

Reconstruction of Three-Dimensional Nerve Fiber Orientations from Histological Brain Sections in Three-Dimensional Polarized Light Imaging

Dissertation zur Erlangung des Doktorgrades:

Dr. rer. nat.

Fakultät für Mathematik und Naturwissenschaften
Bergische Universität Wuppertal

vorgelegt von:

Daniel Schmitz

März 2020

Prof. Dr. Dr. Thomas Lippert

Dr. Markus Axer

The PhD thesis can be quoted as follows:

urn:nbn:de:hbz:468-20200921-090438-8

[<http://nbn-resolving.de/urn/resolver.pl?urn=urn%3Anbn%3Ade%3A468-20200921-090438-8>]

DOI: 10.25926/15kc-t577

[<https://doi.org/10.25926/15kc-t577>]

Abstract

Three-Dimensional Polarized Light Imaging enables to reconstruct nerve fiber orientations from polarimetric measurements of histological brain sections based on the birefringence of myelinated nerve fibers. Measurements of the brain sample from oblique views facilitate the retrieval of the three-dimensional orientation of the nerve fibers as well as the birefringence strength. The enhancement of this noise sensitive reconstruction represents the major motivation for this thesis. First, a novel and fast reconstruction algorithm based on a least squares approach is introduced. At the mesoscale an unprecedented comprehensive view of the human brain's nerve fiber tracts is obtained based on the new reconstruction algorithm. Next, the reconstruction is further improved for very low signals utilizing a Bayesian estimator. Furthermore, as the parameter estimation is sensitive to noise, the uncertainty of the obtained parameter maps is studied. This results in the first nerve fiber orientation confidence estimates for Three-Dimensional Polarized Light Imaging. Finally, the developed analysis is applied at the microscale. A validation against higher-resolved volumetric measurements shows that individual fiber bundles can be reconstructed with high accuracy using the developed algorithms.

Zusammenfassung

Das bildgebende Verfahren Three-Dimensional Polarized Light Imaging (3D-PLI) ermöglicht die Rekonstruktion von Nervenfaserorientierungen mithilfe polarimetrischer Messungen histologischer Gehirnschnitte basierend auf der Doppelbrechung myelinisierter Nervenfasern. Aufnahmen des Hirnschnitts von schrägen Perspektiven erlauben die Abschätzung der dreidimensionalen Faserorientierung sowie der Doppelbrechungsstärke. Die Verbesserung dieser rauschsensitiven Abschätzung ist das Ziel der vorliegenden Arbeit. Zuerst wird ein neuartiger Rekonstruktionsalgorithmus basierend auf dem Verfahren der kleinsten Quadrate entwickelt. Dieser ermöglicht auf der Mesoskala eine umfassende Darstellung der menschlichen Nervenfaserarchitektur. Zur Verbesserung der Rekonstruktion für sehr kleine Signale wird ein bayesscher Schätzer entwickelt. Weiterhin wird die Unsicherheit der rekonstruierten Gewebeeigenschaften untersucht. Dies ermöglicht zum ersten Mal die Abschätzung der Verlässlichkeit der in 3D-PLI ermittelten Nervenfaserorientierung. Schließlich werden die entwickelten Analysemethoden auf mikroskopische 3D-PLI Messungen mit schrägem Lichteinfall angewandt. Ein Vergleich mit höher aufgelösten volumetrischen Messungen zeigt, dass die Orientierung einzelner Nervenfaserbündel mit den entwickelten Algorithmen mit hoher Genauigkeit rekonstruiert werden können.

Contents

1	The challenge of investigating the human brain's nerve fiber tracts	1
2	Foundations of Three-Dimensional Polarized Light Imaging	5
2.1	Polarization and birefringence	5
2.2	Optical properties of brain tissue	7
2.3	Preparation of histological brain sections	9
2.4	Polarimetric setups	9
2.5	The biophysical model of 3D-PLI	11
2.6	Established signal analysis techniques	13
2.7	A review of existing approaches for the analysis of birefringence experiments from oblique views	17
3	Foundations of Bayesian data analysis	21
3.1	Parameter estimation based on Bayes' theorem	21
3.2	Uncertainty quantification	25
4	A Least Squares Approach for the derivation of Nerve Fiber Orientations	31
4.1	The Robust Orientation Fitting via Least Squares algorithm	32
4.2	Evaluation for synthetic data	35
4.3	Evaluation for experimental data	40
4.4	Development of a high performance implementation of the fitting algorithm	48
5	Maximum A Posteriori Estimation of Nerve Fiber Orientations	55
5.1	The Posterior Orientation Solver	55
5.2	Evaluation on synthetic data	59
5.3	Evaluation on experimental data	61
6	Uncertainty estimation in 3D-PLI	67
6.1	Planar 3D-PLI measurement	68
6.1.1	Theory and implementation	68
6.1.2	Validation on synthetic data	69
6.1.3	Validation on experimental data	71
6.2	Oblique 3D-PLI measurement	74
6.2.1	Theory and implementation	74
6.2.2	Validation on synthetic data	76
6.2.3	Validation on experimental data	79

7	Measurements from Oblique Views in Microscopic 3D-PLI	87
7.1	Experimental characterization of the LMP3D	87
7.2	Application of the Bayesian framework at the microscale	89
7.3	Validation of 3D-PLI fiber orientations based on TPFM measurements .	90
8	Conclusion and Outlook	97
	Acknowledgments	101
	Appendix A Mathematical derivations for the ROFL algorithm	103
A.1	Rotation matrices of the tilted measurements	103
A.2	Error propagation from camera noise to the normalized light intensity .	103
A.3	Gradient of the objective function	104
	Appendix B Volumetric reconstruction of the brain sections	107
	Appendix C System configuration for the runtime benchmarks	111
	Appendix D Optimization experiments for the LOriE algorithm	113
	Appendix E Complementary data	115
	List of Symbols	119
	List of Abbreviations	121
	List of Figures	123
	List of Tables	127
	References	129

1

The challenge of investigating the human brain's nerve fiber tracts

“Moore’s Law-based technology is so much easier than neuroscience. The brain works in such a different way from the way a computer does.”

– Paul Allen

Advancing our understanding of the human brain is of high significance for progress not only in fundamental and clinical neuroscience but also in technical fields. As pointed out by the World Health Organization neurological disorders pose a public health risk [1]. Despite substantial efforts of the scientific community only symptomatic treatments are available so far for many neurological disorders such as Parkinson’s disease. Outside of classical neuroscience, adapting principles of neural learning mechanisms lead to great advances in the field of machine learning [2, 3] and fuels the development of neuromorphic computing technology [4, 5]. Therefore the importance of fundamental research into the working principles of the brain cannot be overestimated.

In recent years, large research consortia have been formed around the globe to address this challenge. Besides the European Union funded *Human Brain Project* [6] six countries have initiated similar programs: Australia (*Australian Brain Alliance*) [7], Canada (*Canadian Brain Research Strategy*) [8], China (*China Brain Project*) [9], Japan (*Brain/MINDS*) [10], South Korea (*Korean Brain Initiative*) [11] and the US (*Brain Initiative*) [12]. An overview of the different programs and strategies to advance neuroscience worldwide is given in [13].

This study is a part of the efforts for large scale *human brain mapping*. Brain mapping describes the attempt to divide the brain into regions of distinct functional and structural features. At the microscale, the brain consists of cells (*neurons*) which exchange information via small electrical signals. The network of nerve fibers connecting these cells is called *connectome* [14]. It is well known that specific areas of the brain are more active during certain tasks than others. Studying the connections between these areas, the connectome, can therefore give further insights into the brain’s function. A bundle of nerve fibers which connects a larger number of neurons is called *fiber tract*. In a clinical context, the investigation of the connectome is also highly relevant as knowledge about major fiber tract courses is critical for surgery planning [15]. The study of the connectome though represents a multiscale challenge from the scale of single nerve fibers with a diameter below 1 μm to the whole human brain at the scale of several cm. As

a consequence, not one single imaging technique is sufficient for a coherent mapping of the connectome alone. Instead, multiple imaging techniques which each have their own strengths and pitfalls have been developed to map the connectome at different resolutions.

So far, the only way to study nerve fibers in vivo is via diffusion magnetic resonance imaging (DMRI) based on probing the diffusion of water with strong magnetic gradient fields [16, 17, 18]. The most common model for the diffusion process, Diffusion Tensor Imaging, yields a tensor description of the water diffusion in every voxel (the volumetric analogue of the image pixel) [16]. Presuming that water primarily diffuses along the dominant nerve fiber direction, this results in a three-dimensional map of the fiber pathways of the brain. As the derived fiber orientation gives the tangential orientation of a fiber tract, spatial integration over the fiber orientations results in reconstructions of whole fiber tracts [16]. The reconstruction of fiber tracts has been coined *tractography*. More sophisticated descriptions of the diffusion process enable to extract a statistical distribution of fiber orientations [19] and to infer microstructural tissue properties such as an apparent fiber density [20]. In clinical applications DMRI has been successfully used for the diagnosis of stroke, monitoring of treatment response and treatment planning [21]. Some of the concepts adopted in this thesis, such as using Bayesian inference to analyze nerve fiber orientations, have already been applied to DMRI [22]. Still the resolution of DMRI is limited to millimeters in clinical practice and $200\text{-}300\text{ }\mu\text{m}^3$ in post mortem brains due to scan time restrictions [23, 24, 25] and therefore only able to capture the "brain's highways". An influential recent study by Hein et al. [26] proved that it is a priori impossible to derive correct fiber tracts from DMRI studies purely based on local fiber orientation information. The major drawback of current DMRI tractography is the high number of reconstructed fiber tracts which in fact do not exist [26]. Thus, for deeper insights into the microstructure of the brain such as individual nerve fibers bundles a higher resolution is necessary.

Revealing the composition of individual nerve cells and neural circuits requires a resolution below $1\text{ }\mu\text{m}$ which can only be provided by microscopic techniques. Especially electron microscopy [27], traditional fluorescence microscopy [28] and in recent years X-ray diffraction [29] have been utilized to investigate the brain at this scale. An obvious drawback of these methods is that the brain has to be fixated, sectioned and then reconstructed to a three-dimensional volume as they do not provide depth resolved scanning of the tissue. Another difficulty are excessive scanning times as the microscopic field of view is very small compared to an entire human brain section covering an area up to 150 cm^2 . With respect to the reconstruction of fiber pathways, histological imaging techniques suffer from a lack of information in the direction perpendicular to the sectioning plane as only the two dimensional projection of the three dimensional structures are captured by the camera. While the projection of the fiber courses into the sectioning plane is directly visible at high resolution, the out-of-plane component of the fiber orientation is typically not accessible. Hence, mapping the connectome at high resolution necessitates even more sophisticated measurements and data analytics.

An established technique to investigate fiber tracts in three dimensions at highest resolution is neural tracing based on the injection of a tracer molecule and the later imaging of the location of the tracer molecule [30]. This comes with the disadvantage that the tracer injection procedure and the propagation of the tracer inside the tissue vary strongly over subjects [30]. Another method which emerged in recent years relies on image processing to derive three-dimensional fiber pathways: if nerve fibers are easily identifiable in an

image stack (for example after staining), calculating the three-dimensional gradient of the volume using the structure tensor yields an approximate orientation of the nerve fibers [31, 32]. In contrast to DMRI, this approach does not utilize biophysical models anymore but image processing heuristics which stresses its limitations. A coherent gradient calculation requires a well-registered image stack from the original histological sections which is in general still an unsolved problem in image registration [33] (*image registration* refers to the matching of one image onto another) and also strongly depends on the chosen kernel size and pre-processing steps [32].

Sectioning a human brain into histological sections results in 5,000 - 10,000 sections depending on the sectioning plane and the section thickness. As automated sectioning of human brains could not be achieved so far due to the fragility of the sections the sectioning itself already represents a very labour intensive undertaking even before imaging of the individual sections. To overcome this limitation or to at least reduce the amount of manual work, depth resolved microscopic techniques such as Light Sheet Microscopy [34, 35], Optical Coherence Tomography [36, 37] and Two-Photon Fluorescence Microscopy (*TPFM*) [38, 39] were developed. Still these techniques are not able to scan into the depth of a whole brain due to physical limitations of microscopic lenses, but increase the theoretically possible section thickness to the order of a centimeter [40]. The actual possible section thickness though depends on attenuation and scattering effects of the investigated tissue and the employed setup [41]. For small brains such as the drosophila fly the full connectome could for example be reconstructed from TPFM measurements [42] resulting in high-resolved three-dimensional optical properties of the brain tissue. Another approach to avoid histological sectioning is to chemically make brain tissue transparent with the CLARITY protocol [43], thereby making the structures of interest more accessible. Still, reconstruction of the connectome at large scale which is required for human brains has not been accomplished so far using these depth resolving techniques.

In this thesis the imaging technique Three-Dimensional Polarized Light Imaging (*3D-PLI*) introduced in [44, 45] is studied. 3D-PLI utilizes polarimetric measurements through unstained brain sections to derive the spatial orientation of nerve fibers based on their birefringence. The theoretical background and experimental realization of 3D-PLI are presented in chap. 2. As a microscopic technique, the lateral resolution of 3D-PLI is theoretically limited by the illumination wavelength while the z -resolution is determined by the section thickness. Compared to the microscopic techniques discussed before, 3D-PLI does not primarily employ image processing to derive its connectome but a biophysical model. In that sense, 3D-PLI data analytics are more similar to DMRI data analytics than for example staining methods and structure tensor analysis. The biophysical model behind 3D-PLI presumes a nerve fiber orientation vector in every pixel which is subsequently derived. While the in-plane fiber orientation can easily be calculated based on this model, inferring the out-of-plane orientation is more challenging. Using additional information gained from measurements from oblique views, the out-of-plane orientation can in principle be derived as shown for the first time in [46]. In the context of [46] also a new microscopic setup was developed which facilitates measurements from oblique views at micrometer resolution which was only possible at mesoscopic resolution before. As mentioned before, the ability to derive three-dimensional fiber tract information is the pitfall of histological imaging techniques concerning the connectome. Thus, 3D-PLI has great potential for large scale connectome mapping at high resolution. So far, 3D-PLI has been employed for qualitative neuroanatomical studies of avian, rodent, monkey and human brains [47, 48, 49]. However, further quantitative analysis of the

connectome based on 3D-PLI data is still challenging due to the difficulties of inferring the three-dimensional fiber orientation and the required volumetric reconstruction of the individual sections. This study seeks to advance 3D-PLI by enhancing the derivation of the fiber orientation.

Improving the nerve fiber reconstruction in 3D-PLI requires a careful statistical analysis of the oblique measurements. Without the additional oblique measurements, 3D-PLI measurements can be analyzed based on analytical calculus to derive the model parameters. On the other hand, the derivation of the out-of-plane orientation exploits small variations in the signal captured from different oblique angles and is therefore prone to noise effects. In [46] an algorithm was devised which numerically fits the model to the measurement data in every image pixel. As the model is highly nonlinear this comes with the mathematical difficulties arising from fitting nonlinear models and the cost of excessive computation time as the fitting procedure has to be carried out in millions of pixels. Hence, 3D-PLI faces big data problems due to the sheer amount of necessary computations and limitations in the analysis such as a lack of reliability measures for the model parameters. These challenges arise in other fields of physics such as high energy particle physics and astronomy as well [50, 51] and led, for instance, to the establishment of astrophysics as a new discipline in astronomy [52, 53]. In recent years, a clear trend for tackling these problems in physics as well as in neuroscience is to apply Bayesian approaches [54, 55, 56] which was also adopted in this study. An introduction to the Bayesian methods relevant for this thesis is given in chap. 3.

The estimation of nerve fiber orientations in 3D-PLI represents a nonlinear regression problem. Although a huge number of parameter estimation techniques for nonlinear regression have been developed, the most common technique used is still least squares fitting due to its simplicity [57]. In chap. 4 a least squares estimator for 3D-PLI is developed and evaluated on synthetic data and experimental data of a human brain at the macroscale. In a next step, the novel algorithm is accelerated using high performance computing (*HPC*) resources to scale to big data.

One advantage of Bayesian approaches is that they allow to incorporate prior information. Especially, in cases in which the observed data does not allow a precise determination of the model parameters, prior information has a strong effect. This is applied in chap. 5 to 3D-PLI to improve the reconstruction for regions of low birefringent signals and evaluated on the same datasets as the least squares approach. So far, only the most likely fiber orientation was derived in 3D-PLI. As especially the determination of the out-of-plane orientation is sensitive to noise, it is necessary to assess the reliability of the orientation estimation. Thus, in chap. 6 orientation confidence measures are derived for 3D-PLI based on a Bayesian approach and evaluated based on synthetic data and selected experimental datasets.

In [46], it was shown in a proof of concept experiment that three-dimensional fiber orientations can be derived from measurements with oblique views using the novel prototypical microscopic setup. Still, no phantom has been developed for 3D-PLI so far which limits the possibilities for validation. Another possible validation is the comparison to higher resolved measurements which enable to directly visualize nerve fibers in 3D. In chap. 7 the fiber orientations calculated from 3D-PLI with the developed algorithms at microscopic resolution are validated against a ground truth obtained from TPFM measurements.

A conclusion and outlook for future directions are given in chap. 8.

2

Foundations of Three-Dimensional Polarized Light Imaging

“Physics is really nothing more than a search for ultimate simplicity, but so far all we have is a kind of elegant messiness.”

– Bill Bryson

In this chapter the fundamental principles and measurement techniques of the imaging technique 3D-Polarized Light Imaging are presented. After a review of the theoretical background of polarization and birefringence their connection to nerve fibers is presented. Based on these foundations, the preparation of brain sections and their measurement in 3D-PLI is illustrated. Next, the current state of the art for the analysis of 3D-PLI measurements is explained. The difficulties of the current analysis are finally reviewed which leads to the new methods developed in this work.

2.1 Polarization and birefringence

Polarization characterizes the oscillation plane of transversal waves. Natural light, for example, oscillates randomly in all directions, it is therefore *unpolarized*. Consider the electric field vector of an electromagnetic wave propagating in z -direction [58]:

$$\mathbf{E}(t) = \mathbf{E}_0 e^{i(kz - \omega t)} = \begin{pmatrix} E_x e^{i\Psi_x} \\ E_y e^{i\Psi_y} \\ 0 \end{pmatrix} e^{i(kz - \omega t)} . \quad (2.1)$$

where k denotes the wave number and ω the frequency of the wave. The magnitudes of the field components E_x and E_y and the phases Ψ_x and Ψ_y now determine the polarization state of the wave. The vector $\mathbf{E}_0$ is called *Jones vector*. A wave for which the Jones vector does not change its direction is called *linearly polarized*. Linear polarization occurs if both components of the electrical field oscillate in phase ($\Psi_x = \Psi_y$). The case that $\mathbf{E}_0$ describes a circle is called *circular polarization* and arises for $E_x = E_y$, $\Psi_x - \Psi_y = \frac{\pi}{2}$. The most general case is called *elliptical polarization*: the electric field vector describes an ellipse ($E_x \neq E_y$, $\Psi_x - \Psi_y = \frac{\pi}{2}$).

Polarization microscopy techniques are based on the interaction of polarized light with the specimen, in the case of 3D-PLI brain tissue. Different materials interact differently with incoming light due to their structure at the molecular level which results in different

forces interacting with the incoming photons. In general, incoming light can be absorbed, scattered and refracted. One immediate consequence is that in a material light does not propagate with the same speed as in vacuum. The phase velocity in a material v depends on the material specific *refractive index* n by $v = c/n$ with the speed of light in vacuum c . At the boundary of two ideal isotropic materials with refractive indices n_1 and n_2 , a light ray with an incident angle θ_1 is refracted according to Snell's law: $n_1 \sin(\theta_1) = n_2 \sin(\theta_2)$.

Optically anisotropic media display a more complex refraction behaviour. If the molecular structure of a material results in anisotropic forces interacting with incoming light, the refractive index becomes anisotropic as well. A prominent example are crystals which regular lattice structure results in strongly orientation dependent atomic forces. Depending on the oscillation plane of incoming light, its polarization, it then experiences different interactions. Thus, polarized light can be utilized to investigate the inner structure that cause the optical anisotropy of the material.

Geometrically, the anisotropy of the refractive index can be represented by an ellipsoid which describes the refractive index for any propagation direction (n_x, n_y, n_z) as a function of the principal refractive indices of the material:

$$\frac{n_x^2}{n_1^2} + \frac{n_y^2}{n_2^2} + \frac{n_z^2}{n_3^2} = 1 . \quad (2.2)$$

Two types of optical anisotropy are distinguished: uniaxial birefringence for the case that $n_1 = n_2 \neq n_3$ and biaxial birefringence in case that $n_1 \neq n_2 \neq n_3$. As biaxial birefringence is not relevant for 3D-PLI, we will focus on uniaxial birefringence. The two equal indices are called *ordinary* refractive index $n_o = n_1 = n_2$, while the third index is called the *extraordinary* refractive index $n_E = n_3$. Their difference $\Delta n = n_E - n_o$ is often termed *birefringence* and is a measure of the amount of optical anisotropy.

For a uniaxial birefringent material the refractive index ellipsoid becomes a rotational ellipsoid with a principal axis called *optic axis*. The rotational ellipsoid is also called *indicatrix*. The polarization state and the propagation direction of the incoming light relative to the optic axis now determine the refractive index experienced by the wave (cf. Fig. 2.1). A wave which is polarized perpendicularly to the plane spanned by optic axis and wave vector (denoted by the blue ellipse) behaves as if it was interacting with an optically isotropic material: it obeys Snell's law with n_o as the material's refractive index independent of the wave's direction of propagation. It is therefore denoted as the *ordinary wave*. In case that the incoming light is not polarized perpendicularly to the optic axis, it experiences a refractive index which depends on the amount of birefringence and on the propagation direction of the incoming light. From Fig. 2.1 it can be seen that the cross section of the plane perpendicular to the wave vector and the ellipsoid define an ellipse whose half axes are given by n_o and the extraordinary refractive index $n_e(\alpha)$ ¹. n_e becomes maximal/minimal if the optic axis is oriented perpendicularly to the incoming light ($\alpha = 0^\circ$) for which then $n_e = n_E$. In case that the incident light and optic axis are parallel ($|\alpha| = 90^\circ$), the birefringence vanishes as $\Delta n = n_e(\alpha) - n_o = n_o - n_o = 0$. From the ellipsoidal geometry, the extraordinary refractive index n_e can be derived as [58]

$$n_e(\alpha) = \frac{1}{\sqrt{\frac{\sin^2(\alpha)}{n_o^2} + \frac{\cos^2(\alpha)}{n_E^2}}} . \quad (2.3)$$

¹The extraordinary refractive index of the material is denoted as n_E , while the orientation dependent index is denoted as n_e .




FIGURE 2.1: Indicatrix of a uniaxial positive birefringent material. The intersection plane of the incoming wave light vector $\mathbf{k}$ and the optic axis of the material is shaded in red. A wave polarized perpendicular to the optic axis experiences the refractive index n_o . For any other polarization state, the refractive index $n_e(\alpha)$ depends on the angle between $\mathbf{k}$ and the optic axis.

As ordinary and extraordinary rays propagate with different refractive indices, their optical paths vary from each other. Their phases after passing a birefringent medium of thickness t are given by $\Psi_o = \frac{2\pi n_o t}{\lambda}$ and $\Psi_e = \frac{2\pi n_e t}{\lambda}$. Hence, after passing the birefringent medium, a relative phase retardance δ between extraordinary and ordinary wave of

$$\delta = \Psi_e - \Psi_o = \frac{2\pi \Delta n t}{\lambda} . \quad (2.4)$$

occurs. This is exploited by *wave retarders* to induce a defined phase shift between both waves. For example a retarder which induces a 90° phase shift (*quarter wave retarder*) transforms linearly polarized light into circular polarized light if the polarization axis of the retarder makes an angle of $\pm 45^\circ$ with the linear polarizer. For very small birefringence values Δn , eq. (2.4) can be approximated by

$$\delta \approx \frac{2\pi t \Delta n}{\lambda} \cos^2(\alpha) . \quad (2.5)$$

This approximation has been verified based on simulation studies [59] and used in all 3D-PLI studies so far (e. g. [44, 45, 60]). The derivation is based on a Taylor Expansion of eq. (2.3) [61, 46].

Now that the concepts of polarization and birefringence are introduced, in the next section their relationship to brain tissue is explored after a short overview of relevant neuroanatomy.

2.2 Optical properties of brain tissue

On a large scale, the human brain can be divided into the cerebellum and the cerebrum which consists of two hemispheres. It consists of up to 100 billion nerve cells (*neurons*)




FIGURE 2.2: Myelinated axon. The birefringent myelin sheath produced by oligodendrocyte cells serves as an insulating layer. Modified from [62].

which are connected to 1000-10000 other cells [63, 64]. Some of the connections are up to 1 m long and extend from the brain until the spinal cord. As the cells send time critical information to each other via electrical signals, communication speed along these connections (*axons*) is key for a healthy brain. To increase the conduction speed, an insulating layer of lipids called *myelin sheath* wraps around a large number of axons. The myelin sheath is made out of concentric layers of lipids and originates from oligodendrocyte cells (cf. Fig. 2.2). Gaps in the myelin sheath, the *nodes of Ranvier*, allow to manipulate the transported electric signals in between cells.

The birefringence of nerve fibers was first described by Göthlin in [65]. Various studies have investigated the source of this effect. One source of birefringence are neurofilaments and microtubuli in the myelin sheath [66]. Yet, the strongest contributor to the nerve fiber's birefringence is the highly ordered structure of lipids in the myelin sheath [67]. On a microscale, the myelin sheath itself can be modelled as birefringent material with radially oriented optic axes. Simulation studies showed that this radial birefringence results in a negative birefringence at the macroscopic level whose axis is oriented into the direction of the nerve fiber bundle [59]. Experimental studies confirm this simulative result [68, 69]. The magnitude of the birefringence of the myelin sheath is still unknown. General studies about the birefringence of biologic matter resulted in birefringence strengths of $|\Delta n| = 10^{-3} - 10^{-2}$ [61, 70]. Compared to the refractive index of brain tissue which lies between $n_b = 1.3 - 1.5$ the amount of birefringence is so low that the dominant birefringent effect of nerve fibers is the phase retardation between ordinary and extraordinary ray. Based on these findings, the assumption that myelinated nerve fibers can be modeled as negative uniaxial birefringent material whose optic axis is oriented along the dominant nerve fiber orientation builds the core of the biophysical model employed by 3D-PLI to analyze the polarimetric measurements.

2.3 Preparation of histological brain sections

The treatment and the sectioning of the postmortem brains are subject of ongoing experiments of the INM-1 laboratory team and not identical for all investigated brains. The pipeline described here was used for the brains studied in this thesis.

To avoid degeneration of the brain, it must be removed from the skull as soon as possible after death. After extraction the brains are fixed with a 4 % solution of buffered formaldehyde. As cryoprotection, the brain is then immersed in a 20 % solution of glycerin and Dimethyl sulfoide (DMSO). After dipping the brain into isopentane for one hour, the tissue is frozen in a plastic bag at -80°C . The frozen brain is kept in the freezer until it is designated for sectioning. Before each sectioning procedure, a photo of the brain block, called *blockface image*, is taken to serve as a reference for the 3D reconstruction of the individual sections². ARTag markers serve as references for the individual blockface images [71]. The brain is sectioned by a cryostat microtome (*Polycut CM 3500, Leica, Germany*) at a temperature of about -50°C . So far, a steady quality of the obtained sections could be achieved for a sectioning thickness of $70\text{ }\mu\text{m}$. Each section is then mounted on cooled glass slides and embedded in a 20 % glycerin solution. Next, the tissue is covered by a cover glass and sealed with lacquer. To avoid the development of air bubbles, the sections are additionally weighted for several hours.

2.4 Polarimetric setups

After a short introduction about the general capabilities and limitations of the available measurement setups all of them are presented in more detail. Two custom made polarimeters are available for 3D-PLI routine measurements: the Large-Area-Polarimeter (*LAP*) which enables single-shot acquisition of whole brain sections at a minimal pixel size of $21 \times 21\text{ }\mu\text{m}^2$ and the Large Metripol (*LMP*)³ with a pixel size of $1.33 \times 1.33\text{ }\mu\text{m}^2$ and field-of-view of $2.7 \times 2.7\text{ mm}^2$ requiring time expensive tile-wise measurement of a whole brain section. Besides the big resolution difference, the polarimeters depicted in Fig. 2.3 differ in other aspects as well: the LAP is equipped with a tiltable specimen stage facilitating measurements of the sample from oblique views. Furthermore, in the LAP both polarization filters and the retarder are individually rotatable and removable from the light path which enables complementary measurements such as diattenuation measurements⁴ [61]. As the LMP is primarily designed for high throughput measurements, the required robustness and compactness to achieve a long lifetime of the hardware does not allow this feature: the filters are neither individually rotatable nor removable.

Combining the high resolution of the LMP with the tilting capability of the LAP was realized in a polarizing microscope with the ability of oblique illumination [46]. Oblique illumination is achieved by a movable aperture diaphragm which is placed off axis before the specimen. So far, a prototypic setup of this new microscope (*LMP3D*, Large Metripol 3D) is available at an optical bench. As only the rotating polarization filter is motorized and a scanning stage not feasible on an optical bench, measurements are still limited to individual field of views.

²Details about the blockface setup can be found in B which describes the 3D reconstruction of the histological sections based on the blockface images.

³The name is a tribute to the metripol technique which utilized polarization microscopy with a rotating polarizer to investigate crystals [72].

⁴Diattenuation describes the polarization dependent attenuation of light.




FIGURE 2.3: Employed polarimetric setups. For reasons of simplicity, all optical elements except the objective lenses and the aperture diaphragm for the LMP3D are neglected. At mesoscopic resolution, the LAP enables single-shot imaging of full brain sections. The tilting stage facilitates measurements from oblique views. The polarimetric components are reversed for the microscopic setups compared to the LAP. At microscopic scale, the movable aperture diaphragm enables measurements from oblique views which achieves the same effect as tilting the sample at mesoscopic scale. Image courtesy: N. Gales.

The measurement principle is the same for all setups: the rotatable filters (for LMP and LMP3D only the first polarizer) are rotated in equidistant steps (typically 10°) to sample the polarization angle $\rho \in [0^\circ, 180^\circ]$. For each rotation angle, the camera records an image which yields an image series.

Large Area Polarimeter The light source of the LAP is a LED panel (NSPG 510S manufactured by Nichia corporation) consisting of an array of 36×36 single LEDs illuminating an area of $300 \times 300 \text{ mm}^2$ at a wavelength of $\lambda = 529 \text{ nm}$. The LED panel is powered by a constant current source while its operating temperature is maintained by water cooling system to establish a constant brightness. Diffuser plates above the LED array are employed to homogenize the illumination field. The polarization filters produced by Jos. Schneider Optische Werke GmbH are made of polymer foils and cover the whole field of view with a diameter of 240 mm. After passing the linear polarizer and the quarter-wave retarder which makes an angle of 45° with the polarizer the light becomes circularly polarized and passes the brain tissue. The tissue sample is placed on a specimen stage which is tiltable by up to 8° around two orthogonal axis. Regarding the polarization of the incoming light, the brain tissue acts as an inhomogeneous wave retarder and changes the polarization state to elliptical. The second polarizer which is oriented orthogonally to the first polarizer then serves as an analyzer. Finally, the transmitted light is captured by a CCD color camera with 14 bit depth (AxioCam HRc, Zeiss, Germany). A variety of different objectives is available to be placed in front of the camera to optimize the setup for brain sections of different species (rodent brain

sections require a far smaller field of view than human brain sections) [61]. For the brain analyzed in this thesis the Lametar 2.8/25 objective (JENOP-TIK-Laser, Optik, Systeme GmbH) with a focal length of $f = 25$ mm which enables measurements with a field of view of 18×18 cm² and a pixel size of 64×64 μm^2 was employed.

Large metripol In the LMP (manufactured by Taorad GmbH), the sample is illuminated by light emitted from a white LED (IntraLED 2020+, Volti AG) whose wavelength spectrum is shifted by a band pass filter to $\lambda = 550 \pm 9$ nm to match the optimal working wavelength of the retarder. The light field is further homogenized by a Koehler illumination to ensure a uniform illumination [73]. The employed polarization filters are typical high quality polarization filters⁵. Only the polarizer is rotatable. Above the specimen, a microscopic objective (TU Plan Fluor EPI P, Nikon) is employed. After passing the specimen the light is analyzed by a quarter-wave retarder and a perpendicular polarization filter, which together act as a circular analyzer. The outgoing light is captured by a monochrome camera (Monochrome Retiga 4000R, QImaging) with a bit depth of 12. The field of view is 2.7×2.7 mm². For a tilewise scan of the sample neighbouring tiles are measured with an overlap to enable the stitching of the individual tiles as a postprocessing step.

Large metripol 3D The LMP3D was designed to enable measurements from oblique views at microscopic resolution. Compared to the LMP, the polarimetric components are set up in the same order but additional optical components are employed to enable oblique illumination. For an extensive description of the optical design the reader is referred to [46] chap. 5, here an overview is provided.

The light source (Prizmatix UHP-Mic-LED-520) emits light at a wavelength of 520 ± 36 nm. A windowed uniform diffuser (Thorlabs ED1-C20) is employed to create a divergent light source. To match the retarder, the wavelength spectrum is then shifted by a bandpass filter (Thorlabs FL532-10) to 532 ± 36 nm. The employed linear polarization filters (Thorlabs LPVISE100-A) were chosen for their high performance at the illumination wavelength. Before the specimen, a movable aperture diaphragm connected to a x-y-stage enables to choose the oblique illumination angle. Test on the optical bench revealed a maximum possible offset of the diaphragm of 3 mm while still achieving sufficient illumination of the field of view which corresponds to a maximal oblique illumination angle of 5.7° . Behind the specimen, the light is captured by an objective lens (Nikon Plan Apo $4 \times 0.2\text{NA}$) and an appropriate tube lens (Nikon CFI60). As in the LMP, the light is analyzed by a circular analyzer consisting of a linear polarizer and a quarter wave retarder (Newport 10RP34-532). Finally, light is recorded by a monochrome camera (SVS-vistek evo 4070 MFLGEC) with a sensor size of 15.2×15.2 mm² and 2048×2048 pixels. From the sensor size and the magnification of the objective lens the imaging parameters result in a field of view of 3.8×3.8 mm² and a pixel size of 1.8×1.8 μm^2 [46].

2.5 The biophysical model of 3D-PLI

As myelinated nerve fibers are uniaxially negative birefringent at the resolution relevant for 3D-PLI, they can be described as wave retarders in a first approximation. Thus, in 3D-PLI every image pixel is modelled as a wave retarder whose optic axis gives the

⁵Details about the employed filters are property of Taorad GmbH.

dominant fiber orientation. Inferring the optic axis from the 3D-PLI measurement is the main objective of this thesis.

Mathematically, changes in polarization can be described by the Jones calculus in matrix form [74]. The Jones calculus is limited to the description of fully polarized light. While this assumption holds for the LMP and the LMP3D due to the employed high quality polarization filters, it is not completely fulfilled by the LAP which consists of custom made filters of far greater size than typically manufactured sizes and a quarter-wave retarder whose wavelength does not exactly match the illumination wavelength. These limitations have been investigated in [46, 61, 75] based on the more sophisticated Müller-Stokes calculus [76, 77], yet a pixelwise characterization of all imperfections of the employed hardware is still standing out. Therefore, in this thesis, the Jones calculus was employed. In the following, the Jones calculus is introduced and applied to 3D-PLI to describe the light intensity recorded in the measurement. The derivation follows [59]⁶.

An electromagnetic wave propagating in z -direction can be written as

$$\mathbf{E}(t) = \begin{pmatrix} E_x(t) \\ E_y(t) \\ 0 \end{pmatrix} = \begin{pmatrix} E_x e^{i\psi_x} \\ E_y e^{i\psi_y} \\ 0 \end{pmatrix} e^{i(kz - \omega t)} \quad (2.6)$$

The *Jones vector* $\mathbf{J}$ as introduced in sec. 2.1 is then defined as the vector which describes the polarization state:

$$\mathbf{J} = \frac{1}{|\mathbf{E}|} \begin{pmatrix} E_x e^{i\psi_x} \\ E_y e^{i\psi_y} \end{pmatrix} \quad (2.7)$$

Polarization filters can now be written as 2×2 matrices. Linear polarizers in x and y direction only transmit the x or y component of the electrical field. Thus, their Jones matrices $\mathbf{P}_x$ and $\mathbf{P}_y$ are given by

$$\mathbf{P}_x = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad \text{and} \quad \mathbf{P}_y = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}. \quad (2.8)$$

The Jones matrix of a wave retarder rotated by an angle γ in counterclockwise direction which induces a phase shift δ along the fast axis $\mathbf{M}_\delta$ is given by

$$\begin{aligned} \mathbf{M}_\delta(\gamma) &= \mathbf{R}(\gamma) \cdot \mathbf{M}_\delta \cdot \mathbf{R}(-\gamma) \\ &= \begin{pmatrix} \cos(\gamma) & -\sin(\gamma) \\ \sin(\gamma) & \cos(\gamma) \end{pmatrix} \cdot \begin{pmatrix} e^{i\delta/2} & 0 \\ 0 & e^{-i\delta/2} \end{pmatrix} \cdot \begin{pmatrix} \cos(\gamma) & \sin(\gamma) \\ -\sin(\gamma) & \cos(\gamma) \end{pmatrix} \end{aligned} \quad (2.9)$$

where $\mathbf{R}$ denotes a rotation matrix to change the coordinate system from the laboratory frame to the retarder frame. The quarter-wave retarder of the LAP as a special case of the general wave retarder is rotated by -45° with respect to the axis of the first linear polarizer. Its Jones matrix $\mathbf{M}_{\lambda/4}$ is then given by

$$\mathbf{M}_{\lambda/4} = \mathbf{M}_{\frac{\pi}{2}} \left(-\frac{\pi}{4} \right) = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i \\ -i & 1 \end{pmatrix}. \quad (2.10)$$

In the 3D-PLI measurement, the polarizers and the quarter wave retarder are rotated by ρ . Here, for simplicity, the equivalent case that the brain tissue rotates by $-\rho$ is considered. In this convention, the brain tissue can also be described as a wave retarder

⁶The presented derivation is valid for the LAP. The resulting light intensity profile is equal for the LMP. For a detailed derivation it is referred to [78].

which is rotated by $\varphi - \rho$ with respect to the polarizer. ρ as the in-plane component of the optic axis equates to the in-plane nerve fiber orientation. The combined Jones matrix of brain tissue $\mathbf{M}_b$ is given by

$$\begin{aligned} \mathbf{M}_b(\rho, \varphi, \delta) &= \mathbf{M}_\delta(\varphi - \rho) \\ &= \begin{pmatrix} \cos(\varphi - \rho) & -\sin(\varphi - \rho) \\ \sin(\varphi - \rho) & \cos(\varphi - \rho) \end{pmatrix} \begin{pmatrix} e^{i\delta/2} & 0 \\ 0 & e^{-i\delta/2} \end{pmatrix} \begin{pmatrix} \cos(\varphi - \rho) & \sin(\varphi - \rho) \\ -\sin(\varphi - \rho) & \cos(\varphi - \rho) \end{pmatrix} \end{aligned} \quad (2.11)$$

The electrical field vector at the camera after passing the polarization filters and the brain tissue $\mathbf{E}_t$ is then obtained by multiplying all Jones matrices with the incoming electrical field which yields

$$\mathbf{E}_t(\rho, \varphi, \delta) = \mathbf{P}_y \cdot \mathbf{M}_b(\rho, \varphi, \delta) \cdot \mathbf{M}_{\lambda/4} \cdot \mathbf{P}_x \cdot \mathbf{E}. \quad (2.12)$$

Digital cameras capture the magnitude of the electrical field vector. The resulting sinusoidal light intensity profile is given by

$$I(I_T, \rho, \varphi, \delta) \propto |\mathbf{E}|^2 = \frac{I_T}{2} (1 + \sin(2(\varphi - \rho)) \sin(\delta)). \quad (2.13)$$

The new parameter I_T (*transmittance*) is a measure of light absorption. The other parameters hold the information about the nerve fiber modelled as a wave retarder: the direction angle $\varphi \in [0, \pi]$ represents the in-plane nerve fiber orientation while the phase retardation δ is a measure of birefringence and the out-of-plane orientation (cf. eq. (2.5)). The utilized coordinate system is depicted in Fig. 4.1. The light intensity profile from eq. (2.13) is called the *planar model* throughout the rest of this thesis to avoid confusion with the oblique model introduced in chap. 4.

2.6 Established signal analysis techniques

The 3D-PLI measurement records a light intensity profile for every pixel of each measured brain section. For the LAP, there are additional light intensity profiles for the tilted measurements. After the measurement, these light intensity profiles are analyzed pixelwise to generate maps of the model parameters. These parameter maps then have to be further processed to retain a three-dimensional brain volume.

Calibration As a first step, all data are calibrated. This is necessary to correct for multiplicative errors in the measurement and inhomogeneities of the illumination over the field of view. For that purpose, a series of typically 60 images without tissue sample are recorded per rotation angle. From these 60 images, an average image is computed. The measurement data with a sample are later divided by this average image and multiplied by the mode value of all calibration images [79].

Fourier analysis In a next step, the parameters direction, retardation and transmittance need to be estimated from the measurement data. While it would be possible to perform model fitting via least squares in every image pixel, this approach would be computationally very intensive. As the model is sinusoidal, a




FIGURE 2.4: 3D-PLI coordinate system with fiber orientation vector $\mathbf{r}$, direction angle φ and inclination angle α .

Fourier analysis enables to fit the parameters analytically. Consider a set of N light intensities I_i with index i denoting the rotation angle. Then the Fourier coefficients up to first order are given by

$$\begin{aligned} a_0 &= \frac{1}{N} \sum_{i=0}^{N-1} I_i = \frac{I_T}{2}, \\ a_1 &= \frac{2}{N} \sum_{i=0}^{N-1} I_i \cos(\rho_i) = -\frac{I_T}{2} \sin(\delta) \sin(2\varphi), \\ b_1 &= \frac{2}{N} \sum_{i=0}^{N-1} I_i \sin(\rho_i) = \frac{I_T}{2} \sin(\delta) \cos(2\varphi). \end{aligned} \quad (2.14)$$

Rearranging these leads to the following formulas for the model parameters:

$$\begin{aligned} I_T &= 2a_0, \\ \varphi &= \frac{1}{2} \arctan\left(-\frac{a_1}{b_1}\right), \\ \sin(\delta) &= \frac{\sqrt{a_1^2 + b_1^2}}{a_0}. \end{aligned} \quad (2.15)$$

The light intensities of one exemplary pixel and the model fit are depicted in Fig. 2.5. In the sinusoidal model, the transmittance parameter captures the average light intensity while the retardation is given by the relative amplitude of the sinusoidal signal. The physical interpretation of the direction angle parameter becomes clear from the connection between the light intensity curve and the indicatrix: for incident light which is polarized by an angle φ in the xy -plane the light oscillates perpendicular to the optic axis and experiences no phase shift. In the sinusoidal light intensity the phase shift manifests as the offset of the sine curve.

Derivation of the inclination angle While the computation of the direction angle is straight forward based on this calculus, the derivation of the inclination angle is more difficult as it is concealed in the retardation formula (cf. eq. 2.5). As only the combined effect of section thickness, birefringence strength and wavelength are relevant for 3D-PLI, Axer et al. [45] introduced the *relative section thickness* $t_{\text{rel}} = 4t\Delta n/\lambda$ which simplifies the retardation formula to

$$\delta = \frac{\pi}{2} t_{\text{rel}} \cos(\alpha)^2. \quad (2.16)$$

Without a priori information the birefringence parameter t_{rel} and inclination are mapped onto one value and cannot be separated from each other. The first approaches to this problem assumed a constant t_{rel} value over the whole brain section [44]. Then the inclination can be calculated by inverting eq. (2.16). While the assumption of a constant t_{rel} value can be justified for homogeneous regions of the brain, it is certainly not valid for whole brain sections. Therefore, additional measurement information is necessary for the estimation of the inclination angle. Experimentally, additional information is available from additional measurements with the tilting stage. The interpretation of these additional data are discussed in sec. 2.7.




FIGURE 2.5: Analytical fit of the 3D-PLI model to the light intensities measured in one pixel. The transmittance parameter is two times the average light intensity, the direction angle φ the phase and the retardation the relative amplitude of the sinusoidal light intensity curve.




FIGURE 2.6: 3D-PLI parameter maps of a coronal human brain section measured in the LAP. **Top left:** Transmittance map which encodes light absorption by the tissue. **Top right:** retardation as a measure of phase retardance caused by birefringence. **Bottom left:** direction angle map indicating the orientation in the sectioning plane. **Bottom right:** absolute inclination map indicating the out-of-plane orientation.

Parameter maps The resulting model parameter maps of transmittance, retardation, direction angle and inclination angle are depicted in Fig. 2.6. The transmittance map (cf. Fig. 2.6 top left) closely resembles images obtained from transmission microscopy as it indicates light absorption: white matter with myelinated tissue absorbs and scatters more light, hence the transmittance map appears darker in white matter. The retardation map indicates the phase retardation induced by the brain section. At the resolution of the LAP of $64 \times 64 \mu\text{m}$, individual nerve fibers are not observable which results in vanishing retardation values for the cortex. The direction and inclination angle maps hold the information about the three-dimensional fiber orientation. The inclination map shows a dominant out-of-plane component in cortical regions while white matter regions also contain lower inclination angles. This result is not a reflection of the real nerve fiber orientations in the brain section. Instead it is caused by the inclination angle derivation which assumes a constant value of t_{rel} over the whole section and the low resolution of the LAP which cannot distinguish individual nerve fibers in the cortex.

Once direction angle and inclination angle maps are available, the three-dimensional fiber orientation can be computed using spherical coordinates as

$$\begin{aligned} x &= \cos(\alpha) \cos(\varphi) \\ y &= \cos(\alpha) \sin(\varphi) \\ z &= \sin(\alpha) . \end{aligned} \tag{2.17}$$

Visualization of three dimensional vector fields is typically done via color coding of the vector components. For 3D-PLI, two color encodings are employed. The RGB color encoding assigns a color to every vector as (red, green, blue) = $(|x|, |y|, |z|)$ similar to the color encoding used in DMRI. As only the absolute values of the vector components are considered, there is an inherent ambiguity in this color coding. This ambiguity is avoided using the other color encoding which additionally utilizes the brightness of the color space to indicate the z -component of the vector. In the HSV color space, the color encoding is given by $(H, S, V) = (2\varphi, 1, 1 - 2|\alpha|/\pi)$. The maps showing these color encoded fiber orientations are called *Fiber Orientation Maps (FOMs)*. In the example shown in Fig. 2.7 the difference between the color encodings are clearly visible in the corpus callosum, a brain region connecting both hemispheres, whose sides are both colored in red in the RGB FOM while in the HSV FOM the different in-plane orientations can easily be distinguished. The dominant out-of-plane nerve fiber orientation resulting from the analysis manifests as a strong blue respective black component in the FOMs.

Volumetric reconstruction After the sectionwise analysis, the individual brain sections need to be reconstructed to a coherent volume for a meaning analysis of brain regions. Here, the blockface images serve as a reference for the histological sections. As a first step, the blockface images are reconstructed using the ARTag reference markers [80, 71]. Then each histological section is linearly registered (linear registration allows translation, rotation and shearing) onto its corresponding blockface image using the *ITK* and *elastix* registration frameworks [81, 82]. Finally, the nonlinear distortions caused by the sectioning procedure need to be reversed using nonlinear image registration techniques. So far, nonlinear registration is achieved using an approach which utilizes three-dimensional information about the full image stack for small local deformations [83]. An example of a partly reconstructed human brain volume is depicted in Fig. 2.8.




FIGURE 2.7: 3D-PLI fiber orientation maps. **Left:** RGB color encoding. **Right:** HSV color encoding. cc: corpus callosum.

2.7 A review of existing approaches for the analysis of birefringence experiments from oblique views

Relative thickness and inclination are related to the retardation by (cf. eq. (2.16))

$$\sin(\delta) = \sin\left(\frac{\pi}{2} t_{\text{rel}} \cos(\alpha)^2\right). \quad (2.18)$$

This equation has severe implications for the derivation of the inclination angle. First of all, birefringence and inclination cannot be estimated simultaneously as both parameters are mapped onto one value. Secondly, even if t_{rel} was available, only the absolute inclination angle could be recovered by inverting the equation while the inclination sign would remain unknown. Furthermore, the inclination can only be unambiguously




FIGURE 2.8: Volumetric reconstruction in 3D-PLI. **A:** blockface volume. **B:** registered transmittance volume. The green arrows highlight reconstructed blood vessels while the blue arrows point out smooth transitions between white and grey matter regions. **C:** boundaries of the reconstructed volume in the full blockface volume. Taken from [84].

derived for $t_{\text{rel}} \leq 1$ as otherwise several inclination values are possible for one retardation value due to the periodicity of the outer sin-function.

Additional measurement information can be obtained by either varying the wavelength of the incident light to induce a predefined change of $t_{\text{rel}} = \frac{4t\Delta n}{\lambda}$ or by measuring the sample from oblique views to induce a predefined rotation to the modelled fiber orientation vector. As the polarization efficiency of polarization filters, especially wave retarders, is wavelength dependent, varying the wavelength is not feasible for 3D-PLI without further modification of the experimental setup. This makes oblique views the only feasible option. Experimentally, measurements from oblique views are realized by rotating the sample in the LAP at mesoscopic resolution and by oblique illumination in the LMP3D at microscopic resolution. For further analysis, the different tilted LAP measurements are coregistered onto the planar measurement based on a perspective transform computed using the Scale Invariant Feature Transform [85].

While the exact mathematical calculus of the oblique views is introduced in chap. 4, the geometric implications of an oblique measurement are already sufficient for a qualitative understanding of how the additional measurements enable the separation of inclination angle and t_{rel} . The biophysical model predicts a sinusoidal light intensity variation I during rotation of the polarization filter for a modelled fiber orientation vector $\mathbf{r}$ of $I(\mathbf{r})$. For an oblique measurement the modelled fiber orientation $\mathbf{r}_{\text{ob}}$ vector is slightly rotated compared to $\mathbf{r}$ which results in a different predicted light intensity course for the oblique measurement $I(\mathbf{r}_{\text{ob}})$. Both light intensity curves predict different retardation values for the planar measurement $|\sin \delta|$ and the oblique measurement $|\sin \delta|_{\text{ob}}$. Without the additional measurements, two different combinations of t_{rel} and α can predict the same retardation value. Now that the additional oblique measurement is available t_{rel} and α need to fit two different retardation values. Writing the rotated fiber inclination $\alpha_{\text{ob}} = f(\alpha)$ the two retardations yield a nonlinear system of equations:

$$\begin{aligned} |\sin(\delta)| &= \sin\left(\frac{\pi}{2} t_{\text{rel}} \cos(\alpha)^2\right) \\ |\sin(\delta)|_{\text{ob}} &= \sin\left(\frac{\pi}{2} t_{\text{rel}} \cos(f(\alpha))^2\right) . \end{aligned} \quad (2.19)$$

For each oblique measurement, a different oblique retardation is obtained which adds another equation to this equation system. The solution of this system of equations then returns values of t_{rel} and inclination.

The first attempts to analyze the oblique measurements in 3D-PLI focused on the determination of the inclination sign. This can also be understood qualitatively from the retardation model: the inclination of a nerve fiber with positive inclination increases for oblique views from a certain direction and decreases for oblique views from a different direction. Due to the direct relation between retardation and inclination, analyzing the changes of the retardation signal enables the determination of the inclination sign. As the retardation changes are small resulting in potentially noisy inclination sign maps, different regularization techniques such as a Markov Random Field approach [86] and total variation [87] were utilized to further improve the results. Still, these approaches were limited to the reconstruction of the inclination sign.

The first approach to determine the inclination and relative thickness from the oblique measurements was presented in [88, 46] where it was shown that the retardation variation with respect to different oblique measurements (cf. (2.19)) is sinusoidal. Based on

a Fourier analysis inclination angle and t_{rel} can then be derived analytically. In the following, this approach will be abbreviated as *DFT* as it is based on a discrete Fourier transform. Due to its analytical nature, the DFT based algorithm is computationally very efficient but suffers from noise instabilities especially for highly inclined nerve fibers and in-plane oriented nerve fibers [46]. In [46], a second approach based on maximum likelihood estimation⁷ termed *Likelihood Orientation Estimation (LOriE)* was developed to overcome the noise instability. This approach takes into account not only changes in the retardation but also changes in the direction angle of different oblique measurements (rotation of the fiber orientation vector also changes the in-plane orientation). As the parameter estimation is based on pixelwise numerical optimization to find the most plausible parameters given the measurement data, the computational demands increase to 60-80 core hours per brain section.

While the LOriE algorithm was shown to be more accurate for weak signals than the DFT algorithm [46], it still induces artifacts for in-plane oriented nerve fibers. In Fig. 2.9 a region of interest (*ROI*) in a coronal human brain section of mostly in-plane oriented nerve fibers processed with the DFT and LOriE algorithms is shown. The histograms (2.9 bottom) display very different results: the DFT algorithm results in almost no in-plane oriented nerve fibers with $\alpha \approx 0^\circ$ while the LOriE algorithm yields a very high number of these orientations. In the inclination map computed by the LOriE algorithm these orientations are indicated in red. Physically, this singular high number of in-plane nerve fiber orientations is not plausible which means that this behaviour of the LOriE algorithm has to be classified as an artifact⁸. Hence, new approaches are necessary to improve the analysis of the oblique measurements.

In chap. 4 and 5 two novel algorithms are developed for a more precise reconstruction of tissue parameters from the oblique measurements. The general statistical principles behind them are presented in the next chapter.

⁷The Maximum Likelihood Estimation method aims to find the most probable parameters of a model given a set of observations. A more extensive introduction is given in chap. 3.

⁸The discontinuity of the inclination histogram obtained from LOriE was already observed in [46] p. 87, yet the classification of this finding as an artifact originating from the algorithm was stated here.




FIGURE 2.9: Analysis of inclination angles calculated by the DFT and LOrIE algorithms for in-plane oriented nerve fibers. **Top left:** position of the analyzed ROI in the brain section (top) and delineation of the white matter region from which the histogram was calculated (bottom). **Top right:** inclination map computed by the LOrIE algorithm clipped to the range $[-30^\circ, 30^\circ]$. Red squares indicate pixels with very low absolute inclinations of $|\alpha| < 0.1^\circ$. **Bottom:** inclination histograms obtained from DFT and LOrIE for the white matter ROI delineated in the top left. While the DFT algorithm results in almost no in-plane inclinations, the LOrIE algorithm estimates an extraordinary high ratio of in-plane oriented fibers with $\alpha \approx 0^\circ$ which appears as a large peak in the otherwise smooth histogram.

3

Foundations of Bayesian data analysis

“The only statistics you can trust are those you falsified yourself.”

– Winston Churchill

While a detailed discussion of Bayesian statistics and computation is beyond the scope of this thesis, a brief introduction to motivate the later application of these methods is presented in this chapter. The introduction focuses on Bayesian methods for regression problems as the problem of interest in this thesis is a regression problem: inferring birefringence parameters from the 3D-PLI measurement.

Although Bayes’ theorem dates back to 1763 [89], Bayesian methods have found widespread use only in recent years in many fields due to their natural way to quantify uncertainty. The main reason for the growing adoption of Bayesian approaches are advances in hardware and software technology as well as novel algorithms that make the high computational demands of these approaches which were prohibitive several decades ago feasible [55]. The increasing interest in Bayesian methods manifested in the development of so called probabilistic programming languages which aim to simplify the usage of these methods not just by publicly funded researchers but also by leading software companies. Prominent examples for publicly funded software are *Stan* [90] and *PyMC3* [91] while Facebook’s *HackPPL* [92], Google’s *Edward* [93], Microsoft’s *InferNet* [94] and Uber’s *Pyro* [95] represent the most popular frameworks developed by the private sector.

The chapter starts with an introduction into the terminology of Bayesian statistics and parameter estimation using priors. Then the Bayesian computation of uncertainty estimates is reviewed. Note that this chapter does not offer a rigorous mathematical introduction to Bayesian statistics but instead tries to convey an intuitive understanding of the principles of Bayesian data analysis.

3.1 Parameter estimation based on Bayes’ theorem

In a Bayesian setting, model parameters are treated as random variables whose probability distribution we seek to determine. Bayes’ theorem states that the probability for a parameter set θ of model M given the data D is given by

$$p(\theta|D, M) = \frac{p(D|\theta, M)p(\theta|M)}{p(D|M)} . \quad (3.1)$$




FIGURE 3.1: Example data for linear regression problem

$p(\theta|M)$ is referred to as *prior* and describes the a-priori knowledge about the parameters. $p(D|\theta, M)$ is called the *likelihood* and represents the probability to observe the measurement data given a specific parameter set of the model. This probability requires a noise model for the specific measurement. $p(\theta|D, M)$ is referred to as *posterior* and represents our a-posteriori knowledge about the parameters incorporating our prior knowledge and the measurement data. The denominator $p(D|M)$, the probability to observe the data given the model, is commonly termed the *model evidence* and can be computed as

$$p(D|M) = \int_{\Sigma} p(D|\theta, M)p(\theta|M)d\theta . \quad (3.2)$$

where Σ is the parameter space. Combining eq. (3.1) and eq. (3.2) it becomes clear that the denominator serves as a normalization factor so that the posterior becomes a proper probability density whose probability over the whole parameter space integrates to 1.

The posterior is the essential quantity used by a Bayesian analysis. The best fit parameters are the ones which maximize the posterior yielding the *maximum-a-posteriori estimator* (MAP). In case of uniform prior and evidence this reduces to *Maximum Likelihood Estimation* (MLE) as only the likelihood depends on the model parameters.

The application of Bayes' theorem is best understood from a simple example such as one-dimensional linear regression. Suppose we observe n data pairs and call them $\{(x_i, y_i), i = 1, \dots, n\}$. Linear regression aims to estimate the parameters of the model $f(x) = \alpha + \beta x$. Let us assume that our observations are independent and normally distributed with known variances: $y_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$. While in many cases the variances of the individual observables are unknown, they were obtained from noise measurements for 3D-PLI and utilized in the analysis for the first time in [46]. For the synthetic data depicted in Fig. 3.1 the variances were sampled from a uniform distribution $\sigma \sim \mathcal{U}(0, 8)$. The true model parameters are $\alpha_{true} = 0.5$ and $\beta_{true} = 1$ ¹.

Now, the likelihood to observe a specific y_i for a specific model parameter set $p(y_i|\alpha, \beta)$ is given by inserting the response predicted by the model $f(x)$ into the probability distribution of the observations as its expected value: $p(y_i|\alpha, \beta) \sim \mathcal{N}(\alpha + \beta x_i, \sigma_i^2)$. Combining all n observations, the joint likelihood $p(y_1, \dots, y_n|\alpha, \beta)$ is then the product

¹All plots in this chapter were created with the *matplotlib* and *seaborn* packages [96, 97].




FIGURE 3.2: Synthetic data and weighted least squares estimate

of all individual likelihoods:

$$p(y_1, \dots, y_n | \alpha, \beta) \sim \prod_{i=1}^n \mathcal{N}(\alpha + \beta x_i, \sigma_i^2) . \quad (3.3)$$

This is already the likelihood required for Bayes' theorem. As we have no prior knowledge, we choose uniform priors over α and β . The evidence can also be chosen as uniform as no measurement seems any more likely than another purely based on the model. In consequence, the posterior is equal to the likelihood and given by:

$$p(\alpha, \beta | y_1, \dots, y_n) = p(y_1, \dots, y_n | \alpha, \beta) \quad (3.4)$$

$$= \prod_{i=1}^n \mathcal{N}(\alpha + \beta x_i, \sigma_i^2) \quad (3.5)$$

$$= \prod_{i=1}^n \frac{1}{\sqrt{2\pi}\sigma_i} \exp\left(-\frac{(y_i - (\alpha + \beta x_i))^2}{2\sigma_i^2}\right) . \quad (3.6)$$

Instead of maximizing the expression given in eq. (3.6) it is common to minimize the negative logarithmic likelihood function $l = \ln(p)$ which yields (neglecting constant terms)

$$-l(\alpha, \beta | y_1, \dots, y_n) = \frac{1}{2} \sum_i^n \ln(\sigma_i^2) + \frac{(y_i - (\alpha + \beta x_i))^2}{\sigma_i^2} . \quad (3.7)$$

This result, the log-likelihood for a normal distribution, becomes important in the chapters 5 and 6 as the noise of the 3D-PLI measurement can be modeled by a Normal distribution. For a different problem, only the expected value in eq. (3.7) (here: $f(x) = \mu = \alpha + \beta x$) has to be replaced which results in the general log-likelihood for normally distributed observables:

$$l = -\frac{1}{2} \sum_i^n \ln(\sigma_i^2) + \frac{(y_i - \mu_i)^2}{\sigma_i^2} . \quad (3.8)$$

The optimization problem for the straight fit line problem can then be stated as

$$\operatorname{argmin}_{\alpha, \beta} \frac{1}{2} \sum_i^n \ln(\sigma_i^2) + \frac{(y_i - (\alpha + \beta x_i))^2}{\sigma_i^2} . \quad (3.9)$$




FIGURE 3.3: Synthetic data with outliers. True model: black, weighted least squares estimate (WLS): blue.

As in our case the variances do not depend on the model parameters and have in consequence no effect on the estimation, this expression can further be simplified to

$$\operatorname{argmin}_{\alpha, \beta} \frac{1}{2} \sum_i^n \frac{(y_i - (\alpha + \beta x_i))^2}{\sigma_i^2} \quad (3.10)$$

which is one case of the *Weighted Least Squares Estimator (WLS)*. The name stems from the fact that the weighted squared error between the model predictions and the observations has to be minimized. A closed form solution for the WLS estimator for linear regression is available [98] and yields a satisfactory fit for our problem: $\alpha_{WLS} = 0.519$ and $\beta_{WLS} = 0.986$ (cf. Fig. 3.2). While this result would be possible to obtain without the Bayesian approach, the usefulness of priors becomes more evident in the presence of outliers.

For the dataset with two outliers, the WLS estimate deviates strongly from the true model (cf. Fig. 3.3): $\alpha_{WLS} = 3.51$ and $\beta_{WLS} = 0.48$ compared to $\alpha_{true} = 0.5$ and $\beta_{true} = 1$. The first choice here would be to perform outlier detection and perform the fit only on the inlier data points [98]. Here, for demonstration purposes, the problem is approached using priors. An observer who does not know the true model parameters will likely judge that the slope of the regression line of the WLS estimate is too small and the intercept too high as consequence of the outliers. This impression, our prior knowledge about the model parameters, can be incorporated into the estimation process by a prior. One choice for the priors are normal priors with expected value of $\mu_\beta = 1$ for the slope and $\mu_\alpha = 0$ for the intercept: $p(\beta) \sim \mathcal{N}(\mu_\beta, \sigma_\mu^2)$ and $p(\alpha) \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha^2)$. Taking the logarithm of Bayes' theorem yields

$$\underbrace{\ln p(\theta|D, M)}_{\text{Log-Posterior}} \propto \underbrace{\ln p(D|\theta, M)}_{\text{Log-likelihood}} + \underbrace{\ln p(\theta|M)}_{\text{Log-Prior}}. \quad (3.11)$$

Combining the log-likelihood from eq. (3.10) and the logarithmic priors then gives the logarithmic posterior l_{post} :

$$-l_{post} = \frac{1}{2} \left(\frac{(\alpha - \mu_\alpha)^2}{\sigma_\mu^2} + \frac{(\beta - \mu_\beta)^2}{\sigma_\beta^2} + \ln(\sigma_\mu^2 \sigma_\beta^2) + \sum_i^n \frac{(y_i - (\alpha + \beta x_i))^2}{\sigma_i^2} \right). \quad (3.12)$$




FIGURE 3.4: Synthetic data with outliers. True model: black, weighted least squares estimate (WLS): blue, maximum-a-posteriori estimators with decreasing variances over the normal priors (MAP): red, green, yellow

Again, the variance term $\ln(\sigma_\mu^2 \sigma_\beta^2)$ represents a constant offset and can be neglected for the optimization. The prior variances σ_β^2 and σ_μ^2 measure how strong our belief in our prior knowledge is: the smaller they are, the narrower become the prior probabilities and the stronger they effect the posterior. The result of the minimization which was carried out numerically for convenience, is depicted in Fig. 3.4 for three different variances of the priors (for simplicity, the variances were the same for both paramters). We can conclude that the smaller the prior variance, the stronger the fit line tends towards the prior and the better it fits the true model as the prior assumptions are close to the ground truth.

While the best fit parameters can now be estimated based on the prior and the likelihood, no quantitative statement about the goodness of the fit can be obtained solely on the MAP estimate in general. As all derivations in a Bayesian setting are based on probabilities, Bayesian statistics offer an intuitive way to measure the confidence in the obtained results. This is explored in the next section.

3.2 Uncertainty quantification

Like the best fit parameters, the uncertainty of the parameters is also derived from the posterior. Graphically, we just have to determine how much of the probability mass lies within a certain volume. E.g, for a one dimensional distribution, a 2σ confidence interval is given by the interval that contains 95% of the probability mass. In Bayesian statistics, *credible intervals* are the analogue of confidence intervals from frequentist statistics. Typically, credible intervals are obtained from *high posterior density (HPD)* intervals, the shortest intervals that contain a certain amount of probability. In a multi-variate case, a credible interval for a specific parameter θ_i is obtained from the *marginal*

distribution $P(\theta_i|D, M)$. This marginal distribution is obtained by integrating over all other parameters:

$$P(\theta_i|D, M) = \int_{\Sigma_{-i}} P(\theta|D, M) d\theta_{-i} \quad (3.13)$$

The marginal distribution is a one-dimensional distribution from which HPD intervals can be derived. Visually, integrating over all parameters just means that we only restrict the parameter space of the variable of interest, we *marginalize* over all other parameters.

While conceptually elegant, computing credible intervals is hard as it requires integration over potentially high-dimensional parameter spaces which are not tractable analytically. Also, most of the parameter space will contribute almost nothing to the integral making standard numerical integration very inefficient. To make these computations feasible, *Markov Chain Monte Carlo (MCMC)* methods were developed. MCMC offers an efficient way to generate samples from a distribution and thereby perform the integration implicitly. The difference to standard Monte Carlo integration is that the samples are correlated.

Informally speaking, MCMC methods perform a (more or less depending on the algorithm) random walk through the parameter space biased towards high probability regions. This is achieved by constructing a Markov Chain whose equilibrium distribution is the target posterior. Essentially, from a position x_i with posterior probability $p(x_i)$, a next position x_{i+1} with posterior $p(x_{i+1})$ is chosen using a transition kernel. The trick in MCMC is now that the probability to accept x_{i+1} as a new position for the chain is given by a Metropolis-Hastings criterion proportional to the division of both probabilities $p(x_{i+1})/p(x_i)$. This way the computation of the evidence in eq. (3.2) is no longer required. As an example, the simplest MCMC algorithm, the random walk Metropolis-Hastings (RWMH) algorithm with Normal proposal distribution, will be discussed in more detail [99, 100]. The pseudocode of the algorithm is depicted for the one-dimensional case in Alg. 1.

```

Given: target probability density  $p$ , proposal distribution  $g$ ;
Initialization: pick a starting point  $x_0$ ;
while  $i < \text{number of iterations}$  do
    Proposal generation:  $x_p \sim g(x_i) = \mathcal{N}(x_i, \sigma^2)$ ;
    Acceptance ratio:  $\alpha = \min \left( 1, \frac{p(x_p) g(x_i|x_p)}{p(x_i) g(x_p|x_i)} \right) = \min \left( 1, \frac{p(x_p)}{p(x_i)} \right)$ ;
    Random number generation:  $u = \mathcal{U}[0, 1]$ ;
    Accept or reject:
    if  $u \leq \alpha$ ;
    then
        | Accept proposal:  $x_{i+1} = x_p$ ;
    else
        | Reject proposal:  $x_{i+1} = x_i$ ;
    end
end

```

Algorithm 1: Pseudocode of the Random Walk Metropolis-Hastings algorithm

The RWMH algorithm works in two steps²: proposals are drawn from a normal distribution with the the current position of the chain as its mean. Next, the acceptance

²A very helpful animation of MCMC algorithms can be found at <http://chi-feng.github.io/mcmc-demo/>, last accessed: 23.12.2019

ratio α is calculated as the ratio of the target probability of the proposed and the current state multiplied by the inverse ratio of the proposal probabilities. In the case of the normal proposal distribution where $g(y|x) = \mathcal{N}(x, \sigma^2)$ the proposal is symmetric as $g(x_i|x_p) = g(x_p|x_i)$, so the acceptance ratio reduces to $\alpha = \min(1, p(x_p)/p(x_i))$. The acceptance ratio is compared to a random number u drawn from a uniform distribution in $[0, 1]$. If the probability of the proposed state is higher than the current one, it will always be accepted as $\alpha = 1$ and $u \leq 1$. If the proposed state has a lower probability than the current state ($\alpha < 1$), it will sometimes be rejected and sometimes accepted (the lower the "decrease" in probability, the higher the acceptance probability). This ensures that the chain moves into regions of lower probability, instead of only ascending to higher probability regions. After running the Markov chain for a "sufficient" number of iterations, the chain will explore the whole parameter space and the empirical distribution of the samples converge to the posterior distribution. In practice, this procedure offers a way to determine the posterior in a reasonable amount of time. As the calculation of the evidence is not required, the nominator of the logarithmic posterior from eq. (3.11) can be utilized for uncertainty estimation via MCMC just as for finding the MAP estimate.




FIGURE 3.5: Scatterplot matrix of posterior samples generated by MCMC for the straight line fit problem with outliers and prior variances $\sigma_\alpha = \sigma_\beta = 1$. Top: histogram of α values of the samples, right: histogram of β values of the samples. Straight lines: median values. dashed lines: boundaries of the high posterior density intervals containing 95% of all samples marginalized over the other variable. Middle: scatter plot in which the number of samples within a bin is highlighted by the blackness.

Although theoretical convergence of the RWMH algorithm is guaranteed, its convergence time is impedingly high for many cases. For example, for a multivariate case a correlated distribution will be sampled very ineffectively if the proposal distribution stays symmetric. This can only be coped with by manually adjusting the covariance matrix of the proposal distribution. In general the covariance matrix always needs to be tuned for a specific problem or adjusted during the sampling: if the proposed steps

are too small, the chain will take very long to explore the parameter space, if they are too big, many proposals will be rejected which also results in a extensive convergence time. Furthermore, the RWMH algorithm also suffers strongly from the famous *curse of dimensionality*: with increasing number of dimensions it becomes increasingly difficult to generate "good" proposals from the normally distributed proposal as the available space increases exponentially with the number of dimensions. For the application in an imaging technique such as 3D-PLI, another important aspect besides convergence speed is that an MCMC algorithm should not require extensive manual tuning as the data from different pixels of the brain are very different from each other. For example, a grey matter pixel without birefringent signal will have a much worse goodness of fit than a white matter pixel.

One of the more advanced approaches to MCMC is ensemble sampling [101]. The ensemble sampler runs several Markov chains in parallel and generates proposals by inter- and extrapolating between them resulting in "informed" proposals. The idea behind this can be understood in the following way: as all chains will in the long run converge to the same stationary distribution, moving into the direction of another chain is a better proposal than picking a random point like RWMH does. Furthermore, due to its affine invariance the ensemble sampler can sample linearly correlated densities as efficiently as uncorrelated densities and nonlinearly correlated densities much more efficiently than RWMH. The only parameters which have to be set manually for the ensemble sampler are the number of parallel chains and their starting positions.

The application of MCMC sampling is demonstrated using the same straight line fit as for the parameter estimation. The posterior is again given by eq. (3.12) and fed into the emcee package's Python implementation of the ensemble sampler [102]. The number of chains was set to 50 and the chains were initialized randomly close to the MAP estimate as recommended by the authors of the ensemble sampler. A scatterplot matrix of the resulting samples is depicted in Fig. 3.5. Intercept and slope are clearly correlated. This can be interpreted in a way that if we believe in a higher intercept, we also have to believe in a lower slope. The credible intervals can now be obtained from the samples by calculating HPD intervals for both variables individually resulting in $\alpha = 1.94^{+0.91}_{-0.92}$ and $\beta = 0.74^{+0.16}_{-0.16}$. This result means that we have a higher degree of belief in our estimate of the slope than the intercept.

The question whether a MCMC algorithm has converged after a certain number of iterations is subject of ongoing research. A simple method to evaluate convergence is to run the same chain several times and compare the results. Recently, Vats et al. [103] published a minimal number of effective samples, the number of uncorrelated samples of the chain, required to achieve a specific Monte Carlo Error depending on the dimension of the problem. Unfortunately, estimating the effective sample size from a correlated Markov Chain is very challenging for short chains [104, 103]. For the purposes of 3D-PLI, keeping the number of samples low is key to achieve acceptable computation times. Therefore in the latter parts of this thesis (cf. chap. 6) the convergence of the applied MCMC algorithms is evaluated based on synthetic data for which ground truth HPD intervals are available.

For an interpretation, it is also appealing to visualize the prediction of the posterior samples. Calculating the values predicted by the model for all samples yields the *posterior predictive distribution (PPD)*. As for every predictor value of x now exists a distribution of possible model outcomes, summary statistics such as percentiles and mean value of




FIGURE 3.6: Posterior predictive distribution for the straight line fit problem. Blue: mean value. Grey shaded: 95% credible region.

the PPD can be computed and plotted. In Fig. 3.6 the mean value of the posterior predictive distribution of the straight line fit problem is plotted in blue. The 95% prediction interval obtained as the HPD of the predicted values is shaded in grey.

Here, we merely scratched the surface of Bayesian statistics. A more rigorous and in depth discussion is given in [104, 105]. For the purposes of this thesis though, the concepts presented here are sufficient. The different estimators introduced here are applied to 3D-PLI in the following chapters. Chap. 4 introduces a WLS estimator to 3D-PLI, chap. 5 a MAP estimator. Finally, in chap. 6 MCMC sampling is utilized to quantify uncertainty in 3D-PLI.

A Least Squares Approach for the derivation of Nerve Fiber Orientations

“It is a capital mistake to theorize before one has data.”

– Sherlock Holmes

In sec. 2.7 the two algorithmic approaches developed in [46] to determine three-dimensional fiber orientations from oblique 3D-PLI measurements were reviewed. Both represent a great step forward to an unambiguous estimation of the out-of-plane nerve fiber orientation compared to previous studies but suffer from instabilities for in-plane fibers. Compared to each other, the MLE based LOriE algorithm was shown to be more reliable for in-plane and out-of-plane fibers than the analytical DFT algorithm but leads to artifacts for flat fibers for a significant number of pixels. Another drawback is its high demand in computation time of 60-80 core hours per brain section.

In this chapter a new approach is developed to overcome the pitfalls of the LOriE algorithm. As analytical approaches such as the analytical DFT algorithm are prone to suffer from noise instabilities, the new algorithm must also be based on a model fitting paradigm. Instead of a MLE based approach we will utilize a simpler weighted least squares estimator. As the objective function of WLS estimators is a weighted sum of quadratic differences between the model function and the observed data, it is not as complex as the objective function resulting from an MLE based approach which for a gaussian noise model also contains logarithmic terms (cf. chapter 3). This simplifies the optimization problem and thereby the numerical stability of the estimator.

This chapter is structured as follows: first a new algorithmic framework is developed and a first implementation described. It is then evaluated and compared to the DFT algorithm on synthetic data. The plausibility of the obtained fiber orientations are then evaluated for human brain data. Finally, strategies to reduce the computation time of the developed algorithm by utilizing the computational power of GPUs and classical CPU clusters are explored. Large parts of this chapter were published in [106, 84, 107].




FIGURE 4.1: Tilting coordinate system. Light grey: specimen stage in the planar position without tilt, dark grey: tilted specimen stage. ψ : tilting direction angle, τ : tilting angle. Adapted from [84].

4.1 The Robust Orientation Fitting via Least Squares algorithm

In the following derivation, the coordinate system developed in [46] is utilized (cf. Fig. 4.1). It describes the tilting position by a tilting direction ψ and the angle by which the stage is tilted with respect to the planar plane τ . With N_{Tilt} as the number of tilting positions the tilting directions equidistantly spanning the space are in general given by $\psi_k = \frac{2\pi}{N_{Tilt}}(j - 1)$, $j \in [1, 2, \dots, N_{Tilt}]$ with the index k indicating the tilting position. The standard measurements with four tilting positions is then described by the tilting directions $\psi = 0^\circ, 90^\circ, 180^\circ, 270^\circ$ and a tilting angle of $\tau = 8^\circ$. In this notation, the planar measurement is simply a tilt with tilting angle of $\tau = 0^\circ$. The essential information lies in the change of the light intensity with respect to the tilting position. For a given fiber orientation $\mathbf{r}$ the polarimetric model predicts a specific light intensity I . For a tilting position k the predicted light intensity I_k is then obtained by rotating the fiber orientation accordingly and calculating the light intensity for the rotated vector $\mathbf{r}_k = \mathbf{R}_k \mathbf{r}$ with rotation matrix $\mathbf{R}_k$. These rotation matrices were derived in [46]: the full rotation is obtained by first rotating around the z-axis by $-\psi$, then rotating around the y-axis by τ and then rotating back around the z-axis by ψ which can be written as

$$\mathbf{R}(\psi, \tau) = \mathbf{R}^z(\psi) \mathbf{R}^y(\tau) \mathbf{R}^z(-\psi) . \quad (4.1)$$

The resulting matrices are given in App. A.1. Denoting the transformation from the planar coordinate system into the tilted coordinate system with γ , the rotated vector's direction angle can be expressed as $\varphi_k = \gamma_k(\varphi)$ and its inclination angle as $\alpha_k = \gamma_k(\alpha)$.

As already noted in [46], the relative thickness t_{rel} also needs to be adjusted for a tilted measurement due to the elongated optical path of the light through the tissue as $t_{relk} = \gamma_k(t_{rel}) = t_{rel} / \cos(\tau)$. The tilting angle τ itself is affected by refraction which leads to a tilting angle lower than the incident "tilt". According to Snell's law the internal tilting angle τ_{int} can be calculated from the experimentally employed tilting angle of the tilting stage τ_{exp} and the tissue's refractive index n_{sample} as

$$\tau_{int} = \arcsin(\sin(\tau_{exp}) / n_{sample}) \quad (4.2)$$

We introduce the index j for the measurements, including 0 for the planar measurement and $1, \dots, N_{Tilt}$ for the tilted measurements. Furthermore, i will index the rotation angle of the polarization filters. Additionally, we denote the total number of measurements

(tilted measurements and the planar measurement) as $N_T = N_{Tilt} + 1$ and the number of polarization states acquired as N_P . In this notation the intensity curve of a measurement can then be expressed as

$$I_{ji}(\rho_i, \varphi, \alpha, t_{rel}) = \frac{I_{j,T}}{2} (1 + \sin(2(\rho_i - \varphi_j)) \sin\left(\frac{\pi}{2} t_{rel_j} \cos^2(\alpha_j)\right)) \quad (4.3)$$

where $I_{j,T}$ denotes the transmittance of measurement j . As the transmittances are easily available via the Fourier analysis of the individual measurements and are affected by additional absorption and refraction effects in a tilted measurements, they are not included in the further parameter estimation. Therefore we introduce the normalized intensity

$$I_{N_{ji}} = \frac{2I_{ji}}{I_{j,T}} - 1 \quad (4.4)$$

which is limited to the range $[-1, 1]$. From the variance of the individual light intensity $\sigma_{I_{ji}}^2 = gI_{ji}$ the variance of the normalized light intensity can be derived using error propagation as (see Appendix A.2)

$$\sigma_{N_{ji}}^2 = \frac{4gI_{ji}}{I_{j,T}^2} \left(1 + \frac{2I_{ji}}{N_P I_{j,T}}\right) \quad (4.5)$$

The normalized light intensities predicted by the 3D-PLI model f_{ji} are given by inserting eq. (4.3) into eq. (4.4) which yields

$$f_{ji}(\rho_i, \varphi, \alpha, t_{rel}) = \sin(2(\rho_i - \varphi_j)) \cdot \sin\left(\frac{\pi}{2} t_{rel_j} \cos^2(\alpha_j)\right). \quad (4.6)$$

Now f can be fitted to the normalized intensities $I_{N_{ji}}$. With the weights $w_{ij} = \sigma_{N_{ji}}^{-2}$ the optimization problem is stated as

$$\underset{\varphi, \alpha, t_{rel}}{\operatorname{argmin}} \chi^2 = \underset{\varphi, \alpha, t_{rel}}{\operatorname{argmin}} \sum_{j=0}^{N_T} \sum_{i=0}^{N_P} ((f_{ji}(\rho_i, \varphi_j, \alpha_j, t_{rel_j})) - I_{N_{ji}}) \cdot w_{ji}^2. \quad (4.7)$$

Minimizing this function means that the modelled normalized light intensity curves are fitted to the measured normalized light intensities of all oblique measurements simultaneously. The weights of the objective function yield an interesting observation: as the camera gain factor is a constant multiplicative factor for all weights, it has no effect on the optimization process. Physically, this means that the exact knowledge about the camera gain factor is not considered as only the amounts of the weights relative to each other are taken into account. This fact distinguishes the WLS estimator from the MLE which can utilize the information about the gain factor. In the objective function of the LOrE algorithm (see. [46] p. 48), the logarithmic terms also contain the gain factor.

The minimization strategy uses the approach developed for the LORIE algorithm: a grid search followed by a local minimization. A good guess for the direction angle is given by the result of the Fourier analysis of the planar measurement φ_0 . For inclination and relative thickness, a 6×6 grid equidistantly spanning the parameter space is evaluated. From the best minimum found via this brute-force minimization, the Levenberg-Marquardt algorithm [108, 109] is employed to converge to a local minimum.

The Levenberg-Marquardt algorithm is an adaptive gradient based optimizer¹ but not inherently capable of dealing with bounds on the parameters. Here, the boundaries

¹The gradient is given in App. A.3

for inclination and direction do not pose a problem to the optimization process due to the symmetry of the parameter space [46]. In case that the optimizer converges to a minimum outside of the parameter space with an unbounded orientation (φ_u, α_u) , the result can be "mirrored" back into the desired space by the relation

$$\alpha = \left(\left(\alpha_u + \frac{\pi}{2} \right) \bmod \pi - \frac{\pi}{2} \right) \operatorname{sgn} \left(\frac{1}{2} - \left\lfloor \frac{\varphi_u}{\pi} \bmod 2 \right\rfloor \right) \quad (4.8)$$

$$\varphi = \varphi_u \bmod \pi \quad (4.9)$$

where then $\varphi \in [0, \pi)$ and $\alpha \in [-\pi/2, \pi/2)$. The cost function is symmetric with respect to the sign of the relative thickness so it is sufficient to take the absolute value after the optimization $t_{\text{rel}} = |t_{\text{rel}}|$. The fitting routine will be denoted as *Robust Orientation Fitting via Least Squares (ROFL)* algorithm from now on.

Data: Maps of φ_0 , $I_{j,T}$, I_{ji} , $j \in \{0, \dots, N_T\}$, $i \in \{0, \dots, N_P\}$

Result: Maps of φ , α , t_{rel} , χ^2

```

for all image pixels do in parallel
    // calculate normalized intensities and their standard deviations
    for  $j = 0 : N_T$  do
        for  $i = 0 : N_P$  do
             $I_{N_{ji}} = \frac{I_{ji}}{2I_{j,T}} - 1$ 
             $\sigma_{N_{ji}}$  = error prop. of  $\sigma_{I_{j,i}}$   $\sigma_{I_{j,T}}$  to  $I_{N_{j,i}}$  via eq. (4.5)
        end
    end
    // brute force minimization
    for  $\alpha_l, t_{\text{rel}_k} \in \text{bruteforce grid}$  do
         $\chi_{l,k}^2 = \chi^2(\varphi_0, \alpha_l, t_{\text{rel}_k}, (I_{N_{00}}, \sigma_{N_{00}}, \dots))$ 
    end
     $\alpha_0, t_{\text{rel}_0} = \underset{\alpha, t_{\text{rel}}}{\operatorname{argmin}} \chi_{l,k}^2$ 
    // local optimization
    Levenberg-Marquardt optimization of  $\chi^2$  with initial point
     $(\varphi_0, \alpha_0, t_{\text{rel}_0}) \rightarrow \varphi, \alpha, t_{\text{rel}}, \chi^2$ 
    transform  $\varphi, \alpha, t_{\text{rel}}$  into 3D-PLI parameter space via eq. (4.9)
end

```

Algorithm 2: Pseudocode of the ROFL algorithm. Adapted from [84]

The ROFL algorithm was implemented in Python, utilizing *numpy* [110] for the calculation of the objective function and its gradient and *scipy's* *leastsq* function [111, 112, 113] for the Levenberg-Marquardt optimization. As all pixels can be processed independently of each other, the computations are easily parallelized using *mpi4py* [114]. A pseudocode of the algorithm is given in Alg. 2.

For illustration purposes, the working principle of the ROFL algorithm is demonstrated for a single pixel in Fig. 4.2. The pixel is located in the stratum sagittale (cf. Fig. 4.2 A), therefore the fiber orientation is expected to be perpendicular to the coronal sectioning plane. The light intensity curves of the planar measurement and two tilting positions after the registration onto the planar measurement are clearly distinguishable in Fig. 4.2 B. These light intensities are then normalized and fitted by the model. Finally, the normalized light intensities predicted by the estimated model parameters are compared




FIGURE 4.2: Working principle of the ROFL algorithm demonstrated for a single pixel. **A**: Transmittance map. The red circle points out the position of the analyzed pixel. **B**: Measured light intensities of the planar measurement and tilted measurements to west and east after calibration and registration onto the planar measurement. **C**: Normalized light intensities of the planar measurement and tilted measurements to west and east. The dashed lines depict their best fit curves according to the ROFL algorithm. Fit result: $\varphi = 101^\circ$, $\alpha = -60^\circ$, $t_{\text{rel}} = 0.5$, $R^2 = 0.97$. Taken from [84].

to the normalized measured light intensities in Fig. 4.2 C: as the differences are small, the model is able to describe the observed data accurately in this pixel. The estimated fiber orientation is inclined to the sectioning plane by 60° which conforms to our anatomical expectation.

4.2 Evaluation for synthetic data

Methods

The robustness of the developed approach against measurement noise has to be tested via simulations. First, a suitable procedure to generate synthetic data has to be developed.

Generation of synthetic data

Here, we follow the approach presented in [46]. First, a ground truth fiber with the parameters $(\varphi, \alpha, t_{\text{rel}})$ is defined. This vector is then rotated according to the rotations for four tilted measurements with a tilting angle of $\tau = 5.51^\circ$, the presumed internal tilting angle for human brain sections. This results in four "tilted" vectors. For these vectors and the original "untilted" vector artificial 3D-PLI measurements are constructed by calculating a light intensity profile I_{sim} according to the 3D-PLI model from eq. (4.3). The transmittance was set to 5000, a typical value for human brain sections.

In [46] these light intensities were then distorted according to a Poisson distribution. The Poisson noise model presumes equal variance and expected value for the light intensity.

As shown in [46], this assumption is not valid for the LAP where variance σ^2 and expected value μ are related by a gain factor of 3: $\sigma^2 = 3\mu$. Instead, a Negative binomial distribution is chosen as it is defined for natural numbers (light intensities are positive counts) and allows arbitrary expected value and variance. In a negative binomial model parametrized by the expected value and variance, the occurrence of k is given by

$$P(k|\mu, \sigma) = \binom{k-1 + \frac{\mu^2}{\sigma^2 - \mu}}{k} \left(\frac{\sigma^2 - \mu}{\sigma^2} \right)^k \left(\frac{\mu}{\sigma^2} \right)^{\frac{\mu^2}{\sigma^2 - \mu}}. \quad (4.10)$$

Inserting $\sigma^2 = g\mu$, the likelihood to observe the light intensity I for expected value μ and gain factor g can then be expressed as

$$P(I|\mu, g) = \binom{I-1 + \frac{\mu}{g-1}}{I} \left(\frac{g-1}{g} \right)^I g^{-\frac{\mu}{g-1}}. \quad (4.11)$$

Noisy synthetic light intensities I_{noisy} can now be obtained by drawing from a negative binomial distribution whose expected value is given by the undistorted synthetic light intensity: $I_{\text{noisy}} \sim \mathcal{NB}(I_{\text{sim}}, gI_{\text{sim}})$.

Reconstruction accuracy evaluation

Now, that synthetic data can be generated, the question arises which concrete questions the simulations can evaluate. The major concern is the accuracy of the results obtained from ROFL. For that purpose, a high number of artificial noisy datasets with the same ground truth has to be statistically analyzed. Here, the reconstruction accuracy of the fiber orientation and the relative section thickness, is of interest. Therefore, for all combinations of $t_{\text{rel}} = 0.01, 0.02, \dots, 0.9$, $\alpha = 0^\circ, 1^\circ, \dots, 89^\circ$ and direction $\varphi = 45^\circ$ 100.000 samples of noisy light intensities each were generated. These parameter configurations represent a vast possibility of birefringence strengths given by t_{rel} and inclinations given by α . Negative inclinations are not necessary as the inclination sign does not have an effect of the reconstruction precision. Only one direction angle is simulated as the direction angle has only a small effect on the reconstruction accuracy [46]. The chosen direction angle of $\varphi = 45^\circ$ represents the worst case scenario for a measurement with four tilting positions as the angle between the orientation vector and the rotated orientation vector becomes maximal for $\varphi = \psi$ and decreases with $|\varphi - \psi|$.

The synthetic datasets are then analyzed with the ROFL and the DFT algorithms resulting in a reconstructed fiber with fiber orientation $\mathbf{r}_{\text{rec}}$ and relative thickness t_{rec} for each algorithm. A comparison with the LORIE algorithm is not possible due to the excessive computational demands of the LORIE implementation presented in [46]. The reconstructed fibers can then be evaluated against the ground truth. For the orientation, a suitable measure is the angular deviation between the reconstructed $\mathbf{r}_{\text{rec}}$ and the ground truth orientation $\mathbf{r}_{\text{gt}}$. The angular deviation is calculated as the acute angle γ between both vectors:

$$\gamma = \arccos(\mathbf{r}_{\text{rec}} \cdot \mathbf{r}_{\text{gt}}) \quad (4.12)$$

The overall error for a parameter set is then given by the mean angular deviation $\langle \gamma \rangle$ of all 100.000 samples.

As the t_{rel} parameter was shown to carry structural information about the brain tissue not available via the transmittance and retardation, its reconstruction accuracy also needs to be examined. As a measure of reconstruction accuracy the absolute error $\sigma_{t_{\text{rel}}}$

was chosen: $\sigma_{t_{\text{rel}}} = |t_{\text{rec}} - t_{\text{gt}}|$. The average error for a parameter set is given by the mean absolute error of all samples $\langle \sigma_{t_{\text{rel}}} \rangle$.

The performance of the DFT algorithm strongly depends on the gradient of the retardation with respect to the inclination as shown in [46]. Therefore this relationship is studied further for an evaluation of the simulation results. The retardation gradient is investigated by simply plotting it for the whole parameter range of α from 0° to 90° and t_{rel} from 0 to 1. The gradient is calculated as

$$\begin{aligned} \frac{\partial \sin \delta}{\partial \alpha} &= \frac{\partial}{\partial \alpha} \sin \left(\frac{\pi}{2} t_{\text{rel}} \cos(\alpha)^2 \right) \\ &= -\pi t_{\text{rel}} \sin(\alpha) \cos(\alpha) \cos \left(\frac{\pi}{2} t_{\text{rel}} \cos(\alpha)^2 \right). \end{aligned} \quad (4.13)$$

Inclination bias evaluation

Besides a potentially low reconstruction accuracy, another pitfall of reconstruction techniques can be a biased reconstruction. In our case, it is crucial that the full range of possible inclinations can be estimated. The developed algorithm should not prefer some inclinations over others. To test the ROFL and DFT algorithms against biases in their inclination determination, 500,000 vectors uniformly distributed on the unit sphere were computed. From the orientation vectors, inclination and direction angles were derived. For these samples synthetic 3D-PLI datasets were generated as before with a relative thickness of $t_{\text{rel}} = 0.5$. The resulting datasets are again analyzed with the ROFL and DFT algorithms. A biased reconstruction would become apparent in the distribution of the reconstructed inclinations. This is evaluated based on the histograms of ground truth inclinations and reconstructed inclinations.

Results

Reconstruction accuracy

The results of the reconstruction accuracy evaluation are depicted in Fig. 4.3: it shows the mean orientation deviation and mean relative error of t_{rel} as a function of ground truth inclination and t_{rel} . The mean orientation error achieved by the ROFL algorithm follows a clear trend: for $t_{\text{rel}} > 0.06$, it increases with increasing inclination. Its dependency on the relative thickness is also simple: it decreases with increasing t_{rel} for $t_{\text{rel}} > 0.06$. For very low values of t_{rel} and very high inclinations, the accuracy strongly deteriorates. For very low relative thicknesses, the mean orientation error is even higher than 30° . Highly inclined fiber orientations of $\alpha > 80^\circ$ express a mean orientation error of 12° on average. On the other side, for flat fibers with respect to the sectioning plane of $\alpha \approx 0^\circ$, the reconstructed orientation lies very close to the ground truth: the minimal resulting error is 1° for $\alpha = 0^\circ$ and $t_{\text{rel}} = 0.9$. For the DFT algorithm, the mean orientation error is valley-shaped with respect to the inclination. The minimal error is achieved for inclinations of ca. $40^\circ - 60^\circ$ and from there increases strongly for highly inclined fibers and less strongly for in-plane fibers. With respect to the relative thickness, a similar behaviour as for the ROFL algorithm is observed: in general the accuracy increases with increasing t_{rel} . An exception the combination of very high values of t_{rel} and in-plane inclinations for which the accuracy decreases again stands out. For all simulated parameter sets, the ROFL algorithm achieves a lower mean orientation error than the DFT algorithm. For white matter structures with $t_{\text{rel}} \in [0.2, 0.9]$ and $\alpha \in [0^\circ, 80^\circ]$, the orientation reconstruction accuracy achieved by the algorithms are on




FIGURE 4.3: Reconstruction accuracy of ROFL and DFT algorithms evaluated on synthetic data. **Left:** ROFL algorithm. **Right:** DFT algorithm. **Top:** orientation reconstruction error $\langle \gamma \rangle$ as a function of relative thickness t_{rel} and inclination angle α . **Bottom:** absolute reconstruction error of the relative thickness $\langle \sigma_{t_{\text{rel}}} \rangle$ as a function of relative thickness t_{rel} and inclination angle α .

average 2° for ROFL and 4.5° for DFT with maximal values of 9.5° for ROFL and 18.8° for DFT.

The behaviour of the reconstruction error of the relative section thickness t_{rel} is similar to the reconstruction error of the fiber orientation: it increases with increasing inclination and decreasing relative thickness. In contrary to the orientation error, the absolute error for t_{rel} does not take the form of a valley for the DFT algorithm. Again, the ROFL algorithm achieves a reduced error with respect to the DFT algorithm: on average, the relative thickness of white matter fibers with $t_{\text{rel}} \in [0.2, 0.9]$ and $\alpha \in [20^\circ, 90^\circ]$ is determined with an average absolute error of 0.03 compared to 0.05% for the DFT algorithm. For very steep fibers, both algorithms are not able to determine a plausible value as the absolute error exceeds 0.2. For very low relative thicknesses, the error also increases. While the absolute error does amounts to between 0.02 and 0.1, for a ground truth value of $t_{\text{rel}} = 0.01$ this represents a relative error greater than 1.

For a better understanding of the reconstruction accuracy of the DFT algorithm, the retardation gradient is investigated. The resulting plot of the absolute gradient as a function of relative thickness and inclination is depicted in Fig. 4.4. It can be observed that the gradient strongly depends on the inclination angle: from a value of zero for $\alpha = 0^\circ$ it increases to a maximum between ca. 45° and 60° from which it decreases again to zero for $\alpha = 90^\circ$.




FIGURE 4.4: Absolute retardation gradient with respect to inclination $\left| \frac{\partial \sin \delta}{\partial \alpha} \right|$ as a function of inclination and relative thickness t_{rel} .

Reconstruction bias

The histograms obtained from the simulation of 500,000 uniformly distributed orientations on a sphere are depicted in Fig. 4.5. The ground truth inclinations follow a $\cos(\alpha)$ distribution which is validated by a $a \cdot \cos(\alpha)$ fit for a to the histogram. The histogram of the inclinations obtained from the ROFL algorithm displays a good agreement with the ground truth histogram. For the DFT algorithm two inconsistencies between its inclination histogram and the ground truth are observed. For in-plane fibers which are the most probably orientations, the frequency of reconstructed inclinations drops to almost zero for $\alpha \approx 0^\circ$. Instead of decreasing for $|\alpha| > 0^\circ$, the histogram displays two symmetric peaks close to $\alpha = 0^\circ$ which are indicated by black arrows. The other discrepancy between ground truth and DFT inclinations lies in the regime of highly inclined fibers: for these, the frequency of DFT inclinations drops moderately below the ground truth frequency as highlighted by the blue arrows.

Discussion

The simulations prove the working principle of the algorithm. Also, an improved reconstruction accuracy compared to the analytical DFT approach especially for in-plane and out-of-plane orientations, could be shown. The varying reconstruction accuracy of the DFT algorithm can be explained with the gradient strength plotted in Fig. 4.4. For a constant t_{rel} value the retardation gradient becomes maximal for inclinations between 45° and 60° . In this inclination range the orientation reconstruction error of the DFT algorithm becomes minimal (cf. Fig. 4.3). For in-plane fibers and out-of-plane fibers the gradient becomes zero. In consequence the reconstruction accuracy drops sharply.

The inclination bias simulation revealed that the DFT algorithm is in fact not capable of determining in-plane orientations. Here lies the the major advantage of the ROFL algorithm as it is capable of reconstructing these orientations. For uniformly distributed orientations, in-plane fibers are far more likely than highly inclined fibers which emphasizes the importance of the capability to reconstruct in-plane orientations.

The $\cos(\alpha)$ dependency of the frequency of inclinations for uniformly distributed orientations follows from the spherical geometry (cf. Fig. 4.6). For a given inclination, the frequency of orientations is proportional to the circumference of the circle of a sphere




FIGURE 4.5: Inclination bias evaluation. Plotted are inclination histograms for ground truth (black), the results of the ROFL algorithm (green) and the results of the DFT algorithm (blue). Bin width: 0.25° . Cyan: $a \cdot \cos(\alpha)$ fit to ground truth histogram, fit result: $a = 0.0087$. The arrows point out discrepancies between the histograms of the ground truth and the results of the DFT algorithm for in-plane fibers (black) and out-of-plane fibers (blue). Taken from [84]

at the respective inclination. The radius of this circle is exactly $\cos(\alpha)$. With the prior belief that all orientations should be equally likely, the $\cos(\alpha)$ proportionality can serve as a prior for the inclination in a Bayesian analysis. In [22] this fact was utilized for a Bayesian analysis of DMRI measurements. In chap. 5 the prior $p(\alpha) \propto \cos(\alpha)$ will be incorporated into the analysis of 3D-PLI data.

One pitfall of the inclination determination remains: very steeply inclined fibers with $|\alpha| > 80^\circ$ also pose a problem for the ROFL algorithm. As the measured sinusoidal signals have an almost vanishing amplitude for these orientations, the light intensity curves resemble a constant function plus random noise. Hence, very steep orientations are inherently challenging to interpret by the polarimetric model which tries to fit a sinusoidal model to the data. Nevertheless, assuming uniformly distributed orientations, the fraction of highly inclined fibers with $|\alpha| > 80^\circ$ adds up to a fraction of 1.5% of all orientations². Therefore we can conclude that for the vast majority of possible orientations in white matter regions with $t_{\text{rel}} > 0.2$, the ROFL algorithm achieves a high accuracy of 2° on synthetic data. Still, some brain regions will consist of large bundles of highly inclined fibers which remain hard to explain by 3D-PLI.

4.3 Evaluation for experimental data

The simulations performed in the last section proved the working principle of the ROFL algorithm on synthetic data. Nevertheless, for a true validation of the developed analysis it must be applied on experimental data.

²The fraction is calculated via

$$\eta = \frac{2 \int_{\frac{80\pi}{180}}^{\frac{\pi}{2}} \cos(\alpha) d\alpha}{\int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\alpha) d\alpha} \approx 1.5\% .$$




FIGURE 4.6: Geometrical explanation for the $\cos(\alpha)$ dependency of the inclination distribution of uniformly distributed orientations. The infinitesimal probability for an orientation vector to be inclined by α is proportional to the circumference of the circle of a sphere at the respective inclination. The radius of this circle is given by $\cos \alpha$.

Methods

An ideal experimental validation would be a phantom measurement with a known ground truth which is not available for 3D-PLI. At LAP resolution the derived properties of the nerve fiber tissue can therefore only be assessed qualitatively. At the micrometer scale, the results can be compared to higher resolved volumetric measurement techniques such as TPFM which is discussed in chap. 7. These measurements have the disadvantage of small field of views which do not allow to analyze major fiber tracts through the brain volume in contrast to LAP datasets which are analyzed in this section.

As tissue sample the occipital lobe of a human right hemisphere consisting of 843 coronal sections was utilized. Data was acquired at a pixel size of $64 \times 64 \mu\text{m}$ in the LAP with the routine measurement protocol of 18 polarization angles and four tilting positions. All sections were processed with the DFT and ROFL algorithms. To enable an assessment of the three-dimensional fiber orientations, all but the first 75 sections were reconstructed to a 3D volume. For the reconstruction, no new image registration techniques were developed for 3D-PLI in the scope of this thesis, instead existing methods and tools were utilized. Details on the registration procedure are given in App. B.

Theoretically, the reconstructed vector field of the brain volume would allow a tractography study for a qualitative validation of the derived fiber orientations. While tractography is a standard procedure in DMRI, it is substantially more difficult to apply to 3D-PLI. The main pitfall are registration errors which would induce strong artifacts. From a computational viewpoint, handling tractograms becomes challenging already for DMRI datasets due to the data size and visualization difficulties [115]. This issue becomes even more severe for 3D-PLI datasets which exceed DMRI dataset sizes by at least two orders of magnitude. For these reasons, the derived fiber orientations were evaluated based on visual inspection of the fiber orientation vector field as in [46]. A major fiber tract which runs almost perpendicularly to the coronal sectioning plane is the stratum sagittale (*StS*). To assess the continuity of the fiber orientation, the reconstructed volume was virtually resliced to visualize the horizontal plane. For one horizontal slice, the projection of the vector fields into the horizontal plane were then plotted for selected ROIs located in the *StS*. Based on the synthetic datasets, the ROFL algorithm outperforms the DFT algorithm for highly inclined fiber orientations. This can be evaluated for the human brain data based on the plotted vector fields.

The simulations revealed a distinct difference between the reconstructed orientations

for in-plane oriented fibers computed by the ROFL and the DFT algorithm. This difference was examined for experimental data by analyzing the inclination distribution for a region with dominant in-plane fiber orientations. As the inclination histograms are almost unaffected by the 3D registration process, the same ROI which was analyzed to showcase the difficulties of DFT and LOriE algorithms (cf. Fig. 2.9) was designated for the analysis. This way all three algorithms can be compared which was not possible for the simulation study. Another region for which differences between ROFL and DFT are likely to be observed is the cortex due to the low myelination. Therefore, for one ROI at the boundary of white and grey matter the relative thickness and fiber orientation maps resulting from both algorithms were compared.

A quantitative measure for the accuracy of the model parameters is given by the mean squared error (MSE) between the model prediction and the measurement data. In [46], the MSE between the predicted and measured retardation values was utilized to compare DFT and LOriE for individual pixels. Here, the MSE between the measured light intensities after registration onto the planar measurement and the light intensities predicted by the oblique 3D-PLI model (cf. eq. (4.3)) is utilized as the retardation values neglect the changes in the direction angles of the different tilting positions. For one exemplary brain section, the MSE was computed for the results of both ROFL and DFT and were compared based on their histogram.

Results

A horizontal view through the reconstructed brain volume is provided in Fig. 4.7. In the retardation map, the boundary of the stratum sagittale is easily identifiable. The vectorfield shown in Fig. 4.7 C shows that in the StS a large number of orientation vectors is oriented along the course of the bundle³. Four regions of interests were then further analyzed in detail (cf. Fig. 4.8).

In the first ROI, one very dominant fiber orientation inclined by app. 45° with respect to the coronal sectioning plane was estimated. This orientation fits the direction of the tract from the boundaries visible in the retardation map (cf. Fig. 4.7). ROFL and DFT results are barely distinguishable.

The second ROI displays two differently oriented structures: one runs in the horizontal plane while to the left another fiber orientation running in the coronal sectioning plane can be observed. In a direct comparison of the perpendicular bundle (green colored vectors), the DFT algorithm estimates a low number of fiber orientations which do not follow the outer structure of the bundle (indicated by white arrows). The orientation of the right bundle also agrees with the overall orientation of the bundle based on the retardation map (cf. Fig. 4.7). For the in-plane bundle (blue vectors), ROFL generally estimates a stronger in-plane component than DFT (indicated by red arrows).

For the third ROI, no dominant fiber orientation is recognizable. Both algorithms reconstruct an inhomogeneous vector field. Again the ROFL vector field expresses a stronger in-plane component.

The fourth ROI shows a region of a strongly inclined nerve fiber bundle. Strongly inclined fiber orientations agree with the impression that the bundle runs almost perpendicularly to the sectioning plane (cf. Fig. 4.7). For individual pixels, differences

³For better visibility, the coloring of the vectors was changed compared to typical 3D-PLI by switching green and blue.




FIGURE 4.7: Overview of the analyzed human brain volume. **A**: Blockface volume. The white plane indicates the analyzed plane. **B**: Retardation map of the plane whose vector field is investigated. **C**: Estimated fiber orientation vector field in the region of the stratum sagittale. Every 6th vector is mapped. The ROIs indicated by the four rectangles are shown in detail in Fig. 4.8.

between the algorithms can be observed. While for some of these pixels ROFL seems to derive stronger inclined fibers (white arrows), for other pixels it results in less inclined nerve fibers (red arrows).

The inclination histograms obtained from a region of interest are depicted in Fig. 4.9. For the DFT algorithm, the relative frequency of inclinations strongly decreases for $|\alpha| \rightarrow 0^\circ$. High absolute inclinations are much more likely than for the ROFL and LOriE algorithms. The histogram of the inclinations reconstructed by the LOriE algorithm displays a strong peak for $\alpha \approx 0^\circ$ which disrupts the otherwise smooth distribution. For the ROFL algorithm, no distinct features are observable for $\alpha \approx 0^\circ$: the histogram is continuous.

The agreement between model and data is evaluated based on the distributions of the mean squared error metric (*MSE*) for an exemplary brain section depicted in Fig. 4.10. The *MSE* values obtained with the parameters estimated by ROFL are significantly lower than the *MSE* values obtained with the DFT estimates. The results of the tilting analysis for a cortical region are shown in Fig. 4.11. While the majority of grey matter pixels express very low relative thicknesses, both DFT and ROFL result in a significant number of unconnected pixels with high values of t_{rel} in the cortex. The FOMs display a strong variation in the cortex as well. Of special interest is also the measurement artifact pointed out by the white arrow: it originates from a dust particle. For the DFT algorithm it has severe consequences as the estimated values of t_{rel} strongly disagree with surrounding pixels. In the maps computed by ROFL the artifacts are not recognizable. The bottom row of Fig. 4.11 depicts the vector field on top of the relative thickness map. Both vector fields agree strongly in white matter and disagree in grey matter.

Discussion

The analysis of tilted LAP measurements was evaluated utilizing a large human brain volume. The reconstructed brain volume represents the first large scale fiber model obtained from 3D-PLI. Based on the three-dimensional fiber orientation vector fields, it could be shown that even fiber tracts which run almost perpendicular to the sectioning plane such as the stratum sagittale, can be restored from the individual sections. Still, the derived fiber orientations do not match the apparent orientation of the fiber bundle visible in the retardation map everywhere.

Examples for a strong agreement between expected geometric orientation of the bundle and the fiber orientation obtained from 3D-PLI are the vector fields displayed in Fig. 4.8 A and D. In contrast the vector field from Fig. 4.8 B expresses two differently oriented bundles, one running in the sagittal and one in the coronal plane, which was already observed for another ROI of the stratum sagittale in [46]. One explanation for this could be that the sagittally oriented bundle is part of the internal stratum sagittale and the coronally oriented bundle part of the external stratum sagittale. In fact, the stratum sagittale does not consist of only one fiber bundle but at least three bundles [116]. Another possible explanation is that this region contains fiber crossings which first become visible at a resolution provided by 3D-PLI. Fiber crossings are not correctly described by the current model which presumes one fiber orientation vector. Further analysis of the obtained fiber orientations with help of neuroanatomical experts might enable to find the correct explanation. The vector fields shown in Fig. 4.8 C express random fiber orientations while the bundle seems to be inclined by an angle of app. 60° . As the reconstruction for synthetic fibers inclined by 60° was quite robust and




FIGURE 4.8: Comparison of vectorfields derived from the ROFL and DFT algorithms. **A**: ROI consisting of one dominant fiber orientation approximately inclined by 50° with respect to the sectioning plane. **B**: ROI consisting of two differently oriented fiber bundles. **C**: a ROI which mainly consists of nerve fiber crossings. **D**: ROI of highly inclined nerve fibers with respect to the sectioning plane. White arrows indicate fiber orientations for which ROFL estimates a larger inclination angle than DFT, red arrows indicate fiber orientations for which ROFL estimates a smaller inclination angle than DFT.




FIGURE 4.9: Inclination histograms for one ROI of in-plane oriented nerve fibers obtained from the ROFL, DFT and LOriE algorithms.




FIGURE 4.10: Histograms and kernel density estimates (full lines) of the mean squared errors between the measured light intensities and the light intensities predicted by the oblique 3D-PLI model for one exemplary brain section. Blue: 3D-PLI parameters estimated by the DFT algorithm. Orange: 3D-PLI parameters estimated by the ROFL algorithm. The mean squared errors are clearly lower for the ROFL algorithm.

the higher inclined vector field from Fig. 4.8 D was reconstructed very well, the random fiber orientations most likely originate from fiber crossings.

For synthetic data, the ROFL algorithm clearly outperformed the analytical DFT algorithm. In the vector fields, differences are not that obvious to spot. Still, especially for the steep fiber bundles shown in Fig. 4.8 B and D the vector fields obtained from ROFL appear more coherent and fit the general orientation of the bundle slightly better. Differences were also observed for flat fibers with respect to the sectioning plane. For these, the investigation of the inclination histograms showed that ROFL indeed enables a reliable derivation of the inclination angle in contrast to the DFT and LOriE algorithms. In the vector fields, discrepancies for flat fibers occur for the flat fiber bundle from Fig. 4.8 B. However, as discussed in the last paragraph it is not clear if this particular ROI actually contains two distinct bundles or fiber crossings. Thus, it cannot be evaluated which interpretation is more realistic. From a statistical point of view, the parameters derived by ROFL provide the better prediction due to the lower mean squared errors.




FIGURE 4.11: Analysis of tilted 3D-PLI measurements for the cortex. The investigated ROI is indicated by the red rectangle in the full section view on the left. **Middle:** DFT results. **Right:** ROFL results. **From top to bottom:** relative thickness maps, fiber orientation maps, relative thickness maps overlaid with the vector field of every third vector. The white arrows indicate a dust particle which becomes apparent as an measurement artifact. While the ROFL result is almost unaffected by it, the DFT algorithm propagates the measurement artifact into the model parameters. In the cortex, both relative thickness and fiber orientation maps contain noise.

As ROFL is designed to find the optimal parameters through optimization and the DFT algorithm based on analytical calculus, this result is not surprising.

The performance of the ROFL and DFT algorithms for very low signals and in presence of measurement artifacts was investigated for one ROI at the boundary of the cortex. For the dust particle, it was found that ROFL is still able to provide a comparable interpretation compared to surrounding pixels while the DFT result does not fit the surrounding pixels at all. Due to the analytical nature of the DFT algorithm, any measurement artifact is directly propagated into the results while ROFL still tries to find the parameters that describe all measurement data the best. In future, the influence of outliers could be further reduced utilizing automatic outlier detection techniques similar to algorithms available for DMRI [117].

Regarding the cortex it was shown that both algorithms result in noisy maps for relative thickness and fiber orientation. Again, the maps obtained from ROFL appear slightly more coherent. Still this observation might be subjective. Just like for the synthetic data, the random parameters obtained from the analysis can be explained by the vanishing sinusoidal signal for unmyelinated tissue such as the cortex. As the cortex does

not contain large fiber bundles but mostly thin or even individual nerve fibers, it is probably impossible to derive reliable fiber orientations at a pixel size of $64 \times 64 \mu\text{m}$. The same does not apply for the relative thickness parameter: as it is proportional to the birefringence strength, it is expected to vanish for the cortex at LAP resolution. Based on this assumption, the t_{rel} map obtained from ROFL is more likely to represent a realistic interpretation while it still contains a large amount of noise. However, the only possibility to overcome the random parameters estimated for the cortex is the utilization of prior information which is pursued in chap. 5.

4.4 Development of a high performance implementation of the fitting algorithm

The computation time for the ROFL algorithm for one of the brain sections analyzed in the last section with app. 1 Mio. pixels amounts to 4 Core hours using the pure Python implementation described in section 4.1. Based on these numbers, the computation time would add up to more than 10.000 Core hours for a whole brain measured with the LAP. While in principle this could still be accepted, the computational demands become infeasible for microscopical data. An image of a coronal brain section measured with the LMP consists of up to 500 Mio. pixels which would result in a computation time of 2.000 core hours for a single section measured with the LMP3D. Analyzing whole brains with the current implementation of ROFL then becomes clearly computationally prohibitive.

In this section, two strategies to reduce the computation time are explored. The first one utilizes the partial compilation of the Python code while still using CPUs. The second approach exploits the massively parallel computational power of GPUs. The results of the different implementations and their computation times are then evaluated for experimental data. The developed GPU implementation was presented by Jan-Oliver Kropp at NVIDIA GPU Technology Conference 2019 Washington [107].

Methods

CPU based optimization

Python's interpreted nature makes it very popular for rapid prototyping. Compared to compiled languages interpreted languages are unfortunately typically slow as the program is not compiled to efficient machine code. For example, in a loop the Python interpreter has to check the data types in every iteration. The scientific ecosystem of python is centered around NumPy which provides fast array computations and relies on compiled C code for speed. Many other packages such as scipy [111] provide wrappers of C++ or Fortran routines or are written in Cython, a language to write C code in Python [118]. The same applies to the ROFL algorithm: the cost function and its gradient are written in NumPy while the Levenberg-Marquardt optimization from scipy calls a Fortran routine of the MINPACK library [113].

In recent years, a new approach to speed up Python programs which is pursued by packages such as theano [119], tensorflow [120] and pytorch [121] is to compile the Python code to efficient machine code on the CPU or even graphic cards via computational graphs. While in principle promising for general purpose high performance computing in Python, these libraries were primarily designed for deep learning applications on graphics cards and are not typically used for CPU intensive tasks on non-shared memory

parallel systems such as supercomputers. Furthermore, using one of these frameworks would require a complete rewrite of ROFL and none of them includes the Levenberg-Marquardt algorithm, therefore a different strategy was chosen⁴.

The first step for speeding up numerical computer programs is to identify the computational bottlenecks. Optimization processes typically spend most of their time calculating the cost function and its gradient [122]. Speeding up this part of the code can then possibly result in a significant speedup of the whole optimization process. One possibility would be to write the cost function and its gradient in a compiled language and call it from the main Python program. A simpler alternative is to compile the python code to machine code using the *numba* library [123]. Numba enables just-in-time compilation of Python and NumPy code based on the LLVM compiler infrastructure [124]⁵. Just-in-time compilation works in the following way: when a function is executed for the first time, the compiler inspects the types of all inputs and outputs and compiles it to machine code for the inferred types. In subsequent calls of the function, the function is replaced by its compiled version. Compared to the frameworks discussed in the last paragraph, numba also generates efficient machine code but allows a very convenient integration into existing numpy and scipy based code without the need to rewrite substantial parts of the existing program.

Together with Felix Matuschke who uses ROFL to analyze simulated 3D-PLI measurements, the cost function of the ROFL algorithm and its gradient were adapted so that they could be compiled by Numba which only required minimal code changes. All other parts of the code including the parallelization were left untouched. As only parts of the code were replaced by compiled versions, all functions return the same results up to floating point arithmetic errors.

For a performance analysis, I/O operations and preprocessing were not included in the runtime measurements as they differ from the GPU implementation described in the next section and are not fully standardized yet. Preprocessing includes the exclusion of background pixels from the analysis if the user provides a mask of the brain section and the reshaping of the data into the shapes required by scipy's *leastsq* optimization routine. The runtime measurement starts at distributing the data to the different processes and ends after gathering the results from the individual processes before writing the parameter maps. As the key point is the difference between interpreted and partially compiled code and not the scaling behavior for a large number of CPUs, the runtime was measured for 1 CPU and 16 CPUs. The final runtimes were obtained by averaging over 10 individual runtime measurements⁶.

GPU based optimization

Originally, graphic cards or graphics processing units (GPUs) were developed for computer graphics applications, especially rendering. Due to the massive parallelism of rendering problems, they were designed differently than CPUs. Instead of one processor with a very high clock rate in the GHz range like CPUs GPUs typically consist

⁴While other optimization algorithms can be employed, the Levenberg-Marquardt algorithm is the de-facto standard algorithm for least squares problems due to its efficiency.

⁵LLVM is the abbreviation for Low Level Virtual Machine, the original name of the project. As the scope of the project was extended beyond the typical understanding of virtual machines in software engineering, this name was officially dropped as stated by Chris Clattner, one of the original initiators: <http://lists.llvm.org/pipermail/llvm-dev/2011-December/046445.html>, accessed: 31.03.2020.

⁶Details about the employed hardware are given in App. C.

of thousands of processors with a clock rate of several MHz. The total possible arithmetic throughput of one GPU is therefore higher than the combined throughput of several CPUs for massively parallel problems. Hence, graphic cards offer an appealing alternative to CPU clusters for high performance computing. In recent years, GPU programming has been adopted in many disciplines of science and achieved enormous speedups over traditional cluster computing [125, 126, 127]. For our purposes, a NVIDIA GTX 1080 is available. The most common framework for programming NVIDIA GPUs is CUDA [128] which enables to write C and C++ code with additional GPU specific instructions.

The parameter estimation problem for the ROFL algorithm is massively parallel due to the pixelwise parallelism which makes it suitable for GPUs. The same computational problem, solving a large number of small independent optimization problems, exists in other fields such as MRI, computer vision and microscopic applications and has successfully been addressed using GPUs. For example, the microstructural diffusion parameter estimation could be accelerated by two orders of magnitude in [129, 130], Smith et al. [131] achieved real-time MRI image reconstruction and Bruce et al. [132] accomplished real-time 3D deconvolution using GPU hardware. One illustrative example for the usage of least squares fitting on GPUs is real-time relighting and reshading of videos based on a target scene [133].

For a GPU version of the ROFL algorithm the main difficulty lies in a robust implementation of the Levenberg-Marquardt algorithm as the CUDA standard libraries do not contain nonlinear optimization algorithms. Przybylski et al. [134] developed a well-documented general purpose curve-fitting library (*gpufit*) in CUDA based on the Levenberg-Marquardt algorithm for the parallel fitting of millions of independent datasets. A developer only needs to implement a new model and *gpufit* automatically handles the parallelization over all datasets. Furthermore, *gpufit* provides APIs⁷ to various programming languages such as Python, Matlab [135] and Java and manages the data transfer to the GPU and back to RAM after fitting. All these features make it a convenient candidate for the purpose of a GPU accelerated ROFL algorithm.

The GPU Optimized ROFL algorithm (*GOROFL*) was implemented prototypically by a student, Oliver Kropp, under my supervision. Input and output operations (I/O) were kept in Python but all computations transferred to the GPU. The centering of the light intensity profiles and the brute force minimization utilize the *pyCUDA* framework [136]. PyCUDA enables to write CUDA kernels and execute them from Python via a simple API.

Step by step the GOROFL algorithm works in the following way (cf. Fig. 4.12): the calibrated light intensities are read into system RAM. Next, the data are preprocessed. This includes the exclusion of background pixels from the analysis if the user provides a mask of the brain section using simple numpy array routines and the rearranging of the data into the shapes required by *gpufit*. Then, all data are transferred to the GPU via PyCUDA. After calculating the centered light intensities and the weights the initial parameters are found via the grid search. The centered light intensities, weights and initial parameters are then copied to system memory. Finally, the Levenberg-Marquardt optimization is executed using *gpufit*'s Python binding and the best fit parameters are written to disk.

⁷In software engineering, API stands for Application programming interface and typically describes an interface between different software packages




FIGURE 4.12: Diagram of the GPU implementation of ROFL. Dashed lines depict data transfer between GPU and RAM. The measured light intensities are read into system memory via Python. Next, they are transferred to the GPU where PyCUDA calculates centered light intensities and executes the grid search. The centered light intensities, weights and brute force results are copied back to RAM and then passed to gpufit which carries out the Levenberg-Marquardt optimization on the GPU. The best fit parameter maps are then written to disk from Python. L-M optim.: Levenberg-Marquardt optimization. Image courtesy: O. Kropp

As gpufit's implementation of the Levenberg-Marquardt algorithm is not the same as MINPACK's which is employed in the CPU version, the algorithms might yield different results. Most importantly, gpufit executes the calculations in single precision while MINPACK utilizes double precision. CPU and GPU versions are compared based on the resulting parameter maps of fiber orientation and relative thickness of both implementations using their default optimization stopping criteria. Another criterion to evaluate the algorithms is the χ^2 measure which both algorithms seek to optimize. Especially in the cortex the algorithms need to be compared as finding the optimum of the cost function is substantially more challenging for grey matter than for white matter. Additionally to a runtime measurement of the whole program the computation times of the different steps of the algorithm are measured for a later analysis.

Results

Comparison of CPU and GPU results

The FOMs of the whole analyzed brain sections computed by ROFL and GOROFL are shown in Fig. 4.13 A. By eye, both maps are barely distinguishable in this view. The histograms of the χ^2 measure after the optimization process depicted in Fig. 4.13 B are also very similar. In Fig. 4.14 a ROI in the cortex is analyzed further. Here, differences are observable especially for the t_{rel} map. GOROFL estimates higher t_{rel} values than ROFL in grey matter pixels. The FOMs also express different fiber orientations for a significant number of grey matter pixels. In the χ^2 map however, no systematic difference can be found. For both algorithms, the map appears noisy in grey matter.




FIGURE 4.13: Comparison of ROFL and GOROFL results: global view. **A:** in a direct comparison of the FOMs both algorithms yield very similar results. The white rectangle points out the position of the ROI which is shown again in Fig. 4.14. **B:** the histograms of the χ^2 values of both algorithms are almost identical.

Performance analysis

The total runtimes of the original Python implementation, the Numba compiled Python implementation and the GPU implementation are listed in tab. 4.1. The compilation of cost function and its gradient via Numba results in a speedup of one order of magnitude. For the analyzed brain section the computation time reduced from 4 Core hours to 20 min which is only slightly higher than the noncompiled Python version running on 16 CPUs. On 16 CPUs, the total runtime reduces to ca. 1.5 minutes using the Numba accelerated implementation. The GPU implementation achieves another speedup of two orders of magnitude compared to the single threaded partially compiled Python version. The processing time for one brain section reduces to several seconds compared to 4 hours of the original serial Python implementation.

Implementation	Runtime
Python: 1 CPU	4 h
Python: 16 CPUs	20 min
Python/Numba: 1 CPU	18 min
Python/Numba: 16 CPUs	1.5 min
GPU: NVIDIA GTX 1080	5 sec

TABLE 4.1: Runtimes of different implementations of the ROFL algorithm for a brain section of app. 1 Mio. pixels.

In Fig. 4.15 the runtime shares of the different steps of the GOROFL algorithm are depicted. The full runtime amounts to 16 seconds out of which I/O operations make up more than 50%. The actual computations on the GPU (centering, brute force minimization and nonlinear optimization) amount to roughly one third of the total runtime.

Discussion

The objective of this section to reduce the computation time of the ROFL algorithm was achieved. The simple partial compilation of the Python code via Numba resulted




FIGURE 4.14: Comparison of ROFL and GOROFL results: local analysis in the cortex. The position of the ROI is indicated by the white rectangle in Fig. 4.13 A. While the results are barely distinguishable for white matter, in the cortex GOROFL estimates higher t_{rel} values.

in a 10 times faster execution with identical results. Using 16 CPUs, the total runtime for a section of the analyzed human brain dataset reduces to lower than two minutes. This enables the processing of all 843 sections on the institute’s server within several hours. The high speedup achieved by Numba’s just-in-time compilation coupled with its convenient usage were found so convincing that it was utilized in all other implementations presented in this thesis. Despite the ten fold faster execution accomplished by compiling the main bottleneck, an implementation in a compiled language would reduce the runtime even further. For the purposes of this thesis though Python coupled with Numba compilation was sufficient.

An even higher speedup was accomplished using GPU resources which reduce the total computation time to several seconds. This potentially eliminates the requirement of supercomputer resources for LAP datasets. Compared to the CPU implementation, the results of the GPU accelerated algorithm are almost equal for white matter but different for grey matter. As mentioned before fitting the 3D-PLI model is challenging for the cortex due to the very low retardation values. In consequence the lower floating point precision utilized by gpufit affects the results stronger than in white matter which is the most likely reason for the observed differences. Another reason might be differing stopping criteria for the optimization. However, based on the values of the optimized cost functions neither of both implementations has a distinct advantage. For vanishing signals which occur in the cortex the χ^2 measure might not be the best evaluation criterion though. Simulative studies could provide a better alternative. From a neuroanatomical point of view, the t_{rel} map computed by ROFL is more likely to represent the reality about the birefringence strength as at the resolution of the LAP individual nerve fibers in the cortex cannot be distinguished which is therefore expected to yield vanishing values of t_{rel} . Concerning the reconstruction accuracy of GOROFL, the precision issue can potentially be solved if future releases of the gpufit library support double precision.




FIGURE 4.15: Total runtime shares of the different steps of the GPU implementation of the ROFL algorithm for one brain section. LM: Levenberg-Marquardt.

Another approach could be to change the optimization strategy of the algorithm itself by employing global optimization algorithms or to try out several initial positions for the local optimization at the cost of computation time.

Using a NVIDIA GTX 1080, the actual computations on the GPU take a few seconds and less than I/O operations. On more powerful GPUs, the computation time would potentially reduce to less than one second which means that the main bottleneck for the whole program is reading the data and not the actual algorithm. Future optimizations would therefore be related to the data structures and the streaming to the GPU. Currently, the data handling is not fully optimized: GOROFL copies the data first once to the GPU for the grid search, back to system memory and then again to the GPU for the optimization. As copying data to the GPU is slow, eliminating the Python layer represents a clear starting point for future optimizations. Data structures and I/O operations are subject to standardization to ensure a stable workflow for the whole 3D-PLI processing pipeline in future. Using a suitable data format such as HDF5 [137] would for example enable parallel I/O operations. Another simple optimization would be to distribute the computations sectionwise across several GPUs.

It has to be noted that the runtime measurements carried out here do not represent a comparison of CPU and GPU on equal terms. This would require running the same code base on the different hardware architectures. Also, the implementations are not fully optimized due to the Python overhead. In particular, the execution of PyCUDA kernels requires just in time compilation which increases the processing time for centering and grid search. Nevertheless, it could be shown that GPUs offer great potential for pixelwise parallel model fitting problems in 3D-PLI.

5

Maximum A Posteriori Estimation of Nerve Fiber Orientations

“Humans are unbelievably data efficient. You don’t have to drive 1 million miles to drive a car, but the way we teach a self-driving car is have it drive a million miles.”

– Jeff Bezos

In chapter 4 a least squares approach for the analysis of 3D-PLI measurements from oblique views termed Robust Orientation Fitting via Least Squares (ROFL) was introduced and validated for experimental and synthetic data. ROFL outperformed the DFT and LOriE algorithms especially for in-plane nerve fibers. Despite a high reconstruction accuracy for weak to very high signals, very low signals still pose a challenge for the ROFL algorithm. Very low signals arise for either very low birefringence strengths which mainly occur in unmyelinated tissue and very highly inclined nerve fibers with respect to the sectioning plane. On experimental data, this becomes primarily apparent for the cortex where the t_{rel} and fiber orientation maps contain a significant amount of noise.

In this chapter, this problem is addressed by a Bayesian approach which incorporates prior information. In chapter 4 it was shown that a uniform orientation distribution prior leads to the nonuniform inclination prior $p(\alpha) \propto \cos(\alpha)$. This prior is utilized to extend the MLE based LOriE approach to a MAP estimator. As the LOriE framework results in artifacts for in-plane nerve fibers, first its numerical instabilities need to be solved.

The outline of this chapter is as follows: after a review of the LOriE algorithm, a new optimization approach to overcome its numerical difficulties is presented. As a next step, the new MAP based algorithm is developed. The new framework is then evaluated for synthetic and experimental data.

5.1 The Posterior Orientation Solver

A review of the LOriE algorithm

Theory

For the complete derivation of the Likelihood Orientation Estimation algorithm it is referred to [46]. Here, a short review will be given to gain a better understanding of the

MAP estimator which is subsequently introduced. The LORIE algorithm is based on a MLE approach. In contrast to the ROFL algorithm, it does not utilize normalized light intensities but instead the normalized Fourier coefficients of the respective tilted measurements as dependent variables of the model. In principle, Fourier coefficients carry all necessary information about a sinusoidal signal. The normalized Fourier coefficients A and B and their standard deviations σ_A and σ_B approximated by Gaussian error propagation of tilted measurement j are given by

$$A_j = \frac{a_{1,j}}{a_{0,j}} = \frac{2}{N} \frac{1}{\frac{1}{N} \sum_i I_{j,i}} \sum_i I_{j,i} \cos(2\rho_k), \quad (5.1)$$

$$B_j = \frac{b_{1,j}}{a_{0,j}} = \frac{2}{N} \frac{1}{\frac{1}{N} \sum_i I_{j,i}} \sum_i I_{j,i} \sin(2\rho_k) \quad (5.2)$$

$$\sigma_{A_j}^2 = \sum_i \left(\frac{\partial A_j}{\partial I_{j,i}} \sigma_{I_{j,i}} \right)^2 = \frac{4g}{N I_{T,j}} \left(1 - \frac{1}{2} A_j^2 \right) \quad (5.3)$$

$$\sigma_{B_j}^2 = \sum_i \left(\frac{\partial B_j}{\partial I_{j,i}} \sigma_{I_{j,i}} \right)^2 = \frac{4g}{N I_{T,j}} \left(1 - \frac{1}{2} B_j^2 \right) \quad (5.4)$$

According to the 3D-PLI model, the expectation values of these Fourier coefficients relate to the model parameters by

$$\langle A_j \rangle = -\sin\left(\frac{\pi}{2} t_{\text{rel},j} \cos(\alpha_j)^2\right) \sin(2\varphi_j) \quad (5.5)$$

$$\langle B_j \rangle = \sin\left(\frac{\pi}{2} t_{\text{rel},j} \cos(\alpha_j)^2\right) \cos(2\varphi_j) \quad (5.6)$$

where the notation introduced in chap. 4 was used. Assuming a Gaussian noise model, the likelihood to observe the coefficients A_j and B_j for a parameter set $\theta = (\varphi, \alpha, t_{\text{rel}})$ is now given by

$$p(A_j, B_j | \theta) = p(A_j | \varphi, \alpha, t_{\text{rel}}) \cdot p(B_j | \varphi, \alpha, t_{\text{rel}}) \quad (5.7)$$

$$= \frac{1}{\sqrt{2\pi\sigma_A^2(\theta)}} \exp\left(-\frac{(A_j - \langle A_j(\theta) \rangle)^2}{2\sigma_A^2(\theta)}\right) \frac{1}{\sqrt{2\pi\sigma_B^2(\theta)}} \exp\left(-\frac{(B_j - \langle B_j(\theta) \rangle)^2}{2\sigma_B^2(\theta)}\right). \quad (5.8)$$

The combined likelihood to observe all Fourier coefficients for all tilted measurements is then the joint density

$$p(A_0, B_0, \dots, A_{N_T}, B_{N_T} | \varphi, \alpha, t_{\text{rel}}) = \prod_{j=0}^{N_T} p(A_j | \varphi, \alpha, t_{\text{rel}}) \cdot p(B_j | \varphi, \alpha, t_{\text{rel}}). \quad (5.9)$$

The respective negative log-likelihood for all measurements which has to be minimized is given by

$$-\ln(p(\text{Meas.} | \theta)) = \frac{1}{2} \sum_{j=0}^{N_T} \ln(\sigma_{A_j}^2(\theta) \sigma_{B_j}^2(\theta)) + \frac{(A_j - \langle A_j(\theta) \rangle)^2}{2\sigma_{A_j}^2(\theta)} + \frac{(B_j - \langle B_j(\theta) \rangle)^2}{2\sigma_{B_j}^2(\theta)} \quad (5.10)$$

In [46] the log-likelihood was minimized by a grid search followed by local Nelder-Mead optimization [138, 139].

Now that the algorithm is reviewed, the origin of the observed artifacts for in-plane nerve fibers must be found. The derivation of the log-likelihood itself contains no obvious mistakes. While the derivation of the standard deviation of the Fourier coefficients is based on Gaussian error propagation this approximation cannot explain the problems of the approach for in-plane nerve fibers. A simple explanation might be a software bug which could not be found, however. Therefore, the problem must be a numerical issue. As pointed out in [46], the cost function is not convex. Non-convex optimization is challenging as the found local minimum is not necessarily the global minimum. In case of the LOriE algorithm, an optimization algorithm must also be flexible to cope with the different shapes the cost function assumes for different parameter values. As the artifacts only occur for very slightly inclined nerve fibers, the Nelder-Mead algorithm (N-M) just does not converge to the global minimum for this part of the parameter space. Hence, a new optimization approach is necessary to overcome the numerical instability.

A new approach to the optimization problem

Challenging optimization problems are often solved via global optimization techniques which rely on an exhaustive search of the parameter space. As this search requires a far higher number of calculations than a local optimization, their computation time is prohibitive for the purposes of 3D-PLI making local optimization the only viable option. Harms et al. [129] compared three different local optimization algorithms for DMRI models and found significant differences between the results.

Following that idea, the LOriE algorithm could also be tested with various local optimization algorithms. Fortunately, the first tested optimization algorithm, the L-BFGS-B algorithm [140, 141, 142] which is scipy's default optimization algorithm for constraint optimization problems, already proved to solve the problem. In contrast to the derivative-free N-M algorithm, the L-BFGS-B algorithm utilizes gradient information and an approximation of the Hessian¹. Furthermore it is able to deal with bounded parameter spaces which is required by the constraint that $t_{\text{rel}} \in [0, 1)$. The gradient of the log-likelihood was calculated via a finite difference approximation.

The effect of the choice of optimization algorithm is demonstrated by means of the same ROI which was already analyzed for the comparison of ROFL, DFT and LOriE algorithms in chapter 4. The inclination maps obtained from N-M and L-BFGS-B optimization are depicted in Fig. 5.1 top. Pixels with in-plane fibers where $|\alpha| < 0.1^\circ$ are highlighted in red: the L-BFGS-B algorithm estimates almost no such in-plane fibers in contrast to the N-M algorithm. The high number of in-plane fibers also becomes apparent in the inclination histogram in Fig. 5.1 bottom which is not smooth for the N-M algorithm. The histogram obtained from L-BFGS-B optimization yields a smooth distribution and strongly agrees with the ROFL algorithm.

¹The utilization of the second derivatives might be the reason that the L-BFGS-B algorithm is successful in our case. The retardation gradient with respect to the inclination is zero for in-plane oriented nerve fibers (cf. eq. (4.13)). On the other hand, the second derivative given by

$$\begin{aligned} \frac{\partial^2}{\partial \alpha^2} \sin\left(\frac{\pi}{2} t_{\text{rel}} \cos^2(\alpha)\right) &= \frac{\partial}{\partial \alpha} - \pi t_{\text{rel}} \sin(\alpha) \cos(\alpha) \cos\left(\frac{\pi}{2} t_{\text{rel}} \cos^2(\alpha)\right) \\ &= -\pi t_{\text{rel}} (\cos^2(\alpha) - \sin^2(\alpha)) \cos\left(\frac{\pi}{2} t_{\text{rel}} \cos^2(\alpha)\right) \\ &\quad - \pi^2 t_{\text{rel}}^2 \sin^2(\alpha) \cos^2(\alpha) \sin\left(\frac{\pi}{2} t_{\text{rel}} \cos^2(\alpha)\right) \end{aligned} \quad (5.11)$$

is nonzero for $\alpha = 0$: $\frac{\partial^2}{\partial \alpha^2} \sin\left(\frac{\pi}{2} t_{\text{rel}} \cos^2(\alpha)\right) |_{\alpha=0^\circ} = -\pi t_{\text{rel}} \cos\left(\frac{\pi}{2} t_{\text{rel}}\right)$.




FIGURE 5.1: Investigation of the choice of optimization algorithm for the LORIE framework. **Top:** inclination maps from one ROI resulting from Nelder-Mead optimization (left) and L-BFGS-B optimization (right). For clarity the inclination is clipped to the range $[-30^\circ, 30^\circ]$. Pixels in red: $|\alpha| < 0.1^\circ$. **Bottom:** inclination histograms obtained from the ROFL algorithm and the LORIE algorithm with the two optimization algorithms for the WM ROI.

In contrast to before and to the ROFL algorithm, the birefringence strength t_{rel} is now constraint to the range $[0, 1]$ which avoids ambiguities in the model. Adding this constraint proved necessary as otherwise again artifacts occurred for in-plane fibers. A comparison of the results with and without the constraint is given in D.

In principle, a thorough comparison of ROFL and LORIE could be carried out to examine the performance of the fitting algorithms. A similar comparison of MLE and WLS estimator was carried out by Przybylski et al. [134] for Poisson distributed data. As the Poisson distribution is a special case of the negative binomial distribution where variance and expectation value are equal, their results are applicable for the case of 3D-PLI. The obtained result was practically equal performance of the estimators for high signal to noise ratio (SNR). For 3D-PLI, the SNR can be calculated as

$$\text{SNR} = \frac{\text{Signal}}{\text{Standard deviation}} = \frac{I}{\sqrt{gI}} = \sqrt{\frac{I}{g}}. \quad (5.12)$$

Hence, for typical light intensities of $I = 2000 - 5000$ and the gain factor $g = 3$,

the SNR lies between 25 and 40 meaning that LOriE and ROFL are expected to yield almost identical results. Therefore, in the next section the MLE based LOriE framework is upgraded to a MAP estimator.

Incorporation of prior information

So far, the LOriE algorithm makes no use of prior information. Prior information can be crucial for the parameter estimation where the likelihood itself does not provide enough information. For 3D-PLI, this is mainly the case for unmyelinated tissue such as the cortex where the assumptions of the model do not hold which results in noisy maps of birefringence and fiber orientation.

At the resolution of the LAP a legitimate prior assumption about the cortex is that the birefringence strength is very low and that all orientations are equally likely. One possible choice for a prior which follows from this assumption is a normal prior with an expectation value of zero for the relative thickness t_{rel} similar to the linear regression example presented in chap. 3. This would serve as a regularization term for t_{rel} . While this prior can be justified for unmyelinated tissue it is invalid for pixels which contain very highly inclined fibers and consequently also express very low birefringence signals.

Therefore, the uniform orientation prior is chosen which leads to a prior for the inclination: $p(\alpha) \propto \cos(\alpha)$. With the log-likelihood l_{like} given by eq. (5.10) the application of Bayes' theorem (see eq. (3.11)) gives the log-posterior l_{post}

$$-l_{\text{post}} \propto -\ln(p(\alpha)) - l_{\text{like}} \quad (5.13)$$

$$= -\ln(\cos(\alpha)) + \frac{1}{2} \sum_{j=0}^{N_T} \ln(\sigma_{A_j}^2(\theta) \sigma_{B_j}^2(\theta)) + \frac{(A_j - \langle A_j(\theta) \rangle)^2}{2\sigma_{A_j}^2(\theta)} + \frac{(B_j - \langle B_j(\theta) \rangle)^2}{2\sigma_{B_j}^2(\theta)}. \quad (5.14)$$

As the logarithmic prior is a simple smooth and monotonic function, it does not complicate the optimization. Hence the posterior can be minimized in the same way as the likelihood: a grid search followed by L-BFGS-B optimization. This newly developed MAP estimator is termed *Posterior Orientation Solver (POriSo)* from here on. Following the experience with Numba and Python from sec. 4.4, the algorithm was implemented in Python using NumPy and Numba compilation for the cost function and its approximated gradient and mpi4py for pixelwise parallelization.

5.2 Evaluation on synthetic data

Methods

The effect of the inclination prior on the estimation of the parameters needs to be evaluated against a known ground truth. As for the evaluation of the ROFL algorithm, synthetic data can serve as ground truth. In sec. 4.2 the reconstruction accuracy of the ROFL and DFT algorithms was evaluated based on 100.000 samples of synthetic 3D-PLI signals for each parameter configuration. As the effect of priors become the strongest for extremely low signals, here also the cases of perpendicular inclination $\alpha = 90^\circ$ and zero relative thickness $t_{\text{rel}} = 0$ are studied. For each parameter combinations of $t_{\text{rel}} = 0, 0.01, \dots, 0.8$, $\alpha = 0^\circ, 1^\circ, \dots, 90^\circ$ and again $\varphi = 45^\circ$ 1.000 samples of artificial 3D-PLI signals were generated using the method described in sec. 4.2. Analyzing a larger number of samples was not possible due to a lack of computation time. All




FIGURE 5.2: Reconstruction accuracy of ROFL and PORISO algorithms evaluated on synthetic data. **Left:** ROFL algorithm. **Right:** PORISO algorithm. **Top:** orientation reconstruction error $\langle \gamma \rangle$ as a function of relative thickness t_{rel} and inclination angle α . **Bottom:** absolute reconstruction error of the relative thickness $\langle \sigma_{t_{rel}} \rangle$ as a function of relative thickness t_{rel} and inclination angle α . The red arrows point out discrepancies between the results

synthetic datasets were then analyzed with the ROFL and PORISO algorithms. The outcomes were again evaluated based on the mean orientation error and mean absolute error compared to the ground truth.

Results

The mean orientation error and the mean absolute error of t_{rel} achieved by PORISO and ROFL are depicted in Fig. 5.2. For the mean orientation error no significant difference is observed except for very flat fibers and zero relative thickness: in this case, the mean orientation error computed by PORISO amounts to 60° compared to 90° computed by ROFL. Stronger differences occur for the relative thickness. For very low values of t_{rel} , the PORISO algorithm results in a mean absolute error of almost zero compared to 0.1 as a result of the ROFL algorithm which is pointed out by the red arrows. For highly inclined fibers, the PORISO algorithm is slightly less accurate than ROFL: the maximal deviations are 0.65 for PORISO and 0.5 for ROFL.

Discussion

On synthetic data, PORISO achieves a comparable accuracy as the ROFL algorithm for the reconstructed orientations. The observed difference for very low relative thicknesses is insignificant as in practice an average deviation of 60° is as bad as an average deviation of 90° . For the reconstruction of very low relative thicknesses, PORISO clearly outperforms ROFL. For perpendicular fibers with respect to the sectioning plane, both approaches result in high absolute errors for the estimated value of t_{rel} . While the re-

sults of ROFL are slightly better, both algorithms are not capable of finding a plausible estimate.

The simulations revealed that the chosen inclination prior has a strong effect on the estimated relative thickness. This fact can only be explained by a correlation between inclination and t_{rel} . As both parameters together dictate the amplitude of the sine curve, such a correlation is an inherent part of the employed model. According to the retardation formula, higher t_{rel} values require the $\cos^2(\alpha)$ term to be smaller to achieve a similar retardation. As higher inclinations result in lower values of $\cos^2(\alpha)$ this can explain the correlation. In general, this consideration causes a (possibly nonlinear) positive correlation between t_{rel} and the absolute inclination angle $|\alpha|$. For very low signals only the $\cos \alpha$ prior can be optimized which results in an inclination angle close to 0. Due to the correlation, the t_{rel} estimate is then also close to 0. For a ground truth signal with a very low relative thickness, this result is much closer to the ground truth than the WLS estimate. For a ground truth signal of a steep fiber, the $\cos \alpha$ prior does not provide helpful information as the ground truth has the lowest prior probability. In consequence the estimated flat fibers and low relative section thicknesses are farther from the ground truth than the estimate without the prior information. This explains the higher mean error of t_{rel} of POriso compared to ROFL for perpendicular orientations. Regarding the reconstructed orientation, both POriso and ROFL result in an average angular deviation of 90° for perpendicular fiber orientations which means that both algorithms are unable to reconstruct orientations with $\alpha = 90^\circ$.

The major motivation behind the MAP estimator was to improve the reconstruction accuracy for very weak signals which are hard to interpret solely based on the likelihood. For the case of very low relative thickness, this goal was achieved for synthetic data. For fiber orientations perpendicular to the sectioning plane, the reconstruction accuracy remained the same for the orientation but slightly decreased for the relative thickness. Nevertheless, the superior reconstruction for low relative thicknesses is of far greater importance than the slightly inferior reconstruction for perpendicular orientations as even the marginally better ROFL result for perpendicular orientations does not provide plausible values for t_{rel} .

5.3 Evaluation on experimental data

Methods

To evaluate the performance of the POriso algorithm for experimental data, the same human brain datasets as for the ROFL algorithm were reinvestigated (cf. sec. 4.3). The vector fields in the stratum sagittale and at the boundary of white and grey matter were compared to the vector fields obtained from ROFL. Also, the distribution of the mean squared error was compared to ROFL and DFT for the same exemplary brain section already investigated in sec. 4.3.

Results

The vector fields in the stratum sagittale are shown in Fig. 5.3. Differences are observable for individual pixels pointed out by arrows. Overall, for pixels in which differences occur POriso estimates smaller inclinations than ROFL. For the first ROI, the only difference occurs for pixels with vanishing retardation signal. A similar behaviour can be observed in the second ROI. For the third ROI POriso derives almost no highly inclined

fiber orientations compared to ROFL. Instead, all fiber orientations run in the coronal sectioning plane. In the fourth ROI of very steep fiber orientations, no differences can be found.

The agreement of model prediction and measurement data was again evaluated based on the mean squared error metric of an exemplary brain section. The distributions of the MSE values achieved by DFT, POriso and ROFL are depicted in Fig. 5.4. The histograms reveal that POriso results in slightly smaller errors on average compared to ROFL. As this impression might be subjective, additionally the cumulative density functions (*CDF*) were plotted in Fig. E.1: for $4,000 < MSE < 12,000$ a significant difference between the CDFs can be observed. Here, the distribution obtained from POriso expresses a higher density of lower MSE values. However, the differences between the CDFs are small.

For synthetic data, differences in the reconstruction accuracy between POriso and ROFL were particularly found for low relative thicknesses. This was evaluated for synthetic data based on the cortical ROI shown in Fig. 5.5. In white matter both algorithms agree strongly but the interpretation differs strongly in grey matter. From the relative thickness maps, it can be seen that POriso estimates very low values of t_{rel} throughout the whole cortex in contrast to ROFL. Regarding the fiber orientation, POriso results in in-plane fiber orientations for the cortex while ROFL estimates strongly inclined orientations for a large number of pixels. The investigated ROI also contains a measurement artifact. It can be observed that the measurement artifact induces artifacts for the POriso algorithm which are however not as pronounced as for the DFT algorithm (cf. Fig. 4.11).

Discussion

The POriso algorithm was developed to improve the reconstruction accuracy for very low signals which arise for very steep fibers and unmyelinated tissue. For the case of very steep fibers, the reconstructed fiber orientations strongly agree with the results of the ROFL algorithm. For the ROI depicted in Fig. 5.3 B however, the vector field obtained from ROFL seems to be more homogeneous and appears to fit the overall orientation of the bundle better than the POriso vector field. On the other hand, for the vector field in Fig. 5.3 B which is even more inclined the results are barely distinguishable. In general, POriso results in less inclined fiber orientation estimates as a consequence of the prior. This becomes most apparent in the ROI consisting of fiber crossings depicted in Fig. 5.3 C.

As for the synthetic data, a strong effect of the inclination prior on the estimated relative thickness was observed. For the cortex, POriso results in a smooth map of vanishing values of t_{rel} for cortical areas. As argued in chap. 4, at LAP resolution cortical areas are expected to yield vanishing values for t_{rel} while the fiber orientation cannot reliably be estimated. For the relative thickness POriso enables to achieve this expectation utilizing the inclination prior. At the same time, POriso yields no longer random orientation estimates but very flat fibers with respect to the sectioning plane for the cortex. Anatomically, this is still as unlikely as the random orientations computed by ROFL.

The investigated ROI showed that the measurement artifact has a stronger effect on the parameter estimation for POriso than for ROFL. This can most likely be explained by the fact that ROFL utilizes the measured light intensities directly whereas POriso




FIGURE 5.3: Comparison of vectorfields derived from the ROFL and POriso algorithms. **A**: ROI consisting of one dominant fiber orientation approximately inclined by 50° with respect to the sectioning plane. **B**: ROI consisting of two apparently differently oriented fiber bundles. **C**: a ROI which mainly consists of nerve fiber crossings. **D**: ROI of highly inclined nerve fibers with respect to the sectioning plane. Arrows indicate differences between the results.




FIGURE 5.4: Histograms and kernel density estimates (full lines) of the mean squared errors between the measured light intensities and the light intensities predicted by the oblique 3D-PLI model for one exemplary brain section. 3D-PLI parameters were derived using the DFT (blue), ROFL (orange) and POriso (green) algorithms. The mean squared error values obtained from POriso are marginally lower than for ROFL.




FIGURE 5.5: Reassessment of tilted 3D-PLI measurements for the cortex. **Left:** ROFL results. **Right:** POriso results. **From top to bottom:** relative thickness maps, fiber orientation maps, relative thickness maps overlaid with the vector field of every third vector. The white arrows indicate a dust particle (cf. Fig. 4.11). The measurement artifact has a stronger effect on the POriso estimator than on the ROFL estimator. Whereas ROFL contains a significant amount of noise in the cortex, POriso yields smooth maps for both relative thickness and fiber orientation. In a direct comparison, POriso estimates predominantly in-plane oriented nerve fiber orientations in cortical areas.

relies on the analytically calculated Fourier coefficients. A dust particle corrupts one of the 18 measured light intensities per tilting direction. The calculation of the Fourier coefficients is severely affected by such a corrupted light intensities. For ROFL, the relative number of outliers is 5 out of 90 while for POriso all Fourier coefficients are also corrupted. In consequence, ROFL is more robust against such outliers.

One point which was not investigated as thoroughly as for the ROFL algorithm is the computation time. Using Python and Numba, the current implementation of POriso takes roughly 0.7 core hours for a section with 1 Mio. pixels compared to 0.3 core hours for ROFL. In principle though, POriso could also be implemented on a GPU which might reduce the computation time to comparable times as the GPU implementation of ROFL.

Finally, we can conclude that the prior assumption that all fiber orientations are equally likely biases the parameter estimation towards flat fibers with respect to the sectioning plane and low relative section thicknesses. For the relative thickness, this interpretation is more realistic for cortical areas while the fiber orientation is estimated with a similar accuracy as without prior information. From a purely statistical perspective, POriso yields marginally better results due to the lower mean squared errors. Still, as no ground truth measure for the accuracy of the obtained fiber orientations is available, at this point it is not possible to determine if POriso achieves a better reconstruction of the brain's nerve fiber orientations than ROFL.

6

Uncertainty estimation in 3D-PLI

“Nothing in the world causes so much misery as uncertainty.”

– Martin Luther

In chapters 4 and 5 different strategies to infer birefringence parameters from the polarimetric measurements were explored. The estimators were limited to best fit parameters so far. In this chapter measures of the uncertainty of the obtained parameters are examined. Also the correlations between the parameters which were already observed in the results of the POrISO algorithm are studied further. Following the basic principles of MCMC sampling described in chapter 3, this technique is applied to 3D-PLI.

In DMRI, which also aims at inferring local and global nerve fiber tract information, the importance of the propagation of uncertainty has been recognized and addressed in several studies, such as the bootstrapping approaches by [143, 144, 145]. Furthermore, the local fiber orientation uncertainty was propagated into probabilistic tractography frameworks [146]. Behrens et al. [22] were the first to perform full Bayesian inference on diffusion MRI data via MCMC. They exploited the probabilistic Bayesian framework for local uncertainties as well as tractography. This chapter is primarily inspired by their work. Since then, Bayesian statistics have been adopted not only in the DMRI community, but also in functional MRI¹ [148]. In recent years, MCMC methods were utilized to characterize microstructural models for DMRI [130, 129].

This chapter is structured as follows: after a review of the noise model for 3D-PLI, sampling algorithms for the planar and oblique 3D-PLI measurements are developed. Both are first evaluated based on synthetic data and then applied to selected experimental datasets. Parts of this chapter have been accepted for publication and will be published in [?].

Full noise model of the 3D-PLI measurement

An accurate noise model of the data acquisition process is essential for the derivation of the likelihood function. In chap. 4 it was shown that the recorded light intensities I_r follow a negative binomial distribution $I_r \sim \mathcal{NB}(\mu, g\mu)$ with expected value μ and the setup specific gain factor g . The gain factors g strongly depend on the employed camera. For the LAP experimentally a value of $g = 3$ was experimentally determined [46]. For

¹Functional MRI studies the activation of brain regions during specific tasks by measuring the blood flow in the brain [147].

the LMP noise measurements carried out in [149] resulted in a gain factor of $g = 0.125$ which agrees with the information provided by the manufacturer of the CCD sensor. To reduce the computational complexity, we will approximate the negative binomial distribution by a Normal distribution: $\mathcal{NB}(\mu, g\mu) \approx \mathcal{N}(\mu, g\mu)$. This approximation is justified as the observed light intensities are large ($I_r \gg 10$).

The first step of the 3D-PLI processing pipeline is an intensity based calibration as described in chap. 2 which multiplies every measured light intensity I_r by a calibration factor c . As the calibration factors are obtained by averaging over 100 images of the field of view without a probe, their uncertainty is negligible. The calibration process can then be treated as a multiplication of a normal distribution with a constant: $c \cdot \mathcal{N}(\mu, g\mu) = \mathcal{N}(c\mu, c^2g\mu)$. Typically, the uncalibrated raw data is archived and not available for later processing, therefore it is more convenient to use the calibrated light intensities $I = c \cdot I_r$ for computations and derivations. Thus, the calibrated light intensities follow the distribution $I \sim \mathcal{N}(c\mu, c^2g\mu)$.

6.1 Planar 3D-PLI measurement

6.1.1 Theory and implementation

In a 3D-PLI measurement without additional measurements from oblique views, we seek to determine the parameter set $\theta = (I_T, \varphi, r)$ with transmittance I_T , direction angle φ and retardation r , from a set of N light intensities I_i , where the index i denotes the rotation angle. The polarimetric model predicts the following light intensity M for a given polarization angle ρ :

$$M_i(I_T, \varphi, r) = I_T \cdot (1 + r \cdot \sin(2(\rho_i - \varphi))) . \quad (6.1)$$

Thus, the log-likelihood is given by inserting the calibrated model prediction $c_i\mu_i = c_iM_i(\theta)$ and the variance $\sigma_i^2 = gc_i^2M_i(\theta)$ into the log-likelihood of the normal distribution (see. eq. (3.8)) which yields

$$l = -\frac{1}{2} \sum_{i=0}^N \ln(gc_i^2M_i(\theta)) + \frac{(I_i - c_iM_i(\theta))^2}{gc_i^2M_i(\theta)} . \quad (6.2)$$

Non-informative priors were chosen except for the transmittance which has to be positive:

$$\begin{aligned} I_T &\sim \mathcal{U}(0, \infty) \\ \varphi &\sim \mathcal{U}\left[-\frac{\pi}{2}, \frac{3\pi}{2}\right) \\ r &\sim \mathcal{U}(0, 1) \end{aligned} \quad (6.3)$$

The direction angle prior was chosen to comply to the cyclic parameter space as the ensemble sampler is by default not capable of dealing with cyclic parameter spaces. After sampling the direction samples φ are then projected back into the 3D-PLI parameter space by $\varphi = \varphi \% \pi$.

The empirical probability density of the samples needs to be analyzed to obtain credible intervals. For the scalar parameters, HPD intervals, the smallest interval which contains 95% of all samples, were computed to serve as credible intervals. For the direction samples which we are primarily interested in, this approach would not respect the cyclical




FIGURE 6.1: Results of the BAMBI framework for the planar case on synthetic data. **Left:** empirical direction angle credible interval as a function of relative thickness and inclination. **Middle:** MCMC direction angle credible interval obtained from BAMBI as a function of relative thickness and inclination. **Right:** Difference between ground truth and MCMC estimate.

symmetry of orientations. In cyclical and more general spherical statistics, an analogon of the median of a scalar valued distribution is given by the principal axis [150]. For N orientation vectors $v_i = (x_i, y_i)^T$, the orientation tensor is given by $T = \sum_i^N v_i v_i^T$. The principal axis is then the Eigenvector corresponding to the largest Eigenvalue of the orientation tensor. After computing the principal axis, all direction samples were projected into the half circle centered around the direction angle corresponding to the principal axis. Now HPD intervals can be obtained as for the scalar parameters.

Executing the sampling and postprocessing for all pixels of one brain section was implemented in Python, utilizing the emcee package for ensemble sampling. The initial positions of the Markov chains for the ensemble sampler were set randomly close to the parameter vector obtained from the Fourier analysis. It was heuristically found that running 20 chains for 500 iterations yields satisfactory results which is confirmed by the simulation studies in the next section. To speed up the computations, the log-posterior calculation was compiled via numba. The computation time for one pixel then amounts to ca. 0.1 seconds. Due to the pixelwise parallelism of the computational problem, the computations are again easily parallelized using mpi4py. The sampling and postprocessing procedure will be referred to as *B*ayesian *M*ultivariate *B*irefringence *I*nfERENCE (*BAMBI*) framework from here on. Besides the implementation for processing of whole sections, an alternative application for BAMBI was implemented for interactive usage. This implementation aims at processing single pixels and generating useful plots to demonstrate the results.

6.1.2 Validation on synthetic data

Methods

Synthetic data can serve as a ground truth for the true posterior which is estimated by MCMC sampling. Informally speaking, MCMC sampling seeks to reconstruct the posterior from one measurement. On the other hand, synthetic data generates many possible measurements. Calculating the MAP for each of the synthetic measurements

then recovers the posterior of the model parameters. This ground truth posterior can then be compared to the posterior estimated from MCMC. While this procedure cannot prove that the obtained credible intervals for brain tissue represent the true uncertainty, it can show that the MCMC sampling yields correct results from a purely mathematical standpoint.

For the validation of the BAMBI framework the synthetic datasets which were used to compare ROFL and PORiSo in sec. 5.2 were utilized. The dataset consists of 1.000 synthetic signals for a vast variety of model parameters. In case of the planar 3D-PLI measurement, the credible interval of the direction angle is the most important parameter as it is further utilized for orientation analysis. For all artificial 3D-PLI signals the direction angle was estimated by the Fourier analysis. Next, from this distribution the 95% HPD interval was computed which serves as a ground truth credible interval. For a comparison with BAMBI, the sampling procedure was executed for one signal for each parameter configuration yielding a 95% HPD interval.

Results

The empirical ground truth credible interval and the credible interval estimated by BAMBI are plotted as a function of relative thickness and inclination in Fig. 6.1. For white matter structures of $t_{\text{rel}} > 0.1$ and not too highly inclined fibers with $\alpha < 80^\circ$, the empirical credible interval is smaller than 10° . For high relative thicknesses and in-plane fibers, the uncertainty drops to almost 1° . The credible interval computed by BAMBI strongly agrees with the ground truth for weak to strong signals: for $t_{\text{rel}} > 0.2$ and $\alpha < 70^\circ$ they differ by -0.3° on average. For very low signals, especially for very steep fibers, the MCMC estimate differs strongly from the ground truth (cf. Fig. 6.1 right).

Discussion

The simulations proved a strong agreement between the ground truth credible interval and the credible interval computed by BAMBI for the vast majority of parameter configurations. Only for model parameters that generate almost vanishing signals, the MCMC sampling is not able to reconstruct the correct posterior. Due to the nature of MCMC sampling, this finding is unavoidable. As all parameter sets are equally unlikely, it is hard to find good proposals for any MCMC algorithm.² In consequence, most proposals are rejected.

The in-plane orientation confidence decreases with increasing inclination and decreasing relative thickness. This behavior originates from the retardation signal: the lower the relative thickness and the steeper the nerve fiber, the weaker becomes the retardation signal. Naturally, a smaller amplitude of the sinusoidal signal complicates the estimation of the phase resulting in a lower confidence. Apart from very low signals though the in-plane orientation credible interval is small or even very small for strong signals. This means that theoretically the in-plane orientation estimation in 3D-PLI is quite robust in most white matter regions. Furthermore, for white matter these credible intervals can reliably be estimated by the BAMBI framework based on the simulation results.

²Experiments with higher number of samples did not show significantly better results. In a blog post, professor of statistics Prof. Gelman called this the *folk theorem of statistical computing*: *When you have computational problems, often there's a problem with your model* (https://statmodeling.stat.columbia.edu/2008/05/13/the_folk_theorem/, accessed: 23.12.2019).

6.1.3 Validation on experimental data

Despite a successful application on synthetic data, the BAMBI framework is designed for the analysis of experimental 3D-PLI datasets. Especially, the simulations found a distinct difference in performance of BAMBI between white matter and grey matter structures. This is now evaluated based on experimental data. Due to the excessive computation time, the analysis is limited to single brain sections for LAP datasets and regions of interest for LMP datasets.

Methods

Both LAP and LMP are able to carry out the planar 3D-PLI measurement. By analyzing the same tissue sample the effect of the resolution on the obtained direction angle confidence can be investigated. As tissue sample a coronal human brain section was chosen. At LAP resolution, the whole section was processed resulting in a direction angle credible interval map. The marginal distributions of the parameters were plotted for individual pixels to demonstrate the difference between white and grey matter. At LMP resolution, a ROI of 3×3 tiles at the boundary of white and grey matter was examined³. The resulting direction angle confidence map was correlated to the retardation map.

Results

The direction angle credible interval is depicted along side with retardation and direction angle in Fig. 6.2⁴. The majority of white matter pixels express a credible interval smaller than 5° . Especially in areas of very high retardation values such as the corpus callosum the credible interval drops to ca. 1° . Lower confidence occurs in regions of low retardation signals which are highlighted by arrows. At the boundary of white and grey matter, the credible interval continuously increases and even exceeds 32° in the cortex.

For the white matter pixel pointed out by the blue circle in Fig. 6.2 the full posterior is reported in Fig. 6.3. It can be seen that the posterior predictive distribution is able to accurately describe the measured light intensities. The marginal distributions follow a normal distribution. The obtained direction angle credible interval amounts to $\sigma_\varphi = 1.4^\circ$ for a retardation value of $\sin \delta = 0.75 \pm 0.015$. The posterior for the grey matter pixel pointed out by the red circle is given in Fig. 6.4. 100 samples of the posterior predictive distribution are plotted on top of the measurement data. As the predictions differ strongly in the phase of their sinusoidal curve, a large direction angle credible interval can be assumed which is confirmed by the marginal distribution of the direction angle. The credible interval amounts to $\sigma_\varphi = 100^\circ$. The histograms of retardation and direction angle also do not follow a normal distribution compared to the white matter pixel.

The retardation, direction angle and direction angle credible maps obtained from the microscopic measurement are depicted in Fig. 6.5. In the overview of the whole processed data it can be seen that for the white matter fibers in the upper part of the ROI which express a strong retardation signal the direction angle credible interval does not exceed 10° . The direction angle map also appears very homogeneous in this region. In regions of very low retardation values the confidence decreases strongly resulting in high credible intervals greater than 32° . In the cortex (cf. Fig. 6.5 middle), single nerve

³Computations were carried out using 50 compute nodes of the JURECA supercomputer [151].

⁴Instead of the typical 3D-PLI colormaps for the orientation, here the perceptually uniform cyclical colormap *phase* from the *cmocean* package [152] was chosen for the direction angle map.




FIGURE 6.2: Investigation of the in-plane orientation confidence for the planar 3D-PLI measurement based on LAP data. **Top:** retardation map. The red and blue points indicate pixels for which the results are presented in detail in Fig. 6.3 and Fig. 6.4. The rectangle indicates the ROI which was also analyzed with the LMP (cf. Fig. 6.5). **Middle:** direction angle map. **Bottom:** direction angle credible interval map. The credible interval increases with decreasing retardation. Arrows highlight regions of low retardation and high direction angle credible interval.




FIGURE 6.3: Results of the BAMBI framework for a white matter pixel. The position of the pixel is pointed out by the blue circle in Fig. 6.2. **Top:** measured light intensities and the 2σ confidence interval given by the 95% HPD interval of the posterior predictive distribution. **Bottom:** marginal distributions of transmittance, direction angle and retardation and normal distribution fit. Dashed lines indicate the boundaries of the 95% HPD interval.




FIGURE 6.4: Results of the BAMBI framework for a grey matter pixel. The position of the pixel is pointed out by the red circle in Fig. 6.2. **Top:** measured light intensities and 100 samples of the posterior predictive distribution. **Bottom:** marginal distributions of transmittance, direction angle and retardation. Dashed lines indicate the boundaries of the 95% HPD intervals.

fibers which can clearly be distinguished from each other in the retardation map still display direction angle credible intervals smaller than 5° .

Retardation and direction angle credible interval are correlated in Fig. 6.5 C. A clear trend is observable: the higher the retardation, the lower the credible interval. For very low retardation values the credible interval rises to almost 180° . In this regime, a retardation value can also correspond to a vast variety of possible credible intervals, meaning that the correlation is ambiguous.

Discussion

The in-plane orientation confidence was evaluated based on a coronal human brain section which was analyzed fully at a pixel size of $64 \times 64 \mu\text{m}$ with the LAP and for one ROI at the boundary of white and grey matter with the LMP at a pixel size of $1.3 \times 1.3 \mu\text{m}^2$. It could be shown that at both resolutions white matter fibers can be reconstructed with strong confidence. At microscopic resolution even single nerve fibers in the cortex can still be described with high robustness by 3D-PLI. As suggested by the simulation results, the direction confidence is strongly correlated to the retardation. For high retardation values, the credible interval even decreases to 1° . For unmyelinated regions in the cortex, the credible interval exceeds 32° which means that the estimated orientation is basically random and cannot be trusted. As these regions do not contain myelinated nerve fibers, the assumptions behind the physical model are not valid. This raises the question if the model should be applied to pixels which do not contain nerve fibers at all.

The planar 3D-PLI measurement is limited to the reconstruction of the in-plane fiber orientation. In the next section, the confidence in the reconstruction of the three-dimensional fiber orientation from oblique 3D-PLI measurements is investigated.

6.2 Oblique 3D-PLI measurement

6.2.1 Theory and implementation

As shown in chap. 4 and 5, additional measurements from oblique views add crucial information for the determination of inclination angle and birefringence. In general, there are $N_T + 4$ parameters in the model for N_T measured tilting positions: N_T average transmittances from the tilted measurements, one transmittance of the planar measurement, the birefringence strength t_{rel} and the fiber orientation given by direction angle φ and inclination angle α . The parameter vector θ can then be written as $\theta = (I_{0T}, \dots, I_{N_T}, \varphi, \alpha, t_{\text{rel}})$. Denoting the tilting position with index j , the rotated fiber parameters $(\varphi_j, \alpha_j, t_j)$ are given by applying the appropriate rotation R_j to the modelled fiber orientation. For the light intensity curve this yields

$$M_{ji}(I_{jT}, \varphi, t, \alpha) = I_{jT} \cdot (1 + \sin(2(\rho_i - \varphi_j)) \cdot \frac{\pi}{2} t_j \cos(\alpha_j)^2) . \quad (6.4)$$

As for the planar measurement, the log-likelihood is derived by inserting the calibrated model prediction $c_i \mu_{ji} = c_i M_{ji}(\theta)$ and the variance $\sigma_{ji}^2 = g c_i^2 M_{ji}(\theta)$ into the log-likelihood for a normal distribution which results in

$$l = -\frac{1}{2} \sum_j \sum_i \ln(g c_i^2 M_{ji}(\theta)) + \frac{(I_{ji} - c_i M_{ji}(\theta))^2}{g c_i^2 M_{ji}(\theta)} . \quad (6.5)$$




FIGURE 6.5: Results of the BAMBI framework for microscopic 3D-PLI data. Processed data: 3×3 tiles of a coronal human brain section at the boundary of white and grey matter. **Top:** overview of the complete analyzed ROI. Left: retardation, right: direction angle credible interval clipped at 30° . **Middle:** ROI in the cortex where individual nerve fibers are visible. **Bottom:** 2D histogram of retardation values and direction confidence of the full ROI.

The priors were chosen as

$$\begin{aligned}
I_j &\sim \mathcal{U}(0, \infty) \\
t &\sim \mathcal{U}(0, 1] \\
\varphi &\sim \mathcal{U}\left[-\frac{\pi}{2}, \frac{3\pi}{2}\right) \\
\alpha &\sim \begin{cases} \cos(\alpha) & \text{for } \alpha \in [-\pi, \pi) \\ 0 & \text{else .} \end{cases}
\end{aligned} \tag{6.6}$$

The inclination prior again originates from the uniform orientation prior as discussed in chap. 5. As in the planar case, the boundaries of the parameter spaces for the angles were expanded compared to the parameter space of 3D-PLI to enable efficient sampling.

After sampling, the samples are analyzed as in the planar case: for the scalar parameters, HPD intervals containing 95% of the samples are computed. From the direction and inclination samples, orientation vectors are calculated. The principal axis of the orientation vector samples is obtained again as the principal Eigenvector of the orientation tensor. Then all orientation samples are projected into the half sphere centered around the principal orientation. From these orientation samples, direction and inclination angles are then recomputed and HPD intervals derived.

In analogy to the planar case, the sampling and HPD interval computation was implemented in Python based on the emcee's package for ensemble sampling, numpy and numba for the log-posterior calculation and mpi4py for pixelwise parallelization. The chains are initialized close to the MAP estimate given by the result of the Fourier analysis for the transmittances and the POriso algorithm for t_{rel} and orientation. Due to the higher dimensionality and the correlation between inclination and t_{rel} more parallel chains are required for convergence than for the planar measurement. Heuristically, 100 chains running for 500 iterations were found to yield stable results (more samples did not change the results significantly). Processing a single pixel takes 0.5 seconds.

6.2.2 Validation on synthetic data

Like in the planar case, synthetic data can prove that the MCMC sampling is able to recover the true posterior for the oblique 3D-PLI measurement.

Methods

The same dataset as for the planar case was re-investigated. The results of the POriso algorithm which were already obtained in sec. 5.2 were utilized to generate a ground truth distribution of the posterior. This ground truth distribution was then compared to the MCMC analysis of one of the synthetic datasets by means of the 95% HPD interval for each parameter set.

Results

The results of the simulation are depicted in Fig. 6.6. In the top row the results for the direction angle are given. They are very similar to the result for the planar case with the difference that the obtained credible intervals are slightly smaller than in the planar case both for ground truth and MCMC result. The lowest credible interval of 0.75° is achieved for in-plane fibers and the highest simulated t_{rel} value. Empirical and MCMC

generated HPD intervals show a good agreement except for very steeply inclined fibers and very low t_{rel} values.

With respect to the inclination angle confidence, a subtly different behavior can be observed. The lowest ground truth credible interval of $\sigma_\alpha = 5.8^\circ$ is obtained not for $\alpha = 0^\circ$ but for $\alpha = 65^\circ$. In general, the credible intervals are higher than for the direction angle. For very low relative thicknesses and very high inclinations, the credible interval again exceeds 64° . The credible interval calculated by BAMBI agrees with the ground truth for moderate to high t_{rel} values and inclinations $\alpha < 80^\circ$: the difference amounts to $0.5 \pm 1.1^\circ$. Also, the credible interval map obtained from MCMC is not as smooth as for the direction angle.

The relative thickness expresses a different characteristic regarding its reconstruction confidence. The credible interval depends mostly on the inclination: for in-plane fiber orientations the ground truth credible interval amounts to $\sigma_{t_{\text{rel}}} \approx 0.02$ for $t_{\text{rel}} \in [0.1, 0.7]$. For very low values of t_{rel} , the credible interval increases to $\sigma_{t_{\text{rel}}} \approx 0.1$. With increasing inclination, the credible interval increases exponentially and exceeds 0.5 for inclinations of $\alpha > 80^\circ$. The credible intervals obtained from MCMC sampling agree with the ground truth for most of the parameter space. For $\alpha < 80^\circ$ and $t_{\text{rel}} > 0.2$ the average difference amounts to $\langle \sigma_{t_{\text{rel}}} \rangle = 0 \pm 0.02$. Greater differences are again observed for very high inclinations and very low relative thicknesses. In these cases BAMBI estimates very small credible intervals compared to the ground truth.

Discussion

The simulations carried out here were the first ones to assess the uncertainty of the obtained inclination and direction angles for 3D-PLI. One result is that there is inherently a lower confidence in the obtained inclination angle than in the direction angle. Since the direction angle can directly be measured from one polarimetric measurement and the inclination requires information from only slightly differing additional measurements, this result is not surprising. Intuitively it is also clear that it from an optical measurement of the transmitted light of a thin tissue sample it is easier to derive in-plane than out-of-plane information. While the oblique measurements add complementary information, 3D-PLI does not capture three-dimensional information as DMRI does where gradient fields can be applied in arbitrary orientations.

Regarding the maximally possible confidence, it was observed that the highest confidence for the in-plane orientation occurs for the highest simulated t_{rel} value and in-plane fibers. This parameter combination yields the highest possible retardation value. As in the planar case, the confidence strongly depends on the retardation value. The obtained credible intervals are lower than in the planar case as more measurement information is available and the employed model is more complex.

The minimal credible interval for the inclination angle was observed for the highest t_{rel} value and $\alpha = 65^\circ$. As the inclination is derived from the differences between the oblique measurements, its confidence does not solely depend on the raw signal strength but also on the signal characteristic. One hint could be the retardation gradient (cf. Fig. 4.4) which becomes maximal for inclination angles between 50° and 60° . This implicates that for not too small values of t_{rel} highly inclined fibers can be determined with the same or even higher confidence as in-plane fibers.

Besides the orientation confidence, the confidence of the t_{rel} parameter which is propor-




FIGURE 6.6: Simulation results of the BAMBI framework for the oblique 3D-PLI measurement. **Top row:** examination of the direction angle uncertainty. **Middle row:** examination of the inclination angle uncertainty. **Bottom row:** examination of the relative thickness uncertainty. **Left column:** empirical ground truth HPD intervals. **Middle column:** MCMC HPD intervals computed by BAMBI. **Right column:** Difference between empirical ground truth and MCMC result. All results are plotted as functions of relative thickness t_{rel} and inclination angle α .

tional to the birefringence was investigated. It was found that the maximally possible accuracy of its estimation strongly depends on the inclination: while for in-plane fibers t_{rel} can be reconstructed with a negligible credible interval, for highly inclined fibers the credible interval is greater than 0.5 meaning that the computed relative thickness cannot be trusted. This implies that for steep nerve fiber bundles with respect to the sectioning plane the relative thickness map should be interpreted very carefully.

The major objective of the simulation study was to validate the sampling procedure. For weak to strong signals, the MCMC result could be validated based on the empirical ground truth. For perpendicular fibers and very low relative thicknesses, greater differences were observed as in the planar case. In particular the credible interval for the relative thickness is strongly underestimated by BAMBI. These discrepancies can again be explained by the inherent difficulty to fit a sinusoidal model to a dataset with an almost vanishing amplitude.

The agreement between MCMC result and ground truth is slightly worse than for the planar case. One reason might be the generation of the ground truth distribution. Not all parameters are simultaneously estimated while BAMBI samples all parameter simultaneously. The ground truth parameters were obtained in the following way: first, the transmittance, direction and retardation values were estimated using the Fourier analysis. Then, direction, inclination and t_{rel} were computed by the optimization procedure of the POrIso algorithm. A simultaneous optimization of all parameters including the transmittances and not only direction, inclination and t_{rel} would lead to marginally different results. The thorough analysis carried out here is likely sensitive to these subtle differences. An estimation procedure which simultaneously optimizes all eight parameters would suffer from the big disadvantage of excessive computation time. Also, it might yield slightly different transmittance estimates which are however not needed as the already available transmittance estimates provide enough information.

6.2.3 Validation on experimental data

Methods

For the planar 3D-PLI measurement, a full coronal section measured with the LAP was analyzed with the BAMBI framework. As this particular section does not contain a large fiber bundle oriented out of the sectioning plane, here instead one section of the brain already analyzed in chap. 4 and 5 was chosen. The new analysis provides three new maps to 3D-PLI: credible intervals of direction angle, inclination angle and relative thickness. These maps are reported alongside the fiber orientation and t_{rel} map for the whole section as well as in a larger view of a region of interest of the stratum sagittale. For individual pixels scatterplot matrices are given to showcase the full posteriors and the correlations between the parameters.

As for the planar case, the direction credible interval map is correlated with the retardation map. Additionally, inclination credible interval and relative thickness are correlated.

Results

The maps obtained for the whole processed section are depicted in Fig. 6.7. In white matter, direction angle credible intervals lower than 1° can be observed. Overall, the direction confidence map is very similar to the planar case with the difference that the




FIGURE 6.7: Results of the BAMBI framework for a tilted measurement. **Top from left to right:** fiber orientation map, direction angle credible interval map, inclination angle credible interval map. Note that the angular credible interval maps use different color bars. The rectangle indicates the region of interested shown in Fig. 6.8. White and blue circles indicate the positions of pixels whose posteriors are provided in Fig. 6.9 and Fig. 6.10. Blue arrows indicate regions at the boundary of white and grey matter with moderate direction angle credible interval and high inclination angle credible interval. **Bottom from left to right:** retardation map, relative thickness map, relative thickness credible interval map. The red arrow points out a region of very low relative thickness credible interval and in-plane oriented nerve fibers.

credible intervals are smaller. The inclination credible intervals are greater than the direction credible intervals for the whole brain section. The lowest observed credible intervals amount to app. 8° for the inclination angle. Lower values can be attributed to crossing regions which can be seen from the region of interest in the stratum sagittale in Fig. 6.8 where crossing are indicated by white arrows. For the cortex and regions of vanishing retardation signal, again the orientation confidence measures decrease strongly. Still, the direction credible interval stays lower than 20° in some regions of the boundary between white and grey matter while the inclination credible interval exceeds 32° (cf. blue arrows in Fig. 6.7).

The map of the relative thickness credible interval shows a more complex pattern: the credible interval ranges between $\sigma_{t_{rel}} = 0.01$ and $\sigma_{t_{rel}} > 0.32$ in white matter. The region of very low credible interval pointed out by the red arrow corresponds to a region of very flat fibers with respect to the sectioning plane. The indicated gyrus was used to compare the inclination histograms of in-plane fibers obtained from the different algorithms in chapters 4 and 5 (cf. Fig. 4.9). Especially for high t_{rel} estimates the credible interval increases strongly. For the cortex, the t_{rel} credible interval map displays a great variety of possible values.

In Fig. 6.8 the same maps are shown again for a region of interest in the stratum sagittale. It can be seen that for the nerve fibers inclined between 65° and 75° in the stratum sagittale (values were extracted from the corresponding inclination map) both direction and inclination angle credible interval do not exceed 10° . In fact, the inclination confidence in this region is one of the highest of the whole section (cf. Fig. 6.7). While the inclination confidence is high, the relative thickness confidence deteriorates to $\sigma_{t_{rel}} > 0.2$ in many pixels of the stratum sagittale.

Regions of fiber crossings were identified by their vanishing retardation signal and a diverging fiber orientation map (cf. white arrows in Fig. 6.8). In these regions, even the direction angle credible interval exceeds 32° . On the other hand, the inclination angle credible intervals are afflicted with a less homogeneous appearance. For the crossing pointed out by the right arrow, the right side of the crossing expresses a very low credible interval of $\sigma_\alpha < 10^\circ$ while after an abrupt change in the middle the left side displays a credible interval of $\sigma_\alpha > 64^\circ$. The inverse behavior can be observed for the relative thickness: here the left side of the crossing displays a much lower credible interval than the right side. This discontinuity is actually also observable for the best fit parameters: on the left side of the crossing the estimated t_{rel} values are almost zero and the fiber orientation very flat with respect to the sectioning plane. On the right side of the crossing the t_{rel} estimates reach up to almost 1 while the estimated fiber orientation is perpendicular to the sectioning plane.

The posterior of a pixel located in the stratum sagittale is depicted in Fig. 6.9 and serves as an example for a white matter pixel⁵. The estimated fiber orientation is highly inclined with respect to the sectioning plane. A clear nonlinear correlation between t_{rel} and inclination can be observed: higher t_{rel} values correlate to a steeper inclination. The HPD intervals amount to 11° for the direction angle, 10° for the inclination angle and 0.43 for t_{rel} which conforms to the general observations about the stratum sagittale (cf. Fig. 6.8). In the projections of the angular samples onto the sphere, the orientation samples yield an ellipsoidal distribution (cf. . 6.9 right).

An example for a grey matter pixel is shown in Fig. 6.10. The orientation confidence

⁵Scatterplot matrices were created with the corner package [153].




FIGURE 6.8: Results of the BAMBI framework for a tilted measurement zoomed into the stratum sagittale. **Top from left to right:** fiber orientation map, direction angle credible interval map, inclination angle credible interval map. **Bottom from left to right:** retardation map, relative thickness map, relative thickness credible interval map. StS: stratum sagittale. White arrows point out fiber crossings. In the stratum sagittale the inclination confidence is high while the t_{rel} confidence is low. In crossing regions direction, inclination and t_{rel} confidence deteriorate strongly.




FIGURE 6.9: Results of the BAMBI framework for an oblique 3D-PLI measurement of a white matter pixel. **Left:** scatterplot matrix. **Right:** orientation histogram. The position of the pixel is indicated by a white circle in Fig. 6.7. Full lines: best fit parameters estimated by the POrIso algorithm. Dashed lines: boundaries of the 95% HPD intervals. Fit results: $\varphi = 120.5^{+5.5}_{-5.5}$, $\alpha = -77^{+6}_{-4}$, $t_{\text{rel}} = 0.54^{+0.26}_{-0.17}$. The marginal distributions of inclination and relative thickness are not symmetric and nonlinearly correlated.

is basically zero as almost the entire hemisphere has a non negligible probability due to a very high inclination HPD interval. Even the direction credible interval amounts to almost 40° .

To evaluate the distributions of the angular confidences, the cumulative density functions of all white matter pixels is depicted in Fig. 6.11. It can be observed that more than 90 % of white matter express a direction angle credible interval of $\sigma_\varphi < 10^\circ$. For the inclination angle only a negligible amount of pixels expresses such a small credible interval. Instead, for almost 80 % of all white matter pixels the inclination credible interval lies between 10° and 25° . The cumulative density for the whole brain section is shown in Fig. E.3: now app. 55 % of all pixels express a direction confidence smaller than 10° and in 40 % of all pixels the inclination confidence is smaller than 25° .

For the planar measurement a clear correlation between retardation and direction credible interval was observed. The same behavior can be observed for the tilted measurement in the two-dimensional histogram of the obtained maps given in Fig. E.2. While for very low retardation values of $\sin \delta < 0.1$ almost arbitrary credible intervals occur, a clear negative correlation exists for stronger signals: the credible interval decreases with increasing retardation. The correlation of relative thickness and inclination credible interval is depicted in Fig. 6.12. Here, the inclination credible interval decreases with increasing t_{rel} . For $t_{\text{rel}} < 0.2$ the vast majority of inclination credible intervals is greater than 20° . For very low relative thickness values, the resulting inclination credible interval appear basically arbitrary (cf. arrow in Fig. 6.12) as was observed for very low retardation values for the planar case.




FIGURE 6.10: Results of the BAMBI framework for an oblique 3D-PLI measurement of a grey matter pixel. **Left:** scatterplot matrix. **Right:** orientation histogram. Full lines: best fit parameters estimated by the POriso algorithm. The position of the pixel is indicated by a black circle in Fig. 6.7. Dashed lines: boundaries of the 95% HPD intervals. The derived orientation is basically random.

Discussion

The derivation of the three-dimensional nerve fiber orientation from 3D-PLI data relies on the small differences between the light intensity curves recorded from different oblique views during rotation of the filters. The effect of measurement noise on the uncertainty of the estimated fiber orientation was quantified using the developed Bayesian framework based on human brain data. For white matter regions, a strong confidence in the derived in-plane orientation and a strong to moderate confidence in the estimated out-of-plane orientation was found. At the same time, the uncertainty of the microstructural parameter t_{rel} differs from a very low uncertainty for regions with dominant in-plane fiber orientation and a very high uncertainty for regions with dominant out-of-plane fiber orientation. Inclination and t_{rel} are strongly non linearly correlated as already observed in chap. 5. Here, for the first time the nonlinear correlation was explicitly computed as the full posterior became available. For a white matter pixel, it was found that high inclinations correlate to a higher t_{rel} value. As discussed in chap. 5, this can be explained by the relationship of both parameters to the retardation. From the indicatrix model of nerve fibers, it is clear that inclination and birefringence strength will always be correlated solely based on the polarimetric measurements.




FIGURE 6.11: Cumulative densities of angular credible intervals of white matter pixels. Blue: direction angle. Yellow: inclination angle.




FIGURE 6.12: Correlation of relative thickness t_{rel} and inclination angle credible interval σ_α . The plot displays a two-dimensional histogram of the relative thickness and inclination angle credible interval maps shown in Fig. 6.7. Besides an ambiguity for $t_{\text{rel}} \approx 0$ (cf. black arrow) a clear negative correlation can be observed.

Measurements from Oblique Views in Microscopic 3D-PLI

“We see past time in a telescope and present time in a microscope. Hence the apparent enormities of the present.”

– Victor Hugo

This chapter investigates the application of the developed algorithms at microscopic resolution. In [46] a polarization microscope capable of performing measurements with oblique illumination, the LMP3D, was developed (cf. 2.4). Oblique illumination achieves a similar same effect as tilting the sample in the LAP. Due to the limited depth of field simply tilting the sample is not possible in a microscopic setting. In a proof of principle experiment, it was shown that the same analysis as for the LAP can be applied for microscopic data. Still, the evaluation was limited to a qualitative assessment of the fiber orientation map at high resolution and a comparison to LAP data after downsampling.

As the new microscope is only built on a breadboard system so far, several properties such as noise characterization and possible filter misalignment have to be determined experimentally before the developed algorithms can be applied. This is described in sec. 7.1. Using this information, the new methods are used in sec. 7.2 to process the same datasets already investigated in [46]. While a phantom for 3D-PLI has not been developed yet, a volumetric measurement of the same brain section at high resolution which contrasts nerve fibers from the surrounding tissue can serve as the best available ground truth validation for 3D-PLI. An imaging technique which fulfills these requirements is two-Photon microscopy [154] (TPFM). In sec. 7.3, the results of the developed analysis tools for oblique 3D-PLI measurements are validated based on the ground truth provided by TPFM. All LMP3D measurements shown in this chapter were performed by Philipp Schlömer.

7.1 Experimental characterization of the LMP3D

Quantification of camera noise Within the BAMBI framework camera noise essentially determines the uncertainty of the nerve fiber orientation derived from 3D-PLI. The photon detection noise of the camera employed in the LMP3D has therefore to be determined.




FIGURE 7.1: Noise measurement of the LMP3D. The plot depicts a two-dimensional histogram of the pixelwise mean light intensities $\langle I \rangle$ and light intensity variances σ_I^2 . The fit line shows that the variance can be approximated by $\sigma_I^2 = 1.6 \cdot \langle I \rangle$.

In [46], the noise measurement of the LAP was carried out by measuring the same brain section 100 times to sample the recorded light intensity distribution pixelwise. In a microscopic environment, recording the same tissue sample repeatedly might lead to very different results due to edge effects. Another point which has to be considered for a noise measurement is that a broad light intensity distribution over the field of view has to be captured. Recording the out-of-focus image of an edge enables to avoid the edge effects and results in the desired broad intensity range.

For the noise measurement of the LMP3D, a small part of the USAF chart which includes an edge was imaged far out of focus 100 times and analyzed pixelwise for its mean value $\langle I \rangle$ and samples variance σ_I^2 as estimators of the light intensity expectation value and variance. The resulting noise relationship between variance σ_I^2 and light intensity I is given by $\sigma_I^2 = g \cdot \langle I \rangle = 1.6 \cdot \langle I \rangle$ (cf. Fig. 7.1). Hence, for the PORISO algorithm and the BAMBI framework a gain factor of $g = 1.6$ is used¹.

Determination of the polarization filter offset The Jones calculus as introduced in chap. 2 presumes that the axis of the rotating polarization filter is aligned with the camera axis. Due to the prototypical nature of the setup of the LMP3D at the optical bench this assumption has to be verified experimentally.

Using a reference polarization filter, a possible offset of the polarizer can easily be measured. The reference filter (*Edmund Optics polarization tester 37-699*) was inserted into the sample holder aligned with the camera axis. Rotation of the polarizer now yields a sinusoidal light intensity profile. The rotation angle of minimal transmittance is then the offset between camera axis and polarizer axis. The polarizer rotations were sampled in steps of 1° . All images were averaged over the full field of view and analyzed according to the Fourier analysis from sec. 2.6. The resulting offset ρ_o is $\rho_o = -67^\circ$.

¹According to the manufacturer of the camera sensor, the gain factor should be $g = 0.125$. For the LMP which employs the same sensor this value was validated by noise measurements. For the LMP3D, the discrepancy likely originates from the software used to operate the camera. While the sensor has a depth of 12 bit, the images which are written to disk have a dynamic range of 16 bit which is not the case for the PM. This additional conversion is a potential source of error for the additional noise.

Estimator	Computation time
ROFL	≈ 1.9 Core hours
GOROFL	30 seconds
POriSo	≈ 3 Core hours

TABLE 7.1: Computation times of the developed estimators for one field of view of the LMP3D

While this offset can simply be added to the direction angle for a planar 3D-PLI measurement, it complicates the analysis of measurements from oblique views as it relies on the differences between the modelled three-dimensional optical axis. The offset can be incorporated into the model for the light intensity curve of an oblique measurement by replacing the polarization angle ρ by $\rho + \rho_o$ (see eq. (4.3)) which yields ²

$$I_{ji}(\rho_i, \rho_o, \varphi, \alpha, t_{\text{rel}}) = \frac{I_{j,T}}{2} \left(1 + \sin(2(\rho_i + \rho_o - \varphi_j)) \sin\left(\frac{\pi}{2} t_{\text{rel},j} \cos^2(\alpha_j)\right) \right) . \quad (7.1)$$

This way, the offset can also be incorporated into the ROFL and BAMBI algorithms by the simple replacement $\rho = \rho + \rho_o$. As the POriSo algorithm relies on the phases and amplitudes, it is sufficient to correct the direction angles φ as $\varphi = \varphi + \rho_o$ before.

7.2 Application of the Bayesian framework at the microscale

Methods

In [46], one field of view of a coronal rat brain section which contains in-plane oriented nerve fibers as well as out-of-plane fiber bundles was measured to analyze the oblique measurements. The same sample was reassessed with the algorithms developed in this thesis. From the offset of the aperture diaphragm the oblique illumination direction can be calculated geometrically as shown in [46]. The internal oblique illumination angle was assumed as 3.9° following the arguments in [46]. Using the oblique illumination directions and the oblique illumination angle the same coordinate system as for the tilted measurements can be utilized which means that the developed algorithms can directly be applied. To investigate differences between the parameter estimation algorithms, the section was processed with the ROFL, GOROFL and POriSo algorithms resulting in maps of fiber orientation and t_{rel} . Additionally, the computation times were measured for an extrapolation of the computation times to fully measured brain sections. Finally, the BAMBI framework was utilized to investigate the uncertainty of fiber orientation and t_{rel} .

Results

The parameter maps obtained from ROFL of the full field of view and one ROI of highly inclined nerve fiber bundles are shown in Fig. 7.2. In the fiber orientation map, in-plane structures such as the corpus callosum display a strong in-plane fiber orientation while individual fiber bundles in the caudate putamen display strong out-of-sectioning plane orientations. Between the fiber bundles the FOM expresses strongly varying fiber orientations. Compared to LAP measurements, the relative thickness map has a high ratio of "extreme" relative thickness values. For out-of-plane fiber bundles almost all pixels display values of $t_{\text{rel}} > 1$. For regions of in-plane nerve fibers this behaviour

²This simple yet crucial adjustment was found by Felix Matuschke.

is not observed. The inclination credible interval does not exceed 16° for most pixels containing nerve fibers: even for highly inclined fiber bundles, the inclination confidence is high. Here, the relative thickness credible interval rises strongly which can be seen in the zoomed in region of the caudate putamen. For the corpus callosum, the obtained relative thickness credible interval almost vanishes.

Discussion

Additionally to the analysis of the same dataset carried out in [46] here the relative thickness measure and the confidence of the parameters was investigated. The local fiber orientations seem not unlikely as for example the corpus callosum runs mostly in the sectioning plane while individual fiber bundles in the caudate putamen express strong out-of-plane orientations. The level of detail with which individual bundles can be separated from each other represents a distinct advantage over mesoscopic measurements at LAP resolution.

In [46] a good agreement between the fiber orientations obtained from the LAP measurement and the LMP3D results after downscaling was found. The same does not apply for the relative thickness: in the investigated ROI, a large number of individual steep fiber bundles result in very high values of $t_{\text{rel}} > 1$. As discussed in chap. 2, in this regime the model is not unambiguous anymore as two different inclinations can result in the same retardation. For the LAP, only very few pixels show such "extreme" values of t_{rel} .

For the individual fiber bundles, the model is very confident about the found orientation as the inclination credible interval rarely exceeds 16° . At the same time, the confidence in the obtained relative thickness is basically zero. The small confidence in the estimated t_{rel} values for highly inclined fibers was already observed on synthetic data (cf. Fig. 6.6) and stems from the model itself. One possible conclusion which can be drawn from the confidence metrics is that researchers can be relatively confident about the obtained fiber orientation from a statistical viewpoint but should avoid further quantitative analysis of the t_{rel} values at microscopic resolution.

The presented analysis was only qualitative in its nature due to the lack of a known ground truth. Such an investigation is pursued in the next section by validating the obtained fiber orientations with help of a ground truth from TPFM measurements.

7.3 Validation of 3D-PLI fiber orientations based on TPFM measurements

In [155], the first attempt to validate fiber orientations derived from 3D-PLI based on TPFM was conducted. For the direction angle, a strong agreement was found. Regarding the inclination angle, the 3D-PLI data was afflicted with the ambiguities which arise for the planar measurement without the additional information from oblique views (cf. chap. 2). Compared to the TPFM measurements, the resulting inclination angles were overestimated for almost all investigated fiber bundles.

Here, the TPFM measurements are compared to microscopic 3D-PLI measurements from oblique views which allow a bias-free reconstruction of the three-dimensional fiber orientation. While theoretically a comparison between TPFM and LAP measurements would be possible, the resolution gap is too big to carry out a meaningful validation as one LAP pixel contains several nerve fiber bundles.




FIGURE 7.2: Results from measurements from oblique views at the microscale. Left: full field of view. Right: ROI consisting of steep fiber bundles. From top to bottom: fiber orientation map, relative thickness map, inclination credible interval map, relative thickness credible interval map. cc: corpus callosum. CPu: caudate putamen.

Methods

Two-Photon Fluorescence Microscopy

Fluorescence microscopy is a widely used microscopic technique for the investigation of biological tissue. Its contrast principle is based on fluorescence: first, a fluorescent molecule in its ground state absorbs an incoming photon which excites the molecule into a state of higher energy. The absorbed energy is then partially emitted via a photon of lower energy than the incoming photon. As specific molecules are sensitive to specific wavelengths, imaging a sample with several incoming wavelengths can contrast different structures. In Two-Photon Fluorescence Microscopy, the molecule is excited by two photons of lower energy than in typical fluorescence microscopy [154]. This has the advantage of lower photo damage of the imaged tissue which enables deeper penetration of the sample. The high penetration depth makes volumetric imaging of thin tissue samples possible.

For the TPFM measurement, the cover slip of the brain section has to be removed. Experimentally, it was found myelinated axons show strong autofluorescence [61]. While an explanation for this behaviour has not been found so far, the strong autofluorescence enables a simple identification of fiber bundles by eye.

Measurements

As tissue sample a coronal rat section was chosen. The brain was treated according to the standard 3D-PLI preparation (cf. chap. 2) and measured with the LMP3D in the region of the caudate putamen as this particular region contains a large number of individual fiber bundles. The LMP3D measurement used four oblique views like a standard LAP measurement. Subsequently, the brain section was delivered to Florence where it was measured ten days after the 3D-PLI measurement with TPFM at a voxel size of $0.439 \times 0.439 \times 1 \mu\text{m}$ by Irene Constantini using the setup described in [39].

Comparison of 3D-PLI and TPFM

A comparison of the fiber orientations requires the identification of the same fiber bundles in the TPFM volume and the 3D-PLI dataset. Therefore, David Gräkel performed an affine registration of the 3D-PLI data to the TPFM volume using hand written Matlab software [135]. The coregistration now enables an easy identification of the same fiber bundles in TPFM and 3D-PLI.

The orientation of individual fiber bundles was extracted geometrically from the TPFM volume (cf. Fig. 7.3): the upper and lower boundaries in the bundles were identified in the image stack by eye. These boundaries of the fiber bundle were then delineated manually on both ends. The vector between the center of of the delineations then serves as the ground truth fiber bundle orientation provided by TPFM $\mathbf{r}_{TPFM}$. The same nerve fiber bundle was consequently identified in the coregistered 3D-PLI image and also delineated manually. The delineation was projected onto the original 3D-PLI data by applying the inverse affine transformation. In the original data set, a fiber bundle orientation was then computed as the principal orientation of all orientations of the bundle $\langle \mathbf{r}_{PLI} \rangle$. In total a number of $N = 213$ fiber bundles were evaluated.

3D-PLI and TPFM fiber bundle orientations can now be compared based on their angular deviation ε . To evaluate the different fiber orientation estimation algorithms for




FIGURE 7.3: Derivation of fiber bundle orientations for a comparison of 3D-PLI and TPFM. In the TPFM volume (left), the top and bottom boundaries of the fiber bundle are manually delineated. The vector between the centers of both boundaries yields the ground truth nerve fiber orientation $\mathbf{r}_{TPFM}$. The same fiber bundle is identified in the 3D-PLI image (right). The principal orientation of all orientations in the bundle $\langle \mathbf{r}_{PLI} \rangle$ serves as the 3D-PLI orientation for the comparison. Modified from [155].

3D-PLI quantitatively, a suitable measure is given by the mean angular deviation between 3D-PLI and TPFM orientation of all fiber bundles $\langle \varepsilon \rangle$. Besides the mean deviation the distribution of deviations are investigated to validate the consistency of the different 3D-PLI estimators. To investigate the agreement of 3D-PLI and TPFM direction and inclination angles, they were directly correlated with each other. Besides the principal inclination and direction, the shortest interval containing 67 % of the angular samples were computed to indicate the amount of orientation variance in one fiber bundle in 3D-PLI. In contrast to the analysis of the pixelwise uncertainty (cf. chap. 6) a 67 % interval was chosen here as the manual delineations of the fiber bundles are afflicted with errors which increase the overall variance.

Results

The mean angular deviations between 3D-PLI orientation and TPFM ground truth are given in Tab. 7.2. All three different algorithms yield very similar results of $\langle \varepsilon \rangle \approx 10^\circ$: the CPU and GPU implementations of ROFL achieve almost equal outcome while the mean deviation is slightly higher for the POriso algorithm³. More information about the individual results is available in the distributions depicted in Fig. 7.4⁴. The distributions for ROFL and POriso are almost indistinguishable while the GOROFL distribution expresses a slightly larger fraction of high angular deviations.

The agreement of the individual in and out-of-plane angles can serve as an indicator of the source of error for the angular deviation. Direction and inclination angles are correlated in Fig. 7.5. As the results of ROFL, GOROFL and POriso are very similar, only the ROFL results are plotted. The vast majority of the direction angle pairs are close to the ideal bisecting line. The agreement is weaker for the inclination angle: the

³Standard errors for the sample means are not reported as the angular deviations are not normally distributed.

⁴The histograms utilize Knuth's rule [156] for the bin size from the astropy package [157]. Knuth's rule estimates the optimal bin size from a likelihood based approach to provide the best approximation of the data distribution.

Estimator	Mean angular deviation $\langle \epsilon \rangle$ [°]
ROFL	9.45
GOROFL	9.35
POriSo	10.03

TABLE 7.2: Mean angular deviations between TPFM fiber bundle orientations and 3D-PLI orientations computed by ROFL, POriSo and GOROFL




FIGURE 7.4: Histograms of angular deviations between TPFM fiber bundle orientations and 3D-PLI fiber bundle orientations computed by ROFL, POriSo and GOROFL.

data pairs are far more spread out. Especially for $\alpha_{TPFM} > 60^\circ$, the difference becomes higher.

The variance of the individual angles can be assessed in Fig. E.4 and E.5. The 67% direction angle density interval is smaller than 30° for almost all bundles. On the other hand, the density intervals for the inclination angle are 61° on average. In the scatterplot (cf. E.5), the bigger variance is clearly observable.

Discussion

The comparison of 3D-PLI and TPFM measurements showed a strong agreement of the in-plane fiber bundle orientations. For the out-of-plane orientation, an improved agreement compared to 3D-PLI measurements without oblique views was found. On average, the angular error of the 3D-PLI fiber bundle orientation amounts to 10° . For individual fiber bundles though, deviations of 30° - 40° occurred. The high deviations originate from wrongly estimated inclination angles in 3D-PLI. In the investigated bundles, no fiber bundle expressed a principal orientation with an inclination greater than 60° . Overall, the estimated inclination angles vary much more strongly than the estimated direction angles in individual nerve fiber bundles which again indicates that the employed model is not sufficient for a coherent data interpretation.

The validation revealed no distinct differences between the three algorithms for the fiber orientation estimation. All result in very similar deviations from the ground truth orientation. While the GOROFL algorithm achieves the smallest average deviation, it also results in a distribution of deviations with a longer tail than the CPU based algorithms (cf. Fig. 7.4). As the angular deviations are not normally distributed, evaluating the algorithms solely based on their mean angular deviations is misleading.




FIGURE 7.5: Comparison of 3D-PLI and TPFM ground truth orientation angles. **Left:** direction angle correlation. **Right:** inclination angle correlation. 3D-PLI orientations were obtained from the GOROFL algorithm. 3D-PLI inclination angles differ more strongly from the ground truth than the direction angles.

Based on the angular deviations, the ROFL and POriso algorithms are therefore slightly superior to GOROFL. From a computational viewpoint though, the GOROFL algorithm has a clear advantage as processing a tile with app. 4 million pixels takes 30 seconds compared to 2-3 core hours for ROFL and POriso.

In [46], perspective shifts were identified as a source of error. The theoretical model presumes that the rotation of the illumination can be represented by a rotation of the optic axis. In reality though, the light rays from different oblique view reach different pixels on the camera which is not accounted for by the polarimetric model. At LAP resolution, neighbouring pixels, especially in white matter, can be expected to be very similar in terms of birefringence strength and fiber orientation. At microscopic resolution, this expectation no longer holds, e. g. consider that one pixel is outside of a nerve fiber bundle while its neighbour still contains the bundle.

Conclusion and Outlook

“Discussion and argument are essential parts of science; the greatest talent is the ability to strip a theory until the simple basic idea emerges with clarity.”

– Albert Einstein

In this study several limitations of 3D-PLI data analytics were pushed. For the first time, an artifact-free derivation of the nerve fiber inclination was accomplished based on a least squares fitting routine (ROFL). At the same time, runtime constraints could be removed using GPU hardware. Although it represented more an organizational than a technical challenge, the first large-scale brain volume at mesoscopic resolution was reconstructed in the scope of this thesis. Combined with the developed sectionwise reconstruction of the three-dimensional fiber orientation continuous fiber tracts could be restored from the individual histological sections. The reconstructed dataset of the occipital lobe might enable the first volumetric tractography studies on 3D-PLI data. Besides reliable fiber orientations an equally important prerequisite for tractography is high registration quality. The brain registration in 3D-PLI might benefit from polyaffine and polyrigid registration methods as they would allow to register deformed gyri separately from the rest of the brain section [158, 159], which represents a problem for the employed simultaneous registration approach.

While ROFL obtains robust fiber orientation estimates even for weak birefringent signals, it results in unreliable parameter maps for very low signals which arise predominantly in the cortex. To improve the reconstruction for these cases a maximum-a-posteriori estimator which appoints every fiber orientation the same probability was introduced to 3D-PLI (POriSo). On LAP data, it could be shown that this approach yields more accurate estimates for the relative thickness in cortical areas which appear noisy for ROFL. In the cortex, the algorithms also compute different fiber orientations as they are biased into the sectioning plane for POriSo by construction. Still, the fiber orientation determination is unreliable for very small signals even on synthetic data, rendering these differences negligible. Currently, one obvious disadvantage of POriSo compared to ROFL is its computation time which could however be compensated by a similar GPU based implementation.

So far, 3D-PLI data interpretation was limited to estimations of the best fit parameters. This constraint was removed in this thesis. Based on MCMC sampling confidence measures for the obtained orientation and the relative thickness were obtained. As expected, the out-of-plane fiber orientation is afflicted with higher uncertainty than the

in-plane fiber orientation. It could be shown that for most white matter regions outside of crossing areas the fiber orientation can be reconstructed with high confidence using the LAP. Increasing the confidence further would require a less noise sensitive camera. Yet as the measurement setup requires a camera with a high depth of field compared to other typical applications in order to capture the tilted measurements, there are limited options for suitable cameras for the LAP.

Application of the BAMBI framework to full brains is computationally prohibitive with the current implementation. In DMRI, GPU implementations have been utilized to reduce the computation time [130, 129] but even a reduction of the computation time by a factor of 100 would still yield a computation time of several hours for one section measured with the LMP3D. Another option would be to exploit the relationships between the credible intervals and the best fit parameters. It was shown that the direction angle credible interval decreases with the retardation and the inclination angle credible interval decreases with the relative thickness. These relationships can be fitted by a model for a small number of pixels which then predicts the credible interval for all other pixels. That way, the computationally expensive sampling procedure only has to be executed for a small fraction of the data.

While the oblique measurements could only be assessed qualitatively at LAP resolution, the comparison with higher resolved TPFM measurements represented the first validation experiments for 3D-PLI. For a high number of fiber bundles, a good agreement with the ground truth fiber orientation was found. Especially the derived inclination angles are a much better match of the ground truth than the inclinations calculated with the planar model. Still, particularly for highly inclined fiber bundles larger deviations occurred. The most likely reason for this observation is that the employed model does not account for perspective shifts present in the measurement. A famous quote about scientific modeling states that "All models are wrong but some are useful" [160]. This seems applicable here as the model does certainly not fully describe all necessary effects but is still capable of reconstructing plausible fiber orientations for a considerably large number of fiber bundles. On the other hand it is important to note that the comparison with TPFM does not validate the fiber orientations pixelwise. The insufficiency of the model for example manifests in the strong variation of the inclination angles in individual fiber bundles. More coherent fiber orientations in neighbouring pixels could be enforced by a prior, yet this would lead to strong smoothing effects at the edges of fiber bundles at microscopic scale.

In future studies, the Bayesian framework introduced here could be extended further. Simulative and experimental results showed that highly inclined nerve fibers result in a lower transmittance value [61]. This behaviour could be incorporated into a prior for the inclination angle. Formulating a probabilistic description for this case likely requires a more sophisticated approach including hyperpriors for the inclination dependency of the transmittance. Another approach could be to allow imperfect polarization filters in a model based on Müller-Stokes calculus as presented in [46, 61, 75]. As full Bayesian inference provides more information than the fitting methods utilized in these studies this might enable a pixelwise characterization of the hardware limitations. In this context it might be useful to also characterize the camera detection noise for each pixel individually as the manufacturing process and aging effects result in individual noise characteristics of pixels in CCD cameras. An even more elaborate model could include the noise gain factor as a model parameter like the polarization filter effects. Another parameter which was treated as constant here is in fact not precisely known: the oblique illumination angle in

the brain section. As the refractive index of brain tissue varies between 1.3 and 1.5 [61], refraction results in possible internal oblique illumination angles of $\tau = 5.32 - 6.15^\circ$ for the LAP and $\tau = 3.80 - 4.38^\circ$ for the LMP3D. Taking the refractive index into account might make a determination of the refraction properties of the brain section possible.

In this work, the theoretical limitations of 3D-PLI measurements for low signals were evaluated based on synthetic data. With the Cramer-Rao bound, Bayesian statistics provide a mathematically rigorous method to derive a theoretical lower bound for the possible accuracy of the model parameters [161, 162]. Future studies could employ the Cramer-Rao bound to validate if the estimators developed in this study reach the theoretical limits. For the planning of future 3D-PLI setups, the Cramer-Rao bound could be utilized to theoretically assess the system's restrictions, e. g. if a certain orientation confidence even for highly inclined nerve fibers is required.

From a mathematical perspective, it would be more elegant to reformulate the 3D-PLI parameter estimation problem using modern manifold methods to exploit the spherical symmetry of the parameter space. So far, the parameter space for direction, inclination and t_{rel} was implicitly given by $[0, \pi] \times [-\pi/2, \pi/2] \times (0, 1]$. Instead of separating direction and inclination angles, the parameter space could be formulated based on an orientation on a sphere. This would circumvent the need for the inclination prior utilized in this thesis to assign every point on the sphere the same prior probability. On manifolds, gradient based optimization and MCMC sampling algorithms can be applied in the same way as in Euclidean space once the differential geometry of the manifold is taken into account [163, 164, 165]. For example, exploiting the geometry of the parameter space by using the "natural gradient", the gradient on the manifold, has been successfully applied to train neural networks [166] and for blind source separation via Independent Component Analysis in 3D-PLI [79]. Employing gradient based sampling algorithms such as the self tuning No-U-Turn Sampler [167, 168] on the sphere could potentially drastically reduce the number of samples required to quickly and reliably explore the parameter space compared to the ensemble sampler utilized here.

In a similar way, the field of directional statistics offers tools for a more extensive description of the fiber orientation per pixel. Here, only the principal orientation and HPD intervals for inclination and direction angles were derived. Projected on a sphere, the samples of inclination and direction yield an ellipse. The empirical distribution of the orientation samples could be fitted by directional distributions such as the Kent distribution [169] or the angular central gaussian distribution [170], which represent analogues of the Normal distribution on the sphere. Sampling from these distributions could then represent a variety of possible local fiber orientations in probabilistic tractography frameworks similar to approaches employed in the field of DMRI [171, 22].

The suggestions introduced in the previous paragraphs still utilize the same physical model for the interaction of light with the brain tissue and will suffer from the same pitfalls for regions of crossing fibers. Future studies should investigate if it is possible to extend the polarimetric model to account for several fiber orientations per pixel. One possibility could be to model two crossing fibers as two individual wave retarders. Mathematically these would be represented by two Jones matrices which are multiplied one after the other with the electric field vector. Physically this implies that in this model both fiber populations would be stacked on top of each other. Although this model is again a simplification of crossing nerve fibers it could still be more accurate than the current model. However, it is not clear if this model can be solved unambiguously. Here, a Bayesian approach might be helpful: as investigated in [46], the current model

reconstructs the dominant nerve fiber in the scenario of two crossing fibers. This result could then serve as a prior for the first fiber in the model of two fibers.

While computational aspects were not the major driving force behind this thesis, the GOROFL algorithm once more proved the supremacy of GPUs for massively parallel problems such as they arise in the context of 3D-PLI. Recent efforts of GPU vendors might in future enable to circumvent the RAM and CPU completely based on technology such as *GPUDirectStorage* and stream data directly to GPU memory [172]. For the processing of whole brains at microscopic resolution, adopting such strategies likely represents the most promising way to keep the computational demands as small as possible. At the same time, it would be useful to develop a common framework which can execute both parameter fitting and sampling for 3D-PLI on the GPU similar to frameworks employed by the DMRI community [129, 130].

The most pressing question for 3D-PLI in future is the interpretation of the oblique microscopic measurements. With the GPU accelerated fitting algorithm developed here, a first step towards an automated and in terms of required computation power acceptable analysis has been taken. As argued before, the current model provides plausible results after averaging over individual fiber bundles but is not sufficient at the pixel level due to perspective shifts between the individual measurements. At microscopic resolution, coregistration of the individual oblique measurements is not advisable as this would produce severe interpolation effects. Instead, a physical model for the interaction of light with the brain tissue which includes polarization and optical effects needs to be developed. Future studies should also repeat and extend the validation of 3D-PLI based on other imaging techniques such as TPFM. One significant drawback of the analysis carried out in this study is the lack of a higher number of steep fibers in the chosen tissue sample: only eight of the evaluated fiber bundles are inclined by more than 60° . Another interesting candidate for validation is OCT as it also exploits the birefringence of myelinated nerve fibers and enables volumetric measurements. Also, the ideas developed for the analysis of oblique views in 3D-PLI could be applied to OCT as well.

As pointed out in the the first chapter of this work, the analysis of the connectome at high resolution is challenging due to the problem of reconstructing three-dimensional nerve fibers from histological brain sections. This thesis showed that 3D-PLI indeed enables the reconstruction of the connectome at different scales from individual fibers bundles to large scale brain models. The novel algorithms pave the way for further quantitative studies of nerve fiber models based on 3D-PLI, especially tractography.

Acknowledgements

This work would not have been possible without many other people whom I want to show my sincere gratitude. First of all, I would like to take the opportunity to thank my supervisor Dr. Markus Axer and my professors Prof. Dr. Dr. Thomas Lippert, Prof. Dr. Uwe Pietrzyk and Prof. Dr. Katrin Amunts for their advice and leadership during my PhD project. I also thank Prof. Dr. Bolten for joining my PhD examination committee.

At the same time, I am grateful to many colleagues without whom this thesis would not have been possible. In particular, I thank Felix Matuschke for the many discussions about oblique 3D-PLI measurements, programming and politics we had and maintaining the coffee machine for the rest of us. For carrying out the LMP3D measurements at the bread board system and many hardware related issues, a special thanks goes to Philipp Schlömer. A big thank you goes to Miriam Menzel who made the the comparison between 3D-PLI and TPFM possible. Despite being only described in the appendix, the 3D reconstruction of the studied LAP dataset represented a significant challenge which was mostly faced by Martin Schober: thank you. For error checking the blockface volume another thank you goes to David Gräfel. Installing the simultaneous registration software would have been impossible without the help of Jan Reuter and Marcel Huysegoms. In this context, Nicole Schubert was a great help and colleague for all visualization related work. For implementing the GOROF algorithm without prior knowledge of GPU programming, a big thank you goes to Jan-Oliver Kropp. Outside of the PLI group, I would like to thank Susanne Wenzel for very helpful discussions about MCMC sampling and Andrea Brandstetter and Daniel Zachlod for pointing out how to compare 3D-PLI and DMRI from a neuroanatomical viewpoint.

For the tedious manual work of segmenting the LAP sections and the fiber bundles for the TPFM comparison, I would like to thank all colleagues who participated. Especially Heidi Mellenthin put many hours into a fine segmentation of the LAP sections, thank you.

For proofreading the manuscript, special thanks go to my mother and my sister. Last but not least another thank you goes to the colleagues in the institute who provided opportunities for distractions from the PC work, in particular Daniel Krötz and Sarah Oliveira.

This project has received funding from the European Union’s Horizon 2020 Research and Innovation Programme under Grant Agreement No. 7202070 (HBP SGA1) and Grant Agreement No. 785907 (HBP SGA2). The author gratefully acknowledges the computing time granted by the JARA-HPC Vergabegremium and provided on the JARA-HPC

Partition part of the supercomputer JURECA at Forschungszentrum Jülich.

A

Mathematical derivations for the ROFL algorithm

A.1 Rotation matrices of the tilted measurements

For the analysis of the tilted measurements, a coordinate system which describes the tilted views by a tilting direction ψ and a tilting angle τ is useful (cf. Fig. 4.1). In [46] the rotation matrices for the transformation between this coordinate system and the planar coordinate system without tilt were derived. For the sake of completeness, the rotation matrices are provided here. The full rotation is obtained by first rotating around the z-axis by $-\psi$, then rotating around the y-axis by τ and then rotating back around the z-axis by ψ . In matrix form, this can be written as (see. [46], p. 105)

$$\begin{aligned} \mathbf{R}(\psi, \tau) &= \mathbf{R}^z(\psi) \mathbf{R}^y(\tau) \mathbf{R}^z(-\psi) \\ &= \begin{pmatrix} \cos(\psi) & -\sin(\psi) & 0 \\ \sin(\psi) & \cos(\psi) & 0 \\ 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} \cos(\tau) & 0 & \sin(\tau) \\ 0 & 1 & 0 \\ -\sin(\tau) & 0 & \cos(\tau) \end{pmatrix} \cdot \begin{pmatrix} \cos(\psi) & \sin(\psi) & 0 \\ -\sin(\psi) & \cos(\psi) & 0 \\ 0 & 0 & 1 \end{pmatrix} \\ &= \begin{pmatrix} \cos(\tau) \cos(\psi)^2 + \sin(\psi)^2 & (\cos(\tau) - 1) \sin(\psi) \cos(\psi) & \cos(\psi) \sin(\tau) \\ (\cos(\tau) - 1) \sin(\psi) \cos(\psi) & \cos(\tau) \sin(\psi)^2 + \cos(\psi)^2 & \sin(\psi) \sin(\tau) \\ -\cos(\psi) \sin(\tau) & -\sin(\psi) \sin(\tau) & \cos(\tau) \end{pmatrix}. \end{aligned} \tag{A.1}$$

A.2 Error propagation from camera noise to the normalized light intensity

In this section the error propagation from the measured light intensity to the normalized intensities is calculated. For n light intensities I_i with $i = 1, \dots, n$ whose variance is given by $\sigma_i^2 = gI_i$ with gain factor g the transmittance I_T and the normalized light intensity I_N are defined as

$$I_T = \frac{2}{n} \sum_{i=1}^n I_i \tag{A.2}$$

$$I_{N_i} = \frac{2I_i}{I_T} - 1 \tag{A.3}$$

First, the variance of the transmittance $\sigma_{I_T}^2$ is derived as

$$\sigma_{I_T}^2 = \sum_{i=1}^n \left(\frac{\partial I_T}{\partial I_i} \right)^2 \sigma_i^2 \quad (\text{A.4})$$

$$= \sum_{i=1}^n \frac{4}{n^2} g I_i \quad (\text{A.5})$$

$$= \frac{2g I_T}{n} \quad (\text{A.6})$$

Now, the variance of the normalized intensity $\sigma_{N_i}^2$ can be calculated as

$$\sigma_{N_i}^2 = \left(\frac{\partial I_{N_i}}{\partial I_i} \right)^2 \sigma_i^2 + \left(\frac{\partial I_{N_i}}{\partial I_T} \right)^2 \sigma_{I_T}^2 \quad (\text{A.7})$$

$$= \frac{4}{I_T^2} g I_i + \frac{4I_i^2}{I_T^4} \frac{2g I_T}{n} \quad (\text{A.8})$$

$$= \frac{4g I_i}{I_T^2} \left(1 + \frac{2I_i}{n I_T} \right) \quad (\text{A.9})$$

A.3 Gradient of the objective function

In this section, the gradient of the χ^2 function which has to be minimized for the ROFL algorithm, is derived. While in recent years automatic differentiation has been crucial for e.g. deep learning applications, `gprof` which is utilized by the GPU implementation in sec. 4.4 currently still requires analytical gradients. Also, as the automatic computation of gradients follows simple logical rules, its performance is not necessarily optimal. Therefore, for the ROFL algorithm the gradient was derived and implemented manually for the CPU and the GPU implementation.

The objective function is (see. eq. (A.10))

$$\underset{\varphi, \alpha, t_{\text{rel}}}{\text{argmin}} \chi^2 = \underset{\varphi, \alpha, t_{\text{rel}}}{\text{argmin}} \sum_{j=0}^{N_T} \sum_{i=0}^{N_P} \left((f_{ji}(\rho_i, \varphi_j, \alpha_j, t_{\text{rel},j})) - I_{N_{ji}} \right) \cdot w_{ji}^2. \quad (\text{A.10})$$

with the model

$$f_{ji}(\rho_i, \varphi, \alpha, t_{\text{rel}}) = \sin(2(\rho_i - \varphi_j)) \cdot \underbrace{\sin\left(\frac{\pi}{2} t_{\text{rel},j} \cos(\alpha_j)^2\right)}_{=:r}. \quad (\text{A.11})$$

where index j stands for the application of the rotation matrix of tilting position j from eq. A.1 to the orientation vector. As `scipy's` `leastsq` function and `gprof` only require the gradient of the model function, the sum over all observables is neglected here. The partial derivatives of f are

$$\begin{aligned} \frac{\partial f_{ji}}{\partial \varphi} &= -2 \cos(2\rho_i - 2\varphi_j) \frac{\partial \varphi_j}{\partial \varphi} r \\ \frac{\partial f_{ji}}{\partial \alpha} &= -\sin(2(\rho_i - \varphi_j) \cos(r) \pi t_{\text{rel},j} \sin(\alpha_j) \cos(\alpha_j) \frac{\partial \alpha_j}{\partial \alpha} \\ \frac{\partial f_{ji}}{\partial t_{\text{rel}}} &= \sin(2(\rho_i - \varphi_j) \cos(r) \frac{\pi}{2} \cos(\alpha_j)^2 \frac{\partial t_{\text{rel},j}}{\partial t_{\text{rel}}} \end{aligned} \quad (\text{A.12})$$

For the tilted relative thickness the gradient is given by

$$\frac{\partial t_{\text{rel}j}}{\partial t_{\text{rel}}} = \begin{cases} \frac{\partial}{\partial t_{\text{rel}}} \frac{t_{\text{rel}}}{\cos(\tau)} = \frac{1}{\cos(\tau)} \approx 1 & \text{for } j \neq 0 \\ \frac{\partial}{\partial t_{\text{rel}}} t_{\text{rel}} = 1 & \text{for } j = 0 \end{cases}$$

The partial derivatives of the rotated angles are harder to obtain: calculating the derivative of the multiplication of the rotation matrices from eq. (A.1) with a vector in spherical coordinates would result in very tedious calculations. Still, as the differences between the rotated and planar vector are small, the gradients can be approximated by 1. Therefore, all partial derivatives of the rotated parameters are set to 1:

$$\frac{\partial \varphi_j}{\partial \varphi} = \frac{\partial \alpha_j}{\partial \alpha} = \frac{\partial t_{\text{rel}j}}{\partial t_{\text{rel}}} \approx 1 \quad (\text{A.13})$$

The gradient given by eq. (A.12) with this approximation was validated by a comparison with finite differences.

B

Volumetric reconstruction of the brain sections

In this section, the 3D reconstruction of the analyzed brain volume is presented. Blockface reconstruction and sectionwise registration until the rigid registration step were carried out by Martin Schober. The outer boundaries of the histological images were delineated manually in a joint effort of the 3D-PLI team. Subsequently the delineations were revised by Heidi Mellenthin to correct for holes still present in interior parts of the sections.

Dataset description The whole brain volume consists of a right human hemisphere. The organ was obtained in accordance with legal and ethical requirements. Not the full volume was sectioned so far: 843 sections of $70\ \mu\text{m}$ thickness were produced by the INM-1 lab. All but the first 75 sections were designated for 3D reconstruction.

Blockface reconstruction In a first step, the individual blockface sections have to be stacked together to form a coherent reference volume for the histological sections. The blockface setup consists of the *AVT OscarF810C* camera (sensor size $s = 8.8 \times 6.6$ mm, number of pixels: $n_s = 3288 \times 2470$ pixels) equipped with an objective of a focal length of $f = 35$ mm (*MeVis-C 35*). Blockface images were captured at a fixed camera position while the brain was sectioned which results in an increasing pixel size with each section. As a first step, this pixel size needs to be determined which was carried out by Philipp Schlömer. The measures field of view FOV , free working distance FWD (the distance between the imaged object and the camera), focal length f and sensorsize s are related by $FOV = \frac{s \cdot FWD}{f}$. The field of view is also directly related to the pixel size d by $FOV = n_s d$. Putting both equations together yields a simple formula for the pixel size in object space:

$$d = \frac{s \cdot FWD}{n_s \cdot f} . \quad (\text{B.1})$$

This equation means that only the distance to the camera determines the pixel size in our case which makes intuitively sense. From the known physical size of the ARTag markers, the FWD of the markers at the bottom of the brain block was calculated as 85 cm. To determine the FWD of section 843 the height of the brain block was measured. The FWD is then 85 cm minus the measured height which amounts to 77.5 cm. For the lower section numbers the FWD then reduces by $70\ \mu\text{m}$, the section thickness, per section. The resulting pixel sizes are $54.7 \times 54.7\ \mu\text{m}$ for section 1 and $59.3 \times 59.3\ \mu\text{m}$ for section 843.


In image registration, the moving image is typically resampled at the pixel size of the template image. To avoid an interpolation of the histological sections at a higher resolution, the blockface images were therefore artificially resampled at a pixel size of $64 \times 64 \mu\text{m}$ to mimic the pixel size of the LAP. The resampled images were then reconstructed to a volume using the ARTag markers as described in [80]. The resulting volume was checked for anatomical correctness by David Gräfel and shifts still present in the volume removed by section-to-section registration.

Sectionwise registration The registration of the histological sections onto their corresponding blockface sections follows a stepwise approach. First, the center of the masses of the binary masks of both modalities were aligned. The resulting translation was then applied to the transmittance image of the histological section. Next, the transmittance image was rotated to achieve a first rough alignment. Therefore the transmittance was rotated in steps of 1° . The rotation angle which maximizes the normalized mutual information as similarity metric [173, 174] of transmittance and blockface image provides the best rotation. After this rough alignment in the next step a rigid registration of the transmittance images onto the blockface image which allows both translations and rotations was carried out based on the registration software *elastix* [82] using the normalized mutual information metric. Additionally, for 47 sections landmarks were manually set to guide the rigid registration. The sectionwise registration was completed by an affine registration step which allows shearing and scaling of the images again based on *elastix* and the normalized mutual information metric. The rotational component of the affine transformations was used to rotate the estimated fiber orientations. Nearest neighbour interpolation was applied for the transformations.

Three-dimensional registration The sectioning procedure induces severe stress on the brain tissue which results in strong deformations of the individual sections. These deformations cannot be reversed by a global affine transformation of the whole image but require small local deformations. One approach to this problem is *simultaneous registration* [83]. It utilizes a bspline transformation model [175] informed not only by the blockface template but also by neighbouring sections, thereby registering the histology image not only onto its reference but also onto the neighbouring sections which are of the same modality. Registration of the same modalities allows more direct similarity metrics instead of the purely statistical mutual information metric which is employed between histology and blockface. Intuitively, the bspline registration works in the following way: a fixed amount of grid points is positioned on each histological image. Each of these grid points is then moved until the best displacement is found. As a registration of larger volumes requires the optimization of tens of thousands of displacements and the similarity metric has to be calculated for all sections, this approach is computationally very challenging. In the INM-1, Marcel Huysegoms developed a GPU implementation of the simultaneous registration [176] which was utilized here.

The registration parameters were set as 15×21 grid points per section guided by mutual information between transmittance and blockface template and normalized cross correlation [177] between the histological sections as similarity metrics. To comply to GPU memory restrictions, all images were scaled by 0.7 in both dimensions for the computations. Again the resulting deformations were applied using nearest neighbour interpolation.

Visualization The transmittance and retardation modalities of 3D-PLI strongly depend on the time between embedding of the tissue and the measurement [61] and the employed embedding solution. This results in severe intensity variations between subsequent sections when a registered volume is rendered in 3D. For visualization purposes, the grey values of the individual sections were therefore matched using the *Histogram-Match* function of the *Advanced Normalization Tools* software [178]. The volumes were rendered using the softwares *Paraview* [179] and *vaa3d* [180].



System configuration for the runtime benchmarks

In this section the system configuration of the server used for the runtime benchmarks in sec. 4.4 is given.

CPU	16 × Intel Xeon E5-2690: 2.9 GHz, 2 threads
RAM	256 GiB
GPU	NVIDIA GTX 1080: 8 GiB GPU RAM, 2560 CUDA cores a 1607 MHz

TABLE C.1: Computational environment for benchmarks

D

Optimization experiments for the LOriE algorithm

In 5.1 a new optimization approach for the LOriE algorithm which utilizes the L-BFGS-B algorithm was developed. Here it is shown that adding the constraint $t_{\text{rel}} \in [0, 1)$ is necessary for a stable solution. In Fig. D.1 the results of optimization without constraints (left) and with the constraint are depicted. For the unconstrained optimization the BFGS algorithm [181], the unbounded variant of the L-BFGS-B algorithm, was applied. Without the restriction artifacts are observable.




FIGURE D.1: Comparison of optimization approaches for the LOriE algorithm. Left: unbounded, right: restriction $t_{\text{rel}} \in [0, 1)$

E

Complementary data




FIGURE E.1: Cumulative distribution functions of mean squared errors between measured and predicted light intensities. 3D-PLI parameters were estimated by ROFL (orange) and POriso (green).




FIGURE E.2: Correlation of retardation $\sin \delta$ and direction angle credible interval σ_φ for a tilted 3D-PLI measurement. The plot displays a two-dimensional histogram of the retardation and direction angle credible interval maps shown in Fig. 6.7. The credible interval decreases with increasing retardation values.




FIGURE E.3: Cumulative densities of angular credible intervals for the whole analyzed brain section. Blue: direction angle. Orange: inclination angle.




FIGURE E.4: Comparison of 3D-PLI and TPFM ground truth directions including 3D-PLI variability analysis. 3D-PLI directions were obtained from the GOROFLL algorithm. The errorbars indicate the boundaries of the shortest interval which contains 67% of all samples. For reasons of clarity, only every third fiber bundle is plotted.




FIGURE E.5: Comparison of 3D-PLI and TPFM ground truth inclinations including 3D-PLI variability analysis. 3D-PLI inclinations were obtained from the GOROFLL algorithm. The errorbars indicate the boundaries of the shortest interval which contains 67% of all samples. For reasons of clarity, only every third fiber bundle is plotted.

List of symbols

δ	Phase retardation
Δn	Birefringence
φ	Direction angle
α	Inclination angle
$\mathbf{E}$	Electric field vector
n	Refractive index
n_e	Orientation dependent refractive index
n_E	Extraordinary refractive index
k	Wave number
ω	Angular velocity
$\mathbf{R}$	Rotation matrix
$\mathbf{r}$	Fiber orientation vector
$\mathbf{J}$	Jones vector
$\mathbf{M}$	Jones matrix of a wave retarder
$\mathbf{P}$	Jones matrix of an ideal polarizer
ρ	Polarization angle
t_{rel}	Relative section thickness
p	Probability
μ	Expected value
σ	Standard deviation
σ^2	Variance
$\mathcal{N}$	Normal Distribution
I	Light intensity
$\mathcal{NB}$	Negative Binomial Distribution
ψ	Tilting direction angle
τ	Tilting angle
g	Camera noise gain factor

List of abbreviations

3D-PLI	Three-Dimensional Polarized Light Imaging
API	Application Programming Interface
BAMBI	BAYesian Multivariate Birefringence Inference
CDF	Cumulative Density Function
CPU	Central Processing Unit
CUDA	Compute Unified Device Architecture
DFT	Discrete Fourier Transform
DMRI	Diffusion MRI
GOROFL	GPU Optimized Robust Orientation Fitting via Least Squares algorithm
GPU	Graphics Processing Unit
HPD	High Posterior Density
I/O	Input/Output
LAP	Large Area Polarimeter
L-BFGS-B	Limited memory Broyden–Fletcher–Goldfarb–Shanno algorithm
LLVM	Low Level Virtual Machine
LOrIE	Likelihood Orientation Estimation algorithm
LMP	Large Metripol
LMP3D	Large Metripol 3D
MAP	Maximum-a-posteriori
MCMC	Markov Chain Monte Carlo
MLE	Maximum Likelihood Estimation
MSE	Mean Squared Error
OCT	Optical Coherence Tomography
PPD	Posterior Predictive Distribution
POriSo	Posterior Orientation Solver
RAM	Random Access Memory
ROFL	Robust Orientation Fitting via Least Squares algorithm
RWMH	Random Walk Metropolis Hastings algorithm
SNR	Signal-to-Noise Ratio
TPFM	Two Photon Fluorescence Microscopy
WLS	Weighted Least Squares

List of Figures

2.1	Indicatrix of a uniaxial positive birefringent material	7
2.2	Myelinated axon	8
2.3	Employed polarimetric setups	10
2.4	3D-PLI coordinate system	13
2.5	Analytical fit of the 3D-PLI model to the light intensities measured in one pixel	15
2.6	3D-PLI parameter maps of a coronal human brain section measured in the LAP	15
2.7	3D-PLI fiber orientation maps	17
2.8	Volumetric reconstruction in 3D-PLI	17
2.9	Analysis of inclination angles calculated by the DFT and LOriE algorithms for in-plane oriented nerve fibers	20
3.1	Example data for linear regression problem	22
3.2	Synthetic data and weighted least squares estimate	23
3.3	Synthetic data with outliers. Comparison of WLS estimate and true model	24
3.4	Synthetic data with outliers. Comparison of true model, WLS estimate and MAP estimates	25
3.5	Scatterplot matrix of posterior samples generated by MCMC for the straight line fit problem with outliers and prior variances $\sigma_\alpha = \sigma_\beta = 1$	27
3.6	Posterior predictive distribution for the straight line fit problem	29
4.1	Tilting coordinate system	32
4.2	Working principle of the ROFL algorithm	35
4.3	Reconstruction accuracy of ROFL and DFT algorithms evaluated on synthetic data	38
4.4	Retardation gradient with respect to inclination	39
4.5	Evaluation of inclination bias on synthetic data	40
4.6	Geometrical explanation for the $\cos(\alpha)$ dependency of the inclination distribution of uniformly distributed orientations	41
4.7	Overview of the analyzed human brain volume	43
4.8	Comparison of vectorfields derived from the ROFL and DFT algorithms	45
4.9	Inclination histograms for one ROI of in-plane oriented nerve fibers obtained from the ROFL, DFT and LOriE algorithms.	46
4.10	Mean squared error histograms for ROFL and DFT algorithms	46
4.11	Analysis of tilted 3D-PLI measurements for the cortex	47
4.12	Diagram of GPU implementation of ROFL	51

4.13	Comparison of ROFL and GOROFL results: full section view.	52
4.14	Comparison of ROFL and GOROFL results: local analysis in the cortex	53
4.15	Total runtime shares of the different steps of the GPU implementation of the ROFL algorithm	54
5.1	Comparison of the Nelder-Mead and L-BFGS-B optimization algorithms applied to the LOriE algorithm	58
5.2	Reconstruction accuracy of ROFL and PORISO algorithms evaluated on synthetic data	60
5.3	Comparison of vectorfields derived from the ROFL and POriso algorithms	63
5.4	Mean squared error histograms for DFT, ROFL and POriso algorithms	64
5.5	Reassessment of tilted 3D-PLI measurements for the cortex: comparison of ROFL and POriso	64
6.1	Results of the BAMBI framework for the planar case on synthetic data	69
6.2	Investigation of the in-plane orientation confidence for the planar 3D-PLI measurement based on LAP data	72
6.3	Results of the BAMBI framework for a white matter pixel	73
6.4	Results of the BAMBI framework for a grey matter pixel	73
6.5	Microscopic results for BAMBI	75
6.6	Simulation results of the BAMBI framework for the oblique 3D-PLI measurement	78
6.7	Results of the BAMBI framework for a tilted measurement: whole section	80
6.8	Results of the BAMBI framework for a tilted measurement: region of interest in the stratum sagittale	82
6.9	Results of the BAMBI framework for an oblique 3D-PLI measurement of a white matter pixel	83
6.10	Results of the BAMBI framework for an oblique 3D-PLI measurement of a grey matter pixel	84
6.11	Cumulative densities of angular credible intervals of white matter pixels. Blue: direction angle. Yellow: inclination angle.	85
6.12	Correlation of relative thickness t_{rel} and inclination angle credible interval σ_{α}	85
7.1	Noise measurement of the LMP3D	88
7.2	Results from measurements from oblique views at the microscale	91
7.3	Derivation of fiber bundle orientations for a comparison of 3D-PLI and TPFM	93
7.4	Histograms of angular deviations between TPFM fiber bundle orientations and 3D-PLI fiber bundle orientations computed by ROFL, POriso and GOROFL.	94
7.5	Comparison of 3D-PLI and TPFM ground truth orientation angles	95
D.1	Comparison of constraint and unconstraint optimization approaches for the LOriE algorithm	113
E.1	Comparison of CDFs of Mean Squared Errors achieved by POriso and ROFL	115
E.2	Correlation of retardation $\sin \delta$ and direction angle credible interval σ_{φ} for a tilted 3D-PLI measurement	116

E.3	Cumulative densities of angular credible intervals for the whole analyzed brain section. Blue: direction angle. Orange: inclination angle.	116
E.4	Comparison of 3D-PLI and TPFM ground truth directions including 3D-PLI variability analysis	117
E.5	Comparison of 3D-PLI and TPFM ground truth inclinations including 3D-PLI variability analysis	117

List of Tables

4.1	Runtimes of different implementations of the ROFL algorithm for a brain section of app. 1 Mio. pixels.	52
7.1	Computation times of the developed estimators for one field of view of the LMP3D	89
7.2	Mean angular deviations between TPFM fiber bundle orientations and 3D-PLI orientations computed by ROFL, POriso and GOROF	94
C.1	Computational environment for benchmarks	111

References

- [1] World Health Organization. *Neurological disorders : public health challenges*. 2006.
- [2] Deep learning for visual understanding: A review. *Neurocomputing*, 187:27 – 48, 2016. Recent Developments on Deep Big Vision.
- [3] Gao Huang, Guang-Bin Huang, Shiji Song, and Keyou You. Trends in extreme learning machines: A review. *Neural Networks*, 61:32 – 48, 2015.
- [4] Catherine Schuman, Thomas Potok, Robert Patton, J. Birdwell, Mark Dean, Garrett Rose, and James Plank. A survey of neuromorphic computing and neural networks in hardware. 05 2017.
- [5] R. A. Nawrocki, R. M. Voyles, and S. E. Shaheen. A mini review of neuromorphic architectures and implementations. *IEEE Transactions on Electron Devices*, 63(10):3819–3829, Oct 2016.
- [6] Katrin Amunts, Christoph Ebell, Jeff Muller, Martin Telefont, Alois Knoll, and Thomas Lippert. The human brain project: Creating a european research infrastructure to decode the human brain. *Neuron*, 92:574–581, 11 2016.
- [7] Australian brain alliance. *Neuron*, 92(3):597 – 600, 2016.
- [8] Judy Illes, Samuel Weiss, Jaideep Bains, Jennifer A. Chandler, Patricia Conrod, Yves De Koninck, Lesley K. Fellows, Deanna Groetzinger, Eric Racine, Julie M. Robillard, and Marla B. Sokolowski. A neuroethics backbone for the evolving canadian brain research strategy. *Neuron*, 101(3):370 – 374, 2019.
- [9] Ling Wang. Mu-ming poo: China brain project and future of chinese neuroscience. *National Science Review*, 4, 02 2017.
- [10] Hideyuki Okano, Atsushi Miyawaki, and Kiyoto Kasai. Brain/minds: brain-mapping project in japan. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 370, 05 2015.
- [11] Sung-Jin Jeong, In Young Lee, Bang Ook Jun, Young-Joon Ryu, Jeong woo Sohn, Sung-Phil Kim, Choong-Wan Woo, Ja Wook Koo, Il-Joo Cho, Uhtaek Oh, Kyungjin Kim, and Pann-Ghill Suh. Korea brain initiative: Emerging issues and institutionalization of neuroethics. *Neuron*, 101(3):390 – 393, 2019.
- [12] Meghan C. Mott, Joshua A. Gordon, and Walter J. Koroshetz. The nih brain initiative: Advancing neurotechnologies, integrating disciplines. *PLOS Biology*, 16:1–5, 11 2018.

- [13] Sten Grillner, Nancy Ip, Christof Koch, Walter Koroshetz, Hideyuki Okano, Miri Polachek, Poo M-m, and Terrence Sejnowski. Worldwide initiative to advance brain research. *Nature Neuroscience*, 19:1118–1122, 08 2016.
- [14] Olaf Sporns, Giulio Tononi, and Rolf Kötter. The human connectome: A structural description of the human brain, 2005.
- [15] Walid Essayed, Fan Zhang, Prashin Unadkat, G. Cosgrove, Alexandra Golby, and Lauren O'Donnell. White matter tractography for neurosurgical planning: A topography-based review of the current state of the art. *NeuroImage: Clinical*, 15, 06 2017.
- [16] P.J. Basser, J. Mattiello, and D. LeBihan. Mr diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66(1):259 – 267, 1994.
- [17] S. Mori and Jacques-Donald Tournier. *Introduction to diffusion tensor imaging: And higher order models: Second edition*. 09 2013.
- [18] H. Johansen-Berg and T. Behrens. *Diffusion MRI*. Elsevier, 2009.
- [19] Jacques-Donald Tournier, Susumu Mori, and Alexander Leemans. Diffusion tensor imaging and beyond. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 65:1532–56, 06 2011.
- [20] Daniel C. Alexander, Tim B. Dyrby, Markus Nilsson, and Hui Zhang. Imaging brain microstructure with diffusion mri: practicality and applications. *NMR in Biomedicine*, 32(4):e3841.
- [21] Marta Drake-Pérez, Jose Boto, Aikaterini Fitsiori, Karl Lovblad, and Maria Vargas. Clinical applications of diffusion weighted imaging in neuroradiology. *Insights into Imaging*, 9, 05 2018.
- [22] 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. *Magnetic Resonance in Medicine*, 50(5):1077–1088, 2003.
- [23] Alard Roebroeck, Karla Miller, and Manisha Aggarwal. Ex vivo diffusion mri of the human brain: Technical challenges and recent advances. *NMR in Biomedicine*, 32:e3941, 06 2018.
- [24] F.J. Fritz, S. Sengupta, R.L. Harms, D.H. Tse, B.A. Poser, and A. Roebroeck. Ultra-high resolution and multi-shell diffusion mri of intact ex vivo human brains using kt-dsteam at 9.4t. *NeuroImage*, 202:116087, 2019.
- [25] Justine Beaujoin, Nicola Palomero-Gallagher, Fawzi Boumezbeur, Markus Axer, Jeremy Bernard, Fabrice Poupon, Daniel Schmitz, Jean-François Mangin, and Cyril Poupon. Post-mortem inference of the human hippocampal connectivity and microstructure using ultra-high field diffusion mri at 11.7 t. *Brain Structure and Function*, 223(5):2157–2179, Jun 2018.
- [26] Klaus Maier-Hein, Peter Neher, Jean-Christophe Houde, Marc-Alexandre Côté, Eleftherios Garyfallidis, Jidan Zhong, Maxime Chamberland, Fang-Cheng Yeh, Ying-Chia Lin, Qing Ji, Wilburn Reddick, John Glass, David Chen, Feng Yuan-jing, Chengfeng Gao, Wu Ye, Jieyan Ma, H. Renjie, Qiang Li, and Maxime De-

- scoteaux. The challenge of mapping the human connectome based on diffusion tractography. *Nature Communications*, 8, 12 2017.
- [27] Yoshiyuki Kubota, Jaerin Sohn, and Yasuo Kawaguchi. Large volume electron microscopy and neural microcircuit analysis. *Frontiers in Neural Circuits*, 12:98, 2018.
 - [28] Ennio Pannese. Il contributo di camillo golgi alla conoscenza della struttura del sistema nervoso. *Rendiconti Lincei. Scienze fisiche e naturali*, 18, 06 2007.
 - [29] Clément Hémonnot and Sarah Köster. Imaging of biological materials and cells by x-ray scattering and diffraction. *ACS Nano*, 11, 08 2017.
 - [30] A half century of experimental neuroanatomical tracing. *Journal of Chemical Neuroanatomy*, 42(3):157 – 183, 2011.
 - [31] Matthew Budde and Jacopo Annese. Quantification of anisotropy and fiber orientation in human brain histological sections. *Frontiers in Integrative Neuroscience*, 7:3, 2013.
 - [32] Ahmad Khan, Anda Cornea, Lindsey Leigland, Steven Kohama, Sune Jespersen, and Christopher Kroenke. 3d structure tensor analysis of light microscopy data for validating diffusion mri. *NeuroImage*, 111, 02 2015.
 - [33] Jonas Pichat, Juan Eugenio Iglesias, Tarek Yousry, Sébastien Ourselin, and Marc Modat. A survey of methods for 3d histology reconstruction. *Medical Image Analysis*, 46:73 – 105, 2018.
 - [34] Tomoyuki Mano, Alexandre Albanese, Hans-Ulrich Dodt, Ali Erturk, Viviana Gradinaru, Jennifer B. Treweek, Atsushi Miyawaki, Kwanghun Chung, and Hiroki R. Ueda. Whole-brain analysis of cells and circuits by tissue clearing and light-sheet microscopy. *Journal of Neuroscience*, 38(44):9330–9337, 2018.
 - [35] Stella Corsetti, Frank Gunn-Moore, and Kishan Dholakia. Light sheet fluorescence microscopy for neuroscience. *Journal of Neuroscience Methods*, 319:16 – 27, 2019. Methods and Models in Alzheimer’s Disease Research.
 - [36] Hui Wang, Adam J. Black, Junfeng Zhu, Tyler W. Stigen, Muhammad K. Al-Qaisi, Theoden I. Netoff, Aviva Abosch, and Taner Akkin. Reconstructing micrometer-scale fiber pathways in the brain: Multi-contrast optical coherence tomography based tractography. *NeuroImage*, 58(4):984 – 992, 2011.
 - [37] Serial optical coherence scanner for large-scale brain imaging at microscopic resolution. *NeuroImage*, 84:1007 – 1017, 2014.
 - [38] Ke Wang, Nicholas G. Horton, and Chris Xu. Going deep: Brain imaging with multi-photon microscopy. *Opt. Photon. News*, 24(11):32–39, Nov 2013.
 - [39] Ludovico Silvestri, Anna Letizia Allegra Mascaro, Irene Costantini, Leonardo Sacconi, and Francesco Saverio Pavone. Correlative two-photon and light sheet microscopy. *Methods*, 66(2):268 – 272, 2014. Advanced Light Microscopy.
 - [40] Patrick Theer, Mazahir T. Hasan, and Winfried Denk. Two-photon imaging to a depth of 1000 μm in living brains by use of a ti:al₂o₃ regenerative amplifier. *Opt. Lett.*, 28(12):1022–1024, Jun 2003.

- [41] Martin Oheim, Emmanuel Beaurepaire, Emmanuelle Chaigneau, Jerome Mertz, and Serge Charpak. Two-photon microscopy in brain tissue: parameters influencing the imaging depth. *Journal of Neuroscience Methods*, 111(1):29 – 37, 2001.
- [42] Kuo-Jen Hsu, Yen-Yin Lin, Ann-Shyn Chiang, and Shi-Wei Chu. Optical properties of adult drosophila brains in one-, two-, and three-photon microscopy. *Biomed. Opt. Express*, 10(4):1627–1637, Apr 2019.
- [43] Kwanghun Chung and Karl Deisseroth. Clarity for mapping the nervous system (vol 10, pg 508, 2013). *Nature Methods*, 10:1035–1035, 10 2013.
- [44] Markus Axer, Katrin Amunts, David Gräßel, Christoph Palm, Jürgen Dammers, Hubertus Axer, Uwe Pietrzyk, and Karl Zilles. A novel approach to the human connectome: Ultra-high resolution mapping of fiber tracts in the brain. *NeuroImage*, 54(2):1091 – 1101, 2011.
- [45] Markus Axer, David Gräßel, Melanie Kleiner, Jürgen Dammers, Timo Dickscheid, Julia Reckfort, Tim Hütz, Bjoern Eiben, Uwe Pietrzyk, Karl Zilles, and Katrin Amunts. High-resolution fiber tract reconstruction in the human brain by means of three-dimensional polarized light imaging. *Frontiers in neuroinformatics*, 5:34, 12 2011.
- [46] H. Wiese. *Enhancing the Signal Interpretation and Microscopical Hardware Concept of 3D Polarized Light Imaging*. PhD thesis, Bergische Universität Wuppertal, 2017.
- [47] Karl Zilles, Nicola Palomero-Gallagher, David Gräßel, Philipp Schlömer, Markus Cremer, Roger Woods, Katrin Amunts, and Markus Axer. *High-Resolution Fiber and Fiber Tract Imaging Using Polarized Light Microscopy in the Human, Monkey, Rat, and Mouse Brain*, pages 369–389. 12 2016.
- [48] Michael Zeineh, Nicola Palomero-Gallagher, Markus Axer, David Gräßel, Maged Goubran, Andreas Wree, Roger Woods, Katrin Amunts, and Karl Zilles. Direct visualization and mapping of the spatial course of fiber tracts at microscopic resolution in the human hippocampus. *Cerebral Cortex*, 27:bhw010, 02 2016.
- [49] Christina Herold, Philipp Schlömer, Isabelle Mafoppa-Fomat, Julia Mehlhorn, Katrin Amunts, and Markus Axer. The hippocampus of birds in a view of evolutionary connectomics. *Cortex*, 118:165 – 187, 2019. The Evolution of the Mind and the Brain.
- [50] Johannes Albrecht, Antonio Augusto Alves, Guilherme Amadio, Giuseppe Andronico, Nguyen Anh-Ky, Laurent Aphecetche, John Apostolakis, Makoto Asai, Luca Atzori, and et al. A roadmap for hep software and computing rd for the 2020s. *Computing and Software for Big Science*, 3(1), Mar 2019.
- [51] Yanxia Zhang and Yongheng Zhao. Astronomy in the big data era. *Data Science Journal*, 14:1–9, 05 2015.
- [52] Massimo Brescia and Giuseppe Longo. Astroinformatics, data mining and the future of astronomical research. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 720, 01 2012.
- [53] A. Siemiginowska, G. Eadie, I. Czekala, E. Feigelson, E. B. Ford, V. Kashyap, M. Kuhn, T. Loredo, M. Ntampaka, A. Stevens, A. Avelino, K. Borne, T. Bu-

- davari, B. Burkhart, J. Cisewski-Kehe, F. Civano, I. Chilingarian, D. A. van Dyk, G. Fabbiano, D. P. Finkbeiner, D. Foreman-Mackey, P. Freeman, A. Fruscione, A. A. Goodman, M. Graham, H. M. Guenther, J. Hakkila, L. Hernquist, D. Huppenkothen, D. J. James, C. Law, J. Lazio, T. Lee, M. López-Morales, A. A. Mahabal, K. Mandel, X. L. Meng, J. Moustakas, D. Muna, J. E. G. Peek, G. Richards, S. K. N. Portillo, J. Scargle, R. S. de Souza, J. S. Speagle, K. G. Stassun, D. C. Stenning, S. R. Taylor, G. R. Tremblay, V. Trimble, P. A. Yanamandra-Fisher, and C. A. Young. Astro2020 science white paper: The next decade of astroinformatics and astrostatistics, 2019.
- [54] Udo Toussaint. Bayesian inference in physics. *Reviews of Modern Physics - REV MOD PHYS*, 83:943–999, 09 2011.
 - [55] Sanjib Sharma. Markov chain monte carlo methods for bayesian data analysis in astronomy. *Annual Review of Astronomy and Astrophysics*, 55(1):213–259, 2017.
 - [56] Thomas Parr, Geraint Rees, and Karl J. Friston. Computational neuropsychology and bayesian inference. *Frontiers in Human Neuroscience*, 12:61, 2018.
 - [57] H J Motulsky and L A Ransnas. Fitting curves to data using nonlinear regression: a practical and nonmathematical review. *The FASEB Journal*, 1(5):365–374, 1987. PMID: 3315805.
 - [58] E. Hecht. *Optics*. Addison-Wesley, 4th edition, 1998.
 - [59] M. Menzel, K. Michielsen, H. De Raedt, J. Reckfort, K. Amunts, and M. Axer. A jones matrix formalism for simulating three-dimensional polarized light imaging of brain tissue. *Journal of The Royal Society Interface*, 12(111):20150734, 2015.
 - [60] Luiza Larsen, Lewis D. Griffin, David GRäkel, Otto W. Witte, and Hubertus Axer. Polarized light imaging of white matter architecture. *Microscopy Research and Technique*, 70(10):851–863.
 - [61] Miriam Menzel. *Finite-Difference Time-Domain Simulations Assisting to Reconstruct the Brain’s Nerve Fiber Architecture by 3D Polarized Light Imaging*. PhD thesis, 11 2018.
 - [62] Wikimedia Commons. Plaques of lambda phages on e. coli xll-blue mrf, 2007.
 - [63] Suzana Herculano-Houzel. The human brain in numbers: a linearly scaled-up primate brain. *Frontiers in Human Neuroscience*, 3:31, 2009.
 - [64] Contributors. In Larry R. Squire, Darwin Berg, Floyd E. Bloom, Sascha du Lac, Anirvan Ghosh, and Nicholas C. Spitzer, editors, *Fundamental Neuroscience (Fourth Edition)*, pages xxiii – xxiv. Academic Press, San Diego, fourth edition edition, 2013.
 - [65] G. F. Goethlin. Die doppelbrechenden eigenschaften des nervengewebes - ihre ursachen und ihre biologischen konsequenzen. *Kungl. Svenska Vetenskapsakad Handl*, 51, 1913.
 - [66] Xiang-Run Huang and Robert W. Knighton. Microtubules Contribute to the Birefringence of the Retinal Nerve Fiber Layer. *Investigative Ophthalmology Visual Science*, 46(12):4588–4593, 12 2005.
 - [67] FRANCIS O. SCHMITT and RICHARD S. BEAR. The ultrastructure of the nerve axon sheath. *Biological Reviews*, 14(1):27–50.

- [68] Richard S. Bear and Francis O. Schmitt. The optics of nerve myelin. *J. Opt. Soc. Am.*, 26(5):206–212, May 1936.
- [69] Priscilla Chinn and Francis O. Schmitt. On the birefringence of nerve sheaths as studied in cross sections. *Journal of Cellular and Comparative Physiology*, 9(2):289–296, 1937.
- [70] Nirmalya Ghosh and Alex I. Vitkin. Tissue polarimetry: concepts, challenges, applications, and outlook. *Journal of Biomedical Optics*, 16(11):1 – 30, 2011.
- [71] Daniel Wagner and Dieter Schmalstieg. Artoolkitplus for pose tracking on mobile devices, 2007.
- [72] A. Echalié, R. L. Glazer, V. Fülöp, and M. A. Geday. Assessing crystallization droplets using birefringence. *Acta Crystallographica Section D*, 60(4):696–702, Apr 2004.
- [73] A. Koehler. Ein neues beleuchtungsverfahren fuer mikrophotographische zwecke. *Zeitschrift fuer wissenschaftliche Mikroskopie und fuer mikroskopische Technik*, 1893.
- [74] R. Clark Jones. A new calculus for the treatment of optical systemsi. description and discussion of the calculus. *J. Opt. Soc. Am.*, 31(7):488–493, Jul 1941.
- [75] Julia Reckfort, Hendrik Wiese, Uwe Pietrzyk, Karl Zilles, Katrin Amunts, and Markus Axer. A multiscale approach for the reconstruction of the fiber architecture of the human brain based on 3d-pli. *Frontiers in Neuroanatomy*, 9:118, 2015.
- [76] Hans Müller. Memorandum on the polarization optics of the photo-elastic shutter. In *Report Number 2 of the OSRD Project OEMsr*, 1943.
- [77] George Gabriel Stokes. *On the Composition and Resolution of Streams of Polarized Light from different Sources*, volume 3 of *Cambridge Library Collection - Mathematics*, page 233–258. Cambridge University Press, 2009.
- [78] Julia Reckfort. *New Approaches to the Interpretation of 3D-Polarized Light Imaging Signals for an Advanced Extraction of Fiber Orientation*. PhD thesis, Bergische Universität Wuppertal, 2016.
- [79] Jürgen Dammers, Markus Axer, David Gräfel, Christoph Palm, Karl Zilles, Katrin Amunts, and Uwe Pietrzyk. Signal enhancement in polarized light imaging by means of independent component analysis. *NeuroImage*, 49(2):1241 – 1248, 2010.
- [80] Martin Schober, Philipp Schlömer, Markus Cremer, Hartmut Mohlberg, Anh-Minh Huynh, Nicole Schubert, Mehmet E. Kirlangic, and Katrin Amunts. Reference volume generation for subsequent 3d reconstruction of histological sections. In Heinz Handels, Thomas Martin Deserno, Hans-Peter Meinzer, and Thomas Tolxdorff, editors, *Bildverarbeitung für die Medizin 2015*, pages 143–148, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg.
- [81] Brian B. Avants, Nicholas J. Tustison, Michael Stauffer, Gang Song, Baohua Wu, and James C. Gee. The insight toolkit image registration framework. *Frontiers in Neuroinformatics*, 8:44, 2014.
- [82] S. Klein, M. Staring, K. Murphy, M. A. Viergever, and J. P. W. Pluim. elastix: A toolbox for intensity-based medical image registration. *IEEE Transactions on Medical Imaging*, 29(1):196–205, Jan 2010.

- [83] Marco Feuerstein, Tim Heibel, José Gardiazabal, Nassir Navab, and Martin Groher. Reconstruction of 3-d histology images by simultaneous deformable registration. volume 14, pages 582–589, 09 2011.
- [84] Daniel Schmitz, Sascha E. A. Muenzing, Martin Schober, Nicole Schubert, Martina Minnerop, Thomas Lippert, Katrin Amunts, and Markus Axer. Derivation of fiber orientations from oblique views through human brain sections in 3d-polarized light imaging. *Frontiers in Neuroanatomy*, 12:75, 2018.
- [85] D. G. Lowe. Object recognition from local scale-invariant features. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 2, pages 1150–1157 vol.2, Sep. 1999.
- [86] Melanie Kleiner, Markus Axer, David Gräßel, Julia Reckfort, Uwe Pietrzyk, Katrin Amunts, and Timo Dickscheid. Classification of ambiguous nerve fiber orientations in 3d polarized light imaging. In Nicholas Ayache, Hervé Delingette, Polina Golland, and Kensaku Mori, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2012*, pages 206–213, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [87] Abib O. Y. Alimi, Marco Pizzolato, Rutger H.J. Fick, and Rachid Deriche. Solving the Inclination Sign Ambiguity in Three Dimensional Polarized Light Imaging with a PDE-Based Method. In *ISBI 2017 – IEEE International Symposium on Biomedical Imaging*, Melbourne, Australia, April 2017.
- [88] Hendrik Wiese, David Gräßel, Uwe Pietrzyk, Katrin Amunts, and Markus Axer. Polarized light imaging of the human brain: A new approach to the data analysis of tilted sections. volume 9099, page 90990U, 05 2014.
- [89] T. Bayes. An essay towards solving a problem in the doctrine of chances. *Phil. Trans. of the Royal Soc. of London*, 53:370–418, 1763.
- [90] Andrew Gelman, Daniel Lee, and Jiqiang Guo. Stan: A probabilistic programming language for bayesian inference and optimization. *Journal of Educational and Behavioral Statistics*, 40(5):530–543, 2015.
- [91] John Salvatier, Thomas V. Wiecki, and Christopher Fonnesbeck. Probabilistic programming in python using PyMC3. *PeerJ Computer Science*, 2:e55, apr 2016.
- [92] Jessica Ai, Nimar S. Arora, ﻿Ning Dong, ﻿Beliz Gokkaya, Thomas Jiang, ﻿Anitha Kubendran, ﻿Arun Kumar, Michael Tingley, and ﻿Narjes Torabi. Hackppl: A universal probabilistic programming language. In *Proceedings of the 3rd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*, MAPL 2019, pages 20–28, New York, NY, USA, 2019. ACM.
- [93] Dustin Tran, Alp Kucukelbir, Adji B. Dieng, Maja Rudolph, Dawen Liang, and David M. Blei. Edward: A library for probabilistic modeling, inference, and criticism. *arXiv preprint arXiv:1610.09787*, 2016.
- [94] T. Minka et al. Infer.net. <https://github.com/dotnet/infer>, 2019.
- [95] Eli Bingham, Jonathan P. Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall, and Noah D. Goodman. Pyro: Deep universal probabilistic programming. *J. Mach. Learn. Res.*, 20(1):973–978, January 2019.

- [96] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.
- [97] Michael Waskom et al. mwaskom/seaborn: v0.8.1 (september 2017), September 2017.
- [98] David Hogg, Jo Bovy, and Dustin Lang. Data analysis recipes: Fitting a model to data. *ArXiv e-prints*, 08 2010.
- [99] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- [100] W Keith Hastings. Monte carlo sampling methods using markov chains and their application. *Biometrika*, 57, 04 1970.
- [101] Jonathan Goodman and Jonathan Weare. Ensemble samplers with affine invariance. *Commun. Appl. Math. Comput. Sci.*, 5(1):65–80, 2010.
- [102] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. emcee: The mcmc hammer. *PASP*, 125:306–312, 2013.
- [103] Dootika Vats, James M Flegal, and Galin L Jones. Multivariate output analysis for Markov chain Monte Carlo. *Biometrika*, 106(2):321–337, 04 2019.
- [104] Andrew Gelman, John B. Carlin, Hal S. Stern, and Donald B. Rubin. *Bayesian Data Analysis*. Chapman and Hall/CRC, 2nd ed. edition, 2004.
- [105] John K. Kruschke. *Doing Bayesian Data Analysis: A Tutorial with R and BUGS*. Academic Press, Inc., Orlando, FL, USA, 1st edition, 2010.
- [106] D. Schmitz, K. Amunts, T. Lippert, and M. Axer. A least squares approach for the reconstruction of nerve fiber orientations from tilttable specimen experiments in 3d-pli. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 132–135, April 2018.
- [107] Daniel Schmitz Markus Axer Jan-Oliver Kropp, Alexander Kobusch and Katrin Amunts. <https://www.nvidia.com/content/dam/en-zz/Solutions/gtc/conference-posters/nvidia-gtc19-dc-poster-resizing-web-1920x1607-JanOliver-Kropp.jpg>, 2019.
- [108] K. Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168, 1944.
- [109] Donald W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.
- [110] Stéfan van der Walt, S. Chris Colbert, and Gaël Varoquaux. The numpy array: A structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22–30, 2011.
- [111] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. [Online; accessed <today>].
- [112] J. J. Moré. The Levenberg-Marquardt algorithm: Implementation and theory. *Lecture Notes in Mathematics*, Berlin Springer Verlag, 630:105–116, 1978.

- [113] J J Moré, B S Garbow, and K E Hillstom. User guide for MINPACK-1. Technical Report ANL-80-74, Argonne Nat. Lab., Argonne, IL, Aug 1980.
- [114] Lisandro Dalcín, Rodrigo Paz, and Mario Storti. Mpi for python. *Journal of Parallel and Distributed Computing*, 65(9):1108 – 1115, 2005.
- [115] Francois Rheault, Jean-Christophe Houde, and Maxime Descoteaux. Visualization, interaction and tractometry: Dealing with millions of streamlines from diffusion mri tractography. *Frontiers in Neuroinformatics*, 11:42, 2017.
- [116] Davide Tiziano Di Carlo, Nicola Benedetto, Hugues Duffau, Federico Cagnazzo, Alessandro Weiss, Maura Castagna, Mirco Cosottini, and Paolo Perrini. Microsurgical anatomy of the sagittal stratum. *Acta Neurochirurgica*, 161(11):2319–2327, Nov 2019.
- [117] Lin-Ching Chang, Lindsay Walker, and Carlo Pierpaoli. Informed restore: A method for robust estimation of diffusion tensor from low redundancy datasets in the presence of physiological noise artifacts. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 68:1654–63, 11 2012.
- [118] Stefan Behnel, Robert Bradshaw, Craig Citro, Lisandro Dalcin, Dag Sverre Seljebotn, and Kurt Smith. Cython: The best of both worlds. *Computing in Science and Engg.*, 13(2):31–39, March 2011.
- [119] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.02688, May 2016.
- [120] Martín Abadi et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [121] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [122] A. J. Booker, J. E. Dennis, P. D. Frank, D. B. Serafini, V. Torczon, and M. W. Trosset. A rigorous framework for optimization of expensive functions by surrogates. *Structural optimization*, 17(1):1–13, Feb 1999.
- [123] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. Numba: A llvm-based python jit compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, LLVM ’15, pages 7:1–7:6, New York, NY, USA, 2015. ACM.
- [124] Chris Lattner and Vikram Adve. Llvm: A compilation framework for lifelong program analysis & transformation. In *Proceedings of the International Symposium on Code Generation and Optimization: Feedback-directed and Runtime Optimization*, CGO ’04, pages 75–, Washington, DC, USA, 2004. IEEE Computer Society.
- [125] John Owens, David Luebke, Naga Govindaraju, Mark Harris, Jens Krüger, Aaron Lefohn, and Timothy Purcell. A survey of general-purpose computation on graphics hardware. *Computer Graphics Forum*, 26:80 – 113, 03 2007.
- [126] James Fung and Steve Mann. Using multiple graphics cards as a general purpose parallel computer : Applications to computer vision. volume 1, pages 805–808, 01 2004.

- [127] Alberto Cano. A survey on graphic processing unit computing for large-scale data mining. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(1):e1232, 2018.
- [128] John Nickolls, Ian Buck, Michael Garland, and Kevin Skadron. Scalable parallel programming with cuda. *Queue*, 6(2):40–53, March 2008.
- [129] R.L. Harms, F.J. Fritz, A. Tobisch, R. Goebel, and A. Roebroek. Robust and fast nonlinear optimization of diffusion mri microstructure models. *NeuroImage*, 155:82 – 96, 2017.
- [130] Moises Hernandez-Fernandez, Istvan Reguly, Saad Jbabdi, Mike Giles, Stephen Smith, and Stamatis N. Sotiropoulos. Using gpus to accelerate computational diffusion mri: From microstructure estimation to tractography and connectomes. *NeuroImage*, 188:598 – 615, 2019.
- [131] David Smith, John C Gore, Thomas E Yankeelov, and E Welch. Real-time compressive sensing mri reconstruction using gpu computing and split bregman methods. *International journal of biomedical imaging*, 2012:864827, 02 2012.
- [132] Marc Bruce and Manish J Butte. Real-time gpu-based 3d deconvolution. *Optics express*, 21:4766–4773, 02 2013.
- [133] Zachary Devito, Michael Mara, Michael Zollhöfer, Gilbert Bernstein, Jonathan Ragan-Kelley, Christian Theobalt, Pat Hanrahan, Matthew Fisher, and Matthias Niessner. Opt: A domain specific language for non-linear least squares optimization in graphics and imaging. *ACM Trans. Graph.*, 36(5):171:1–171:27, October 2017.
- [134] Adrian Przybylski, Björn Thiel, Jan Keller, Bernd Stock, and Mark Bates. Gpufit: An open-source toolkit for gpu-accelerated curve fitting. *Scientific Reports*, 7, 12 2017.
- [135] The Mathworks, Inc., Natick, Massachusetts. *MATLAB version 9.3.0.713579 (R2017b)*, 2017.
- [136] Andreas Klöckner, Nicolas Pinto, Yunsup Lee, Bryan Catanzaro, Paul Ivanov, and Ahmed Fasih. Pycuda and pyopencl: A scripting-based approach to gpu run-time code generation. *Parallel Comput.*, 38(3):157–174, March 2012.
- [137] The HDF Group. Hierarchical data format version 5, 2000-2010.
- [138] John A. Nelder and Roger Mead. A simplex method for function minimization. *Computer Journal*, 7:308–313, 1965.
- [139] Fuchang Gao and Lixing Han. Implementing the nelder-mead simplex algorithm with adaptive parameters. *Computational Optimization and Applications*, 51(1):259–277, Jan 2012.
- [140] R. Byrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208, 1995.
- [141] Ciyou Zhu, Richard H. Byrd, Peihuang Lu, and Jorge Nocedal. Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Trans. Math. Softw.*, 23(4):550–560, December 1997.

- [142] José Luis Morales and Jorge Nocedal. Remark on “algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound constrained optimization”. *ACM Trans. Math. Softw.*, 38(1):7:1–7:4, December 2011.
- [143] Derek K. Jones. Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor mri. *Magnetic Resonance in Medicine*, 49(1):7–12, 2003.
- [144] Brandon Whitcher, David S. Tuch, Jonathan J. Wisco, A. Gregory Sorensen, and Liqun Wang. Using the wild bootstrap to quantify uncertainty in diffusion tensor imaging. *Human Brain Mapping*, 29(3):346–362, 2008.
- [145] Pew-Thian Yap, Hongyu An, Yasheng Chen, and Dinggang Shen. The non-local bootstrap – estimation of uncertainty in diffusion mri. In James C. Gee, Sarang Joshi, Kilian M. Pohl, William M. Wells, and Lilla Zöllei, editors, *Information Processing in Medical Imaging*, pages 390–401, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [146] Jennifer S. W. Campbell, Parya MomayyezSiahkal, Peter Savadjiev, Ilana R. Lepert, Kaleem Siddiqi, and G. Bruce Pike. Beyond crossing fibers: Bootstrap probabilistic tractography using complex subvoxel fiber geometries. *Frontiers in Neurology*, 5:216, 2014.
- [147] Maggie Chow, Sharon Wu, Sarah Webb, Katie Gluskin, and David Yew. Functional magnetic resonance imaging and the brain: A brief review. *World Journal of Radiology*, 9:5, 01 2017.
- [148] Bayesian analysis of neuroimaging data in fsl. *NeuroImage*, 45(1, Supplement 1):S173 – S186, 2009. Mathematics in Brain Imaging.
- [149] Patrick Derichs. Optimierung des datenflusses polarisationsmikroskopischer messungen mit automatisierter qualitätssicherung. Master’s thesis, FH Aachen, 2018.
- [150] N. I. Fisher, T. Lewis, and B. J. J. Embleton. *Statistical analysis of spherical data*. 1987.
- [151] Dorian Krause and Philipp Thörnig. Jureca: General-purpose supercomputer at jülich supercomputing centre. *Journal of large-scale research facilities JLSRF*, 2, 03 2016.
- [152] College Station TX USA Kristen M. Thyng | Department of Oceanography, Texas AM University, TX USA Chad A. Greene | The University of Texas at Austin Institute for Geophysics, Austin, College Station TX USA Robert D. Hetland | Department of Oceanography, Texas AM University, College Station TX USA Heather M. Zimmerle | Department of Oceanography, Texas AM University, Steven F. DiMarco | Department of Oceanography, Geochemical, and College Station TX USA Environmental Research Group, Texas AM University. True colors of oceanography: Guidelines for effective and accurate colormap selection. *Oceanography*, 29, September 2016.
- [153] Daniel Foreman-Mackey. corner.py: Scatterplot matrices in python. *The Journal of Open Source Software*, 24, 2016.
- [154] Peter TC So. *Two-photon Fluorescence Light Microscopy*. American Cancer Society, 2001.

- [155] Irene Costantini, Miriam Menzel, Ludovico Silvestri, Nicole Schubert, Markus Axer, Katrin Amunts, and Francesco Pavone. Polarized light imaging and two-photon fluorescence microscopy correlative approach for 3d reconstruction of the orientation of myelinated fibers. 04 2017.
- [156] Kevin H. Knuth. Optimal data-based binning for histograms and histogram-based probability density models. *Digital Signal Processing*, 95:102581, 2019.
- [157] Astropy Collaboration, A. M. Price-Whelan, B. M. Sipőcz, H. M. Günther, P. L. Lim, S. M. Crawford, S. Conseil, D. L. Shupe, M. W. Craig, N. Dencheva, A. Ginsburg, J. T. VanderPlas, L. D. Bradley, D. Pérez-Suárez, M. de Val-Borro, T. L. Aldcroft, K. L. Cruz, T. P. Robitaille, E. J. Tollerud, C. Ardelean, T. Babej, Y. P. Bach, M. Bachetti, A. V. Bakanov, S. P. Bamford, G. Barentsen, P. Barmby, A. Baumbach, K. L. Berry, F. Biscani, M. Boquien, K. A. Bostroem, L. G. Bouma, G. B. Brammer, E. M. Bray, H. Breytenbach, H. Buddelmeijer, D. J. Burke, G. Calderone, J. L. Cano Rodríguez, M. Cara, J. V. M. Cardoso, S. Cheedella, Y. Copin, L. Corrales, D. Crichton, D. D’Avella, C. Deil, É. Depagne, J. P. Dietrich, A. Donath, M. Droettboom, N. Earl, T. Erben, S. Fabbro, L. A. Ferreira, T. Finethy, R. T. Fox, L. H. Garrison, S. L. J. Gibbons, D. A. Goldstein, R. Gommers, J. P. Greco, P. Greenfield, A. M. Groener, F. Grollier, A. Hagen, P. Hirst, D. Homeier, A. J. Horton, G. Hosseinzadeh, L. Hu, J. S. Hunkeler, Ž. Ivezić, A. Jain, T. Jenness, G. Kanarek, S. Kendrew, N. S. Kern, W. E. Kerzendorf, A. Khvalko, J. King, D. Kirkby, A. M. Kulkarni, A. Kumar, A. Lee, D. Lenz, S. P. Littlefair, Z. Ma, D. M. Macleod, M. Mastropietro, C. McCully, S. Montagnac, B. M. Morris, M. Mueller, S. J. Mumford, D. Muna, N. A. Murphy, S. Nelson, G. H. Nguyen, J. P. Ninan, M. Nöthe, S. Ogaz, S. Oh, J. K. Parejko, N. Parley, S. Pascual, R. Patil, A. A. Patil, A. L. Plunkett, J. X. Prochaska, T. Rastogi, V. Reddy Janga, J. Sabater, P. Sakurikar, M. Seifert, L. E. Sherbert, H. Sherwood-Taylor, A. Y. Shih, J. Sick, M. T. Silbiger, S. Singanamalla, L. P. Singer, P. H. Sladen, K. A. Sooley, S. Sornarajah, O. Streicher, P. Teuben, S. W. Thomas, G. R. Tremblay, J. E. H. Turner, V. Terrón, M. H. van Kerkwijk, A. de la Vega, L. L. Watkins, B. A. Weaver, J. B. Whitmore, J. Woillez, V. Zabalza, and Astropy Contributors. The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package. , 156(3):123, Sep 2018.
- [158] Vincent Arsigny, Olivier Commowick, Nicholas Ayache, and Xavier Pennec. A fast and log-euclidean polyaffine framework for locally linear registration. *J. Math. Imaging Vis.*, 33(2):222–238, February 2009.
- [159] Vincent Arsigny, Xavier Pennec, and Nicholas Ayache. Polyrigid and polyaffine transformations: A new class of diffeomorphisms for locally rigid or affine registration. In Randy E. Ellis and Terry M. Peters, editors, *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2003*, pages 829–837, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.
- [160] George E. P. Box. Science and statistics. *Journal of the American Statistical Association*, 71(356):791–799, 1976.
- [161] Harald Cramér. *Mathematical Methods of Statistics*. Princeton University Press, 1946.
- [162] C. Radhakrishna Rao. *Information and the Accuracy Attainable in the Estimation of Statistical Parameters*, pages 235–247. Springer New York, New York, NY, 1992.

- [163] Nicolas Boumal, Bamdev Mishra, P.-A. Absil, and Rodolphe Sepulchre. Manopt, a matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15:1455–1459, 2014.
- [164] James Townsend, Niklas Koep, and Sebastian Weichwald. Pymanopt: A python toolbox for optimization on manifolds using automatic differentiation. *Journal of Machine Learning Research*, 17(137):1–5, 2016.
- [165] Mark Girolami and Ben Calderhead. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.
- [166] Shun-Ichi Amari. Natural gradient works efficiently in learning. *Neural Comput.*, 10(2):251–276, February 1998.
- [167] Matthew D. Hoffman and Andrew Gelman. The no-u-turn sampler: Adaptively setting path lengths in hamiltonian monte carlo. *Journal of Machine Learning Research*, 15:1593–1623, 2014.
- [168] M. J. Betancourt. Generalizing the no-u-turn sampler to riemannian manifolds, 2013.
- [169] John T. Kent. The fisher-bingham distribution on the sphere. *Journal of the Royal Statistical Society. Series B (Methodological)*, 44(1):71–80, 1982.
- [170] David E. Tyler. Statistical analysis for the angular central gaussian distribution on the sphere. *Biometrika*, 74(3):579–589, 1987.
- [171] Pierre Fillard, Cyril Poupon, and Jean-François Mangin. A novel global tractography algorithm based on an adaptive spin glass model. In Guang-Zhong Yang, David Hawkes, Daniel Rueckert, Alison Noble, and Chris Taylor, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2009*, pages 927–934, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
- [172] Adam Thompson and Chris J. Newburn. Gpudirect storage: A direct path between storage and gpu memory. <https://devblogs.nvidia.com/gpudirect-storage/>, 2019.
- [173] Anton Bardera, Miquel Feixas, and Imma Boada. Normalized similarity measures for medical image registration. *Proc SPIE*, 5370, 05 2004.
- [174] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever. Mutual-information-based registration of medical images: a survey. *IEEE Transactions on Medical Imaging*, 22(8):986–1004, Aug 2003.
- [175] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes. Nonrigid registration using free-form deformations: application to breast mr images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, Aug 1999.
- [176] M. Huysegoms. Using multiple gpus to reconstruct the brain from histological images. NVIDIA GPU Technology Conference, 2015.
- [177] J. N. Sarvaiya, S. Patnaik, and S. Bombaywala. Image registration by template matching using normalized cross-correlation. In *2009 International Conference on Advances in Computing, Control, and Telecommunication Technologies*, pages 819–822, Dec 2009.

- [178] Brian B. Avants, Nicholas J. Tustison, Gang Song, Philip A. Cook, Arno Klein, and James C. Gee. A reproducible evaluation of ants similarity metric performance in brain image registration. *NeuroImage*, 54(3):2033 – 2044, 2011.
- [179] Utkarsh Ayachit. *The ParaView Guide: A Parallel Visualization Application*. Kitware, Inc., USA, 2015.
- [180] Hanchuan Peng, Zongcai Ruan, Fuhui Long, Julie Simpson, and Eugene Myers. V3d enables real-time 3d visualization and quantitative analysis of large-scale biological image data sets. *Nature biotechnology*, 28:348–53, 03 2010.
- [181] Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, New York, NY, USA, second edition, 2006.