *Diffusion MRI: Theory, methods, and application*. Oxford; New York: Oxford University Press, 2010. ISBN: 978-0-19-536977-9 0-19-536977-7.

[186] L. A. Jorgenson, W. T. Newsome, D. J. Anderson, C. I. Bargmann, E. N. Brown, K. Deisseroth, J. P. Donoghue, K. L. Hudson, G. S. F. Ling, P. R. MacLeish, E. Marder, R. A. Normann, J. R. Sanes, M. J. Schnitzer, T. J. Sejnowski, D. W. Tank, R. Y. Tsien, K. Ugurbil, and J. C. Wingfield. "The BRAIN Initiative: developing technology to catalyse neuroscience discovery". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 370.1668 (Mar. 2015), pp. 20140164–20140164. ISSN: 0962-8436, 1471-2970. DOI: `10.1098/rstb.2014.0164`.

[187] J. Jost. *Riemannian geometry and geometric analysis*. Berlin: Springer, 2005. ISBN: 3-540-25907-4 978-3-540-25907-7 3-540-42627-2 978-3-540-42627-1.

[188] A. Jurcoane, F. Keil, A. Szelenyi, W. Pfeilschifter, O. C. Singer, and E. Hattingen. "Directional diffusion of corticospinal tract supports therapy decisions in idiopathic normal-pressure hydrocephalus". In: *Neuroradiology* 56.1 (Oct. 2013), pp. 5–13. ISSN: 0028-3940, 1432-1920. DOI: `10.1007/s00234-013-1289-8`.

[189] B. Jüttler. "Surface fitting using convex tensor-product splines". In: *Journal of Computational and Applied Mathematics* 84.1 (Oct. 1997), pp. 23–44. ISSN: 03770427. DOI: `10.1016/S0377-0427(97)00100-3`.

[190] K.-I. Kanatani. "Stereological determination of structural anisotropy". In: *International Journal of Engineering Science* 22.5 (Jan. 1984), pp. 531–546. ISSN: 00207225. DOI: `10.1016/0020-7225(84)90055-7`.

[191] C. Y. Kao, S. Osher, and J. Qian. "Lax–Friedrichs sweeping scheme for static Hamilton–Jacobi equations". In: *Journal of Computational Physics* 196.1 (2004), pp. 367–391.

[192] C.-Y. Kao, S. Osher, and Y.-H. Tsai. "Fast sweeping methods for static Hamilton–Jacobi equations". In: *SIAM journal on numerical analysis* 42.6 (2005), pp. 2612–2632.

[193] A. Kawaguchi. "Theory of connections in a Kawaguchi space of higher order". In: *Proceedings of the Imperial Academy* 13.7 (1937), pp. 237–240. ISSN: 0369-9846. DOI: `10.3792/pia/1195579892`.

[194] J. B. Keller. "Diffusion at finite speed and random walks". In: *Proceedings of the National Academy of Sciences* 101.5 (Feb. 2004), pp. 1120–1122. ISSN: 0027-8424, 1091-6490. DOI: `10.1073/pnas.0307052101`.

[195] I. Kezele, M. Descoteaux, C. Poupon, F. Poupon, and J.-F. Mangin. "Spherical wavelet transform for ODF sharpening". In: *Medical Image Analysis* 14.3 (June 2010), pp. 332–342. ISSN: 13618415. DOI: `10.1016/j.media.2010.01.002`.

[196] M. Kiderlen and A. Pfrang. "Algorithms to estimate the rose of directions of a spatial fibre system". In: *Journal of microscopy* 219.2 (2005), pp. 50–60.

[197] G. Kindlmann. "Superquadric tensor glyphs". In: *Proceedings of the Sixth Join Eurographics - IEEE TCVG Conference on Visualization*. Konstanz, 2004, pp. 147–154.

[198] G. Kindlmann, R. S. J. Estépar, M. Niethammer, S. Haker, and C.-F. Westin. "Geodesic-loxodromes for diffusion tensor interpolation and difference measurement". In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2007* 10.01 (2007). PMID: 18051037, pp. 1–9.

[199] G. Kindlmann, X. Tricoche, and C.-F. Westin. "Delineating white matter structure in diffusion tensor MRI with anisotropy creases". In: *Medical Image Analysis* 11.5 (Oct. 2007). PMID: 17804278 PMCID: PMC2367700, pp. 492–502. ISSN: 1361-8415. DOI: `10.1016/j.media.2007.07.005`.

[200] V. G. Kiselev. "The cumulant expansion: An overarching mathematical framework for understanding diffusion NMR". In: *Diffusion MRI*. Ed. by D. K. Jones. Oxford University Press, Nov. 2010, pp. 152–168. ISBN: 978-0-19-536977-9.

[201] S. Klein, M. Staring, K. Murphy, M. Viergever, and J. Pluim. "elastix: A toolbox for intensity-based medical image registration". In: *IEEE Transactions on Medical Imaging* 29.1 (Jan. 2010), pp. 196–205. ISSN: 0278-0062, 1558-254X. DOI: `10.1109/TMI.2009.2035616`.

[202] C. G. Koay, J. D. Carew, A. L. Alexander, P. J. Basser, and M. E. Meyerand. "Investigation of anomalous estimates of tensor-derived quantities in diffusion tensor imaging". In: *Magnetic Resonance in Medicine* 55.4 (Apr. 2006), pp. 930–936. ISSN: 0740-3194, 1522-2594. DOI: `10.1002/mrm.20832`.

[203] C. G. Koay, L.-C. Chang, J. D. Carew, C. Pierpaoli, and P. J. Basser. "A unifying theoretical and algorithmic framework for least squares methods of estimation in diffusion tensor imaging". In: *Journal of Magnetic Resonance* 182.1 (Sept. 2006), pp. 115–125. ISSN: 10907807. DOI: `10.1016/j.jmr.2006.06.020`.

[204] B. Kolman and D. R. Hill. *Elementary linear algebra with applications*. Upper Saddle River, N.J.: Pearson Prentice Hall, 2008. ISBN: 0-13-229654-3 978-0-13-229654-0 978-0-13-135063-2 0-13-135063-3.

[205] S. Krein. *Linear differential equation in a Banach space*. 2012.

[206] C. D. Kroenke, J. J. Ackerman, and D. A. Yablonskiy. "On the nature of the NAA diffusion attenuated MR signal in the central nervous system". In: *Magnetic Resonance in Medicine* 52.5 (Nov. 2004), pp. 1052–1059. ISSN: 0740-3194, 1522-2594. DOI: `10.1002/mrm.20260`.

[207] S. Lang. *Differential and Riemannian manifolds*. New York, NY: Springer New York, 1995. ISBN: 978-1-4612-4182-9.

[208] M. C. de Lara. "Geometric and symmetry properties of a nondegenerate diffusion process". In: *The Annals of Probability* 23.4 (Oct. 1995), pp. 1557–1604. ISSN: 0091-1798. DOI: `10.1214/aop/1176987794`.

[209] M. Lazar and A. L. Alexander. "Bootstrap white matter tractography (BOOT-TRAC)". In: *NeuroImage* 24.2 (Jan. 2005), pp. 524–532. ISSN: 10538119. DOI: `10.1016/j.neuroimage.2004.08.050`.

[210]   M. Lazar, J. H. Jensen, L. Xuan, and J. A. Helpern. "Estimation of the orientation distribution function from diffusional kurtosis imaging". In: *Magnetic Resonance in Medicine* 60.4 (Oct. 2008), pp. 774–781. ISSN: 07403194, 15222594. DOI: 10.1002/mrm.21725.

[211]   M. Lazar, D. M. Weinstein, J. S. Tsuruda, K. M. Hasan, K. Arfanakis, M. E. Meyerand, B. Badie, H. A. Rowley, V. Haughton, A. Field, and A. L. Alexander. "White matter tractography using diffusion tensor deflection". In: *Human Brain Mapping* 18.4 (Apr. 2003), pp. 306–321. ISSN: 1065-9471, 1097-0193. DOI: 10.1002/hbm.10102.

[212]   D. Le Bihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis, and M. Laval-Jeantet. "MR imaging of intravoxel incoherent motions: Application to diffusion and perfusion in neurologic disorders". In: *Radiology* 161.2 (1986), p. 401.

[213]   D. Le Bihan, J. F. Mangin, C. Poupon, C. A. Clark, S. Pappata, N. Molko, and H. Chabriat. "Diffusion tensor imaging: Concepts and applications". In: *Journal of magnetic resonance imaging* 13.4 (2001), pp. 534–546.

[214]   D. Le Bihan. "Looking into the functional architecture of the brain with diffusion MRI". In: *International Congress Series* 1290 (June 2006), pp. 1–24. ISSN: 05315131. DOI: 10.1016/j.ics.2006.04.006.

[215]   A. Leemans, B. Jeurissen, J. Sijbers, and D. K. Jones. "ExploreDTI: A graphical toolbox for processing, analyzing, and visualizing diffusion MR data". In: *17th Annual Meeting of Intl Soc Mag Reson Med*. Vol. 209. 2009, p. 3537.

[216]   A. Leemans, J. Sijbers, S. De Backer, E. Vandervliet, and P. Parizel. "Multiscale white matter fiber tract coregistration: A new feature-based approach to align diffusion tensor data". In: *Magnetic Resonance in Medicine* 55.6 (June 2006), pp. 1414–1423. ISSN: 0740-3194, 1522-2594. DOI: 10.1002/mrm.20898.

[217]   A. Leemans and D. K. Jones. "The B-matrix must be rotated when correcting for subject motion in DTI data". In: *Magnetic Resonance in Medicine* 61.6 (2009), pp. 1336–1349.

[218]   C. Lenglet, R. Deriche, and O. Faugeras. "Inferring white matter geometry from diffusion tensor MRI: Application to connectivity mapping". In: *Computer Vision - ECCV 2004*. Ed. by T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, T. Pajdla, and J. Matas. Vol. 3024. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 127–140. ISBN: 978-3-540-21981-1 978-3-540-24673-2.

[219]   C. Lenglet, E. Prados, J.-P. Pons, R. Deriche, and O. Faugeras. "Brain connectivity mapping using Riemannian geometry, control theory, and PDEs". In: *SIAM Journal on Imaging Sciences* 2.2 (Jan. 2009), pp. 285–322. ISSN: 1936-4954. DOI: 10.1137/070710986.

[220]   J. C. Lewis. "Elementary statistical models for vector collision-sequence interference effects with Poisson-distributed collision times". In: *International Journal of Spectroscopy* 2010 (2010), pp. 1–5. ISSN: 1687-9449, 1687-9457. DOI: 10.1155/2010/561697.

[221]   J. D. Lewis, R. J. Theilmann, J. Townsend, and A. C. Evans. "Network efficiency in autism spectrum disorder and its relation to brain overgrowth". In: *Frontiers in Human Neuroscience* 7 (2013). ISSN: 1662-5161. DOI: 10.3389/fnhum.2013.00845.

[222]   C. Liu, R. Bammer, B. Acar, and M. E. Moseley. "Characterizing non-gaussian diffusion by using generalized diffusion tensors". In: *Magnetic Resonance in Medicine* 51.5 (May 2004), pp. 924–937. ISSN: 0740-3194, 1522-2594. DOI: 10.1002/mrm.20071.

[223] C. Liu, S. C. Mang, and M. E. Moseley. "In vivo generalized diffusion tensor imaging (GDTI) using higher-order tensors (HOT)". In: *Magnetic Resonance in Medicine* (2009), NA–NA. ISSN: 07403194, 15222594. DOI: 10.1002/mrm.22192.

[224] A. Magnani, S. Lall, and S. Boyd. "Tractable fitting with convex polynomials via sum-of-squares". In: *Decision and Control, 2005 and 2005 European Control Conference. CDC-ECC'05. 44th IEEE Conference on*. IEEE, 2005, pp. 1672–1677.

[225] K. H. Maier-Hein, P. Neher, J. Houde, M. Côté, E. Garyfallidi, J. Zhong, M. Chamberland, F. Yeh, Y. Lin, Q. Ji, W. E. Reddick, J. O. Glass, D. Q. Chen, Y. Feng, C. Gao, Y. Wu, J. Ma, H. Renjie, Q. Li, C. Westin, S. Deslauriers-Gauthier, J. O. O. González, M. Paquette, S. St-Jean, G. Girard, F. Rheault, J. Sidhu, C. M. Tax, F. Guo, H. Y. Mesri, S. Dávid, M. Froeling, A. M. Heemskerk, A. Leemans, A. Boré, B. Pinsard, C. Bedetti, M. Desrosiers, S. Brambati, J. Doyon, A. Sarica, R. Vasta, A. Cerasa, A. Quattrone, J. Yeatman, A. R. Khan, W. Hodges, S. Alexander, D. Romascano, M. Barakovic, A. Auría, O. Esteban, A. Lemkaddem, J. Thiran, H. E. Cetingul, B. L. Odry, B. Mailhe, M. S. Nadar, F. Pizzagalli, G. Prasad, J. E. Villalon-Reina, J. Galvis, P. M. Thompson, F. D. S. Requejo, P. L. Laguna, L. M. Lacerda, R. Barrett, F. Dell'Acqua, M. Catani, L. Petit, E. Caruyer, A. Daducci, T. B. Dyrby, T. Holland-Letz, C. C. Hilgetag, B. Stieltjes, and M. Descoteaux. "Tractography-based connectomes are dominated by false-positive connections". In: *bioRXiv preprint 084137* (2016).

[226] K. V. Mardia and P. E. Jupp. *Directional statistics*. Wiley series in probability and statistics. Chichester ; New York: J. Wiley, 2000. ISBN: 978-0-471-95333-3.

[227] J. Melonakos, E. Pichon, S. Angenent, and A. Tannenbaum. "Finsler active contours". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 30.3 (Mar. 2008), pp. 412–423. ISSN: 0162-8828. DOI: 10.1109/TPAMI.2007.70713.

[228] J. Melonakos, V. Mohan, M. Niethammer, K. Smith, M. Kubicki, and A. Tannenbaum. "Finsler tractography for white matter connectivity analysis of the cingulum bundle". In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2007*. Springer, 2007, pp. 36–43.

[229] N. S. Metwalli, M. Benatar, G. Nair, S. Usher, X. Hu, and J. D. Carew. "Utility of axial and radial diffusivity from diffusion tensor MRI as markers of neurodegeneration in amyotrophic lateral sclerosis". In: *Brain Research* 1348 (Aug. 2010), pp. 156–164. ISSN: 00068993. DOI: 10.1016/j.brainres.2010.05.067.

[230] M. M. Mielke, O. C. Okonkwo, K. Oishi, S. Mori, S. Tighe, M. I. Miller, C. Ceritoglu, T. Brown, M. Albert, and C. G. Lyketsos. "Fornix integrity and hippocampal volume predict memory decline and progression to Alzheimer's disease". In: *Alzheimer's & Dementia* 8.2 (Mar. 2012), pp. 105–113. ISSN: 15525260. DOI: 10.1016/j.jalz.2011.05.2416.

[231] J.-M. Mirebeau. "Efficient fast marching with Finsler metrics". In: *Numerische Mathematik* 126.3 (2014), pp. 515–557.

[232] H. Mirzaalian, L. Ning, P. Savadjiev, O. Pasternak, S. Bouix, O. Michailovich, G. Grant, C. Marx, R. Morey, L. Flashman, M. George, T. McAllister, N. Andaluz, L. Shutter, R. Coimbra, R. Zafonte, M. Coleman, M. Kubicki, C. Westin, M. Stein, M. Shenton, and Y. Rathi. "Inter-site and inter-scanner diffusion MRI data harmonization". In: *NeuroImage* 135 (July 2016), pp. 311–323. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2016.04.041.

[233] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. San Francisco: W.H. Freeman, 1973. ISBN: 0-7167-0334-3 978-0-7167-0334-1 0-7167-0344-0 978-0-7167-0344-0.

[234] P. P. Mitra and B. I. Halperin. "Effects of finite gradient-pulse widths in pulsed-field-gradient diffusion measurements". In: *Journal of Magnetic Resonance, Series A* 113.1 (Mar. 1995), pp. 94–101. ISSN: 10641858. DOI: 10.1006/jmra.1995.1060.

[235] P. P. Mitra. "Multiple wave-vector extensions of the NMR pulsed-field-gradient spin-echo diffusion measurement". In: *Physical Review B* 51.21 (1995), p. 15074.

[236] P. P. Mitra, P. N. Sen, L. M. Schwartz, and P. Le Doussal. "Diffusion propagator as a probe of the structure of porous media". In: *Physical Review Letters* 68.24 (June 1992), pp. 3555–3558. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.68.3555.

[237] M. Moakher and P. G. Batchelor. "Symmetric positive-definite matrices: From geometry to applications and visualization". In: *Visualization and Processing of Tensor Fields*. Ed. by J. Weickert and H. Hagen. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 285–298. ISBN: 978-3-540-25032-6 978-3-540-31272-7.

[238] M. Moakher and M. Zéraï. "The Riemannian geometry of the space of positive-definite matrices and its application to the regularization of positive-definite matrix-valued data". In: *Journal of Mathematical Imaging and Vision* 40.2 (June 2011), pp. 171–187. ISSN: 0924-9907, 1573-7683. DOI: 10.1007/s10851-010-0255-x.

[239] A. S. Monin. "A statistical interpretation of the scattering of microscopic particles". In: *Theory of Probability & Its Applications* 1.3 (Jan. 1956), pp. 298–311. ISSN: 0040-585X, 1095-7219. DOI: 10.1137/1101024.

[240] J. L. Morgan and J. W. Lichtman. "Why not connectomics?" In: *Nature Methods* 10.6 (May 2013), pp. 494–500. ISSN: 1548-7091, 1548-7105. DOI: 10.1038/nmeth.2480.

[241] S. Mori and B. J. Crain. *MRI atlas of human white matter*. Amsterdam; Boston: Elsevier, 2005. ISBN: 978-0-444-51741-8 0-444-51741-3 0-08-045616-2 978-0-08-045616-4.

[242] S. Mori and P. C. M. van Zijl. "Fiber tracking: Principles and strategies - A technical review". In: *NMR in Biomedicine* 15.7-8 (Nov. 2002), pp. 468–480. ISSN: 0952-3480, 1099-1492. DOI: 10.1002/nbm.781.

[243] P. Mosayebi, D. Cobzas, A. Murtha, and M. Jagersand. "Tumor invasion margin on the Riemannian space of brain fibers". In: *Medical Image Analysis* 16.2 (Feb. 2012). PMID: 22154876, pp. 361–373. ISSN: 1361-8423. DOI: 10.1016/j.media.2011.10.001.

[244] M. Mussel, L. Inzelberg, and U. Nevo. "Insignificance of active flow for neural diffusion weighted imaging: A negative result". In: *Magnetic Resonance in Medicine* (Aug. 2016). ISSN: 07403194. DOI: 10.1002/mrm.26375.

[245] P. F. Neher, M. Descoteaux, J.-C. Houde, B. Stieltjes, and K. H. Maier-Hein. "Strengths and weaknesses of state of the art fiber tractography pipelines – A comprehensive in-vivo and phantom evaluation study using Tractometer". In: *Medical Image Analysis* 26.1 (Dec. 2015), pp. 287–305. ISSN: 13618415. DOI: 10.1016/j.media.2015.10.011.

[246] T. Niendorf, R. M. Dijkhuizen, D. G. Norris, M. van Lookeren Campagne, and K. Nicolay. "Biexponential diffusion attenuation in various states of brain tissue: Implications for diffusion-weighted imaging". In: *Magnetic Resonance in Medicine* 36.6 (Dec. 1996), pp. 847–857. ISSN: 07403194, 15222594. DOI: 10.1002/mrm.1910360607.

[247] D. G. Norris. "The effects of microscopic tissue parameters on the diffusion weighted magnetic resonance imaging experiment". In: *NMR in Biomedicine* 14.2 (Apr. 2001), pp. 77–93. ISSN: 0952-3480, 1099-1492. DOI: 10.1002/nbm.682.

[248] D. S. Novikov, J. H. Jensen, J. A. Helpern, and E. Fieremans. "Revealing mesoscopic structural universality with diffusion". In: *Proceedings of the National Academy of Sciences* 111.14 (Apr. 2014), pp. 5088–5093. ISSN: 0027-8424, 1091-6490. DOI: `10.1073/pnas.1316944111`.

[249] D. S. Novikov, E. Fieremans, J. H. Jensen, and J. A. Helpern. "Random walks with barriers". In: *Nature Physics* 7.6 (Mar. 2011), pp. 508–514. ISSN: 1745-2473, 1745-2481. DOI: `10.1038/nphys1936`.

[250] D. S. Novikov, S. N. Jespersen, V. G. Kiselev, and E. Fieremans. "Quantifying brain microstructure with diffusion MRI: Theory and parameter estimation". In: *arXiv preprint bio-ph/1612.02059v1* (Dec. 2016).

[251] D. S. Novikov and V. G. Kiselev. "Effective medium theory of a diffusion-weighted signal". In: *NMR in Biomedicine* 23.7 (Sept. 2010), pp. 682–697. ISSN: 09523480. DOI: `10.1002/nbm.1584`.

[252] D. S. Novikov, J. Veraart, I. O. Jelescu, and E. Fieremans. "Mapping orientational and microstructural metrics of neuronal integrity with in vivo diffusion MRI". In: *arXiv preprint bio-ph/1609.09144v1* (Sept. 2016).

[253] L. J. O'Donnell, W. M. Wells III, A. J. Golby, and C.-F. Westin. "Unbiased groupwise registration of white matter tractography". In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2012, pp. 123–130.

[254] L. O'Donnell, S. Haker, and C.-F. Westin. "New approaches to estimation of white matter connectivity in diffusion tensor MRI: Elliptic PDEs and geodesics in a tensor-warped space". In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2002*. Ed. by T. Dohi and R. Kikinis. Vol. 2488. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 459–466. ISBN: 978-3-540-44224-0.

[255] E. Özarslan, M. E. Komlosh, M. J. Lizak, F. Horkay, and P. J. Basser. "Double pulsed field gradient (double-PFG) MR imaging (MRI) as a means to measure the size of plant cells: Double-PFG MRI to measure the size of plant cells". In: *Magnetic Resonance in Chemistry* 49 (Dec. 2011), S79–S84. ISSN: 07491581. DOI: `10.1002/mrc.2797`.

[256] E. Özarslan and P. J. Basser. "Microscopic anisotropy revealed by NMR double pulsed field gradient experiments with arbitrary timing parameters". In: *The Journal of Chemical Physics* 128.15 (2008), p. 154511. ISSN: 00219606. DOI: `10.1063/1.2905765`.

[257] E. Özarslan, P. J. Basser, T. M. Shepherd, P. E. Thelwall, B. C. Vemuri, and S. J. Blackband. "Observation of anomalous diffusion in excised tissue by characterizing the diffusion-time dependence of the MR signal". In: *Journal of Magnetic Resonance* 183.2 (Dec. 2006), pp. 315–323. ISSN: 10907807. DOI: `10.1016/j.jmr.2006.08.009`.

[258] E. Özarslan, C. G. Koay, and P. J. Basser. "Remarks on q-space MR propagator in partially restricted, axially-symmetric, and isotropic environments". In: *Magnetic Resonance Imaging* 27.6 (July 2009), pp. 834–844. ISSN: 0730725X. DOI: `10.1016/j.mri.2009.01.005`.

[259] E. Özarslan, C. G. Koay, T. M. Shepherd, M. E. Komlosh, M. O. İrfanoğlu, C. Pierpaoli, and P. J. Basser. "Mean apparent propagator (MAP) MRI: A novel diffusion imaging method for mapping tissue microstructure". In: *NeuroImage* 78 (Sept. 2013), pp. 16–32. ISSN: 10538119. DOI: `10.1016/j.neuroimage.2013.04.016`.

[260] E. Özarslan and T. H. Mareci. "Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution diffusion imaging". In: *Magnetic Resonance in Medicine* 50.5 (2003), pp. 955–965.

[261] E. Özarslan, T. M. Shepherd, B. C. Vemuri, S. J. Blackband, and T. H. Mareci. "Resolution of complex tissue microarchitecture using the diffusion orientation transform (DOT)". In: *NeuroImage* 31.3 (July 2006), pp. 1086–1103. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2006.01.024.

[262] E. Özarslan, B. C. Vemuri, and T. H. Mareci. "Fiber orientation mapping using generalized diffusion tensor imaging". In: *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on*. 2004, pp. 1036–1039.

[263] E. Özarslan, B. C. Vemuri, and T. H. Mareci. "Generalized scalar measures for diffusion MRI using trace, variance, and entropy". In: *Magnetic Resonance in Medicine* 53.4 (Apr. 2005), pp. 866–876. ISSN: 0740-3194, 1522-2594. DOI: 10.1002/mrm.20411.

[264] S. Pajevic and C. Pierpaoli. "Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: application to white matter fiber tract mapping in the human brain". In: *Magnetic Resonance in Medicine* 42.3 (Sept. 1999). PMID: 10467297, pp. 526–540. ISSN: 0740-3194.

[265] E. Panagiotaki, T. Schneider, B. Siow, M. G. Hall, M. F. Lythgoe, and D. C. Alexander. "Compartment models of the diffusion MR signal in brain white matter: A taxonomy and comparison". In: *NeuroImage* 59.3 (Feb. 2012), pp. 2241–2254. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2011.09.081.

[266] G. J. M. Parker, C. A. M. Wheeler-Kingshott, and G. J. Barker. "Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging". In: *Medical Imaging, IEEE Transactions on* 21.5 (2002), pp. 505–512.

[267] O. Pasternak, M. Nilsson, Y. Cohen, E. Özarslan, H. Knutsson, and C.-F. Westin. "Double-PFG filtered diffusion tensors". In: *Proceedings of the 22nd Annual Meeting of the ISMRM*. Milan, 2014, p. 4413.

[268] O. Pasternak, Y. Rathi, M. E. Shenton, and C.-F. Westin. "Estimation of the angle between crossing fibers as a novel structural quantity". In: *Proc. Int. Soc. Magn. Reson. Med*. 2012, p. 1915.

[269] M. Péchaud, M. Descoteaux, and R. Keriven. "Brain connectivity using geodesics in HARDI". In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2009* (2009), pp. 482–489.

[270] X. Pennec, P. Fillard, and N. Ayache. "A Riemannian framework for tensor computing". In: *International Journal of Computer Vision* 66.1 (2006), pp. 41–66.

[271] E. Pichon, C.-F. Westin, and A. R. Tannenbaum. "A Hamilton-Jacobi-Bellman approach to high angular resolution diffusion tractography". In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. 2005, pp. 180–187.

[272] S. Pieper, M. Halle, and R. Kikinis. "3D Slicer". In: *Nano to Macro*. 2004, pp. 632–635.

[273] C. Pierpaoli, P. Jezzard, P. J. Basser, A. Barnett, and G. Di Chiro. "Diffusion tensor MR imaging of the human brain". In: *Radiology* 201.3 (Dec. 1996). PMID: 8939209, pp. 637–648. ISSN: 0033-8419. DOI: 10.1148/radiology.201.3.8939209.

[274] C. Pierpaoli and P. J. Basser. "Toward a quantitative assessment of diffusion anisotropy". In: *Magnetic Resonance in Medicine* 36.6 (Dec. 1996), pp. 893–906. ISSN: 07403194, 15222594. DOI: 10.1002/mrm.1910360612.

[275] M. A. Pinsky. "Isotropic transport process on a Riemannian manifold". In: *Transactions of the American Mathematical Society* 218 (1976), pp. 353–360.

[276] L. Pizarro, B. Burgeth, S. Didas, and J. Weickert. "A generic neighbourhood filtering framework for matrix fields". In: *Computer Vision–ECCV 2008*. Springer, 2008, pp. 521–532.

[277] M. Pizzolato, R. Fick, T. Boutelier, and R. Deriche. "Noise floor removal via phase correction of complex diffusion-weighted images: Influence on DTI and q-space metrics". In: *Computational Diffusion MRI*. 2016.

[278] H. Poincaré. "Sur les intégrales irrégulières des équations linéaires". In: *Acta Mathematica* 8.1 (1886), pp. 295–344.

[279] J. Portegies and R. Duits. "New exact and numerical solutions of the (convection-)diffusion kernels on SE(3)". In: *arXiv preprint math/1604.03843v5* (May 2016).

[280] J. Portegies, G. Sanguinetti, S. Meesters, and R. Duits. "New approximation of a scale space kernel on SE(3) and applications in neuroimaging". In: *Scale Space and Variational Methods in Computer Vision*. Ed. by J.-F. Aujol, M. Nikolova, and N. Papadakis. Vol. 9087. Cham: Springer International Publishing, 2015, pp. 40–52. ISBN: 978-3-319-18460-9 978-3-319-18461-6.

[281] E. Prados, S. Soatto, C. Lenglet, J.-P. Pons, N. Wotawa, R. Deriche, and O. Faugeras. "Control theory and fast marching techniques for brain connectivity mapping". In: *Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. Vol. 1. IEEE, 2006, pp. 1076–1083.

[282] V. Prčkovska, P. R. Rodrigues, R. Duits, B. M. ter Haar Romeny, and A. Vilanova. *Extrapolating fiber crossings from DTI data. Can we gain the same information as HARDI?* Technical Report 10. Eindhoven: Eindhoven University of Technology, 2010.

[283] W. S. Price and O. Soederman. "Some "Reflections" on the effects of finite gradient pulse lengths in PGSE NMR experiments in restricted systems". In: *Israel journal of chemistry* 43.1-2 (2003), pp. 25–32.

[284] D. Raffelt, J.-D. Tournier, J. Fripp, S. Crozier, A. Connelly, and O. Salvado. "Symmetric diffeomorphic registration of fibre orientation distributions". In: *NeuroImage* 56.3 (June 2011), pp. 1171–1180. ISSN: 10538119. DOI: `10.1016/j.neuroimage.2011.02.014`.

[285] D. Raffelt, J.-D. Tournier, S. Rose, G. R. Ridgway, R. Henderson, S. Crozier, O. Salvado, and A. Connelly. "Apparent Fibre Density: A novel measure for the analysis of diffusion-weighted magnetic resonance images". In: *NeuroImage* 59.4 (Feb. 2012), pp. 3976–3994. ISSN: 10538119. DOI: `10.1016/j.neuroimage.2011.10.045`.

[286] M. Reisert, V. G. Kiselev, B. Dihtal, E. Kellner, and D. S. Novikov. "MesoFT: unifying diffusion modelling and fiber tracking". In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2014, pp. 201–208.

[287] M. Reisert, I. Mader, C. Anastasopoulos, M. Weigel, S. Schnell, and V. Kiselev. "Global fiber reconstruction becomes practical". In: *NeuroImage* 54.2 (Jan. 2011), pp. 955–962. ISSN: 10538119. DOI: `10.1016/j.neuroimage.2010.09.016`.

[288] B. Riemann. "Über die Hypothesen, welche der Geometrie zu Grunde liegen". In: *Abhandlungen der Königlichen Gesellschaft der Wissenschaften zu Göttingen* 13 (1868), pp. 133–150.

[289] G. K. Rohde, A. S. Barnett, P. J. Basser, S. Marenco, and C. Pierpaoli. "Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI". In: *Magnetic Resonance in Medicine* 51.1 (Jan. 2004), pp. 103–114. ISSN: 0740-3194, 1522-2594. DOI: `10.1002/mrm.10677`.

[290] S. Roosendaal, J. Geurts, H. Vrenken, H. Hulst, K. Cover, J. Castelijns, P. Pouwels, and F. Barkhof. "Regional DTI differences in multiple sclerosis patients". In: *NeuroImage* 44.4 (Feb. 2009), pp. 1397–1403. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2008.10.026.

[291] M. Rubinov and O. Sporns. "Complex network measures of brain connectivity: Uses and interpretations". In: *NeuroImage* 52.3 (Sept. 2010), pp. 1059–1069. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2009.10.003.

[292] J. Rudie, J. Brown, D. Beck-Pancer, L. Hernandez, E. Dennis, P. Thompson, S. Bookheimer, and M. Dapretto. "Altered functional and structural brain network organization in autism". In: *NeuroImage: Clinical* 2 (2013), pp. 79–94. ISSN: 22131582. DOI: 10.1016/j.nicl.2012.11.006.

[293] H. Rund. *The Hamilton-Jacobi theory in the calculus of variations : Its role in mathematics and physics*. London: D. Van Nostrand Company, 1966.

[294] R. Salvador, A. Peña, D. K. Menon, T. A. Carpenter, J. D. Pickard, and E. T. Bullmore. "Formal characterization and extension of the linearized diffusion tensor model". In: *Human Brain Mapping* 24.2 (Feb. 2005), pp. 144–155. ISSN: 1065-9471, 1097-0193. DOI: 10.1002/hbm.20076.

[295] P. Savadjiev, Y. Rathi, S. Bouix, R. Verma, and C.-F. Westin. "Multi-scale characterization of white matter tract geometry". In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2012, pp. 34–41.

[296] M. Schachter, M. Does, A. Anderson, and J. Gore. "Measurements of restricted diffusion using an oscillating gradient spin-echo sequence". In: *Journal of Magnetic Resonance* 147.2 (Dec. 2000), pp. 232–237. ISSN: 10907807. DOI: 10.1006/jmre.2000.2203.

[297] B. Scherrer, A. Schwartzman, M. Taquet, M. Sahin, S. P. Prabhu, and S. K. Warfield. "Characterizing brain tissue by assessment of the distribution of anisotropic microstructural environments in diffusion-compartment imaging (DIAMOND): Characterizing Brain Tissue with DIAMOND". In: *Magnetic Resonance in Medicine* (Sept. 2015), n/a–n/a. ISSN: 07403194. DOI: 10.1002/mrm.25912.

[298] R. Schneider. *Convex bodies: The Brunn-Minkowski theory*. Second Expanded Edition. Encyclopedia of mathematics and its applications. New York: Cambridge University Press, 2014. ISBN: 978-1-107-60101-7.

[299] M. Schober, N. Kasenburg, A. Feragen, P. Hennig, and S. Hauberg. "Probabilistic shortest path tractography in DTI using Gaussian process ODE solvers". In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014*. Nagoya, 2014.

[300] T. Schultz. "Towards resolving fiber crossings with higher order tensor inpainting". In: *New developments in the visualization and processing of tensor fields*. Ed. by D. H. Laidlaw and A. Vilanova. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 253–265. ISBN: 978-3-642-27342-1 978-3-642-27343-8.

[301] T. Schultz, H. Theisel, and H.-P. Seidel. "Crease surfaces: From theory to extraction and application to diffusion tensor MRI". In: *IEEE transactions on visualization and computer graphics* 16.1 (Feb. 2010). PMID: 19910665, pp. 109–119. ISSN: 1077-2626. DOI: 10.1109/TVCG.2009.44.

[302] T. Schultz, C.-F. Westin, and G. Kindlmann. "Multi-diffusion-tensor fitting via spherical deconvolution: A unifying framework". In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2010* (2010), pp. 674–681.

[303] E. Scola, M. Bozzali, F. Agosta, G. Magnani, M. Franceschi, M. P. Sormani, M. Cercignani, E. Pagani, M. Falautano, M. Filippi, and A. Falini. "A diffusion tensor MRI study of patients with MCI and AD with a 2-year clinical follow-up". In: *Journal of Neurology, Neurosurgery & Psychiatry* 81.7 (July 2010), pp. 798–805. ISSN: 0022-3050. DOI: 10.1136/jnnp.2009.189639.

[304] A. Seehaus, A. Roebroeck, M. Bastiani, L. Fonseca, H. Bratzke, N. Lori, A. Vilanova, R. Goebel, and R. Galuske. "Histological validation of high-resolution DTI in human post mortem tissue". In: *Frontiers in Neuroanatomy* 9 (July 2015). ISSN: 1662-5129. DOI: 10.3389/fnana.2015.00098.

[305] N. Sepasian, J. H. M. ten Thije Boonkkamp, B. M. Ter Haar Romeny, and A. Vilanova. "Multivalued geodesic ray-tracing for computing brain connections using diffusion tensor imaging". In: *SIAM Journal on Imaging Sciences* 5.2 (Jan. 2012), pp. 483–504. ISSN: 1936-4954. DOI: 10.1137/110824395.

[306] N. Sepasian, J. H. ten Thije Boonkkamp, L. J. Astola, M. Breeuwer, and A. Fuster. "Modified geodesic ray-tracing for diffusion tensor imaging". In: *Proceedings of the 2016 IEEE International Symposium on Biomedical Imaging*. Prague: IEEE, Apr. 2016.

[307] N. Sepasian, J. H. ten Thije Boonkkamp, L. M. Florack, B. M. T. H. Romeny, and A. Vilanova. "Riemann-Finsler multi-valued geodesic tractography for HARDI". In: *Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data*. Springer, 2014, pp. 209–225.

[308] K. Setsompop, R. Kimmlingen, E. Eberlein, T. Witzel, J. Cohen-Adad, J. McNab, B. Keil, M. Tisdall, P. Hoecht, P. Dietz, S. Cauley, V. Tountcheva, V. Matschl, V. Lenz, K. Heberlein, A. Potthast, H. Thein, J. Van Horn, A. Toga, F. Schmitt, D. Lehne, B. Rosen, V. Wedeen, and L. Wald. "Pushing the limits of in vivo diffusion MRI for the Human Connectome Project". In: *NeuroImage* 80 (Oct. 2013), pp. 220–233. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2013.05.078.

[309] N. Shemesh, S. N. Jespersen, D. C. Alexander, Y. Cohen, I. Drobnjak, T. B. Dyrby, J. Finsterbusch, M. A. Koch, T. Kuder, F. Laun, M. Lawrenz, H. Lundell, P. P. Mitra, M. Nilsson, E. Özarslan, D. Topgaard, and C.-F. Westin. "Conventions and nomenclature for double diffusion encoding NMR and MRI". In: *Magnetic Resonance in Medicine* 75.1 (Jan. 2016), pp. 82–87. ISSN: 07403194. DOI: 10.1002/mrm.25901.

[310] A. J. Sherbondy, R. F. Dougherty, M. Ben-Shachar, S. Napel, and B. A. Wandell. "ConTrack: Finding the most likely pathways between brain regions using diffusion tractography". In: *Journal of Vision* 8.9 (July 2008), pp. 15–15. ISSN: 1534-7362. DOI: 10.1167/8.9.15.

[311] N. Shu, Y. Liu, K. Li, Y. Duan, J. Wang, C. Yu, H. Dong, J. Ye, and Y. He. "Diffusion tensor tractography reveals disrupted topological efficiency in white matter structural networks in multiple sclerosis". In: *Cerebral Cortex* 21.11 (Nov. 2011), pp. 2565–2577. ISSN: 1047-3211, 1460-2199. DOI: 10.1093/cercor/bhr039.

[312] D. Sinnaeve. "The Stejskal-Tanner equation generalized for any gradient shape-an overview of most pulse sequences measuring free diffusion". In: *Concepts in Magnetic Resonance Part A* 40A.2 (Mar. 2012), pp. 39–65. ISSN: 15466086. DOI: 10.1002/cmr.a.21223.

[313] J. Sjölund, F. Szczepankiewicz, M. Nilsson, D. Topgaard, C.-F. Westin, and H. Knutsson. "Constrained optimization of gradient waveforms for generalized diffusion encoding". In: *Journal of Magnetic Resonance* 261 (Dec. 2015), pp. 157–168. ISSN: 10907807. DOI: 10.1016/j.jmr.2015.10.012.

[314] C. P. Slichter. *Principles of magnetic resonance*. 3rd enl. and updated ed. Springer series in solid-state sciences 1. Berlin ; New York: Springer, 1996. ISBN: 3-540-50157-6.

[315] R. E. Smith, J.-D. Tournier, F. Calamante, and A. Connelly. "Anatomically-constrained tractography: Improved diffusion MRI streamlines tractography through effective use of anatomical information". In: *NeuroImage* 62.3 (Sept. 2012), pp. 1924–1938. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2012.06.005.

[316] S. M. Smith, M. Jenkinson, H. Johansen-Berg, D. Rueckert, T. E. Nichols, C. E. Mackay, K. E. Watkins, O. Ciccarelli, M. Z. Cader, P. M. Matthews, and T. E. Behrens. "Tract-based spatial statistics: Voxelwise analysis of multi-subject diffusion data". In: *NeuroImage* 31.4 (July 2006), pp. 1487–1505. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2006.02.024.

[317] S. N. Sotiropoulos, T. E. Behrens, and S. Jbabdi. "Ball and rackets: Inferring fiber fanning from diffusion-weighted MRI". In: *NeuroImage* 60.2 (Apr. 2012), pp. 1412–1425. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2012.01.056.

[318] S. N. Sotiropoulos, S. Jbabdi, J. Xu, J. L. Andersson, S. Moeller, E. J. Auerbach, M. F. Glasser, M. Hernandez, G. Sapiro, M. Jenkinson, D. A. Feinberg, E. Yacoub, C. Lenglet, D. C. Van Essen, K. Ugurbil, and T. E. Behrens. "Advances in diffusion MRI acquisition and processing in the Human Connectome Project". In: *NeuroImage* 80 (Oct. 2013), pp. 125–143. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2013.05.057.

[319] M. Spivak. *A comprehensive introduction to differential geometry*. Vol. I-V. Boston, Mass. (U.S.A.): Publish or Perish, 2005.

[320] E. Spodarev. "On the rose of intersections of stationary flat processes". In: *Advances in Applied Probability* 33.3 (2001), pp. 584–599.

[321] G. J. Stanisz, G. A. Wright, R. M. Henkelman, and A. Szafer. "An analytical model of restricted diffusion in bovine optic nerve". In: *Magnetic Resonance in Medicine* 37.1 (Jan. 1997), pp. 103–111. ISSN: 07403194, 15222594. DOI: 10.1002/mrm.1910370115.

[322] R. M. Starr. "Quasi-equilibria in markets with non-convex preferences". In: *Econometrica* 37.1 (Jan. 1969), pp. 25–38. ISSN: 00129682. DOI: 10.2307/1909201.

[323] E. O. Stejskal and J. E. Tanner. "Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient". In: *The Journal of Chemical Physics* 42.1 (1965), p. 288. ISSN: 00219606. DOI: 10.1063/1.1695690.

[324] D. Stoyan, W. S. Kendall, and J. Mecke. *Stochastic geometry and its applications*. 2nd ed. Wiley series in probability and statistics. Chichester; New York: Wiley, 1995. ISBN: 978-0-471-95099-8.

[325] K. Tabelow, H. U. Voss, and J. Polzehl. "Modeling the orientation distribution function by mixtures of angular central Gaussian distributions". In: *Journal of Neuroscience Methods* 203.1 (Jan. 2012), pp. 200–211. ISSN: 01650270. DOI: 10.1016/j.jneumeth.2011.09.001.

[326] M. Taquet, B. Scherrer, O. Commowick, J. M. Peters, M. Sahin, B. Macq, and S. K. Warfield. "A mathematical framework for the registration and analysis of multi-fascicle models for population studies of the brain microstructure". In: *IEEE Transactions on Medical Imaging* 33.2 (Feb. 2014), pp. 504–517. ISSN: 0278-0062, 1558-254X. DOI: 10.1109/TMI.2013.2289381.

[327] C. M. W. Tax. *Discussion & debate: "Sheet or no sheet?" - Introduction on the mathematical background of the sheet structure*. Singapore, May 2016.

[328]   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, and A. Leemans. "Towards quantification of the brain's sheet structure in diffusion MRI data". In: *Proceedings of the International Biomedical and Astronomical Signal Processing (BASP) Frontiers workshop*. Villars-sur-Ollon, 2015, p. 74.

[329]   C. M. W. Tax, T. C. J. Dela Haije, A. Fuster, R. Duits, M. A. Viergever, L. M. J. Florack, and A. Leemans. "Towards quantification of the brain's sheet structure: Evaluation of the discrete Lie bracket". In: *Proceedings of the 22nd Annual Meeting of the ISMRM*. Milan, 2014, p. 0975.

[330]   C. M. W. Tax, T. C. J. Dela Haije, A. Fuster, M. A. Viergever, L. M. J. Florack, and A. Leemans. "Considerations on the theory of sheet structure of cerebral pathways". In: Podstrana, Oct. 2013.

[331]   C. M. W. Tax, T. C. J. Dela Haije, A. Fuster, C.-F. Westin, M. A. Viergever, L. M. J. Florack, and A. Leemans. "Mapping the brain's "sheet probability index" (SPI) with diffusion MRI: Sheet happens?!" In: *Proceedings of the 24th Annual Meeting of the ISMRM*. Singapore, 2016, p. 0791.

[332]   C. M. W. Tax, C.-F. Westin, T. C. J. Dela Haije, A. Fuster, M. A. Viergever, E. Calabrese, L. M. J. Florack, and A. Leemans. "Quantifying the brain's sheet structure with normalized convolution". In: *Medical Image Analysis* (2017).

[333]   C. M. W. Tax, C.-F. Westin, T. C. J. Dela Haije, A. Fuster, M. A. Viergever, L. M. J. Florack, and A. Leemans. "Normalized convolution for robust assessment of the brain's sheet structure". In: *Abstract Book of the 8th Annual ISMRM Benelux Chapter Meeting*. Eindhoven, 2016, p039.

[334]   C. M. W. Tax, C.-F. Westin, T. C. J. Dela Haije, A. Fuster, M. A. Viergever, L. M. J. Florack, and A. Leemans. "Robust assessment of the brain's sheet structure using normalized convolution". In: *Proceedings of the 24th Annual Meeting of the ISMRM*. Singapore, 2016, p. 2073.

[335]   C. M. Tax, T. Dela Haije, A. Fuster, C.-F. Westin, M. A. Viergever, L. Florack, and A. Leemans. "Sheet Probability Index (SPI): Characterizing the geometrical organization of the white matter with diffusion MRI". In: *NeuroImage* 142 (Nov. 2016), pp. 260–279. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2016.07.042.

[336]   C. M. Tax, B. Jeurissen, S. B. Vos, M. A. Viergever, and A. Leemans. "Recursive calibration of the fiber response function for spherical deconvolution of diffusion MRI data". In: *NeuroImage* 86 (Feb. 2014), pp. 67–80. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2013.07.067.

[337]   C. M. Tax, W. M. Otte, M. A. Viergever, R. M. Dijkhuizen, and A. Leemans. "REKINDLE: Robust extraction of kurtosis indices with linear estimation". In: *Magnetic Resonance in Medicine* 73.2 (Feb. 2015), pp. 794–808. ISSN: 07403194. DOI: 10.1002/mrm.25165.

[338]   C. Tax, R. Duits, B. ter Haar Romeny, A. Vilanova, and P. Ossenblok. "Tractography of the optic radiation for vision sparing epilepsy surgery". In: IEEE, June 2012, pp. 441–445. ISBN: 978-1-4673-2237-9 978-1-4673-2238-6 978-1-4673-2236-2. DOI: 10.1109/ICInfA.2012.6246846.

[339]   N. Tomczak-Jaegermann. *Banach-Mazur distances and finite-dimensional operator ideals*. Pitman monographs and surveys in pure and applied mathematics 38. Harlow, Essex, England : New York: Longman Scientific & Technical ; J. Wiley, 1989. ISBN: 978-0-470-20982-0.

[340]   J.-D. Tournier, S. Mori, and A. Leemans. "Diffusion tensor imaging and beyond". In: *Magnetic Resonance in Medicine* 65.6 (June 2011), pp. 1532–1556. ISSN: 07403194. DOI: 10.1002/mrm.22924.

[341]   J.-D. Tournier, F. Calamante, and A. Connelly. "MRtrix: Diffusion tractography in crossing fiber regions". In: *International Journal of Imaging Systems and Technology* 22.1 (Mar. 2012), pp. 53–66. ISSN: 08999457. DOI: 10.1002/ima.22005.

[342]   J.-D. Tournier, F. Calamante, and A. Connelly. "Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution". In: *NeuroImage* 35.4 (May 2007), pp. 1459–1472. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2007.02.016.

[343]   J.-D. Tournier, F. Calamante, D. G. Gadian, and A. Connelly. "Diffusion-weighted magnetic resonance imaging fibre tracking using a front evolution algorithm". In: *NeuroImage* 20.1 (Sept. 2003), pp. 276–288. ISSN: 10538119. DOI: 10.1016/S1053-8119(03)00236-2.

[344]   J.-D. Tournier, F. Calamante, D. G. Gadian, and A. Connelly. "Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution". In: *NeuroImage* 23.3 (Nov. 2004), pp. 1176–1185. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2004.07.037.

[345]   A. Tristán-Vega and S. Aja-Fernández. "Joint LMMSE estimation of DWI data for DTI processing". In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2008*. Springer, 2008, pp. 27–34.

[346]   D. S. Tuch, T. G. Reese, M. R. Wiegell, and V. J. Wedeen. "Diffusion MRI of complex neural architecture". In: *Neuron* 40.5 (2003), pp. 885–895.

[347]   D. S. Tuch. "Q-ball imaging". In: *Magnetic Resonance in Medicine* 52.6 (Dec. 2004), pp. 1358–1372. ISSN: 0740-3194, 1522-2594. DOI: 10.1002/mrm.20279.

[348]   D. S. Tuch, T. G. Reese, M. R. Wiegell, N. Makris, J. W. Belliveau, and V. J. Wedeen. "High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity". In: *Magnetic Resonance in Medicine* 48.4 (Oct. 2002), pp. 577–582. ISSN: 0740-3194, 1522-2594. DOI: 10.1002/mrm.10268.

[349]   H. J. V. Tyrrell and K. R. Harris. *Diffusion in liquids: A theoretical and experimental study*. OCLC: 893741022. Burlington: Elsevier Science, 1984. ISBN: 978-1-4831-0089-0.

[350]   A. M. Uluğ and P. C. van Zijl. "Orientation-independent diffusion imaging without tensor diagonalization: Anisotropy definitions based on physical attributes of the diffusion ellipsoid". In: *Journal of Magnetic Resonance Imaging* 9.6 (June 1999), pp. 804–813. ISSN: 1053-1807, 1522-2586. DOI: 10.1002/(SICI)1522-2586(199906)9:6<804::AID-JMRI7>3.0.CO;2-B.

[351]   D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub, and K. Ugurbil. "The WU-Minn Human Connectome Project: An overview". In: *NeuroImage* 80 (Oct. 2013), pp. 62–79. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2013.05.041.

[352]   N. G. Van Kampen. "Derivation of the phenomenological equations from the master equation: I. Even variables only". In: *Physica* 23.6-10 (1957), pp. 707–719.

[353]   S. R. S. Varadhan. "Asymptotic probabilities and differential equations". In: *Communications on Pure and Applied Mathematics* 19.3 (Aug. 1966), pp. 261–286. ISSN: 00103640, 10970312. DOI: 10.1002/cpa.3160190303.

[354] J. Veraart, J. Sijbers, S. Sunaert, A. Leemans, and B. Jeurissen. "Weighted linear least squares estimation of diffusion MRI parameters: Strengths, limitations, and pitfalls". In: *NeuroImage* 81 (Nov. 2013), pp. 335–346. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2013.05.028.

[355] A. Vilanova, G. Berenschot, and C. v. Pul. "DTI visualization with streamsurfaces and evenly-spaced volume seeding". In: (2004). DOI: 10.2312/VisSym/VisSym04/173-182.

[356] A. Vilanova, S. Zhang, G. Kindlmann, and D. Laidlaw. "An introduction to visualization of diffusion tensor imaging and its applications". In: *Visualization and Processing of Tensor Fields*. Springer, 2006, pp. 121–153.

[357] S. Wakana, A. Caprihan, M. M. Panzenboeck, J. H. Fallon, M. Perry, R. L. Gollub, K. Hua, J. Zhang, H. Jiang, and P. Dubey. "Reproducibility of quantitative tractography methods applied to cerebral white matter". In: *Neuroimage* 36.3 (2007), pp. 630–644.

[358] R. Walker. *The human brain project: A report to the European commision*. Tech. rep. Lausanne: HBP-PS Consortium, Apr. 2012.

[359] R. Wang, T. Benner, A. G. Sorensen, and V. J. Wedeen. "Diffusion toolkit: A software package for diffusion imaging data processing and tractography". In: *Proceedings of the 16th Annual Meeting of the ISMRM*. Vol. 15. 2007, p. 3720.

[360] Z. Wang, B. Vemuri, Y. Chen, and T. Mareci. "A constrained variational principle for direct estimation and smoothing of the diffusion tensor field from complex DWI". In: *IEEE Transactions on Medical Imaging* 23.8 (Aug. 2004), pp. 930–939. ISSN: 0278-0062. DOI: 10.1109/TMI.2004.831218.

[361] S. Watanabe and T. Watanabe. "Convergence of isotropic scattering transport process to Brownian motion". In: *Nagoya Mathematical Journal* 40 (1970), pp. 161–171.

[362] T. Watanabe. "Weak convergence of the isotropic scattering transport process with one speed in the plane to Brownian motion". In: *Proceedings of the Japan Academy* 44.7 (1968). MR: MR0236996 Zbl: 0177.45403, pp. 677–680. ISSN: 0021-4280. DOI: 10.3792/pja/1195521091.

[363] V. J. Wedeen. *Discussion & debate: "Sheet or no sheet?" - For*. Singapore, May 2016.

[364] V. J. Wedeen, D. L. Rosene, R. Wang, G. Dai, F. Mortazavi, P. Hagmann, J. H. Kaas, and W.-Y. I. Tseng. "Response to comment on "The geometric structure of the brain fiber pathways"". In: *Science* 337.6102 (Sept. 2012), pp. 1605–1605. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1223493.

[365] V. J. Wedeen, D. L. Rosene, R. Wang, G. Dai, F. Mortazavi, P. Hagmann, J. H. Kaas, and W.-Y. I. Tseng. "The geometric structure of the brain fiber pathways". In: *Science* 335.6076 (Mar. 2012), pp. 1628–1634. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1215280.

[366] V. J. Wedeen, R. P. Wang, J. D. Schmahmann, T. Benner, W. Y. I. Tseng, G. Dai, D. N. Pandya, P. Hagmann, H. D'Arceuil, and A. J. de Crespigny. "Diffusion spectrum magnetic resonance imaging (DSI) tractography of crossing fibers". In: *NeuroImage* 41.4 (July 2008), pp. 1267–1277. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2008.03.036.

[367] V. J. Wedeen, P. Hagmann, W.-Y. I. Tseng, T. G. Reese, and R. M. Weisskoff. "Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging". In: *Magnetic Resonance in Medicine* 54.6 (Dec. 2005), pp. 1377–1386. ISSN: 0740-3194, 1522-2594. DOI: 10.1002/mrm.20642.

[368] V. J. Wedeen, F. Mortazavi, R. Wang, W.-Y. I. Tseng, T. Witzel, A. Nummenmaa, W. Morrison, H. E. Stanley, L. Wald, and D. L. Rosene. "Grid structure of brain pathways – Validation and the character of turns". In: *Proceedings of the 22nd Annual Meeting of the ISMRM*. Milan, 2014, p. 0803.

[369] J. Weickert and M. Welk. "Tensor field interpolation with PDEs". In: *Visualization and Processing of Tensor Fields*. Ed. by J. Weickert and H. Hagen. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 315–325. ISBN: 978-3-540-25032-6 978-3-540-31272-7.

[370] C.-F. Westin, H. Knutsson, O. Pasternak, F. Szczepankiewicz, E. Özarslan, D. van Westen, C. Mattisson, M. Bogren, L. J. O'Donnell, M. Kubicki, D. Topgaard, and M. Nilsson. "Q-space trajectory imaging for multidimensional diffusion MRI of the human brain". In: *NeuroImage* 135 (July 2016), pp. 345–362. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2016.02.039.

[371] C.-F. Westin, F. Szczepankiewicz, O. Pasternak, E. Özarslan, D. Topgaard, H. Knutsson, and M. Nilsson. "Measurement tensors in diffusion MRI: Generalizing the concept of diffusion encoding". In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014*. Ed. by P. Golland, N. Hata, C. Barillot, J. Hornegger, and R. Howe. Vol. 8675. Cham: Springer International Publishing, 2014, pp. 209–216. ISBN: 978-3-319-10442-3 978-3-319-10443-0.

[372] C.-F. Westin, S. E. Maier, H. Mamata, A. Nabavi, F. A. Jolesz, and R. Kikinis. "Processing and visualization for diffusion tensor MRI". In: *Medical image analysis* 6.2 (2002), pp. 93–108.

[373] C.-F. Westin, S. Peled, H. Gudbjartsson, R. Kikinis, and F. A. Jolesz. "Geometrical diffusion measures for MRI from tensor basis analysis". In: *Proceedings of the 5th Scientific Meeting of the ISMRM*. Vancouver, 1997, p. 1742.

[374] R. Wong. *Asymptotic approximations of integrals*. Boston: Academic Press, 1989. ISBN: 0-12-762535-6 978-0-12-762535-5.

[375] D. Xing, N. G. Papadakis, C. L.-H. Huang, V. M. Lee, T. Adrian Carpenter, and L. D. Hall. "Optimised diffusion-weighting for measurement of apparent diffusion coefficient (ADC) in human brain". In: *Magnetic Resonance Imaging* 15.7 (Jan. 1997), pp. 771–784. ISSN: 0730725X. DOI: 10.1016/S0730-725X(97)00037-4.

[376] D. A. Yablonskiy, G. L. Bretthorst, and J. J. Ackerman. "Statistical model for diffusion attenuated MR signal". In: *Magnetic Resonance in Medicine* 50.4 (Oct. 2003), pp. 664–669. ISSN: 0740-3194, 1522-2594. DOI: 10.1002/mrm.10578.

[377] M. Yamashita, K. Fujisawa, K. Nakata, M. Nakata, M. Fukuda, K. Kobayashi, and K. Goto. *A high-performance software package for semidefinite programs: SDPA 7*. Research Report B-460. Tokyo, Japan: Tokyo Institute of Technology, 2010.

[378] F. Yang, Y.-M. Zhu, I. E. Magnin, J.-H. Luo, P. Croisille, and P. B. Kingsley. "Feature-based interpolation of diffusion tensor fields and application to human cardiac DT-MRI". In: *Medical Image Analysis* 16.2 (Feb. 2012), pp. 459–481. ISSN: 13618415. DOI: 10.1016/j.media.2011.11.003.

[379] C. Yolcu and E. Özarslan. "Diffusion-weighted magnetic resonance signal for general gradient waveforms: Multiple correlation function framework, path integrals, and parallels between them". In: *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*. Ed. by I. Hotz and T. Schultz. Cham: Springer International Publishing, 2015, pp. 3–19. ISBN: 978-3-319-15089-5 978-3-319-15090-1.

[380] P. A. Yushkevich, H. Zhang, T. J. Simon, and J. C. Gee. "Structure-specific statistical mapping of white matter tracts". In: *NeuroImage* 41.2 (June 2008), pp. 448–461. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2008.01.013.

[381] H. Zhang, P. Yushkevich, D. Alexander, and J. Gee. "Deformable registration of diffusion tensor MR images with explicit orientation optimization". In: *Medical Image Analysis* 10.5 (Oct. 2006), pp. 764–785. ISSN: 13618415. DOI: 10.1016/j.media.2006.06.004.

[382] H. Zhang, T. Schneider, C. A. Wheeler-Kingshott, and D. C. Alexander. "NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain". In: *NeuroImage* 61.4 (July 2012), pp. 1000–1016. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2012.03.072.

[383] J. Zhang, P. C. van Zijl, and S. Mori. "Image contrast using the secondary and tertiary eigenvectors in diffusion tensor imaging". In: *Magnetic Resonance in Medicine* 55.2 (Feb. 2006), pp. 439–449. ISSN: 0740-3194, 1522-2594. DOI: 10.1002/mrm.20767.

[384] M. Zhang and P. T. Fletcher. "Finite-dimensional Lie algebras for fast diffeomorphic image registration". In: *International Conference on Information Processing in Medical Imaging*. Springer, 2015, pp. 249–260.

[385] S. Zhang, C. Demiralp, and D. H. Laidlaw. "Visualizing diffusion tensor MR images using streamtubes and streamsurfaces". In: *IEEE Transactions on Visualization and Computer Graphics* 9.4 (2003), pp. 454–462.

[386] H. Zhao. "A fast sweeping method for Eikonal equations". In: *Mathematics of Computation* 74.250 (May 2004), pp. 603–628. ISSN: 0025-5718. DOI: 10.1090/S0025-5718-04-01678-3.

Acknowledgements

This thesis is the product of years of work, fun, traveling, speaking, and learning, and I can safely say that it would have looked entirely different without the invaluable support and contributions of my colleagues, friends, and family. Naming everyone who has had an impact on this project is an impossible task, so I apologize in advance for any omissions in this attempt.

Luc and Andrea, it is only natural that you have the honor to be mentioned first in this recital, as without you this thesis would not have existed at all. I thank you for the freedom that I have enjoyed these past years, visiting groups and conferences around the world. I also thank you for your patience and faith in allowing me to pursue my own ideas and interests. Although I was often frustratingly wrong, you have always supported me, and looking back I regret not listening to your advice more often. You have without a doubt made me a better researcher, and I hope our collaboration can continue beyond this PhD project. Luc, before I started I knew you as a great teacher with a keen eye for detail, and I am happy to have had a few more years to learn from you. Andrea, I am not only grateful for your support in the scientific aspects of this PhD, but more so for your friendship. Your company, words of confidence, and constant supply of coffee made the job manageable. Luckily we can still share some cappuccinos in Denmark! Luc and Andrea, it has been a joy to be your student, and I look forward to a reunion with some drinks at the next conference!

I would also like to extend my gratitude to the members of my reading committee—C-F, Rachid, Bart, Evren, and Aasa. Not only did you spend your time combing through my thesis, but it is also your work that has motivated and inspired my own, and I am honored that you were willing to take part in my PhD defense. I sincerely hope that I will have the opportunity to (continue to) work with you on some of the interesting open questions in diffusion MRI. Along the same lines I would like to thank all my collaborators, co-authors, and colleagues from around the world, and express the same sentiment. Alexander, Remco, and everyone else—this thesis builds on your accomplishments.

Next I would like to thank the people that have welcomed me in their labs over the course of my PhD. Karsten and Jörg at WIAS, C-F, Peter, and everyone else at LMI, Max and the people at SCIL, and of course Aasa and my new colleagues at

DIKU. I have learned different things from each of you, and thoroughly enjoyed your hospitality. I would also like to take this opportunity to thank in particular Peter and Preetika, who have made Boston more like a home to me. Peter, I hope that when the paper is finally published we can find an opportunity to have a small celebration!

At the many conferences I had the fortune to visit I have always been surrounded by great people, many of whom I now consider my friends. Of course it helped a lot that my university friends Rutger and Chantal often ended up in the same places. Though we seem to be moving farther and farther away from each other, I look forward to visiting (and hosting) you wherever we may all end up. Rutger, conference parties will for sure become a lot less entertaining without you there! To my paranymph Chantal, I had a lot of fun working with you—thank you for thinking of me when you started looking at sheet structure! Your work ethic and passion are something to aspire to, and I look forward to your next big idea. I learned a lot from you, and I am glad to be your friend.

I would also like to thank the university and all my colleagues and friends there. Playing squash with the CASA group (being beaten by Mark and Upanshu), listening to talks and having beers with an ever-evolving group of interesting PhD students, and being shown the ropes by Enna, Jolijn, and Marèse, the TU/e has been a great place to work. In particular I would like to thank Remco, who was my MSc supervisor and my gateway to geometry. Remco, you are one of the reasons I am currently in science, and I hope we can keep collaborating. Of course I also have to mention my 'students', Maddy and Joep. It was great fun to work with you both, and I look forward to seeing where you will end up. Finally there are my MetaForum mates: Erik, Frank, Gonzalo, Jorg, Michiel, Stefan, and Stephan—the Metaboyz. The many hours of quizzing and comedy, of discussions about mathematics, programming, philosophy, and—on rare occasions—somewhat less serious prompts, will be sorely missed! Best of luck to you, and I hope that we can keep in touch!

Aside from the professional and scientific support of my colleagues, I would also like to mention my friends, who provided an occasionally necessary escape from work. Mies and Jurre, Chantal and Max, Simone, Rik, and our newest—as of this writing still unnamed—member, our Dutch getaways have become a yearly tradition, and I look forward to the next five! Jurre, my other paranymph, we have been best friends since our first day at the TU/e, and your character and friendship have since become a bedrock in my life. There are few people I can be as comfortable around as I can be with you. I love you, man. Jeroen and Rutger, though we don't see each other nearly as often as I'd like, there is never a dull moment with you, and I am glad to know that you are there when I need you. Bart, the same goes for you. You are my oldest friend, and it is great to know that I can always rely on you. Finally Jorre and Floris, you have been two of my best friends since the beginning of high school, and I consider it a rare fortune to have met you when I did. I appreciate that you are always ready to challenge me, intellectually or otherwise. I hope there will

be many more winter sports and trips to Portugal with you, Sacha, and the rest of the Nijmegen group.

Pat Mo, Sai Pa, Fungie, Sau, Jangie, and Emma, thank you for welcoming me into your family, and Pat Mo and Sai Pa, thank you for quite literally providing me the energy to complete this thesis! Finally I would like to thank my parents Harry and Jolanda, brothers Geert and Rob, and sister Tessa. Mom and Dad, you have given me all the tools I need to live the life that I want, and you have always been there for me. You have, both in your own ways, been living examples of the principles that guide me in my work and in my life, and I will always carry that with me. Thank you. Geert, Rob, and Tessa, our endless arguments are one of the things I enjoy most in life, and I am immensely happy that—though we walk very different paths—we remain so close. You truly are some of the best people I know. Denise and Jasmin, naturally you are included in this acknowledgement as well, I could not have wished for a better family!

Lastly I would like to thank May, my amazing girlfriend. You were there for every difficult decision, every long night, and every short weekend. I am immensely happy that we got to share our times as PhD students—in particular for the great trips we got to take together! I know that it could not have been easy dealing with my unorthodox working hours and absentmindedness, but you have always been there when I needed you. Thank you for your patience, your support, and your love. I hope I can be there for you as you have been there for me, in your own research, and in your life. Ngo oi nei!

So that's it—a little over 4 years, a little under 300 pages. Thanks for reading!

Tom Dela Haije
Eindhoven, 8 April 2017

About the Author

Tom Dela Haije was born on 2 February 1988 in Landgraaf, the Netherlands. He studied at the Eindhoven University of Technology, where he received an MSc degree in Biomedical Engineering with honors in 2012. During his master track he spent five months in Paris working at Philips Healthcare on 'Onset Detection in First-Pass Perfusion MRI', and subsequently worked on the enhancement of diffusion-weighted MRI data at the Biomedical Image Analysis group (BMIA), continuing work he started during a previous internship. In September 2012 he started a PhD project titled 'Riemann–Finsler Geometry for Human Brain Connectomics' in the Department of Mathematics and Computer Science at the Eindhoven University of Technology, of which the results are presented in this dissertation. Several months of this project were spent in labs abroad, including visits to the Laboratory of Mathematics in Imaging at the Harvard Medical School (Boston), the Weierstrass Institute for Applied Analysis and Stochastics (Berlin), and the Copenhagen University. Various conference contributions received accolades, and the work in this thesis led to invited talks at a number of workshops and research institutes. Since March 2017 he is employed at the Image Group at the Department of Computer Science of the Copenhagen University.

Publications

Journal articles

T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. *The Finslerian character of the diffusion MRI signal*. **In preparation, preprint available upon request.**

T. C. J. Dela Haije, P. Savadjiev, A. Fuster, A. Tristán-Vega, R. Schultz, R. Verma, L. M. J. Florack, and C.-F. Westin. *Structural connectivity analysis using Finsler geometry*. **In preparation, preprint available upon request.**

C. M. W. Tax, C.-F. Westin, **T. C. J. Dela Haije**, A. Fuster, M. A. Viergever, E. Calabrese, L. M. J. Florack, and A. Leemans. "Quantifying the brain's sheet structure with normalized convolution". In: *Medical Image Analysis* (2017). **Accepted for publication, preprint available upon request.**

R. Duits, A. Ghosh, **T. C. J. Dela Haije**, and A. Mashtakov. "On sub-Riemannian geodesics in SE(3) whose spatial projections do not have cusps". In: *Journal of Dynamical and Control Systems* (July 2016). ISSN: 1079-2724, 1573-8698. DOI: `10.1007/s10883-016-9329-4`.

A. Fuster, **T. Dela Haije**, A. Tristán-Vega, B. Plantinga, C.-F. Westin, and L. Florack. "Adjugate diffusion tensors for geodesic tractography in white matter". In: *Journal of Mathematical Imaging and Vision* 54.1 (2016), pp. 1–14. ISSN: 0924-9907, 1573-7683. DOI: `10.1007/s10851-015-0586-8`.

C. M. Tax, **T. Dela Haije**, A. Fuster, C.-F. Westin, M. A. Viergever, L. Florack, and A. Leemans. "Sheet Probability Index (SPI): Characterizing the geometrical organization of the white matter with diffusion MRI". In: *NeuroImage* 142 (Nov. 2016), pp. 260–279. ISSN: 10538119. DOI: `10.1016/j.neuroimage.2016.07.042`. **Joint first authors Tax and Dela Haije.**

L. Astola, N. Sepasian, **T. Dela Haije**, A. Fuster, and L. Florack. "A simplified algorithm for inverting higher order diffusion tensors". In: *Axioms* 3.4 (Nov. 2014), pp. 369–379. ISSN: 2075-1680. DOI: `10.3390/axioms3040369`.

R. Duits, **T. C. J. Dela Haije**, E. J. Creusen, and A. Ghosh. "Morphological and linear scale spaces for fiber enhancement in DW-MRI". In: *Journal of Mathematical Imaging and Vision* (Nov. 2012). ISSN: 0924-9907, 1573-7683. DOI: `10.1007/s10851-012-0387-2`.

Book chapters

L. Florack, **T. Dela Haije**, and A. Fuster. *Cartan scalars in Finsler-DTI for higher order local brain tissue characterization*. **In preparation, preprint available upon request.**

T. Dela Haije, N. Sepasian, A. Fuster, and L. Florack. "Adaptive enhancement in diffusion MRI through propagator sharpening". In: *Computational Diffusion MRI*. Ed. by A. Fuster, A. Ghosh, E. Kaden, Y. Rathi, and M. Reisert. Cham: Springer International Publishing, 2016, pp. 131–143. ɪsʙɴ: 978-3-319-28586-3 978-3-319-28588-7. **Joint first authors Dela Haije and Sepasian.**

T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. "Finslerian diffusion and the Bloch-Torrey equation". In: *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*. Ed. by I. Hotz and T. Schultz. Cham: Springer International Publishing, 2015, pp. 21–35. ɪsʙɴ: 978-3-319-15089-5 978-3-319-15090-1.

L. Florack, **T. Dela Haije**, and A. Fuster. "Direction-controlled DTI interpolation". In: *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data*. Ed. by I. Hotz and T. Schultz. Cham: Springer International Publishing, 2015, pp. 149–162. ɪsʙɴ: 978-3-319-15089-5 978-3-319-15090-1.

T. C. J. Dela Haije, R. Duits, and C. M. W. Tax. "Sharpening fibers in diffusion weighted MRI via erosion". In: *Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data*. Ed. by C.-F. Westin, A. Vilanova, and B. Burgeth. Mathematics and Visualization. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 97–126. ɪsʙɴ: 978-3-642-54300-5 978-3-642-54301-2.

R. Duits, A. Ghosh, **T. Dela Haije**, and Y. Sachkov. "Cuspless sub-Riemannian geodesics within the Euclidean motion group SE(d)". In: *Neuromathematics of Vision*. Ed. by G. Citti and A. Sarti. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 173–215. ɪsʙɴ: 978-3-642-34443-5 978-3-642-34444-2.

A. Fuster, A. Tristán-Vega, **T. Dela Haije**, C.-F. Westin, and L. Florack. "A novel Riemannian metric for geodesic tractography in DTI". In: *Computational Diffusion MRI and Brain Connectivity*. Springer, 2014, pp. 97–104.

E. J. Creusen, R. Duits, and **T. C. J. Dela Haije**. "Numerical schemes for linear and non-linear enhancement of DW-MRI". In: *Scale Space and Variational Methods in Computer Vision*. Springer, 2012, pp. 14–25.

R. Duits, **T. C. J. Dela Haije**, A. Ghosh, E. J. Creusen, A. Vilanova, and B. M. ter Haar Romeny. "Fiber enhancement in diffusion-weighted MRI". In: *Scale Space and Variational Methods in Computer Vision*. Springer, 2012, pp. 1–13.

Conference contributions

T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. "A new fiber orientation distribution function". In: *Proceedings of the 25th Annual Meeting of the ISMRM*. Honolulu, HI, 2017, p. 3368. **Stipend. Accepted for publication, preprint available upon request.**

M. H. J. Janssen, **T. C. J. Dela Haije**, F. C. Martin, E. J. Bekkers, J. Olivan Bescos, and R. Duits. *The Hessian of axially symmetric functions on* SE(3) *and application in 3D image analysis*. 2017. **Accepted for publication, preprint available upon request.**

T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. "The apparent range of spin movement in diffusion MRI data". In: *Proceedings of the 24th Annual Meeting of the ISMRM*. Singapore, 2016, p. 2084. **Stipend. Selected for presentation at the Diffusion Study Group Session.**

T. C. J. Dela Haije, C. M. W. Tax, A. Fuster, C.-F. Westin, M. A. Viergever, L. M. J. Florack, and A. Leemans. "Does sheet happen? Mapping the brain's "sheet probability index" with diffusion MRI". In: *Abstract Book of the 8th Annual ISMRM Benelux Chapter Meeting*. Eindhoven, 2016, p–040.

C. M. W. Tax, **T. C. J. Dela Haije**, A. Fuster, C.-F. Westin, M. A. Viergever, L. M. J. Florack, and A. Leemans. "Mapping the brain's "sheet probability index" (SPI) with diffusion MRI: Sheet happens?!" In: *Proceedings of the 24th Annual Meeting of the ISMRM*. Singapore, 2016, p. 0791. **Joint first authors Tax and Dela Haije. Magna Cum Laude Award.**

C. M. W. Tax, C.-F. Westin, **T. C. J. Dela Haije**, A. Fuster, M. A. Viergever, L. M. J. Florack, and A. Leemans. "Robust assessment of the brain's sheet structure using normalized convolution". In: *Proceedings of the 24th Annual Meeting of the ISMRM*. Singapore, 2016, p. 2073.

T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. "Reconstruction of convex polynomial diffusion MRI models using semi-definite programming". In: *Abstract Book of the 7th Annual ISMRM Benelux Chapter Meeting*. Ghent, 2015.

T. C. J. Dela Haije, A. Fuster, and L. M. J. Florack. "Reconstruction of convex polynomial diffusion MRI models using semi-definite programming". In: *Proceedings of the 23rd Annual Meeting of the ISMRM*. Toronto, 2015, p. 2821. **Stipend.**

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, and A. Leemans. "Towards quantification of the brain's sheet structure in diffusion MRI data". In: *Proceedings of the International Biomedical and Astronomical Signal Processing (BASP) Frontiers workshop*. Villars-sur-Ollon, 2015, p. 74. **Joint first authors Tax and Dela Haije.**

C. M. W. Tax, **T. C. J. Dela Haije**, A. Fuster, R. Duits, M. A. Viergever, L. M. J. Florack, and A. Leemans. "Towards quantification of the brain's sheet structure: Evaluation of the discrete Lie bracket". In: *Proceedings of the 22nd Annual Meeting of the ISMRM*. Milan, 2014, p. 0975. **Magna Cum Laude Award.**

R. H. J. Fick, R. Duits, E. J. Creusen, **T. C. J. Dela Haije**, P. P. W. Ossenblok, B. M. ter Haar Romeny, and A. Vilanova. "Contextual enhancements on DW-MRI". In: *HARDI reconstruction challenge 2013 proceedings*. San Francisco, Apr. 2013, p. 24.

L. M. J. Florack, A. Fuster, and **T. C. J. Dela Haije**. "Riemann-Finsler geometry and its applications to diffusion magnetic resonance imaging". In: *Proceedings of the International Biomedical and Astronomical Signal Processing (BASP) Frontiers workshop*. Villars-sur-Ollon, Jan. 2013, p. 61.

A. Fuster, **T. C. J. Dela Haije**, and L. M. J. Florack. "On the Riemannian rationale for diffusion tensor imaging". In: *Proceedings of the International Biomedical and Astronomical Signal Processing (BASP) Frontiers workshop*. Villars-sur-Ollon, Jan. 2013, p. 62.

A. Fuster, A. Tristán-Vega, **T. C. J. Dela Haije**, C.-F. Westin, and L. M. J. Florack. "A novel Riemannian metric for geodesic tractography in DTI". In: *CDMRI 2013*. Nagoya, Sept. 2013, pp. 47–54.

C. M. W. Tax, **T. C. J. Dela Haije**, A. Fuster, M. A. Viergever, L. M. J. Florack, and A. Leemans. "Considerations on the theory of sheet structure of cerebral pathways". In: Podstrana, Oct. 2013.

T. C. J. Dela Haije, S. Makram-Ebeid, J. Rouet, M. Fradkin, and C. Ciofolo-Veit. "Onset detection in first-pass perfusion MRI". In: *Nederlandse Vereniging voor Patroonherkenning en Beeldverwerking (NVPHBV) Meeting Fall 2010 - Abstracts*. Eindhoven, Oct. 2010.

Reports and theses

T. C. J. Dela Haije. "Geometrical methods for enhancement and tractography in diffusion MRI". MSc Thesis. Eindhoven: Eindhoven University of Technology, Apr. 2012. **Cum Laude.**

T. C. J. Dela Haije. *Erosion on $\mathbb{R}^3 \rtimes S^2$ and its application to sharpen fibers*. Internship report. Eindhoven: Eindhoven University of Technology, Aug. 2011.

R. Duits, E. J. Creusen, A. Ghosh, and **T. C. J. Dela Haije**. *Diffusion, convection and erosion on $\mathbb{R}^3 \rtimes S^2$ and their application to the enhancement of crossing fibers*. Technical Report 18. Eindhoven: Eindhoven University of Technology, 2011.

T. C. J. Dela Haije. *Onset detection in myocardial first-pass perfusion magnetic resonance imaging*. Internship report. Suresnes: Philips Healthcare, Nov. 2010.