

INAUGURAL-DISSERTATION

zur Erlangung der Doktorwürde der

NATURWISSENSCHAFTLICH-MATHEMATISCHEN
GESAMTFAKULTÄT

der

RUPRECHT-KARLS-UNIVERSITÄT
HEIDELBERG

vorgelegt von
Diplom-Mathematiker

Jan Hering

aus Plzeň (Pilsen, Tschechische Republik)

Tag der mündlichen Prüfung: _____

Robust Motion and Distortion Correction of Diffusion-Weighted MR Images

Gutachter: **Prof. Dr. rer. nat. Bernd Jähne**

Abstract

Effective image-based correction of motion and other acquisition artifacts became an essential step in diffusion-weighted Magnetic Resonance Imaging (MRI) analysis as the micro-structural tissue analysis advances towards higher-order models. These come with increasing demands on the number of acquired images and the diffusion strength (b -value) yielding lower signal-to-noise ratios (SNR) and a higher susceptibility to artifacts. These conditions, however, render the current image-based correction schemes, which act retrospectively on the acquired images through pairwise registration, more and more ineffective. Following the hypothesis, that a more consequent exploitation of the different intensity relationships between the volumes would reduce registration outliers, a novel correction scheme based on memetic search is proposed. This scheme allows for incorporating all single image metrics into a multi-objective optimization approach. To allow a quantitative evaluation of registration precision, realistic synthetic data are constructed by extending a diffusion MRI simulation framework by motion and eddy-currents-caused artifacts. The increased robustness and efficacy of the multi-objective registration method is demonstrated on the synthetic as well as in-vivo datasets at different levels of motion and other acquisition artifacts. In contrast to the state-of-the-art methods, the average target registration error (TRE) remained below the single voxel size also at high b -values ($3000 \text{ s} \cdot \text{mm}^{-2}$) and low signal-to-noise ratio in the moderately artifacted datasets. In the more severely artifacted data, the multi-objective method was able to eliminate most of the registration outliers of the state-of-the-art methods, yielding an average TRE below the double voxel size. In the in-vivo data, the increased precision manifested itself in the scalar measures as well as the fiber orientation derived from the higher-order Neurite Orientation Dispersion and Density Imaging (NODDI) model. For the neuronal fiber tracts reconstructed on the data after correction, the proposed method most closely resembled the ground-truth. The proposed multi-objective method has not only impact on the evaluation of higher-order diffusion models as well as fiber tractography and connectomics, but could also find application to challenging image registration problems in general.

Zusammenfassung

Effektive bildbasierte Korrektur von Bewegungs- und Aufnahmeartefakten ist durch die Weiterentwicklung der mikrostrukturellen Gewebemodelle zu einem notwendigen Vorverarbeitungsschritt in der Analyse von Bilddaten der diffusionsgewichteten Magnetresonanztomographie (MRT) geworden. Für eine robuste Auswertung benötigen diese Modelle höherer Ordnung stets mehr Aufnahmen bei einer höheren Diffusionswichtung (dem b -Wert), die wiederum anfälliger für Aufnahmeartefakte sind und mit einem schlechteren signal-to-noise Verhältnis (SNR) einhergehen. Für die aktuellen Korrekturverfahren, welche mittels von paarweiser Bildregistrierung arbeiten, stellen die dadurch verschlechterten Kontrastverhältnisse eine Herausforderung dar, die bis zur ineffektiven Korrektur führt. Ausgehend von der Hypothese, dass eine konsequente Ausnutzung der unterschiedlichen Signalähnlichkeiten zwischen den Bildvolumina eines diffusionsgewichteten Bildes zur Reduktion von Ausreißern in der Bildregistrierung führen kann, wurde in dieser Arbeit ein neues Korrekturverfahren vorgestellt, das auf dem memetischen Optimierungsprinzip aufbaut und so die einzelnen Kostenfunktionen (objectives) in einem multi-objektiven Ansatz bündelt. Für eine quantitative Auswertung der Registrierungspräzision wurden realistische synthetische Datensätze mit simulierten Bewegungs- sowie Aufnahmeartefakte konstruiert. Die durch das multi-objektive Verfahren erreichte Effizienz und Robustheit wird auf synthetischen und in-vivo Datensätzen mit je unterschiedlicher Ausprägung der Bewegungs- und Aufnahmeartefakten demonstriert. Bei Datensätzen mit mittlerer Artefaktintensität blieb der durchschnittliche target registration error (TRE), im Gegensatz zu state-of-the-art Verfahren, unterhalb der Voxelgröße auch für höhere b -Werte ($3000 \text{ s} \cdot \text{mm}^{-2}$) und niedrigeres SNR. Bei höheren Artefaktintensitäten konnten durch den multi-objektiven Ansatz fast alle Ausreißer der state-of-the-art Methoden bei einem TRE unterhalb der doppelten Voxelgröße eliminiert werden. Für die in-vivo Daten wurde die gewonnene Präzision anhand der Skalargrößen sowie der Hauptfaserrichtung aus dem Neurite Orientation Dispersion and Density Imaging (NODDI) Modell demonstriert. Bei der Rekonstruktion der Nervenfasern hat die vorgestellte Methode die ground-truth am besten wiedergegeben. Neben den Auswirkungen auf Signalmodelle höherer Ordnung in der diffusionsgewichteten MRT, auf Nervenbahnenrekonstruktion und Connectomics könnte die multi-objektive Methode auch bei weiteren herausfordernden Problemen der Bildregistrierung im Allgemeinen einen Vorteil bieten.

Contents

List of Figures	v
List of Abbreviations	vii
1 Introduction	1
1.1 Motivation	2
1.2 Objectives	3
1.3 Outline	4
2 Diffusion-weighted MRI	5
2.1 MR Acquisition Physics	5
2.1.1 MR Imaging	6
2.1.2 Diffusion-weighted Imaging	9
2.1.3 Signal Modeling	11
2.2 Diffusion-weighted Images	14
2.3 Data Processing and Applications	17
2.3.1 Estimating Tissue Parameters	17
2.3.2 Tractography	19
2.3.3 Connectomics	20
2.4 Artifacts in Diffusion MR	21
2.4.1 Patient Motion	21
2.4.2 Eddy Currents	23
2.4.3 Higher Order Distortions	24
2.4.4 Conclusion	25
3 State of the Art	27
3.1 DW MRI Distortion Correction	27
3.1.1 Rejection Approaches	28
3.1.2 Prospective Motion Correction	28
3.1.3 Retrospective Motion Correction	29
3.1.4 Discussion	33
3.2 Image Registration	34
3.2.1 Image Metrics	35
3.2.2 Registration Optimizer	37
3.2.3 Discussion	39

3.3	Evolutionary Optimization	40
3.3.1	Particle Swarm Optimization	41
3.3.2	Multi-objective Approaches	42
3.3.3	Discussion	43
3.4	Memetic Optimization Approaches	44
3.5	Conclusion	45
4	Methods	47
4.1	Retrospective DWI Distortion Correction	48
4.1.1	Objective Functions	48
4.1.2	Optimization Scheme	49
4.1.3	Pyramid Image Registration Scheme	52
4.1.4	Information Flow and Parametrization	54
4.2	Construction of Validation Data	57
4.2.1	Synthetic Validation Data	58
4.2.2	In-Vivo Validation Data	61
4.3	Evaluation Metrics	62
4.3.1	Objective's Behavior	62
4.3.2	Qualitative Evaluation	63
4.4	Evaluation on Synthetic Data	63
4.4.1	Precision	64
4.4.2	Particle History	64
4.4.3	Initial Position Variance	64
4.4.4	Fiber Tracking	65
4.5	Evaluation on In-vivo Data	65
4.5.1	Tissue Parameter Estimation	65
4.5.2	Particle History	66
4.5.3	Fiber Tracking	66
5	Results	67
5.1	Evaluation on Synthetic Data	67
5.1.1	Synthetic Validation Data	67
5.1.2	Precision	70
5.1.3	Particle History	73
5.1.4	Runtime	74
5.1.5	Fiber Tracking	74
5.2	Evaluation on In-vivo Data	76
5.2.1	In-vivo Validation Data	76
5.2.2	Tissue Parameter Estimation	80
5.2.3	Particle History	83
5.2.4	Fiber Tracking	83
6	Discussion	87
6.1	Evaluation Data	88
6.2	Synthetic Experiments	89

6.3	In-vivo Experiments	90
6.4	Local Search Method	91
6.5	Global Optimization Method	94
6.6	Relevance for Diffusion MRI Analysis	96
6.7	Outlook and Further Work	96
7	Conclusion	99
A	Own Publications	103
	Acknowledgements	105
	Bibliography	107

List of Figures

2.1	Magnetization of spins	6
2.2	Schematic depiction of the spin echo signal	7
2.3	Sequence diagram of a typical diffusion MRI EPI acquisition	10
2.4	Illustration of diffusion MRI signal models	12
2.5	Scheme of a multi-shell diffusion-weighted image	15
2.6	Exemplary axial slices from a diffusion-weighted image	16
2.7	Exemplary axial slices of DTI scalar indices maps	18
2.8	Principle of diffusion MR fiber reconstruction	19
2.9	The connectomics principle	20
2.10	Diffusion-weighted image distortions by head motion	22
2.11	Effects of head motion on measured diffusion gradient	23
2.12	Eddy currents in a diffusion MRI EPI sequence	23
2.13	Eddy-currents-caused distortions in diffusion-weighted images	24
3.1	Unweighted reference correction scheme	30
3.2	Extrapolation method correction scheme	31
3.3	Similarity-based correction scheme	33
3.4	Visualization of the Mutual Information metric in ideal case	36
3.5	Image pyramid for multi-resolution registration	38
4.1	Visualization of the fixed image assignment for the objective functions	50
4.2	Schematic depiction of the memetic search scheme	51
4.3	Exemplary particle trace during memetic registration	63
5.1	Exemplary axial slices from the simulated synthetic datasets	68
5.2	Ground-truth estimation with reference registration	69
5.3	Target Registration Error for synthetic data	71
5.4	Target Registration Error for different initial position variance choices	72
5.5	Aggregated history of particles for synthetic experiments	73
5.6	The fornix tracts reconstructed from the simulated dataset	75
5.7	The principle of visual registration validation	77
5.8	Example of the in-vivo validation data	78
5.9	Illustration of the alignment of an in-vivo dataset	78
5.10	Mean error of the corrected raw data at higher b-shell	79
5.11	Visualization of the error in the NODDI Index ICVF	80
5.12	Visualization of the error in the NODDI Index ODI	81

5.13	Reference-based evaluation of the correction schemes on basis of NODDI	82
5.14	Aggregated history of particles for in-vivo experiments	83
5.15	The fornix tracts reconstructed from an in-vivo dataset	84
5.16	The cingulum tracts reconstructed from an in-vivo dataset	85
6.1	Preliminary results illustrating capabilities of different local search methods	92
6.2	Preliminary results of ASGD-based local search method	93
6.3	Preliminary results of the correction performance of a weak local search	95

List of Abbreviations

AD	Axial Diffusivity
ADC	Apparent Diffusion Coefficient
ASGD	Adaptive Stochastic Gradient Descent
CHARMED	Composite Hindered and Restricted Model of Diffusion
CSD	Constrained Spherical Deconvolution
CSF	Cerebro-spinal Fluid
CT	Computed Tomography
dMRI	diffusion MRI
DOF	Degree of Freedom
DTI	Diffusion Tensor Imaging
DWI	diffusion-weighted Image
EA	Evolutionary Algorithm
EPI	Echo-planar Imaging
FA	Fractional Anisotropy
FOV	Field of View
FV	Fitness Value
GA	Genetic Algorithm
HARDI	High Angular Resolution Diffusion Imaging
ICVF	Intra-cellular Volume Fraction
IR	Image Registration
ITK	Insight ToolKit
MB	Model-based reference objective
MD	Mean Diffusivity
MH	meta-heuristic
MI	Mutual Information
MITK	Medical Imaging and Interaction Toolkit
MO	multi-objective
MR	Magnetic Resonance

MRI	Magnetic Resonance Imaging
nMI	normalized Mutual Information
NMR	Nuclear Magnetic Resonance
NN	Nearest-neighbor reference objective
NODDI	Neurite Orientation Dispersion and Density Imaging
NW	Non-weighted reference objective
ODF	Orientation Distribution Function
ODI	Orientation Dispersion Index
PSO	Particle Swarm Optimization
RD	Radial Diffusivity
RF	Radio Frequency
ROI	Region of Interest
RSGD	Regular Step Gradient Descent
SD	Spherical Deconvolution
SH	Spherical Harmonics
SNR	Signal-to-Noise Ratio
SPECT	Single Photon Emission Computed Tomography
TE	Echo Time
TR	Repetition Time
TRE	Target Registration Error
VEGA	Vector-evaluated Genetic Algorithm
VEPSO	Vector-evaluated Particle Swarm Optimization

Chapter 1

Introduction

Diffusion-weighted Magnetic Resonance Imaging (MRI) is an image acquisition technique that allows the mapping of molecular diffusion in biological tissues, reflecting the interaction with obstacles such as myelin sheaths or cell membranes. As patterns of water molecule diffusion can reveal micro-structural tissue properties, this imaging method offers a unique non-invasive insight into the human brain.

With this acquisition technique, an examination of the micro-structural changes of the tissue that can occur as a result of pathological changes due to disease conditions or naturally because of healthy aging is made possible. This promotes the diffusion-weighted MRI to a valuable tool in the search for *biomarkers*, in other words for providing a set of image-based values that allow for a separation of diseased and healthy subjects. It also enables reconstruction of the neuronal pathways (tractography), which offer further insights in multiple facets. Next to application as navigation aid for neurosurgery and tissue classification, the reconstructed pathways also reveal the connections between different anatomical as well as functional areas of the brain and thus enable a global evaluation of the *wiring* (connectomics).

To achieve the mentioned, a typical diffusion MRI (dMRI) acquisition consists of a set of volumes (3D), each of them providing the measure of water diffusion in a pre-defined direction (the gradient direction) and strength (the b -value) and of at least one additional baseline volume acquired without gradient direction and with $b=0 \text{ s} \cdot \text{mm}^{-2}$. The total number of acquired volumes is usually between 30 and 60. To handle this high-dimensional data, a signal model describing the measured anisotropy in each imaging voxel is employed. The most common measures used for micro-structural analysis are derived from plain Diffusion Tensor Imaging (DTI). DTI, however, does not allow a differentiation of multiple tissue compartments and/or fiber directions within a

voxel. To overcome this limitation, higher order models were formulated in more recent methodological developments, which allow the estimation of diffusivity measures for individual tissue compartments within a voxel and thus provide more precise insights for the analysis and understanding of brain development and diseases [82, 37, 85, 119]. Increasingly high angular resolution is provided by recent reconstruction schemes on basis of multiple b -shells and High Angular Resolution Diffusion Imaging (HARDI) acquisitions [1, 51]. These advances in signal modeling directly translate to tractography and, consequently, connectomics [44], that are both known to benefit from HARDI protocols at higher b -values [35, 20].

1.1 Motivation

The high amount of acquired volumes leads inevitably to longer acquisition time, and thus increased difficulty for the scanned subject to hold still during the whole acquisition. As a result of the subject's motion, the individual gradient volumes are not aligned to each other. In addition, the implementation of a dMRI acquisition sequence cannot avoid a further data distortion which occurs mostly as a consequence of the emerging additional magnetic fields, the eddy currents. Robust and successful head motion and artifact correction is thus a critical prerequisite for all the mentioned analysis techniques. The expected sensitivity and specificity with regard to micro-structural effects can only be achieved if the datasets are exactly aligned to each other in the imaging space. Outlier rejection of motion-corrupted diffusion imaging acquisitions [17] can increase the reconstruction quality, especially when combined with optimized gradient schemes for partial datasets [31]. However, the exclusion of motion-corrupted directions decreases the ability of resolving micro-structural features of the tissue [32]. Prospective correction schemes are effective, but negatively impact acquisition time [15, 63]. Therefore, retrospective correction schemes based on image registration are widely adopted in the community nowadays [72, 94]. The pairwise registration in the retrospective approaches is applied to each volume in order to align it with a selected baseline volume of the acquisition.

The role of an effective image-based correction scheme becomes even more important for higher-order models that include larger numbers of free parameters and a generally higher sensitivity to noise. Unfortunately, the retrospective correction of such datasets is not only becoming increasingly important here, but also increasingly difficult to achieve. First, the increasing demands related to q -space sampling (i.e. higher number of b -values and gradient directions) lead to increased acquisition times and thus potentially

increased head motion. Second, acquisitions at higher b -values introduce stronger eddy-currents-caused artifacts and additional challenges as the signal intensity and the signal-to-noise ratio drop proportionally. This leads to larger contrast deviations when comparing the weighted gradient acquisition with the unweighted reference signal. Altogether, the correction scheme using pairwise registration of diffusion-weighted images to a non-weighted ($b=0 \text{ s} \cdot \text{mm}^{-2}$) reference image [72], can easily produce a remarkable amount of outliers in such settings.

In conclusion, effective and robust image-based artifact correction is an essential step in the analysis of diffusion MR images which must be carried out with particular care.

1.2 Objectives

The central objective of this work is to improve the retrospective image-based artifact and motion correction in diffusion-weighted MRI to provide a reliable basis for any further processing and modeling steps. The main hypothesis is that a more consequent exploration of the different intensity relationships between the individual gradient volumes could enhance the general capabilities of solving challenging registration problems that occur in artifact correction schemes of diffusion-weighted images. An optimal method would be able to prevent outliers and local minima by choosing the right combination of different registration metrics and efficiently lead the way towards the global optimum.

The progress towards the main objective involves several additional stages. Firstly, to allow quantitative evaluation of the precision of the constructed method itself and against the state-of-the-art retrospective correction methods, datasets with expressed motion and eddy-currents-caused artifacts as well as with known ground-truth transforms must be constructed. As this cannot be achieved for in-vivo acquisitions, a realistic simulation of the data and the artifacts is necessary. Secondly, an evaluation of the impact of the different correction methods on a higher-order signal model should be conducted on in-vivo datasets at different levels of artifact expression. Finally, the impact on further processing steps (i.e. fiber tractography) in dMRI analysis should be examined.

1.3 Outline

Chapter 2 provides the necessary background on the acquisition physics, image reconstruction and the data analysis pipelines in diffusion-weighted MRI. A special focus is dedicated to the formation of acquisition artifacts and their consequences on the reconstructed images.

Chapter 3 first summarizes the current state-of-the-art in motion and artifact correction techniques in diffusion MRI, then introduces the necessary definitions and current approaches in (medical) image registration and evolutionary optimization with particular focus on multi-objective optimization methods. Lastly, the concept of hybrid optimization—the memetic algorithms—as well as its application to (image) registration problems are described.

The main contributions of this work are described in Chapter 4. Here, a novel, multi-objective image registration approach for the retrospective motion and artifact correction of diffusion-weighted images is presented. Further, the construction of both synthetic and in-vivo validation datasets for the method’s evaluation are introduced. Finally, the experimental settings for both quantitative and qualitative evaluation on the constructed validation data are described.

The results of these experiments are presented in Chapter 5 and discussed in Chapter 6, where also the presented method is analyzed with regard to the experiments in further detail, including its impact and relevance in diffusion-MRI analysis as well as future courses of investigations.

Chapter 2

Diffusion-weighted MRI

In this chapter, the main principles of structural MRI and its extension to diffusion-weighted MRI are presented. In Nuclear Magnetic Resonance (NMR), the signal from atomic nuclei is used for image signal generation and, if no contrast agent is present, the main signal in a human body is due to hydrogen nuclei (^1H). In addition, diffusion MR focuses on the microscopic movement of the nuclei over time which provides an insight into structural properties of different tissues based on the restriction of the nuclei's movement.

For the thesis it is essential to understand the basic physics principles of MR and diffusion-weighted MR along with signal generation (Section 2.1) and signal modeling (Section 2.1.3). The concept of diffusion-weighted images, which is central to this work, is explained in Section 2.2 followed by a summary of common data processing pipelines alongside the main application areas in Section 2.3. The most important aspect for the thesis's main objective—the acquisition artifacts and distortions—are explained in Section 2.4 at the end of this chapter.

2.1 MR Acquisition Physics

In contrast to X-ray and Computed Tomography (CT) imaging, where the absorption of ionizing radiation is used to differentiate between tissues, MRI offers a non-destructive alternative. It is based on the spin, a property of elementary particles, so as for the proton of ^1H nucleus. The proton rotates along its axis and, as a rotating mass with an electric charge that has a *magnetic momentum* (B), behaves like a small magnet. Hence,

the protons get affected by an externally applied magnetic field and, when moving, induce measurable voltage in a receiver coil.

2.1.1 MR Imaging

An MRI scanner operates with multiple magnetic fields: first a strong, static field B_0 and varying Radio Frequency (RF) pulses to achieve excitation from the steady orientation parallel to the static field and second a collection of additional magnetic gradients to allow spatial encoding of the measured signal. After a sufficient long exposure to the external B_0 field, the spins of hydrogen nuclei align with the direction of B_0 and a longitudinal magnetization M_z builds up (see Figure 2.1). Meanwhile, the magnetic moments undergo a precession around the main field and their angular frequency (*Larmor* or *precession frequency* ω_0) is proportional to the field strength and is given by the *Larmor equation*:

$$\omega_0 = \gamma_0 \cdot B_0,$$

with γ_0 the constant, nuclei-specific gyromagnetic ratio and B_0 the strength of the main field in Tesla [T]. A RF pulse with a frequency matching the resonance frequency ω_0 applied for a sufficient time flips the longitudinal magnetization by exactly 90° into the transverse magnetization M_{xy} . At this state, the rotating spins act as an electrical generator inducing an alternating voltage in a receiver coil - the MR signal.




Figure 2.1: Magnetization of spins. (a) With no external field, spins rotate around their axis pointing in random directions. (b) In the presence of an external static field B_0 , the spins align along the field and produce a longitudinal magnetization M_z .

With time, the magnetic moments gradually realign with the z -axis and the transverse magnetization M_{xy} within the xy -plane decreases slowly. The average time for reaching the equilibrium is designated T_1 and the realigning process is called T_1 or *longitudinal relaxation*. The time scale of T_1 relaxation for biological tissues is typically in the range of 1 second at 1.5T field strength [114]. In parallel with the decay of the longitudinal

relaxation another relaxation process takes place. Immediately after the RF pulse, the spins precess synchronously and are said to be *in phase*. At this state, they also yield maximal signal in the receiving coil. However, mainly due to local field inhomogeneities, some spins advance faster than others which leads to a spin decoherence (also called *dephasing*). Dephasing occurs at a time constant of T_2 and is (more or less) independent of the strength of the external field B_0 . An even faster spin dephasing caused by the effects of the static magnetic field inhomogeneities (T_2^* relaxation) can be eliminated by applying a 180° refocusing RF pulse. Such pulse applied at time τ causes a phase reversal and the initially faster precessing spins are placed behind the slower ones and at time of 2τ , all spins are in phase again (c.f. Figure 2.2). The signal induced in the receiver coil after the phase coherence is restored is known as *spin echo*.




Figure 2.2: Schematic depiction of the spin echo signal. (a) The magnetization along B_0 is flipped onto the xy -plane by a 90° RF pulse and the spins remain in phase. (b) The dephasing leads to loss of coherence among the spins. (c) The phase is reversed by an 180° RF pulse and the continuing dephasing alters into re-phasing until (d) the phase coherence is restored. t marks the time of the acquisition relative to a starting time t_0 and the pulse time τ .

Image Contrast

The different times of T_1 and T_2 relaxation are an intrinsic feature of biological tissues and the imaging contrast differs dependent on how these tissue relaxation properties are emphasized in the MR acquisition. One of the acquisition parameters is the *repetition time* (TR), the interval between two successive RF excitations of the same slice. With $TR > T_1$, most of the excited spins rotate back into the z -plane and can be fully excited with the next pulse. With *short* repetition times, the quickly relaxing tissue parts (i.e. with short T_1) will produce a larger signal after the next RF pulse and appear brighter on the image. The tissues with long T_1 recover only partly between the pulses and the next excitation produces only lesser magnetization and they appear darker in the image. Since the resulting image contrast contains mainly the T_1 information,

acquisitions with short TR are called T_1 -weighted. The second important acquisition parameter is the *echo time* (TE) which denotes the time between excitation and signal readout and controls the influence of T_2 time on image contrast. Tissues with short T_2 relaxation time dephase more quickly and their signal becomes dark with longer TE times. Acquisitions with long TE times are called T_2 -weighted.

Spatial Signal Encoding

Since the precession (Larmor) frequency is proportional to the magnetic field strength an additionally applied field gradient will make the frequency dependent on the spin location along that gradient. Hence, with an additional gradient along the main field, each plane is encoded (*phase encoding*) by its specific frequency and will only be excited by the respective RF pulse. The magnitude of an additional gradient and the RF pulse bandwidth control the thickness of the excited slice. The Larmor frequency-field dependency further allows for a spatial encoding of the signal by using two additional gradients (G_x and G_y) along the plane axes of the excited slice.

Image Reconstruction

By means of the localization gradients G_x and G_y each spin precesses with a location dependent frequency and acquires a phase Θ with:

$$\Theta(x, y) = 2\pi(k_x x + k_y y), \quad (2.1)$$

where

$$k_x = \gamma \int G_x(t) dt \quad \text{and} \quad k_y = \gamma \int G_y(t) dt. \quad (2.2)$$

The resulting net magnetization across the excited xy -slice can be represented as a function $f(x, y)$ and for the measured signal $s(t)$ holds:

$$s(t) \propto F(k_x, k_y) = \int f(x, y) e^{i2\pi(k_x x + k_y y)} dx dy \quad (2.3)$$

Equation 2.3 describes a 2D Fourier integral and thus $f(x, y)$ can be reconstructed from the signal by applying an inverse Fourier transformation.

Before applying the 2D inverse, the data collected from the signals is usually stored in the so called k -space. The horizontal axis (k_x) of the k -space represents the frequency information whereas the vertical axis (k_y) denotes the phase information. Typically, the k -space is acquired in lines, each with a fixed phase k_y .

2.1.2 Diffusion-weighted Imaging

In physics, diffusion describes the transport process molecules or particles undergo in different concentration settings. The particle flux \mathbf{J} caused by a concentration gradient is described by Fick's law $\mathbf{J} = -D\nabla n(r, t)$, with D the *diffusion coefficient* and $n(r, t)$ the local concentration of particles. However, diffusion processes also take place when no macroscopic concentration gradient is present. The so called *self-diffusion* happens in microscopic scales and can be accounted to the phenomenon of Brownian motion. Einstein reformulated the diffusion equations in terms of diffusion under probability gradients, introducing the concept of a *diffusion propagator* $P(\mathbf{r}|\mathbf{r}', t)$ which expresses the probability of a particle at position \mathbf{r} moving to \mathbf{r}' in time t . In case of free diffusion (i.e. not restricted by any obstacles) the diffusion propagator is a Gaussian function.

Diffusion-weighted MRI

Diffusion-weighted imaging extends the structural MRI by measuring the characteristics of water diffusion at some spatial location in order to infer the geometry of micro-structures limiting the diffusion at such. In the spin-echo experiment described previously (see Figure 2.2), a refocusing 180° RF pulse is applied at time τ to reverse the process of phase decoherence among the excited protons which started after flipping the magnetization into the xy -plane with a 90° RF pulse. However, the implicit premise of full phase recovery at time 2τ holds only if the spins do not move between application of the two pulses. Thus, under the influence of diffusion at microscopic scales, the spin re-phasing is not complete and the measured MRI signal is reduced. Exposing the imaged tissue to an additional strong gradient will emphasize the phase changes and thus increase the sensitivity to microscopic diffusion. Stejskal and Tanner [104] have developed the application of additional constant gradient by introducing two separate gradient pulses with an encoding time δ (pulse duration) and diffusion time Δ (separation of the two pulses). The first pulse is applied between the excitation and refocusing RF pulse, the second one between the refocusing pulse and the spin echo as illustrated by Figure 2.3. With gradient strength G and the gyro-magnetic ratio γ , the sensitivity to diffusion, the so-called *b-value*, is defined by:

$$b = \gamma^2 G^2 \delta^2 \left(\Delta - \frac{\delta}{3} \right) \quad (2.4)$$

Under the assumption of a negligible diffusion process taking place during the application of the pulses (which holds for small δ), the resulting phase change of a particle at position x induced by the gradient is given by $\phi = -\gamma x$. Hence, after two gradient




Figure 2.3: Sequence diagram of a typical diffusion MRI EPI acquisition. The upper line (RF) shows the applied focusing and re-focusing RF pulses, the middle line (G_x) shows the diffusion gradient activation and the bottom line (Echo) represents the measurable echo signal. The factors forming the diffusion b -value are shown on the gradient line: the separation of two pulses Δ , the pulse duration δ and the pulse strength $|G|$.

pulses (with the refocusing 180° RF pulse in-between them), the net phase change is given by

$$\phi_1 - \phi_2 = -q(x_2 - x_1) \quad (2.5)$$

This means that for stationary particles ($x_1 = x_2$) the net phase change will vanish. However, if the particles underwent a random diffusion, the phase increment caused by the first pulse will not cancel out through the second pulse and the resulting signal will be reduced. The signal attenuation is then calculated by dividing the signal $S(q)$ observed while using diffusion weighting by the signal $S(0)$ without diffusion weighting:

$$E(q) = \frac{S(q)}{S(0)} = \int \rho(x_1) \int P(x_1, x_2, \Delta) \exp[-iq(x_2 - x_1)] dx_2 dx_1, \quad (2.6)$$

with the spin density $\rho(x_1)$ at x_1 at the time of application of the first pulse. In most application $\int \rho(x_1) dx_1$ is set to unity. The second function used, $P(x_1, x_2, \Delta)$ is the diffusion propagator describing the likelihood of a particle at x_1 moving to x_2 after the diffusion time Δ . With free diffusion, the propagator is a Gaussian and the signal attenuation (2.6) reduces to:

$$E(q) = \exp\left[-q^2 D(\Delta - \delta/3)\right] = \exp[-bD] \quad (2.7)$$

After rearranging, D can be estimated from the measured MR signal. The result is usually referred to as Apparent Diffusion Coefficient (ADC).

2.1.3 Signal Modeling

Describing the signal attenuation by a single scalar value, the apparent diffusion coefficient D , is sufficient for samples where the measured diffusivity originates solely from random molecule interaction and is independent of the direction of the diffusion-encoding gradient, like it is in the brain gray matter. In more complex media, where the motion is restricted also by geometrical obstacles, a single coefficient cannot fully describe the orientation-dependent (*anisotropic*) diffusion motion. The ordered tissue in the brain's white matter, formed by bundles of myelinated fibers, will hinder the diffusion in certain directions while allowing more diffusion in others.

The next, more complex model to describe Gaussian diffusion in anisotropic case is the symmetric tensor of second order

$$\mathbf{D} = \begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{pmatrix} \quad (2.8)$$

The *diffusion tensor* \mathbf{D} is often visualized as an ellipsoid, a surface representing constant probability of the distance a molecule will diffuse from origin for a given direction $\mathbf{g} \in \mathbb{R}^3$. The probability is given by the Gaussian propagator discussed above for $P(\mathbf{g}, \mathbf{0}, \Delta)$. Several tensor configurations are shown on Figure 2.4. The diffusion tensor is symmetric and as such has six independent variables. Hence, for its estimation at least six diffusion weighted measurements with different gradient orientations $\mathbf{g}_i \in \mathbb{R}^3$ together with one acquisition without diffusion weighting $S(0)$ is needed. The signal attenuation equation 2.7 extends to

$$E(g_i) = \frac{S(\mathbf{g}_i)}{S_0} = \exp \left[-b \cdot \mathbf{g}_i^T \mathbf{D} \mathbf{g}_i \right]. \quad (2.9)$$

With a sufficient amount of measured gradients, the tensor can be robustly estimated by a linear least squares approach.

Higher Order Models

With the ability of quantifying diffusion anisotropy, the diffusion tensor imaging became a popular technique for measuring micro-structural tissue parameters. It performs well in areas of single fiber bundles, however it fails in the presence of complex fiber configurations, for example at fiber bundle crossings. In such a case, the diffusion propagator cannot be assumed Gaussian anymore. To overcome these limitations in

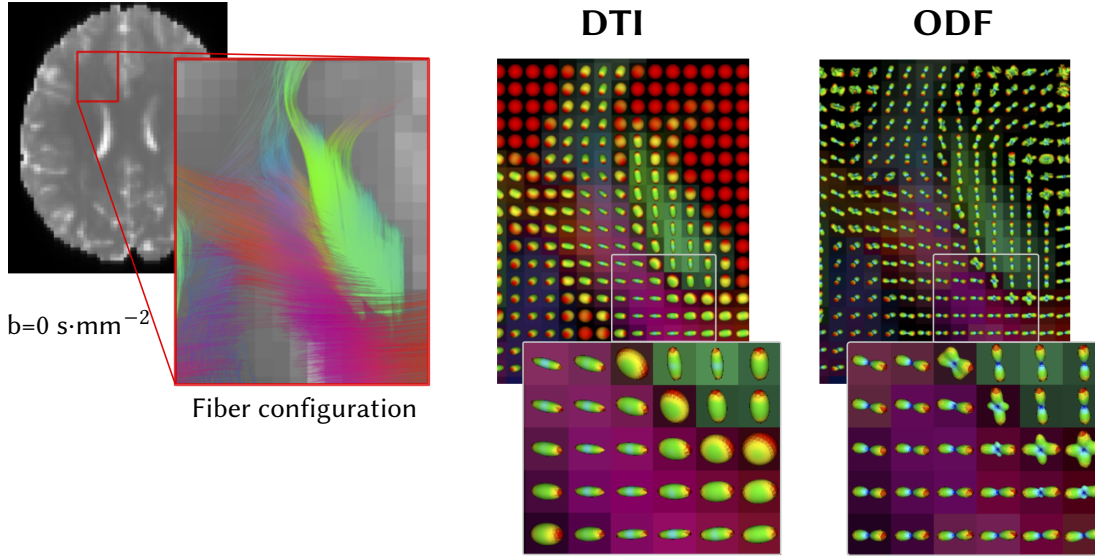


Figure 2.4: Diffusion MRI signal models example for crossing fibers. A crossing fiber region in the area marked on the left image, on the right the reconstructed diffusion tensors (DTI) and the orientation distribution functions (ODF). The DTI model is not capable of resolving the crossing resulting into an oblate shape whereas the ODF shows multiple maxima at such voxels.

the diffusion modeling and to provide a more specific model description of microscopic brain tissue properties, *higher order* models have been proposed.

A straightforward extension of the tensor model covering also for crossing areas is to consider a collection of tensors $\mathbf{D}_1, \dots, \mathbf{D}_N$, each representing a different fiber population. In general, the relation in 2.9 can be lifted to

$$E(q) = \sum_{j=1}^N w_j \exp \left[-b \cdot \mathbf{g}^T \mathbf{D}_j \mathbf{g} \right], \quad \sum_j w_j = 1 \quad (2.10)$$

and for most crossing areas $N = 2$ is sufficient. However, unlike for the single diffusion tensor, the estimation of multiple tensors requires non-linear optimization approaches.

The aim of all diffusion signal models is to reconstruct the (true) fiber Orientation Distribution Function (ODF) which quantifies the fraction of fibers oriented in a given direction. The models mentioned above are restricted to recovering a finite number of predominant fiber orientation directions. A further generalization leads to so-called non-parametric models that reconstruct the fiber ODF with minimal constraints [102]. The most common non-parametric approaches are the Q-ball imaging that reconstructs a diffusion ODF (see Figure 2.4), a diffusion probability rather than the true fiber ODF [1]; and the approach of Spherical Deconvolution (SD) which considers the signal as a linear

combination over the basis of Spherical Harmonics (SH) functions and recovers a more direct estimate of the fiber ODF [109].

Compartment Models

The signal models introduced in previous paragraphs can be seen as a *top-down* approach, starting at macroscopic scales and improving their signal description to capture structures at lower scales. Compartment models, on the contrary, follow the *bottom-up* direction, aiming at precise modeling of the microscopic tissue types (*compartments*) and their influence on water diffusion and hence the dMRI signal at more coarse levels. Three basic compartment types can be identified in brain tissue – the *intra-axonal*, which describes the space bounded by the membrane of neurites, the *extra-axonal*, which refers to the space between the neurons and the *isotropic* Cerebro-spinal Fluid (CSF) compartment where free diffusion is expected. For instance, the DTI signal model describes only a single extra-axonal compartment. For each compartment, several model representations were proposed taking also the previous models into account. An intuitive model choice for the extra-axonal compartment is the DTI ellipsoid, but also other shapes like the *zeppelin* are proposed. A comprehensive taxonomy of different models for the compartments was published by Panagiotaki *et al.* [82], who additionally consider an isotropically restricted compartment. However, a study on human brain tissue by Ferizi *et al.* [36] proved that the three-compartment models are best suited for the modeling.

One of the first models separating the influence of a restricted (intra-axonal) and hindered (extra-axonal) diffusion was the Composite Hindered and Restricted Model of Diffusion (CHARMED) model [6]. Such tissue modeling provides a more precise description of crossing fiber configurations, however the number of model parameters is higher than for the multi-tensor in the $N = 2$ case (2.10) and thus requires far more data at different b -values for stable fitting.

One of the prominent models is the Neurite Orientation Dispersion and Density Imaging (NODDI) approach [119]. NODDI adopts the three compartment model by considering the intra-axonal compartment as a set of *sticks*, the extra-axonal compartment as the anisotropic Gaussian diffusion (the DTI tensor), and the isotropic CSF compartment. The overall signal is formed as a weighted sum of the three contributing compartment signals

$$A = (1 - v_{\text{iso}}) (v_{\text{intra}} \cdot A_{\text{intra}} + (1 - v_{\text{intra}}) \cdot A_{\text{extra}}) + v_{\text{iso}} \cdot A_{\text{iso}}$$

The intra-axonal (A_{intra}) signal is modeled as a set of oriented sticks and their orientation is considered to follow the Watson distribution

$$f(n) = M\left(\frac{1}{2}, \frac{3}{2}, \kappa\right)^{-1} \cdot e^{\kappa(\mu \cdot n)^2} \quad (2.11)$$

for a hypergeometric function M , mean orientation μ and concentration κ [119].

2.2 Diffusion-weighted Images

In the previous sections, the physics of diffusion-weighted MR acquisition along with the approaches to signal modeling were introduced. This section will focus on the concept and the related terminology of diffusion-weighted images which are the central data structure in all following chapters. In the field of (medical) image processing, an image I is commonly defined as a discrete three-dimensional (3D) grid of *voxels*, each of them associated with a gray value:

$$I : \Omega \subset \mathbb{R}^3 \rightarrow \mathbb{R} \quad (2.12)$$

The 3D image can be seen as an ordered stack of two-dimensional (2D) images - the *slices*. Based on the MR imaging principle, the acquired xy -plane is directly associated with the slice (called *axial* slice in radiological terminology) and the slices are ordered along z -axis, the *out-of-plane* axis. The spatial extent of a voxel, and equivalently the distance between midpoints of the neighboring voxels, is referred to as *spacing*, one speaks of *isotropic* voxels in the case of identical spacing in all three directions.

The main principle of diffusion MR acquisition is to probe the target imaging volume (e.g. the head) at different diffusion weightings and different diffusion directions - the gradients. Hence, a diffusion-weighted image (dw-image) I_{dw} with N acquired gradient directions can be seen as a 3D vector image

$$I_{\text{dw}} : \Omega \subset \mathbb{R}^3 \rightarrow \mathbb{R}^N \quad (2.13)$$

Another interpretation is to view a dw-image as a set of 3D volumes

$$I_{\text{dw}} = (I_1, \dots, I_N) : \mathbb{R}^{1+3} \rightarrow \mathbb{R}, \quad (2.14)$$

where each volume is associated with a gradient direction $\mathbf{g}_i, i = 1, \dots, N$. According to the diffusion-weighting the volumes will be referred to as *unweighted* ($b = 0 \text{ s} \cdot \text{mm}^{-2}$)




Figure 2.5: Scheme of a multi-shell diffusion-weighted image. The schematic on the left side shows an exemplary organization of the volumes with different weighting (gray circles) and the unweighted image (white circle). The position on the arc represents the diffusion gradient direction. Corresponding typical dw-images at the different b -values are shown on the right half of the picture. Note that the intensities of the weighted images are displayed with different level window settings.

and *weighted* ($b > 0 \text{ s} \cdot \text{mm}^{-2}$) volumes (also referred to as *images* when a confusion is unlikely).

A common representation of the gradients $(g_i)_{i=1}^N$ is to identify them with points on a sphere with a radius equal to the gradient vector magnitude, which is proportional to the applied diffusion weighting, the b -value. Within this representation, all gradients with equal b -value will be associated with one sphere, often referred to as *shell* (or b -shell). The unweighted image ($g_0 = (0, 0, 0)$) forms the midpoint of each b -shell. One speaks of a *multi-shell* diffusion-weighted image when at least two distinct weighted shells are acquired. The gradient scheme is illustrated in Figure 2.5. With the ordering by b -value, also the intuitive notation *higher* and *lower* shell is commonly used.

The most important characteristics of diffusion-weighted images in the context of methods presented later on is the contrast (signal) behavior which can be simplified by two main points:

1. The global signal level and the contrast decrease at higher shells. As a consequence, the the Signal-to-Noise Ratio (SNR) drops.
2. The global signal level and the SNR are identical for images from one b -shell, however the contrast varies locally in dependence to the diffusion gradient direction.

Exemplary slices from a dMRI dataset illustrating this behavior are shown in Figure 2.6.

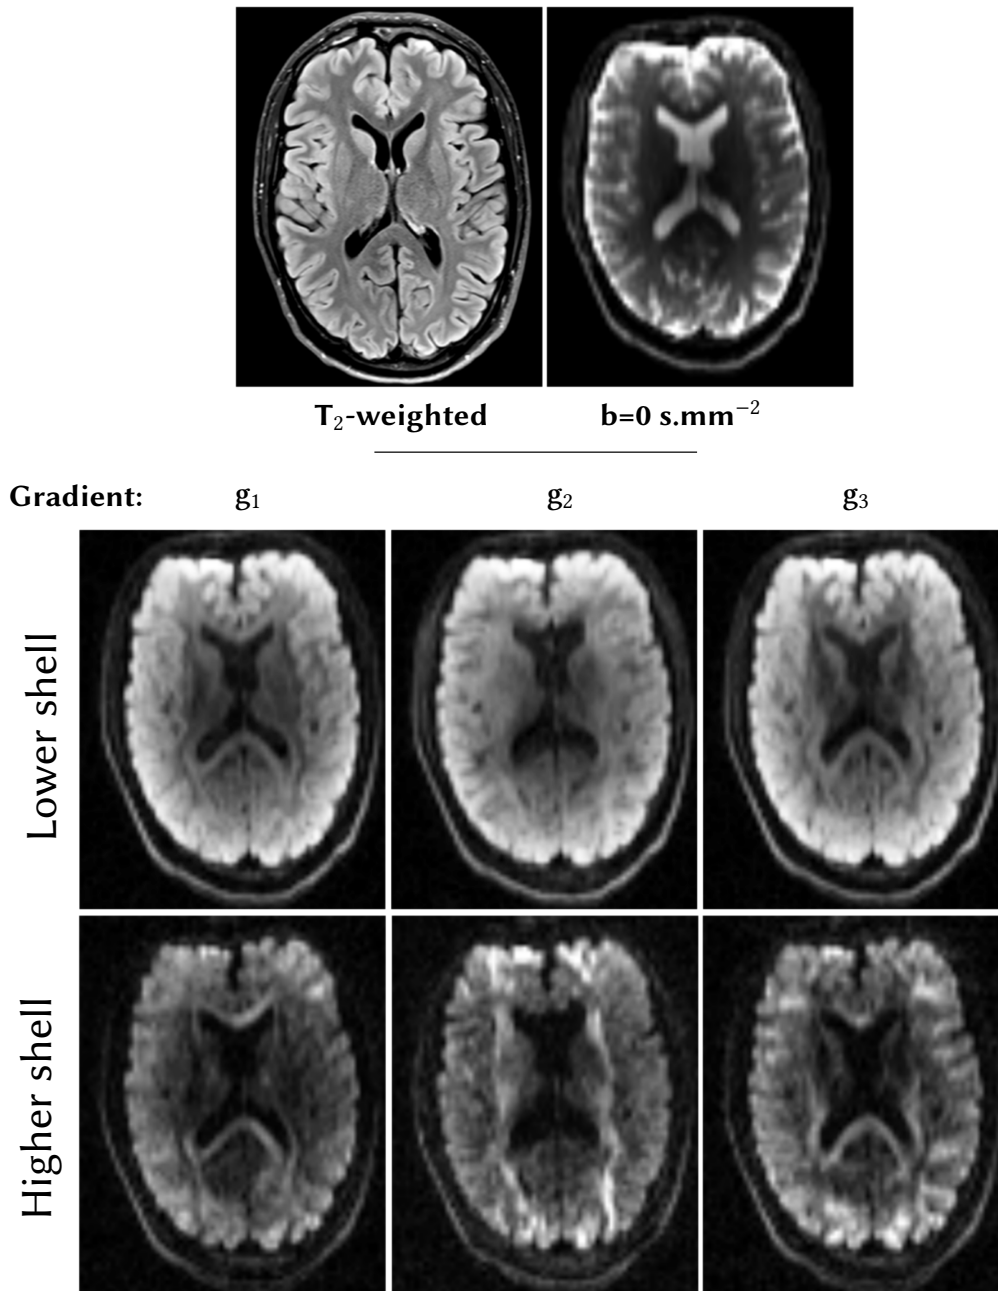


Figure 2.6: Exemplary axial slices from a diffusion-weighted image. The different MR modalities are shown for an identical axial slice from one subject. In the upper row the anatomical images (T_2 -weighted) and the unweighted diffusion image ($b = 0 \text{ s.mm}^{-2}$) are shown. The lower parts show the weighted images for three different diffusion gradient directions acquired at lower ($b=1000 \text{ s.mm}^{-2}$) and higher ($b=3000 \text{ s.mm}^{-2}$) diffusion weighting. Note that the intensities of the weighted images are displayed with different level window settings - the lower (LW: $[0,800]$) and the upper (LW: $[0,200]$).

2.3 Data Processing and Applications

This section summarizes some of the main analysis approaches of the evaluation and quantification of diffusion MRI data. The diffusion MRI was introduced as acquisition technique in the mid-1980s and has enabled the non-invasive insight into the tissue (micro-)structure, and forth on, it became a pillar of modern neuroimaging [64]. To emphasize the wide (clinical) domain of dMRI, the three main analysis principles and their applications are presented starting with the estimation of tissue parameters, followed by diffusion tractography and the graph-based analysis by connectomics.

2.3.1 Estimating Tissue Parameters

One of the main research branches focuses on the estimation and the identification of tissue parameters that differ in diseased and healthy subjects and could thus provide a novel *biomarker* for early diagnosis of neurodegenerative diseases.

Diffusion Tensor Imaging

One of the first approaches to tissue characterization was enabled by diffusion tensor imaging by introducing rotationally invariant measures (the DTI-derived *scalar indices*) of the local anisotropy and diffusivity, reducing the tensor data, which are formed by a 3×3 symmetric matrix (c.f. Section 2.1.3), to a single value per voxel. This reduction is motivated by properties of the tensor ellipsoid, whose axes and their extent correspond to the matrix eigenvectors and -values λ , respectively. The Axial Diffusivity (AD) along the principal axis is defined by the biggest eigenvalue (λ_1), the Radial Diffusivity (RD) is defined as the average of the remaining orthogonal eigenvalues $(\lambda_2 + \lambda_3)/2$. The Mean Diffusivity (MD) is defined as $Tr(\mathbf{D})/3$ via the trace of the tensor, which is the sum of all three eigenvalues $Tr(\mathbf{D}) = \lambda_1 + \lambda_2 + \lambda_3$. The most widely used invariant measure of local anisotropy is the Fractional Anisotropy (FA) [61]:

$$FA = \sqrt{3 \cdot \frac{(\lambda_1 - MD)^2 + (\lambda_2 - MD)^2 + (\lambda_3 - MD)^2}{2(\lambda_1^2 + \lambda_2^2 + \lambda_3^2)}}$$

For illustration, examples of the different scalar indices are shown on Figure 2.7.

Evaluation of FA maps became particularly popular in investigation of neurodegenerative diseases, since the loss of fiber bundle integrity is reflected by a decreased FA value. FA was proved a valuable biomarker in Alzheimer's disease [70], Parkinson




Figure 2.7: Exemplary axial slices of DTI scalar indices maps. All images show the same axial slice from one subject, left the DTI data, color-coded by the tensor principal direction, right the maps of the common scalar indices - the Fractional Anisotropy (FA), Mean Diffusivity (MD), Radial Diffusivity (RD) and Axial Diffusivity (AD).

disease [22], borderline personality disorder [69], epilepsy [93], schizophrenia [87], traumatic brain injury [105] and others [7].

Higher Order Models

Despite its extreme popularity DTI provides only a limited description of the tissue in complex white-matter regions like fiber crossings. One of the higher order models which compensates the DTI shortcoming is the NODDI-model (introduced in Section 2.1.3). Here, similar to the construction of the DTI indices (FA, MD, etc.), several scalar measures are proposed. From the fiber model (see equation (2.11) on p. 14) several parameters with a direct link to the tissue can be derived. The Intra-cellular Volume Fraction (ICVF) v_{ic} measures the neurite density [101] and the Orientation Dispersion Index (ODI)

$$ODI = \frac{2}{\pi} \arctan(1/\kappa) \quad (2.15)$$

describes the integrity of fiber bundles. In fact, v_{ic} and ODI are the major contributing factors of the FA index and thus with NODDI a separate analysis is possible [118]. Especially ODI was found to correlate with the histological tissue measures [24, 101] and to describe the maturation tissue changes better than FA [19].

Parallel to the voxel-based evaluation of scalar measures, the "configuration" of the neuronal pathways in a more global view became of particular interest.

2.3.2 Tractography

Diffusion tractography focuses on the reconstruction of white-matter neuronal pathways in the brain. In fact, it is the only *non-invasive* tool for measuring such pathways *in-vivo*. Myelinated axons, the building blocks of neuronal pathways in the brain, are the main contributors for the measured restricted and hindered diffusion signal. Starting with a model approximation of the fiber orientation distribution function (see Section 2.1.3) in each voxel, fiber tractography aims at reconstructing the fibers that explain the measured signal. The approach of streamline tractography, depicted in Figure 2.8(a,b), provides a good illustration of the reconstruction principle: starting from a seed point, a streamline is reconstructed step-wise in each voxel by following the principle direction provided by the underlying signal model to enter one of the neighboring voxels. This procedure is iterated until a stopping criterion is satisfied or a constraint is violated (f.i. curvature value, direction change). To obtain a dense tracking result multiple seed points in different locations must be considered. The tract reconstruction problem has attracted a lot of attention resulting in a vast amount of different methods, including deterministic and non-deterministic global [78], machine learning-based [77] and other approaches. An overview of some methods can be found in the summary to the annual Tractography Challenge [89]. In a recent evaluation

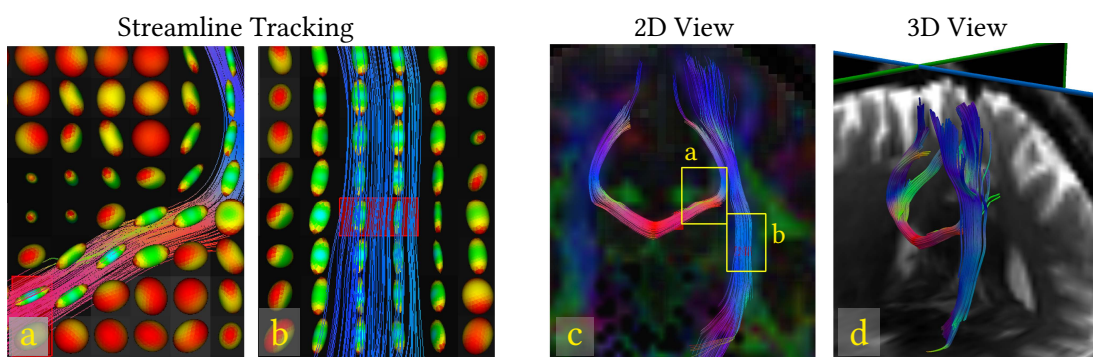


Figure 2.8: Principle of diffusion MR fiber reconstruction. (a,b) Examples of the streamlines following the main fiber orientation as represented by the DTI glyphs starting in the seed voxels marked by red rectangles. (c) A complete view of the cortico-spinal tracts (blue) and lateral projections from the corpus callosum tracts (red). The yellow rectangles mark the locations depicted in (a) and (b) in a higher detail, (d) provides a three-dimensional view of the reconstructed tracts using a tube for visual representation of a tract.

of different signal model reconstructions and different tractography pipelines, the Constrained Spherical Deconvolution (CSD) as fODF reconstruction with a deterministic tracking provided the most reliable results [75].

The reconstructed tracts (the *tractogram*) provide a valuable information for critical neuro-surgical interventions, which let the tractography establish in neurosurgical imaging and surgical guidance systems [115], for example in deep-brain stimulation in Parkinson's disease [95].

2.3.3 Connectomics

Connectomics focuses on the macroscopic scale trying to elucidate the structural (and functional) connections of the brain [39, 53]. To achieve this, the connectomics approach considers the structural analysis in terms of graph theory which is a well-established research field and as such provides a compendium on theoretical results and algorithms.

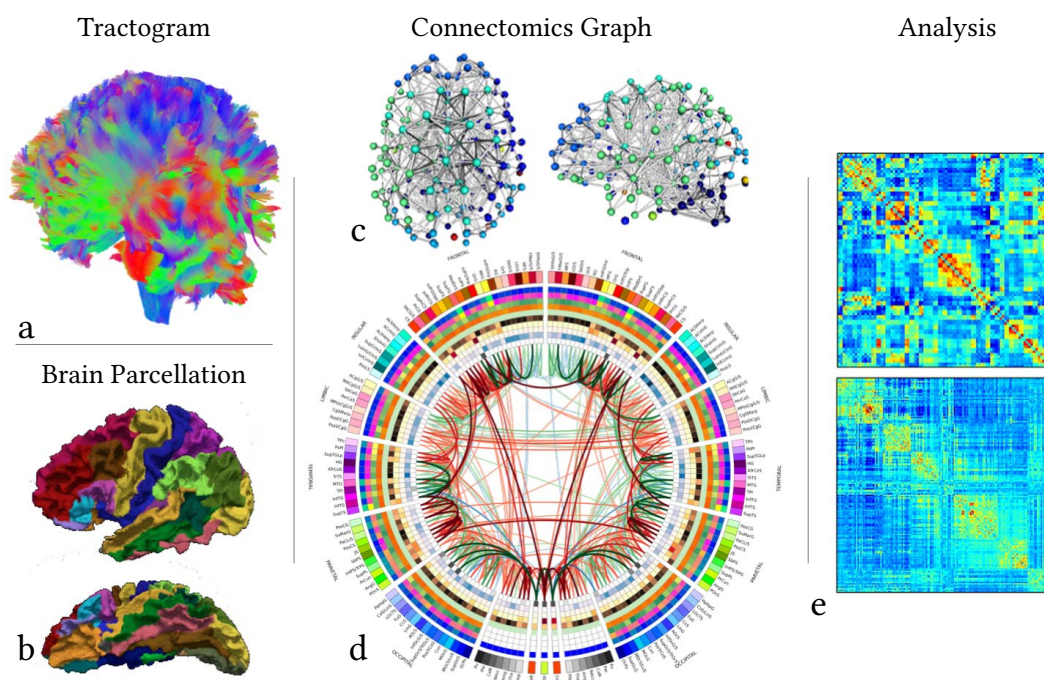


Figure 2.9: The connectomics principle. From a full-brain tractography (a) and the parcellation in anatomical areas (b) the connectomics graph is constructed. The graph itself is visualized in two ways: (c) as a 3D graph with real coordinates of the vertices or (d) in a more abstract way as a *connectogram*. In the analysis, the connectivity matrix (e) describes the connectivity patterns in a structured way. Images (b,c,e) were published under CC license in [66, 113, 14] and (d) *Illustration* CC BY-SA by Lonicula.

To construct a graph $G = (V, E)$, the set of vertices V is formed by the different anatomical areas in the brain (the parcellation) and an edge $e = (v_k, v_j)$ between two nodes is constructed if there exists a connection in the full-brain tractogram between the corresponding areas. Here, also weighted edges can be considered to capture the strength (the "bandwidth") of the connection. The components and the graph construction are illustrated by Figure 2.9.

The constructed graph $G = (V, E)$ enables graph-theoretical measures to be used in analyzing the brain topology. Among the measures, which also provide a meaningful representation in terms of brain function, are the *characteristics path length*, *global network efficiency* or *betweenness centrality*. The results were already used to examine the disease-related changes to the connectome (e.g. in autism spectrum disorder [41], schizophrenia [107], epilepsy [40]) but also the alterations during healthy aging and development [111]).

2.4 Artifacts in Diffusion MR

The precision of the analysis steps described in the preceding section depends strongly on the *quality* of the input image data which can be strongly influenced by different distortion and artifacts in dMR imaging. This section explains the main artifact sources, which are the subject's movement during the scanning and image distortions induced by eddy currents.

2.4.1 Patient Motion

A stable basis for the signal model fitting can only be achieved by sufficient sampling of diffusion directions, which leads to the acquisition of many diffusion-weighted volumes. Commonly, in the clinical protocol a few unweighted images along with ~ 30 weighted volumes are acquired for DTI, the CSD reconstruction with a reasonable SH basis requires a minimum of 45 weighted volumes [108]. Higher order models like NODDI require in addition the acquisition of multiple *b*-shells which pushes the number of weighted images to a count of over 60. All in all, the total acquisition time increases with each volume and this inevitably leads to subject's motion. Furthermore, an effect that cannot be ignored due to the application area of diffusion MRI, subject's age and health conditions can further impair the ability to hold still also during shorter acquisitions.




Figure 2.10: Effects of subject's head motion during acquisition. All images show an axial slice from one subject, the upper row (Steady) comes from an acquisition without motion, the lower row (Motion) is distorted by intentional movement of the subject. (a) the unweighted image, (b,c) weighted images without motion artifacts, (d) the moved unweighted image, (e) a weighted image with moderate motion artifacts, to emphasize the effect of the rotation on the diffusion-weighted signal two areas with changing intensities are marked by the colored ellipses, (f) a weighted image affected by severe motion artifacts that led to signal drop-outs.

The slow, bulk motion of the subject can be considered as a rigid-body motion, i.e. a combination of translation and rotation. The effects of the subject's motion on the reconstructed images depend on the time of occurrence. Slow motion of the head results into a misalignment between the different volumes (c.f. Figure 2.10 (d,f)). Motion taking place during the diffusion gradient encoding phase, on the contrary, will result into signal attenuation in a single slice, or even a complete slice signal deletion (c.f. Figure 2.10 (g)). Such effects happen mostly with rapid and erratic motion of the subject.

A further consequence of head motion is the relative change of the applied diffusion gradient direction. A head rotation by an angle α will change the orientation of the brain tissue relatively to the gradient \mathbf{g} and the measured signal will correspond to a "true" gradient direction \mathbf{g}' rotated by $-\alpha$ as illustrated by Figure 2.11.




Figure 2.11: Effects of head motion on measured diffusion gradient. (A) The intended original gradient direction g , (B) rotation of the head relative to the scanner and (C) the true measured gradient g' in relation to the original g .

2.4.2 Eddy Currents

The gradient coils of an MRI scanner are energized to create magnetic field gradients (for spatial encoding, slice selection, etc.) in the scanner's bore. A change of the field is achieved by switching gradients on/off. This will induce currents in any conductive materials located inside the field, and especially in the coils. In turn, the so induced currents produce an own additional field—the *eddy-currents-induced off-resonance field*. In diffusion imaging, the effects caused by eddy currents are more significant than in "conventional" MR imaging, since the diffusion-encoding gradients have higher amplitude [5].




Figure 2.12: Eddy currents in a diffusion MRI EPI sequence. Turning the diffusion gradient G on/off will induce additional magnetization (blue lines) with slow decay reaching into the readout (echo) phase.




Figure 2.13: Eddy-currents-caused distortions in diffusion-weighted images. (a) An axial slice of an undistorted volume, (b,c) volumes distorted by eddy currents effects showing the zooming effect. The yellow contour shows the brain outline of the undistorted volume. All volumes were acquired with b -value of $3000 \text{ s} \cdot \text{mm}^{-2}$ during one session.

The eddy currents build up during the activation phase of a gradient and decay in the off-gradient phases as illustrated by Figure 2.12. This slow exponential decay of additional field during the readout (echo) phase causes geometrical distortions of the resulting image as illustrated on a diffusion acquisition sequence in Figure 2.13. An eddy current field will add to the spatial-encoding gradient, which will modify the k -space traversal in the signal readout phase and result in geometrical distortion of the reconstructed image. The kind of distortion depends on the eddy-currents-induced gradient direction. For example, an eddy current gradient in the x -direction will result in a *sheared* k -space traversal and the reconstructed image will suffer from a xy -shearing distortion, whereas an additional gradient in the phase-encoding direction (mostly y) will stretch the traversal and lead to zooming effects in the reconstructed image.

In the first EPI sequences for diffusion MRI, the effects were particularly severe, however, new acquisition sequences with reduced eddy currents, like the widely adopted twice refocused spin echo were proposed [91].

2.4.3 Higher Order Distortions

The head motion and eddy-currents-induced distortions are the most relevant sources of image misalignment in diffusion MRI sequences but also other effects and distortions can occur. In the ideal case, the scanner's static field B_0 is taken as homogeneous but this assumption does not hold in real conditions. First, the field itself is prone to minimal

fluctuations, and second, the scanned tissue itself is diamagnetic and the field inside will become inhomogeneous. However, the B_0 inhomogeneities (or *susceptibility* artifacts) affect all diffusion images in the same manner and do not produce misalignment effects between the different volumes, but they have to be accounted for when transferring the dMRI data into the reference frame of structural data (f.e. the T_2 -weighted image).

All previously described artifacts originate from distortions (or inhomogeneities) of the phase-encoding gradients. If the frequency-encoding gradients are affected, they cause a zigzag offset between the lines in the k -space readout and lead successively to so-called ghosting artifacts (a shifted overlay) in the reconstructed image [76].

2.4.4 Conclusion

Clearly, the distortions and artifacts explained in the previous section negatively influence the image data, and if uncorrected, this influence propagates into all subsequent analysis steps. One of the distortion consequences is the misplacement of the voxels in the different weighted volumes, regardless whether the motion or eddy-currents-caused distortion were its origin. This misplacement, in turn, will affect the signal model fitting procedure by reducing its precision and explanatory power.

The impact of imprecise preprocessing of the diffusion-weighted images was evaluated in multiple aspects. An initial alignment error in the data, for example in DTI studies, can influence the resulting scalar maps [50] or can affect the results of fiber tracking and connectomics-based evaluation in HARDI data [32] or result into spurious group differences [117]. The tractography is known to profit from higher b -value acquisitions [35], which directly influences the connectomics analysis [42]. Here the initial misalignment effect would be more pronounced, as the proneness to imaging artifacts increases with higher b -values.

These findings assign a key role to the image preprocessing step which aims at mitigating the distortions to provide a reliable basis for the data analysis steps.

Chapter 3

State of the Art

This chapter aims to give context to the presented new work. To give a sense of the capabilities and limitations this work aims to address, Section 3.1 reviews the current methods in diffusion-weighted MR artifact and motion correction. A background of ongoing research in the field of general image registration (Section 3.2) is followed by a more thorough discussion on evolutionary optimization principles and methods (Section 3.3) as one of the more promising avenues for the motion correction of dMRI scans, with a particular focus on their extension for handling multiple objective functions. In the last section (3.4) hybrid approaches in evolutionary optimization and their application to image registration in both single- and multi-objective settings are introduced as they offer great potential for addressing some of the problems presented in Section 3.3.

3.1 DW MRI Distortion Correction

In the previous chapter, the common acquisition artifacts (Section 2.4) and their impact on typical data analysis approaches in diffusion MRI (Section 2.3) were presented. In conclusion: to achieve a reasonable quality in the subsequent micro-tissue analysis a robust handling of the artifacts that can occur during the acquisition becomes essential. Thus, a proper data preprocessing grew to a separate field of interest within diffusion MRI research. In this section, the state-of-the-art correction approaches are introduced. These can be separated by their point of application into three groups: the *rejection* approaches that aim at detecting and discarding of the corrupted imaging volumes prior to the processing steps (c.f. Section 3.1.1); the *prospective* correction approaches (introduced in Section 3.1.2) that use an additional MRI-independent measurement to

detect and compensate the distortion during acquisition and the *retrospective* correction approaches that employ image registration methods to eliminate the distortion from the reconstructed images after the acquisition (c.f. Section 3.1.3).

3.1.1 Rejection Approaches

A straightforward approach to distortion correction is to discard the affected images from further processing. A combined approach to outlier detection and elimination was presented for the special case of diffusion tensors—the RESTORE estimator for DTI sequences adapts the fitting method of iteratively reweighted least-squares fitting to allow detection of possible outliers and to subsequently exclude them from the model fitting [17]. The authors further improved their approach recently for datasets with low number of acquired diffusion volumes [18], and at the same time, Pannek *et al.* [83] proposed a closely related approach using spherical harmonics for the reconstruction of the fODF and thus enabling reliable outlier detection also for higher b -values. However, the rejection approaches are only of limited use as was pointed out by Elhabian *et al.* [34] in their evaluation of motion correction schemes on the subsequent processing steps. They concluded that volume rejection approaches can be used only if less than 10% of the volumes are removed. Such amounts of corrupted volumes may not be reached in scans of young healthy subjects, however can occur while scanning elderly patients, small children or in consequence of the subjects' health condition. The authors also conclude, that further gradient removal should not result in unbalanced sampling of the q -space to maintain the uniform distribution of the gradients on the sphere.

3.1.2 Prospective Motion Correction

Common to all prospective approaches is the need for an additional measurement of subject's motion independent of the diffusion-weighted MR signal itself. Aksoy *et al.* [2] presented an optical, real-time prospective correction system. It consists of a small camera mounted on the head coil that tracks the movement by capturing a unique checkerboard pattern attached to subject's forehead.

To overcome the limitations posed by an additional, independent and magnetic-field resistant tracking system, the approaches using only the acquired data itself are gaining more and more attention. A correction and re-acquisition scheme was proposed by Benner *et al.*—the diffusion-weighted prospective acquisition correction [13]. Here, the motion is compensated by registering the moved imaging volume to an unmoved

reference volume with the same b -value and in case of signal drop-outs due to severe motion, the volume is reacquired. Further approaches use additional MR sequences, so called *navigators*, instead of the data itself. Kober *et al.* [63] used free induction decay (FID) navigators in each imaging volume in order to detect subject's motion by volume-wise comparison of the navigator signal. If motion is detected, an additional unweighted $b=0 \text{ s} \cdot \text{mm}^{-2}$ image is scanned to provide a stable contrast for subsequent correction performed by retrospective registration of the additional unweighted images. The imaging volume that triggered the motion detection is scanned again after the extra b_0 scan. Alhamud *et al.* [4] introduced the usage of volumetric navigator scans that contain 3D anatomical information for direct computation of the motion parameters which should provide motion estimates that are not affected by the diffusion gradients even at higher b -values, and extended the work very recently also for the compensation of B_0 field inhomogeneities [3].

While minimizing the effects of bulk motion during acquisition, the prospective approaches do not address the artifacts that originate from other sources, in particular the eddy currents caused distortions. In addition, the prospective approaches are limited either by the extra measurement hardware, or by the necessary sequence modification, which are both not easily deployed into clinical practice.

3.1.3 Retrospective Motion Correction

The approach of pairwise image registration was used among the first methods in DWI motion and artifact correction [11, 46]. The first seminal work was presented by Rohde *et al.* [94] who considered a 6 Degree of Freedom (DOF) rigid-body transformation compensating the head motion together with an 8 DOF function describing the effects by eddy currents in a unified cost function. Another seminal work was later presented by Mohammadi *et al.* [72] who proposed to optimize a full-affine transform (12 DOF) under the minimization of normalized Mutual Information (nMI) metric. The authors further demonstrated, that their approach successfully corrects both motion and eddy-currents-induced distortions and hence is equivalent to the solution by Rohde *et al.*

As a consequence of Mohammadi *et al.*'s work, later proposed retrospective correction methods implement their full-affine transform approach. Hence, they differ in the choice of the fixed and moving images in the pairwise registration settings, as well as the choice of image metric (see Section 3.2 for definition). The methods will be explained on a multi-shell DW image and illustrated by alterations of the Figure 2.5 (c.f. Section 2.2, p. 15). The target reference space in the following is the voxel-space

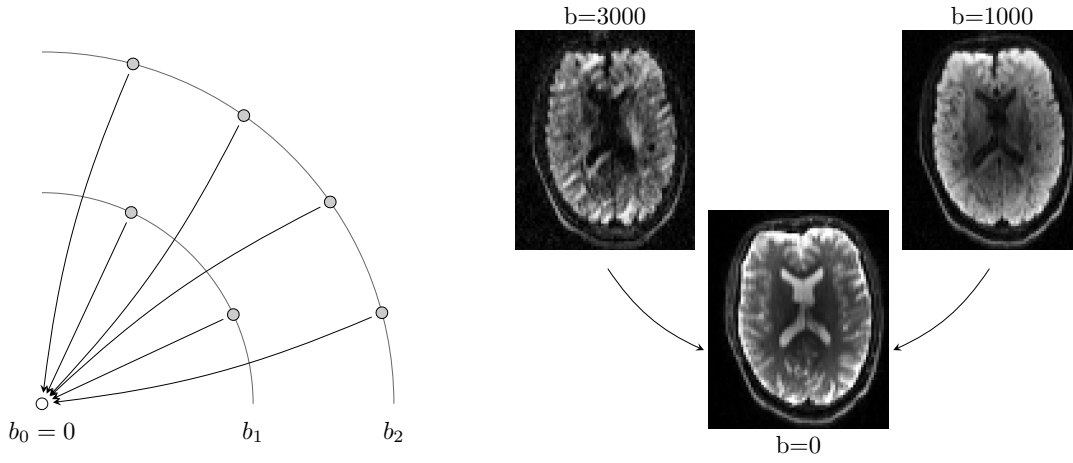


Figure 3.1: Unweighted reference correction scheme. The solid line arrows in the DW Image scheme on the left side indicate the pairwise registration tasks with moving image at tail and reference image at head. On the right side, exemplary images from two different b -shells from an in-vivo dataset are shown.

defined by the unweighted ($b = 0 \text{ s} \cdot \text{mm}^{-2}$) image I_0 . It can be assumed that there is only a single unweighted image, otherwise, if multiple unweighted images are acquired, they are aligned to the first one with respect to the normalized correlation metric and the average of the aligned images is used in further processing.

Unweighted Reference Method

One of the most employed corrections is the method by Mohammadi *et al.* [72], with all weighted images being registered to the unweighted image under the minimization of nMI metric. The method is schematically depicted in Figure 3.1.

The unweighted reference method is well suited for typical DTI acquisition, since it reliably aligns images at lower b -shells, typically acquired at $b = 1000 \text{ s} \cdot \text{mm}^{-2}$. However, it does not provide sufficient robustness for precise registration of weighted images at higher b -values ($b > 2000 \text{ s} \cdot \text{mm}^{-2}$). The application in alignment of HARDI single-shell images acquired typically at $b = 2000 \text{ s} \cdot \text{mm}^{-2}$ is a limit case [33]. With increased b -value, the overall signal in the weighted images and also SNR drop and the gradient direction-dependent signal in fibrous structures is more emphasized, as can be seen on example slices from different shells in Figure 3.1.

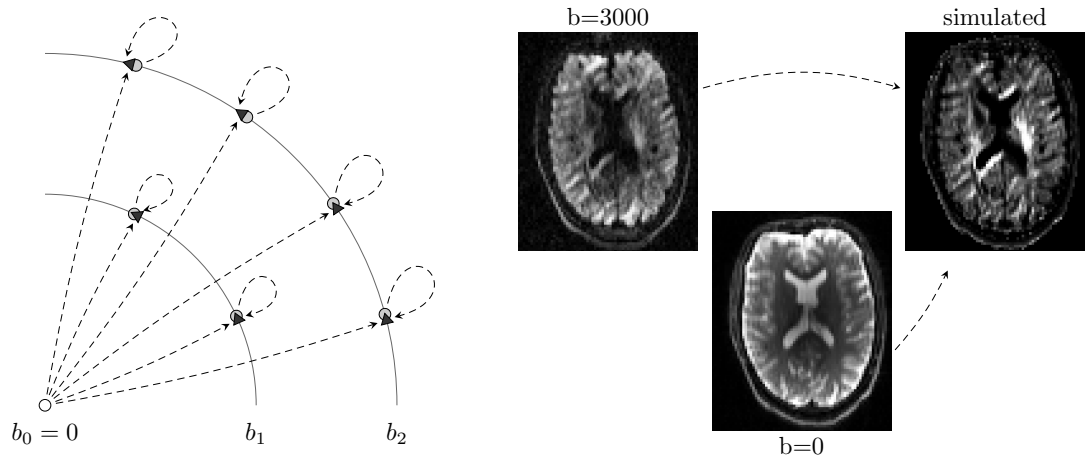


Figure 3.2: Extrapolation method correction scheme. The DW Image scheme on the left side is enhanced by the simulated reference images (black triangles) computed on the basis of a model (dashed line). The corresponding pairwise registration tasks (solid line arrows) aim from the moving image (arrow tail) to the simulated reference image (arrow head). On the right side, an image from a $b=3000 \text{ s}\cdot\text{mm}^{-2}$, the corresponding simulated reference image and the $b=0$ image are shown.

Extrapolation Methods

To overcome the contrast differences between the unweighted and the weighted images for increasing b -values, approaches using reference images that better mimic the signal on higher shells were introduced [10, 12]. Common to the approaches is the usage of a signal model \mathcal{M} for the estimation (extrapolation) of the dMRI contrast, in this case the DTI model (see Section 2.1.3 for definition). In a generalized form, a diffusion-weighted MR signal model \mathcal{M} can be described as a function which depends on the b -value and the diffusion gradient direction \mathbf{g} :

$$\mathcal{M} = M(b, \mathbf{g})$$

and can be fitted to the measured data \mathbf{Y} with some residual error ϵ

$$\mathbf{Y} = \hat{M}(b, \mathbf{g}) + \epsilon$$

The model fit then provides the base for the simulation (extrapolation) of the new target image I_{sim} for a fixed b -value b' and a gradient direction \mathbf{g}' by simply evaluating the fitted model at each voxel position \mathbf{x} :

$$I_{\text{sim}}(b', \mathbf{g}')(\mathbf{x}) = \hat{M}(b', \mathbf{g}')(\mathbf{x}), \quad \mathbf{x} \in \Omega(I_0)$$

Evaluating the model at positions \mathbf{x} from the voxel grid of the unweighted image $\Omega(I_0)$ guarantees the simulated image to align perfectly with the registration target space. The registration of an uncorrected weighted image J is then computed by using the corresponding simulated image $I_{\text{sim}}(b_J, \mathbf{g}_J)$ as *fixed image* as illustrated in Figure 3.2. Here also the full affine transform is optimized while minimizing the NMI Metric between the simulated and the uncorrected moving image. In cases where the b -value shells required for the model-fit are not yet corrected, a pre-alignment of the respective weighted images should be performed, f.i. by the unweighted reference method.

Similarity Methods

A model-free approach for the alignment of higher b -value gradient images was proposed by Zhuang et al. [122]. The idea is to use the contrast similarity of two weighted images from different b -shells with similar diffusion gradient directions \mathbf{g}, \mathbf{g}' . The dMRI signal of white brain matter depends mainly on its orientation relative to the gradient direction \mathbf{g} and will alter proportionally to the angular difference between \mathbf{g} and \mathbf{g}' . As schematically illustrated by Figure 3.3, for an uncorrected weighted image, the weighted image with the smallest gradient angular difference is selected from a lower b -shell as fixed image. Since the eddy currents strength increases with the b -value, a full affine transform minimizing the nMI metric is computed. The method is identical with the unweighted reference method at the lowest b -shell and requires a sequential registration of the different shells to provide an initial alignment of the fixed images and guarantee their alignment with the target registration space of the unweighted image.

Even though both extrapolation- and similarity-based correction approaches improve the insufficient performance of the standard unweighted reference method at higher b -shells, they still can produce a registration outlier, especially in the presence of higher distortions. The similarity-based method relies strongly on a perfect alignment of the lower shells, otherwise an initial alignment error of a selected fixed image will propagate and possibly increase when correcting the higher b -shells.

For the extrapolation-based methods, an approach using improved extrapolation of the two-compartment CHARMED [6] model and thus better simulated reference images was published very recently by Nilsson *et al.* [79]. The authors presented a robust way of fitting the signal fraction of the two tensor compartments of the CHARMED model taking the partial volume effects as well as the CSF compartment signal into account.

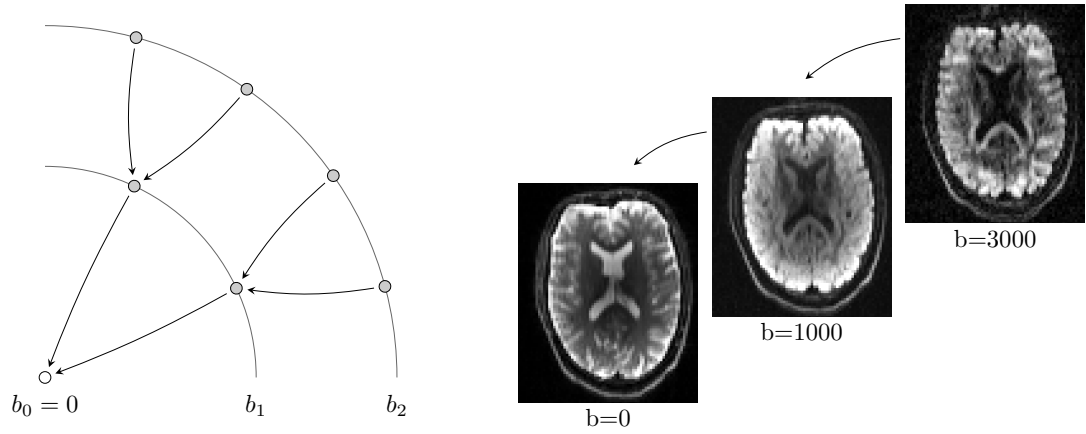


Figure 3.3: Similarity-based correction scheme. The corresponding pairwise registration tasks (solid line arrows) aim from the moving image (arrow tail) to the reference image (arrow head) on a lower b -shell. On the right side, exemplary images at $b=1000, 3000 \text{ s} \cdot \text{mm}^{-2}$ with similar gradient direction and the $b=0$ image are shown.

In a general remark to both retrospective and prospective approaches that aim at estimating and correcting of the motion transform, Leemans *et al.* [65] pointed out the necessity of proper handling of the B -matrix when rotation is applied during motion correction. Since the dMRI data is not rotationally invariant, the diffusion gradient direction \mathbf{g}_i of an image volume must be also corrected with the rotational part of the corresponding transform T_i : $\mathbf{g}_i \rightarrow \text{Rot}(T_i)(\mathbf{g}_i)$.

3.1.4 Discussion

The *rejection* approaches present a simple method for minimizing the influence of distorted gradient volumes on the subsequent data processing. However, this approach is applicable only if a small portion of gradient volumes is rejected and if a balanced sampling of the q -space and hence a uniform distribution of the gradients is maintained afterwards. The requirement for sampling uniformity is fulfilled automatically by the *prospective* approaches, however, the correction is limited by the necessity of either an additional hardware providing the motion measurement or modifications of the acquisition sequences, which are not easily deployed into clinical practice. Furthermore, the prospective approaches address only artifacts caused by bulk motion of the subject and leave other acquisition artifacts (e.g. the ones caused by eddy currents) uncorrected. Due to their ability to effectively compensate both motion and acquisition artifacts, the *retrospective* approaches became the most popular method applied in diffusion-MRI data preprocessing. However, the signal characteristics of higher b -value images that are

required by the more precise higher-order signal models pose a challenging task and the image registration can fail at the correction of such images. This behavior was at first shown for the unweighted reference method and b -values of above $2000 \text{ s} \cdot \text{mm}^{-2}$, but even the extrapolation and the similarity-based retrospective approaches can produce a registration outlier for low SNR conditions or in the presence of acquisition artifacts. The following sections on image registration and optimization in general aim at examining the reasons for this behavior.

3.2 Image Registration

Image Registration (IR) belongs to the fundamental tasks in (medical) image processing as it, in a general formulation, provides a spatial correspondence of two or more images. Such classes of problems arises when a combination or fusion of different data sources is required. The data can be acquired at different times (e.g. longitudinal studies), with different modalities (e.g. CT-MR fusion), with different subjects (e.g. atlas construction for group studies) or from different viewpoints (e.g. in computer-assisted scenarios). The task is – given two images I, J – to find a spatial transformation for one of the images such that the transformed image *matches* the other one. In the usual notation, the transformed image is referred to as *moving* or *template* image, the other one is referred to as *fixed* or *reference* image. Following the notation in [71], an image can be understood as a d -dimensional mapping from a spatial domain set $\Omega \subset \mathbb{R}^d$ to \mathbb{R} . It assigns each spatial point $x \in \Omega$ a *gray value* $b(x)$. The task of image registration is to find a suitable mapping $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that some given (image) distance measure $\mathcal{D} : I \times J \rightarrow \mathbb{R}$ between two images I, J is minimized, i.e.

$$\min_{\varphi} \mathcal{D}(I, J \circ \varphi). \quad (3.1)$$

To reduce the space of possible transforms, φ is typically searched for in some class of transforms which is explicitly restricted by the application. If only rigid movement is possible, for example in the registration of bone CT images, it suffices to search for *rigid transforms*, i.e. $\varphi(x) = Qx + b$ where $Q \in \mathbb{R}^{d \times d}$ is an orthogonal matrix with $\det Q = 1$ describing the rotation and $b \in \mathbb{R}^d$ the translation vector. Additionally, allowing scaling and shearing leads to the class of *linear* (full affine) transforms with $\varphi(x) = Ax + b$ where $A \in \mathbb{R}^{d \times d}$ is a real valued matrix with $\det A > 0$. In case of 3D images, the full affine transformation has 12 Degrees of Freedom. φ can also be restricted implicitly as it is for the class of deformable transforms. For meaningful transforms either smoothness ($\varphi \in C^2(\mathbb{R}^d)$) or even diffeomorphism is required.

The general definition (3.1) combines the two key components of image registration. First, the similarity measure \mathcal{D} , also called *image metric* and second, the *optimizer* finding the metric's minimum $\mathcal{D} \rightarrow \min$ have to be specified. These two components will be explained in detail in the following sections.

3.2.1 Image Metrics

Image registration approaches can be categorized by multiple facets. For the definition of image metrics, the most important differentiation is between *feature-* and *intensity-based* methods. The former extract only some information from the images and define the metric on the basis of a distance measure of the selected features, for example the Euclidean distance. The latter ones take the whole image information into account.

One of the simplest similarity measures is the sum of squared (intensity) differences:

$$\text{SSD}(I, J) = \frac{1}{N} \sum_{x \in \Omega(I, J)} |I(x) - J(x)|^2. \quad (3.2)$$

A more stable measure allowing linear relationship between the intensity values is the *normalized correlation coefficient*

$$\text{CC}(I, J) = \frac{\sum_{x \in \Omega(I, J)} (I(x) - \bar{I}) \cdot (J(x) - \bar{J})}{\left[\sum_{x \in \Omega(I, J)} (I(x) - \bar{I})^2 \cdot \sum_{x \in \Omega(I, J)} (J(x) - \bar{J})^2 \right]^{0.5}} \quad (3.3)$$

Both metrics are only suitable for registration tasks of images from the same imaging source (CT, MR, etc.). In multi-modal medical image settings, the dependencies between the intensities are typically non-linear and the CC cannot reflect this relationships. This leads to the introduction of entropy based image metrics, with Mutual Information (MI) being the most prominent among them.

Mutual Information

The consideration of image registration as a task of maximizing the amount of shared information in two images leads to the usage of *information* as a similarity metric [45]. The most common definition of information measure goes back to Shannon's entropy

$$H = - \sum_i p_i \log p_i,$$




Figure 3.4: Visualization of the Mutual Information metric in ideal case. For two aligned images with similar intensities and low noise levels the values of the MI metric are densely sampled for the translation and rotation along the optimal alignment position.

initially defined in information theory for a set of symbols with associated probabilities p_j . In the context of image registration, the probabilities p_i translate directly into the distribution of different voxel intensities (histogram). Entropy $H(X)$ measures the uncertainty of a random variable X . In the presence of another variable Y the conditional entropy $H(X|Y)$ measures the uncertainty of X when Y is known (to avoid a symbol collision with the information $I(\cdot)$, this paragraph will use the notation X, Y instead of I, J for the images). Hence, the difference

$$I(X, Y) = H(X) - H(X|Y)$$

expresses the amount of *information* Y provides about X . The fact that X and Y can be interchanged leads to the term *mutual information*. An equivalent definition arises when considering the joint entropy of X and Y :

$$H(X, Y) = - \sum_x \sum_y p_{XY}(x, y) \log p_{XY}(x, y)$$

instead of the conditional leading to:

$$I(X, Y) = H(X) + H(Y) - H(X, Y)$$

Both, the conditional and the joint entropy can be computed only for the overlap and the measure as defined above depends on the size and the contents of the overlapping region [88]. To overcome the sensitivity of the basic mutual information metric to the overlap of X and Y , a more robust version, the normalized Mutual Information (nMI)

metric was proposed [106]:

$$I(X, Y) = \frac{H(X) + H(Y)}{2H(X, Y)} \quad (3.4)$$

3.2.2 Registration Optimizer

The optimization task in image registration is given by the equation (3.1) which can be reformulated to

$$\hat{\phi} = \arg \min_{\phi} \mathcal{D}(I, J \circ \phi) \quad (3.5)$$

For non-parametric registration techniques which do not employ a parametric model of the transformation (or deformation), the metric \mathcal{D} must be extended by some regularization term $\mathcal{R}(\cdot)$ to ensure smoothness of $\hat{\phi}$ [60] which reads on an example of the Mutual Information (MI) metric as follows:

$$\mathcal{D}(I, J \circ \phi) = -\text{MI}(I, J \circ \phi) + \omega \mathcal{R}(\phi).$$

In parametric registration, the optimal transform $\hat{\phi}$ is determined through an iterative strategy

$$\phi_{k+1} = \phi_k + a_k \mathbf{d}_k, \quad k = 0, 1, 2, \dots \quad (3.6)$$

with the *search direction* \mathbf{d}_k at iteration k , and a scalar *gain* factor a_k modulating the step size along the search direction. Successful optimization selects the search directions and gain factors such that the sequence $\{\phi_k\}$ converges to a local minimum of the cost function \mathcal{D} . The different methods differ in the kind of modulation of the search directions and the gain parameter.

The class of *gradient descent* optimizers follows the direction of the negative gradient of the cost function

$$\phi_{k+1} = \phi_k - a_k \cdot \nabla \mathcal{D}(\phi_k) \quad (3.7)$$

with a_k being the critical parameter for hitting the minimum. Too small a_k will slow-down the optimization process while too large values will end up in an oscillation behavior and thus a smart strategy for an adaptive choice of a_k is required. The *regular step* gradient descent starts with relatively large a_0 and reduces the value if the gradient direction changes by more than 90° which means that the searched minimum is placed somewhere between the current and the previous position. The *line-search* gradient descent tries to estimate the best value for a_k by performing an iterative search along

the gradient's direction

$$a_k \leftarrow \min_{\alpha} \mathcal{D}(\varphi_k + \alpha d_k).$$

Even more sophisticated adaptive strategy is employed by the *adaptive stochastic* gradient descent [59].

Other classes of optimizers differ from the gradient descent approaches by considering different descent directions d_k . In addition, the *conjugate gradient* optimizer computes the conjugate gradient direction s_k and searches for the minimum of \mathcal{D} in an optimization subspace:

$$d_k = -\nabla \mathcal{D}(\varphi_k) + \beta_k s_{k-1}.$$

The *Newton* optimizer takes the second order Taylor expansion of $f(\varphi_k)$ into account and considers the reciprocal of the quadratic term, i.e. the Hesse-matrix $\mathcal{H}_k = \mathbf{H}\mathcal{D}(\varphi_k)$:

$$d_k = -[\mathcal{H}_k]^{-1} \nabla \mathcal{D}(\varphi_k).$$

The derived *quasi-Newton* methods use an iterative approximation B_k of $[\mathcal{H}_k]^{-1}$ rather than computing a direct estimate of \mathcal{H}_k itself:

$$d_k = -B_k \nabla \mathcal{D}(\varphi_k).$$

Multi-resolution Registration

The common image registration metrics provide a precise guidance for the local search optimization. However, they suffer from a limited attraction radius and an initial misalignment of the images can potentially exceed it. To minimize this possibility, the image registration is carried out in a multi-resolution manner. The so-called *pyramid*




Figure 3.5: Image pyramid for multi-resolution registration. A single axial slice from a T_2 -weighted image is shown at different resolution (pyramid) levels from left to right. The down-sampling factors are shown in the bottom-left corner, the rightmost image shows the full resolution.

image registration approach considers a set (a pyramid) of downsampled copies of the moving and fixed images (example given in Figure 3.5), by choice coupled with Gaussian smoothing [68, 88]. In *multi-scale* approaches, the image pyramid is built only by Gaussian smoothing without downsampling. At the most coarse level, a rough estimate of the registration transform φ can be quickly determined and refined at the finer levels of the image pyramid. In addition, a reasonable initial estimate can reduce the computational effort on finer scales. However, the multi-resolution approach must be reflected by the optimizer by adapting the relevant parameters (e.g. the maximal step size) at each resolution level.

A comprehensive list of related literature can be found in several review studies, either on image registration in general [123], for the field of medical image registration [81, 103], or on a specific class of problems, e.g. mutual-information-based registration [88, 60].

3.2.3 Discussion

Image registration is a local optimization method capable of providing robust image alignment if the “boundary” conditions are fulfilled. These are: a pronounced global minimum of the image registration metric at the position of perfect alignment and by possibility no distinct local minima at other positions, the starting position for the registration located within the attraction area of the metric’s global minimum and a configuration of the optimizer allowing to reach the position of the global minimum. The global minimum condition on the image metric is fulfilled in almost all real-world image registration tasks, the attraction area and the configuration condition can be addressed by employing a pyramid (coarse-to-fine) optimization strategy with adaptive settings at each pyramid level. The local minima condition is the most vulnerable one. In retrospective correction approaches, it can be violated especially at low SNR levels like in higher b -value gradient images, or in cases where the metric cannot fully describe the image characteristics like in the gradient-dependent signal.

3.3 Evolutionary Optimization

An optimization problem is a search for a minimum of a function f in the search (or optimization) space Ω

$$\min_{x \in \Omega} f(x) \quad (3.8)$$

possibly with some additional constraints of the type $g_i(x) \leq b_i$. In general, such problems are difficult to solve, except for special classes of problems, where the optimal solution x^* can be obtained analytically (for example least-squares problems), or at least efficiently (for example convex optimization or linear programming problems). To provide at least a high-quality solution, i.e. not necessarily the optimal one, but near to the true optimum, in case none of the efficient methods is applicable, several *heuristics* and *meta-heuristics* were developed.

A heuristic approach tries to iteratively improve an initial solution guess, either in a constructive manner (greedy algorithms) or by means of a local search in the optimization space (hill-climbing, gradient methods etc.). The traditional image registration optimization belongs to the class of local search approaches. Meta-heuristics can be understood as standing on a level above the heuristic approaches since they are strategies that actively *guide* a local search method, or a heuristics in general. They are often not problem-specific and permit an abstract description. A meta-heuristic can act on a single point during the local search, controlling its trajectory through the search space. This class of *trajectory methods* includes approaches like simulated annealing, tabu search or simple evolutionary strategies. Or a meta-heuristics can employ a population of points in the search space and control their *evolution* during the optimization process. The latter, also called *evolutionary optimization* approaches, will be detailed in the following paragraphs.

Population-based Approaches

The underlying idea common to all population-based approaches is the "survival of the fittest" principle, inspired by the natural selection, which causes a rise in the global fitness of a developing population. The natural process is mirrored in selecting some of the fitter population members and employing *recombination* and *mutation* operators in order to create a new generation of candidate solutions. In optimization, each member of the population represents a solution candidate and has an associated fitness value measured by a fitness function, often the optimization function $f(\cdot)$ itself. Recombination takes place by allowing two selected members to produce an offspring, whereas mutation alters a single candidate in a random manner. Hence, the fundamental

forces in evolutionary optimization are the *selection* that increases the population fitness and the *variation* (recombination, mutation) that maintains the necessary diversity in the whole population. The most prominent population-based evolutionary methods are Genetic Algorithms (GAs), Evolutionary Algorithms (EAs) and Particle Swarm Optimization (PSO) approaches. Some of the applications of EA and GA can be found in review articles [52, 28], the PSO will be introduced in detail in the next paragraph.

3.3.1 Particle Swarm Optimization

Similar to EA and GA the meta-heuristics PSO, originally proposed by Kennedy and Eberhardt [57], uses a population of particles. In contrast to EA, no explicit recombination and mutation operators are used. Using the term swarm is motivated by the nature's example of fish or bird swarms and the dynamics of their collective motion and their social behavior, where the individual members profit from the discoveries and previous experience of all other members [57]. An animal swarm has no designated coordinator, but each individual member interacts with its local neighbors by following their movement while avoiding collisions. To mimic this behavior, the PSO considers along with the particle swarm $\{p_1, \dots, p_N\}$ a set of associated velocities $\{v_1, \dots, v_N\}$ which are both iteratively updated while taking the swarm's experience on the optimization function $f(\cdot)$ into account. Each particle tracks its historical best position $p_{\text{best},j}$ with

$$f(p_{\text{best},j}) \leq f(p_j^{(t)}), t = 0, \dots, T$$

and further the swarm maintains its global best particle p_{gbest} with

$$f(p_{\text{gbest}}) \leq f(p_j^{(t)}), j = 0, \dots, N.$$

At each iteration $t \rightarrow t + 1$ the velocity of each particle p_i is updated according to the *velocity update equation*:

$$\begin{aligned} v_i^{(t+1)} = & \omega v_i^{(t)} + c_1 \cdot r_1 \cdot (p_{\text{best},i} - p_i^{(t)}) \\ & + c_2 \cdot r_2 \cdot (p_{\text{leader}} - p_i^{(t)}) \end{aligned} \quad (3.9)$$

The position of the particle is updated through

$$p_i^{(t+1)} = p_i^{(t)} + v_i^{(t)} \quad (3.10)$$

The velocity equation (3.9) contains several parameters influencing the behavior of PSO significantly. The inertia weight ω controls the impact of the history of previous velocities and can allow faster direction changes. The learning factors c_1, c_2 control the attraction of a particle to its own success and to the neighboring particles, respectively. They are scaled by the uniform random factors $r_1, r_2 \in [0, 1]$. The *cognitive* factor c_1 represents the attraction towards the historical best position, the *social* factor c_2 represents the attraction to the leading particle. The selection of the leading particle p_{leader} modulates the cognition range of each particle, also called the neighborhood *topology* of the PSO [92]. The main approaches consider either all particles connected, i.e. it is $p_{\text{leader}} = p_{\text{gbest}}$, or take only the k -nearest neighbors into account and set $p_{\text{leader}} = p_{\text{localbest},k}$. The leader selection will become even more important in the extension of PSO to multi-objective optimization problems later in this chapter.

The application area for PSOs is vast as documented in the comprehensive survey of different PSO approaches given recently by Zhang *et al.* [120].

3.3.2 Multi-objective Approaches

Real-world optimization problems typically have to satisfy multiple, and often competing cost functions (e.g. project duration and project costs). The problem of finding an equilibrium leads to the class of *multi-objective* optimization approaches. In general, the multi-objective (MO) problem is defined as a minimization (maximization) of

$$F(x) = (f_1(x), \dots, f_k(x))$$

possibly subject to some constraints $g_i(x) \leq b_i$.

A straightforward approach is to reduce the problem dimensionality by *aggregating* the multi-objective function $F(\cdot)$ to a scalar-valued $F'(\cdot)$, the simplest example here is the weighted sum

$$F'(x) = \sum_i w_i f_i(x)$$

of the multiple objectives. The *lexicographic* approach avoids the usage of aggregating functions, by sorting the single objectives $f_i(\cdot)$ by their importance and optimizing them by first finding the optimal solution x_1^* for the most important one and then consecutively optimizing the others while using constraints fixing values for the already optimized $f_1(x) = f_1(x_1^*), \dots, f_j(x) = f_j(x_j^*)$. Aggregating approaches are both easy to implement and understandable while allowing to employ standard search-based algorithms to detect the optimum. However, they require *a priori* knowledge

in determining either appropriate weights w_i or proper ordering, respectively [23]. Population-based meta-heuristics (MHs), on the contrary, seem to be particularly suited for MO problems as their straightforward extensions directly allow realization of the *a posteriori* approaches, where first a search as widespread as possible is performed and the decision process is made afterwards.

One class of prominent *a posteriori* approaches is based on the usage of *sub-populations*, each one assigned to optimization of the single-objective $f_i(\cdot)$. The main challenge in such cases is the information exchange between the sub-populations to progress towards the global optimum. One of the first published approaches, the Vector-evaluated Genetic Algorithm (VEGA)[99] shuffles all sub-populations together to apply the usual GA's crossover and mutation operators.

An important term in the fitness estimation of different solution candidates is the *Pareto dominance*. A candidate solution x is said to *dominate* another candidate y if and only if moving from y to x will improve one objective while staying at least same in the remaining, or formally $f_i(x) \leq f_i(y) \wedge \exists i f_i(x) < f_i(y)$. The global optimum is *per sé* a non-dominated point. The set of all non-dominated points is called the *Pareto set*, the set of objective functions values over the Pareto set is called the *Pareto front*. The Pareto Sampling techniques complete the classification of the *a posteriori* approaches. Inspired by the VEGA, Parsopoulos *et al.* [84] proposed a MO-variant of PSO—the Vector-evaluated Particle Swarm Optimization (VEPSO). The communication strategy between the single-objective swarms is realized through sharing the best experience, the p_{gbest} . However, the best particles are estimated only locally in each swarm, i.e. only with respect to the single objective function. The Pareto-based methods search for a global leader in all non-dominated points across the whole swarm, evaluating all objective functions for each particle. The different variants of the Pareto-based methods are described in the comprehensive overview of both multi-objective Evolutionary Algorithms and Genetic Algorithms by Coello *et al.* [23] or in a more recent one by Zhou *et al.* [121]. Multi-objective PSO approaches were reviewed by Reyes-Sierra and Coello [92] and by Parsopoulos and Vrahatis[84].

3.3.3 Discussion

The main strength of EAs lies in their ability of global search space exploration. In this way, they address the limitations of local methods in the presence of local minima of the optimization function. Their application to image registration problems was examined in several configurations. Das and Bhattacharya [27] showed superiority of an original

PSO method over a GA-based registration in their comparative study for 2D CT-MR registration. The authors used a multi-resolution approach in both settings. However, they presented only a qualitative (visual) comparison of the registration results. In a direct application to intensity image registration, the main disadvantages of basic PSO or other meta-heuristic approaches are that they tend to perform worse compared to straight local search methods, are more difficult to configure and not efficient in computing local optima with high precision [112].

3.4 Memetic Optimization Approaches

To increase the effectiveness and efficiency in locating the Pareto front, hybridized population-based meta-heuristics incorporating generic local search techniques to improve the candidate solutions were proposed. Moscato [73] introduced the notation of *memetic algorithms* for such approaches, inspired by Dawkins' "meme" [29], to express a unit of evolution that can individually adapt the information it transfers. Petalas *et al.* [86] proposed a memetic version of PSO and presented an evaluation on a large set of test problems, including constrained and unconstrained, minmax and integer optimization problems. The authors conclude that, from an optimization point of view, memetic algorithms have been shown both to require orders of magnitude fewer evaluations to find optima (i.e. they are much more efficient) and to identify higher-quality solutions (i.e. they are more effective).

Application to Image Registration

To overcome the limitations of pure EAs in the application to image registration, hybrid approaches have been introduced recently. Lin *et al.* [67] proposed a hybridization of an original PSO through incorporation of two GA concepts—sub-population and cross-over operators—in 3D rigid image registration. The crossover operation is performed only with the two best-ranked (with respect to the MI cost function) sub-population leader. In the feature-based registration Santamaría *et al.* [97] proposed an approach based on *scatter search*, an evolutionary optimization method which is based on the original GA. However, they use a systematic combination between the solutions instead of random one.

Following the idea of memetic algorithms, a class of sequentially hybridized approaches tries to combine the advantages of both local and global approaches, e.g. by adapting a two-step procedure with a global initialization and a local refinement step: Kagadis *et al.* [55] used a GA hybridized in the last optimization stage by a local search

with a multi-dimensional Powell method for volumetric- and surface-based registration of CT and SPECT images and extracted surfaces, respectively. He and Narayana [47] employed a GA in the continuous space in combination with dividing rectangle, a deterministic local method, for the optimization of MI in rigid-body registration of brain MR images. They realized the image registration in multi-resolution manner, applying the GA only at the coarsest level to estimate starting positions for the local method at the finer resolution levels. A similar multi-resolution hybrid approach was presented by Xu *et al.* [116], where a differential GA is used at the coarsest resolution level to provide initial solutions for the Powell's method employed at further resolution pyramid stages. Du *et al.* [30] utilized simulated annealing (SA) at the coarsest resolution level for the estimation of good initial positions for the subsequent computation by Powell's local search. A pure memetic method for (feature-based) image registration was first proposed by Santamaría *et al.* [98] where the authors embed the local search into each EA iteration, instead of using the EA only for global initialization, showing an increased performance when compared to the sequential hybridized EAs.

In traditional image registration, a weighted aggregation approach was already used in the formulation of a multi-metric registration problem by Avants *et al.* [9, 8]. In the proposed case of fixed weights, only a single non-dominated point is optimized per optimization run. First application of population-based meta-heuristics to intensity-based image registration goes back to Wachowiak *et al.* [112]. The authors propose and evaluate a simple PSO with several adaptations in the multi-modal settings using the MI metric and concluding the hybrid PSO using sub-populations to achieve the best results. In their recent review article, Valsecchi *et al.* [110] summarize the EA approaches in the intensity-based image registration, whereas Santamaría *et al.* [96] provided a comparison of EA-based methods in feature-based image registration.

3.5 Conclusion

The three main approaches in diffusion-weighted image preprocessing were introduced in this chapter. The *rejection* approaches that discard corrupted gradient images from further processing, the *prospective correction* approaches, which extend the acquisition set-up by implementing a detection of motion corruption in order to trigger a re-acquisition and the *retrospective correction* approaches that employ image registration methods to correct for the acquisition artifacts afterwards. The latter one can correct for both motion and eddy currents artifacts, while not losing image information essential for precise evaluation of higher-order models through removal of uncorrected images.

For this reason, the retrospective approaches became the most common way in diffusion MRI preprocessing.

However, since the retrospective methods are based on image registration techniques, the correction precision relies on the capabilities of a local search optimization to locate the global optimum. These can be impaired e.g. due to lower SNR levels or due to less precise image metric, as can be seen for the unweighted reference method when matching a gradient image at higher b -value to the unweighted image. The extrapolation reference method aims at improving the image metric performance by mimicking the gradient-dependent contrast with simulated reference images. The similarity reference method achieves a metric improvement by using the signal similarity of two weighted images from different shells and with similar gradient directions. But even these metrics do not remain completely outlier-free. However, due to the differences in the image metrics associated with the three reference methods, the conditions for a registration outlier are likely also to be different. Thus, a multi-objective formulation could provide a robust registration approach for the correction of diffusion-weighted images.

The global optimization approaches like *GAs* or *PSO* firstly provide great capabilities in global space exploration and secondly can be extended to multi-objective problems in a straightforward manner. But, when applied to image registration they are not as precise as the local methods in localizing the global optimum. A local method (e.g. image registration) will most probably reach the global optimum when provided with a starting position near in its proximity.

However, a hybrid optimization approach that combines both global and local methods, can take advantage of the complementary abilities of the global method's search space exploration while allowing a precise localization by the local one.

Chapter 4

Methods

The key objective of this thesis is to develop a robust retrospective registration method for correction of motion and acquisition artifacts in diffusion MRI data. The main target in the method's design is to minimize registration outliers by better exploiting the different intensity similarity relationships between the individual volumes in a diffusion-weighted image. The contribution splits into three main parts:

- (A) Design of a meta-heuristic optimization scheme based on the memetic principle, combining local search phases following multiple objective functions with global PSO updates to overcome problematic areas in the search space.
- (B) Construction of realistic synthetic diffusion-weighted datasets suitable for quantitative evaluation of retrospective correction approaches in diffusion MRI.
- (C) Quantitative and qualitative evaluation of the proposed method (A) on both realistic synthetic (B) and real in-vivo data.

Outline

This chapter is organized as follows: the design of the retrospective correction method (*contribution A*) is presented in Section 4.1. The construction of validation datasets (*contribution B*) is described in Section 4.2, where first the simulation of synthetic datasets with motion and eddy-currents-caused artifacts is described (4.2.1) and second the assembly of evaluation datasets for the in-vivo experiments is explained (4.2.2). The remaining sections address different aspects of the evaluation (*contribution C*). Section 4.3 introduces two specific measures used for performance evaluation of the proposed correction approach and the last two sections provide a full description of

the evaluation data and the experiments settings for the synthetic (Section 4.4) and the in-vivo datasets (Section 4.5).

4.1 Retrospective DWI Distortion Correction

In this section, the multi-objective memetic method for diffusion-weighted Image (DWI) distortion correction is proposed. First, the objective functions used in the multi-objective optimization are defined (4.1.1), followed by the description of the memetic optimization scheme in general (4.1.2) and its adaptation to the specific case of pyramidal image registration (4.1.3) and lastly the application of the optimization scheme to dMRI data correction is introduced (4.1.4).

4.1.1 Objective Functions

This section formalizes the objective functions used in the proposed multi-objective scheme. For the formalization, a fixed registration transform space \mathcal{T} , a moving image \mathcal{J} , a fixed image \mathcal{I} and an image similarity metric $\mathcal{D}(\mathcal{I}, \mathcal{J} \circ \varphi)$ as defined in (3.1) are considered. The diffusion-weighted image is handled as a set of 3D images $\{I_0, I_1, \dots, I_N\}$ as defined in (2.14) with I_0 explicitly denoting the unweighted image. The domain of I_0 defines the reference (i.e. the target registration) space.

The retrospective head motion and artifact correction can be performed on the basis of several alternative registration settings described in Section 3.1.3. These settings differ in the selected function assigning a fixed image $\mathcal{I} \in \{I_0, I_1, \dots, I_N\}$ to each weighted image $\mathcal{J} \in \{I_1, \dots, I_N\}$.

The **unweighted reference** method registers all weighted images to the unweighted one, hence $\mathcal{I} = I_0$ for all gradient images. The corresponding objective function

$$\mathcal{D}_{NW} = \mathcal{D}(I_0, \mathcal{J} \circ \varphi) \quad (4.1)$$

will be further referred to as *non-weighted reference* (NW) objective. The transform estimated by registration under the minimization of \mathcal{D}_{NW} will be denoted as φ_{NW} , i.e. it holds

$$\varphi_{NW} = \arg \min_{\varphi \in \mathcal{T}} \mathcal{D}_{NW}, \quad \mathcal{J} \in I_1, \dots, I_n,$$

The NW objective is depicted in Figure 4.1a.

The **extrapolation method** uses a signal model $M(b, \mathbf{g})$ to compute a simulated fixed image $\mathcal{M}(J) = M(b(J), \mathbf{g}(J))$ for each uncorrected weighted volume J that fits its gradient direction $\mathbf{g}(J)$ and b -value $b(J)$. The associated *model-based reference* (MB) objective function is defined as

$$\mathcal{D}_{\text{MB}} = \mathcal{D}(\mathcal{M}(\mathcal{J}), \mathcal{J} \circ \varphi) \quad (4.2)$$

and the estimated transform is denoted φ_{MB} . Since the model \mathcal{M} resides *per definition* in the I_0 space, the φ_{MB} provides an alignment with the target registration space automatically. The MB objective is depicted on Figure 4.1b.

The **similarity method** selects the fixed image from a lower b -shell which has the most similar gradient direction with \mathcal{J} . This selection can be generalized in form of a neighborhood operator \mathcal{N} acting on the separate gradient images. This, with \mathcal{N} , the definition of the *nearest-neighbor reference* (NN) objective is as follows:

$$\mathcal{D}_{\text{NN}} = \mathcal{D}(\mathcal{N}(\mathcal{J}), \mathcal{J} \circ \varphi). \quad (4.3)$$

The selected fixed image $\mathcal{N}(\mathcal{J})$ does not necessarily align with the target registration space, but provides this alignment over a transform $\varphi_{\mathcal{N}(\mathcal{J})}$ which must be considered in the definition of the estimated final transform φ_{NN} :

$$\varphi_{\text{NN}} = \varphi_{\mathcal{N}(\mathcal{J})} \circ \arg \min_{\varphi \in \mathcal{T}} \mathcal{D}_{\text{NN}}.$$

The transform and the NN objective function are illustrated by Figure 4.1c. All three objective functions will be used in the optimization scheme which is described in the next section.

4.1.2 Optimization Scheme

Memetic methods are hybrid optimization approaches which extend some selected population-based meta-heuristics by a *local search* method to improve the positions provided by the *global* optimization. The proposed memetic optimization scheme alternates a global phase based on a PSO with the local search phase realized by a "standard" pairwise image registration as schematically depicted in Figure 4.2. The search space is defined by the selected transform space \mathcal{T} and each particle represents a position in the search space, i.e. an individual transform $\varphi \in \mathcal{T}$.




Figure 4.1: Visualization of the fixed image assignment for the objective functions on the example of a high shell image. The moving image \mathcal{J} is emphasized in red, the assigned image \mathcal{I} in blue. The solid red arrow marks the pairwise image registration in each sub-figure, φ_{NW} , φ_{MB} and φ_{NN} denote the transforms aligning the selected moving image with the target registration space of the unweighted image. $\mathcal{M}(\mathcal{J})$ denotes the extrapolation of the simulated images and $\varphi_{N(\mathcal{J})}$ the pre-alignment transform of the selected fixed image to the target registration space.

Local Search

Motivated by the sub-populations in the *VEPSO* approach (cf. Section 3.3.2), the particles are organized in separate swarms during the local search phase, one swarm per objective function. The individual progression of each particle is guided by a single-objective

pairwise image registration which is fully defined by the swarm-specific objective function \mathcal{D}_i introduced in the previous section.

Global Update Step

The global update step is designed in a PSO manner. In contrast to the *VEPSO* approach, which shares information across the swarms only through the locally estimated swarm leaders, all particles from all swarms are considered as members of a single swarm during the global update. This design allows a direct comparison of particles' fitness values and thereby selection of the leading particle(s) regardless of their swarm association. In multi-objective registration settings, the image metrics are not competing against each other but rather cooperating to detect the global minimum \mathbf{x}^* . It is common, and without loss of generality, to consider the registration as a minimization problem. A well-designed image metric such as nMI will also have a global minimum at \mathbf{x}^* . However, it is not guaranteed for a local search to reach \mathbf{x}^* , the search can get trapped in some local minimum inside the attraction radius of the global optimum or even leave the attraction area completely. In such cases one speaks of a *registration outlier*.

Let us consider the registrations guided by the single metrics $\mathcal{D}_1, \dots, \mathcal{D}_m$ to stop at points x_1, \dots, x_m respectively. Then each metric \mathcal{D}_i will have at least a local minimum at x_i . The *quality* (the Pareto non-dominance) of the estimated minimum is assessed by cross-evaluating all other objectives $\mathcal{D}_{j \neq i}$ at x_i . A position x_i reached by the local search is more likely only a local minimum if it is dominated by some other position,




Figure 4.2: Schematic depiction of the memetic search scheme for three objectives \mathcal{D}_1 , \mathcal{D}_2 and \mathcal{D}_3 depicted by different line styles. During the local search, there is one swarm per objective function (marked by small ellipses) and each particle follows its corresponding single objective. In the global PSO step, particles are cross-evaluated with multiple objectives and the leading particles are chosen (magenta). In this example, $N = 6$ global particles and $k = 2$ leaders are chosen.

more precisely if there exists a x_j such that

$$\mathcal{D}_i(x_j) < \mathcal{D}_i(x_i), \quad j \neq i.$$

For the global optimum it is

$$\mathcal{D}_j(\mathbf{x}^*) \leq \mathcal{D}_j(x_i), \quad \forall j < m.$$

This consideration leads to the following definition of the Fitness Value (FV) for a particle's position p_j :

$$\text{FV}_j = \text{FV}(p_j) = \|(\mathcal{D}_1(p_j), \dots, \mathcal{D}_m(p_j))\|_1. \quad (4.4)$$

The 1-norm is chosen to allow for an effective outlier detection.

The employed PSO is the standard PSO as defined in Section 3.3.1, using the velocity update equation (3.9) and the position update equation (3.10). One additional parameter k defines the number of leading particles that are considered in the update step. Ranking of the particles is estimated by their FV. Usually, the initial position of the images in a head motion and artifact correction setup are not far away from the global optimum (compared to the size of the search space). This allows to set the starting positions for the particles randomly around the identity transform following the normal distribution

$$p_i^{(0)} = id_{\mathcal{T}} + \mathcal{N}(0, \sigma(N)), \quad i = 1, \dots, N \quad (4.5)$$

with an increasing variance $\sigma(N)$ for increasing number of particles N in order to maintain diversity among the initial positions. A description of the complete optimization scheme is given in Algorithm 1.

4.1.3 Pyramid Image Registration Scheme

In the previous section, the proposed optimization scheme was defined in a general form with an arbitrary registration procedure used for the local search. It remains to specify at which points the global phases are employed during the optimization. For higher robustness, image registration approaches typically follow the multi-resolution, coarse-to-fine principle (also called *pyramid* registration) to increase the attraction radius or to obtain a good initial solution (see Section 3.2). In pyramid registration, the final result on one level $x^{(l)}$ is passed as a starting point for the registration to the next one. Thus, the pyramid level changes provide natural points for performing the global


```

Data:  $N, k, \mathcal{T}, \mathbf{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_m\}$ 
Result:  $\hat{\phi} \in \mathcal{T}$ 

begin
  Initialize particles  $\mathbf{P}^{(0)} = \{p_1, \dots, p_N\}$ 
  Initialize objective swarms  $S_i^{(0)}, i = 1, \dots, m$  with  $\mathbf{P}^{(0)} = S_1^{(0)} \dot{\cup} \dots \dot{\cup} S_m^{(0)}$ 
  while  $n < \text{max iterations}$  do
    /* Perform local search for each objective */
    foreach  $\mathcal{D}_k \in \mathbf{D}$  do
      /* Pairwise registration for each particle in swarm */
      foreach  $p_j \in S_k^{(n)}$  do
         $p'_j \leftarrow \mathcal{D}_k(I, J, p_j)$ 
      end
    end

    /* Perform global update with single global swarm  $\mathbf{P}'$  */
     $\mathbf{P}' = \{p'_1, \dots, p'_N\}$ 
    foreach  $p'_j \in \mathbf{P}'$  do
      /* fitness value */
       $FV_j < -FV(p_j)$  // using equation (4.4)
    end
    Rank particles by  $FV_j$  and select the  $k$  leaders
    foreach  $S_i, i = 1, \dots, m$  do
      foreach  $p_{\text{leader},j} \quad j = 1, \dots, k$  do
         $v_j^{(n)} \leftarrow \text{velocity update with } p_{\text{leader},j}$  // Equation (3.9)
         $p_j^{(n+1)} = p_j^{(n)} + v_j^{(n)}$  // Position update, equation (3.10)
      end
      /* Collect updated particles in objective's swarm */
       $S_i^{(n+1)} \leftarrow p_j^{(n+1)}$ 
    end
     $n \leftarrow n + 1$ 
  end
   $\hat{\phi} = \text{best ranked } (\hat{p} \in \mathbf{P}')$ 
end

```

Algorithm 1: MultiObjectiveMemeticSearch - a pseudo-code of the multi-objective memetic PSO optimization scheme.

update step. The number of resolution levels L must be set so that the coarsest level still provides sufficient information for reliable metric computation. $L = 3$ is typically selected for head MRI registrations, some high-resolution images may allow even the choice of $L = 4$. The adapted memetic optimization scheme to the pyramid image registration is described in Algorithm 2.

```

Input: fixed image  $I$ , moving image  $J$ , objectives  $\mathbf{D}$ 
Output: registered image  $J'$ 

begin
  Estimate maximal pyramid level  $L, l = 1$ 
  Compute image pyramids  $\{I^{(1)}, \dots, I^{(L)}\}, \{J^{(1)}, \dots, J^{(L)}\}$ 
  Initialize particles  $\mathbf{P}^{(0)} = \{p_1, \dots, p_N\}$ 
  Initialize swarms  $S_i^{(0)}, i = 1, \dots, m$ 
  while  $l < L$                                      // For all pyramid levels
  do
    foreach  $\mathcal{D}_k \in \mathbf{D}$                              // Local search for each objective
    do
      Compute objective's fixed image  $\mathcal{I}^{(l)} \leftarrow \mathcal{D}_k[I^{(l)}]$ 
      foreach  $p_j \in S_k^{(n)}$                              // Pairwise registration for each particle
      do
         $p'_j \leftarrow \mathcal{D}_k(\mathcal{I}, J^{(l)}, p_j)$ 
      end
    end
    Rank particles by  $FV_j$  and select the  $k$  leaders           // see Algorithm 1
    Global update  $S_i^{(n)} \rightarrow S_i^{(n+1)}, i = 1, \dots, m$    // see Algorithm 1

     $l = l + 1$                                            // increase pyramid level
    Adapt registration optimizer parameters
  end
   $J' \leftarrow J \circ \hat{\phi}$                                // Apply the estimated transform to the input moving image
end

```

Algorithm 2: MemeticPyramidRegistration - Realization of the multi resolution pyramid image registration as memetic optimization scheme (cf. Algorithm 1).

4.1.4 Information Flow and Parametrization

With the memetic search-based pyramid registration introduced in the previous section the proposed head motion and artifact correction in diffusion-weighted images can be finalized. The diffusion-weighted image I_{DW} will be considered a multi-shell image (as introduced in Section 2.2) with s different weighted shells, i.e.

$$I_{\text{DW}} = I_0 \cup B_1 \cup \dots \cup B_s$$

The target registration space is set to the space of full-affine transforms $\mathcal{T} = \mathcal{Aff}$ in accordance with the correction approach by Mohammadi *et al.* [72] introduced in Section 3.1.3. For the multi-objective formulation all objective functions introduced in

```

Input: Diffusion-weighted image  $I_{\text{DW}}$ 
Result: Corrected diffusion-weighted image  $I_{\text{DW}}^{(\text{C})}$ 
Data:  $N, \mathcal{T} = \mathcal{A}\text{ff}$ 

begin
  Objectives  $\mathbf{D} \leftarrow \{\mathcal{D}_{\text{NW}}, \mathcal{D}_{\text{MB}}, \mathcal{D}_{\text{NN}}\}$ 
  Split image into shells  $I_{\text{DW}} \rightarrow I_0 \cup B_1 \cup \dots \cup B_s$ 
  // Lowest shell pre-alignment
  foreach  $J \in B_1$                                      // register each weighted image
  do
    // see Algorithm 2
     $J' \leftarrow \text{MemeticPyramidRegistration}(I_0, J, \{\mathcal{D}_{\text{NW}}\}, N = 1)$ 
  end
   $B'_1 \leftarrow \{J'_1, \dots, J'_{n_1}\}$ 
   $\mathcal{M} \leftarrow M(b, I_0 \cup B'_1)$                                      // model fit
   $I_{\text{DW}}^{(\text{pre})} \leftarrow I_0 \cup B'_1 \cup \dots \cup B_s$ 

  foreach  $i = 1, \dots, s$                                      // for each weighted shell
  do
    foreach  $J \in B_i$                                      // register each weighted image
    do
       $J' \leftarrow \text{MemeticPyramidRegistration}(I_0, J', \mathbf{D})$        // see Algorithm 2
    end
  end
end

```

Algorithm 3: MultiObjectiveCorrection - pseudo-code of the proposed multi-objective motion and distortion correction approach.

Section 4.1.1 are used:

$$\mathbf{D} \leftarrow \{\mathcal{D}_{\text{NW}}, \mathcal{D}_{\text{MB}}, \mathcal{D}_{\text{NN}}\}.$$

The lowest b -shell B_1 is required to be already corrected for the DTI model-fit by the \mathcal{D}_{MB} . This initial pre-alignment with the reference space of the unweighted I_0 image is performed by a single-objective registration using the non-weighted reference objective function \mathcal{D}_{NW} . The so pre-aligned shell B'_1 is merged with the yet uncorrected shells to

$$I_{\text{DW}}^{(\text{pre})} = I_0 \cup B'_1 \cup B_2 \dots \cup B_s$$

and this pre-aligned dw-image is corrected by applying the multi-objective pyramid registration (Algorithm 2) to each weighted image. The complete approach is summarized in Algorithm 3. Parametrization and implementation details are presented in the next paragraph.

Parametrization

The local search method was realized by a regular step gradient descent optimization with adaptive parametrization. The maximal and minimal optimizer step, the maximum number of iterations and the gradient magnitude tolerance were adapted for each phase according to its corresponding resolution pyramid level. Full listing of the parameter values are given in Table 4.1.

OPTIMIZER PARAMETERS			
Parameter	Lev. 1	Lev. 2	Lev. 3
Max Step Size	4.0	2.0	1.0
Min Step Size	10^{-2}	10^{-3}	10^{-4}
Max Iterations	100	400	800
Grad Tolerance	10^{-4}	10^{-5}	10^{-6}

Table 4.1: Parameter settings for the maximal and minimal step size, the maximum number of iterations and the magnitude tolerance of the regular step gradient descent optimizer used for the local search phase. Each pyramidal level was parametrized differently (Level 1 is the coarsest).

The optimizer scales were set to 10^{-4} for the translational parameters and 10 for all other (affine) transform parameters. All objective functions were implemented using the Mattes Mutual Information Image Metric (v4) as provided by the Insight ToolKit (ITK) library (v4.7.2) [54]. The optimizer was also realized using the implementation provided by ITK. The optimizer scales were also considered for the particles' initialization (Equation 4.5) yielding

$$\mu_j = (id_{\mathcal{T}})_j + \text{scale}_j \cdot \mathcal{N}(0, \sigma(N)), \quad j = 1, \dots, 12 \quad (4.6)$$

for each of the 12 affine parameters μ_j that define the particle's position.

Complexity Analysis

A pairwise image registration following a pyramid scheme with L levels with a maximum number of optimizer steps $\{S_0, S_1, \dots, S_L\}$ will require at most $\sum_{i=0}^L S_i$ evaluations of the image metric, which is the most computationally intensive factor. Let M denote the costs (number of operations) of a single metric evaluation, then the maximal costs of a pyramid scheme registration at L levels can be written as

$$\text{Pyr}(L) = \left[\sum_{i=1}^L S_i \right] \cdot M \quad (4.7)$$

With the number of images \mathcal{I} , the single-objective correction with the *non-weighted reference* objective will require at most

$$\mathcal{I} \cdot \text{Pyr}(L) \quad (4.8)$$

operations. For the *model-based* and *nearest neighbor reference* objectives, this requirement increases by the pre-alignment of the lower shell. Considering the lower b -shell separately with $\mathcal{I} = \mathcal{I}_{\text{low}} + \mathcal{I}_{\text{high}}$, the correction including the pre-alignment will require at most

$$(\mathcal{I}_{\text{low}} + \mathcal{I}) \cdot \text{Pyr}(L) \quad (4.9)$$

operations. In the multi-objective approach, the costs of (4.8) arise for each of the N particles, in addition to the pre-alignment costs and the metric evaluations performed for fitness value estimation in each global update phase. Altogether, the maximum costs in term of metric evaluations are

$$(\mathcal{I}_{\text{low}} + N \cdot \mathcal{I}) \cdot \text{Pyr}(L) + N^2 \cdot \mathcal{I} \cdot L \cdot M \quad (4.10)$$

Although quadratic in the number of particles N , the last term increases the overall complexity only to a small extent, because the main cost factor $\text{Pyr}(L)$ contains the multiple of the single metric evaluation M by all optimizer steps. Thus, with the choice of $N < \mathcal{I}$, all single-objective and the multi-objective approaches have a time complexity of $O(\mathcal{I})$ in the number of image volumes \mathcal{I} . The multi-objective approach has a linear time complexity $O(N)$ also with respect to the number of particles, as long as the sum of optimizer steps on all pyramid levels L exceeds the factor $N \cdot L$.

4.2 Construction of Validation Data

The strongest quality measure for a finished image registration task is the achieved alignment accuracy. There are several options how image registration tasks can be evaluated. From an optimization point of view, the value of the cost function (image metric) provides only a weak indicator of the final alignment, since it strongly depends on both the image contrasts and the cost function itself. Hence, such a measure becomes ineffective when it comes to comparison of registration tasks in different scenarios.

Following the main purpose of image registration, the alignment of images, an error measure based on Euclidean distance between well defined anatomical structures identified in both images—the Target Registration Error (TRE)—is an intuitive choice. To

increase the robustness of the estimation, TRE is usually evaluated over a set of multiple corresponding point pairs (fiducials). Formally, for two sets of points $\{x_1, \dots, x_N\}$ and $\{y_1, \dots, y_N\}$, the mean TRE is defined by

$$\overline{\text{TRE}} = \frac{1}{N} \sum_{i=1}^N \|x_i - y_i\|_2. \quad (4.11)$$

Consequently, the crucial task of the evaluation moves to the precise identification of the corresponding points in both images. When the fiducials are set manually by (medical) experts, one has to account for the fiducial localization error [26, 38], which describes the error of manually targeting a structure in images relative to the real target position. The localization error is not of concern if the ground-truth transform $T_{(\text{gt})}$ for an estimated transform T is known. In such case, the TRE can be evaluated over a set of *virtual* fiducials $\{x_1, \dots, x_N\}$:

$$\overline{\text{TRE}} = \frac{1}{N} \sum_{i=1}^N \|T \circ T_{(\text{gt})}^{-1}(x_i) - x_i\|_2. \quad (4.12)$$

The ground-truth transforms can be obtained in several ways. A simple and generally used approach is to apply the known transform to an image and then register the transformed image back onto the original one [81]. In deformable registration validation, the ground-truth deformation is usually computed by means of simulation to provide realistic tissue deformations [100].

4.2.1 Synthetic Validation Data

In diffusion-weighted MRI, none of the ground-truth constructions above is directly applicable to registration validation. A manual annotation of anatomical landmarks is hardly practicable for the high numbers of acquired volumes per dataset. And because of the rotational dependency of the signal, the only ground-truth transforms that could be applied to a diffusion-weighted image volume without a need of modifying the signal are translation transforms. However, using only translations for the generation of ground-truth does not provide sufficient complexity for a meaningful validation [48].

The problem of missing ground-truth for algorithm validation arises also for fiber tracking methods. The possibilities by using hand-crafted hardware phantoms [90, 74, 25] are restricted in the bundle configuration complexity. Hence, the research focused on the creation of dMRI software phantoms, first based only on a tissue signal

model [21, 16] and therefore with limited simulation complexity, later with realistic simulations of the MR acquisition emerging recently [76, 43].

Construction of dMRI Ground-truth Data

The fully accessible open-source tool *Fiberfox* starts with a full-brain tractogram and a definition of the diffusion MR acquisition settings including the gradient directions, the b -values and the general image geometry. With this input a realistic simulation of a dMRI acquisition is computed by mimicking a single-shot Echo-planar Imaging (EPI) k -space acquisition with a standard Stejskal-Tanner gradient sequence [76]. *Fiberfox* operates in the (Fourier) k -space and can consider up to four different compartments. This allows to model the tissue specific signal relaxation responsible for the tissue contrast in the images as well as to simulate the EPI-specific distortions. The signal with a gradient g at time t is gathered over the compartments c

$$S_v(g, t) = \sum_{c \in \text{Comp.}} f_c \cdot S_c(g) \cdot e^{\text{RELAXATION}(c)}. \quad (4.13)$$

The simulation of a dw-image $I_{\text{DW}} = \{I_0, I_1, \dots, I_M\}$ is computed separately for each gradient image I_j . This allows to apply a known transform to the tractogram between the simulation of two successive images. This corresponds to a movement inside the scanner, transforming the imaged tissue relative to the acquisition frame. The eddy currents effects originate from a residual signal with slow decay (see Section 2.4) and can be simulated with *Fiberfox* exactly in such way:

$$S_v^{(\text{ec})}(g, t) = S_v(g, t) \cdot e^{\text{EDDY-SIGNAL}(b(g), t)}, \quad (4.14)$$

with $S_v(g, t)$ the gradient- and time-dependent voxel signal described in (4.13). The residual signal $\text{EDDY-SIGNAL}(b(g), t)$ depends on the b -value and scales with \sqrt{b} for increasing b .

The simulation provides ground-truth transforms for the simulated rigid motion in each image. The ground-truth deformations for simulated eddy-currents-caused distortions cannot be obtained directly in an analytical way, since the exponential signal-decay was simulated in Fourier space. However, the ground truth can be reliably estimated by non-linear registration of noise-free data sets with and without the eddy-currents-caused distortions. This step was carried out using the B-spline-based deformable registration from *elastix* [58], using the known rigid transforms for initial positioning.

Data Description

Starting with the tractogram of the ISMRM Tractography Challenge¹, 12 DW datasets with 18 gradient directions at $b = 1000$ and $3000 \text{ s} \cdot \text{mm}^{-2}$ each as well as one baseline image at $b = 0 \text{ s} \cdot \text{mm}^{-2}$ were simulated, all with the same diffusion parameters and isotropic voxel size of 2.0 mm , yielding a in-plane slice dimension of 90×108 voxels and 90 axial slices. The datasets differ in the severity of simulated artifacts, motion and signal-to-noise ratio (SNR). For the datasets with moderate artifacts, only eddy currents were simulated. In the simulation of more severe artifacts, the eddy currents were doubled, Gibbs ringing and ghosting artifacts were introduced. For both artifact severities, three distinct levels of motion were added to the datasets by randomly selecting rotation and translation parameters along the three image axes (x, y, z). Finally, for all datasets, Rician noise was added to all images at signal-to-noise ratios of 10 and 20. A comprehensive list of the synthetic data and the corresponding simulation parameters is given in Table 4.2.

SYNTHETIC DATASETS OVERVIEW

Name	Artifacts	Rotation [deg]		Translation [mm]		SNR
		xy	z	xy	z	
Ia₂₀	moderate	± 5	± 10	-	-	20
Ib₂₀				± 10	± 6	
Ic₂₀				± 20	± 12	
Ia₁₀	moderate	± 5	± 10	-	-	10
Ib₁₀				± 10	± 6	
Ic₁₀				± 20	± 12	
IIa₂₀	severe	± 10	± 15	-	-	20
IIb₂₀				± 10	± 6	
IIc₂₀				± 20	± 12	
IIa₁₀	severe	± 10	± 15	-	-	10
IIb₁₀				± 10	± 6	
IIc₁₀				± 20	± 12	

Table 4.2: Overview of the simulated artifacts (I,II), motion distortion (a,b,c) and SNR (subscripts: ₁₀, ₂₀) for the synthetic data. The datasets are referenced in text by the name listed in column (Name).

¹ ISMRM 2015 Tractography Challenge Website: http://www.tractometer.org/ismrm_2015_challenge/

4.2.2 In-Vivo Validation Data

The missing ground-truth problem arises for in-vivo settings as well. Here, an additional MRI-independent measurement for the distortions would be required. This could be provided for the bulk motion using a tracking system identical to the prospective correction approaches (see Section 3.1.2). However, the technical requirements are rather high and in addition the ground truth for the eddy currents-caused distortions is not provided within. For these reasons typically a *reference-based* evaluation is used instead. In such, dataset pairs are evaluated, one is taken for the ground-truth, while the other one is evaluated with respect to it. In diffusion MRI, a signal model or its scalar values are used in the evaluation rather than the "raw" diffusion-weighted image data. In an example, the limits of the unweighted reference method in the correction of reduced Field of View (FOV) data were demonstrated by constructing reduced-FOV datasets by cropping the original dataset and evaluating the difference in DTI scalar maps, caused by the initial misalignment [50]. In another study on the impact of different correction approaches in HARDI acquisitions, Elhabian *et al.* [34] also used a reference-based evaluation, however with one reference image and multiple datasets manifesting different levels of motion.

The in-vivo evaluation in this thesis follows the evaluation data scheme by Elhabian *et al.* A young healthy subject was scanned twice while using the same acquisition settings: *acquisition 1* was acquired while the subject was advised to hold still; during *acquisition 2* the subject was advised to move throughout the complete acquisition. Acquisition 1 is taken as the reference and the evaluation datasets are constructed by replacing selected weighted volumes by the corresponding volumes from the distorted acquisition 2. For the construction, four different levels of motion corruption: 25, 50, 75 and 90% are defined. For each level, the corresponding amount of volumes from acquisition 2 was randomly selected, while the remaining volumes and the unweighted ($b = 0 \text{ s} \cdot \text{mm}^{-2}$) volume were selected from acquisition 1. The selection of the identical I_0 guarantees an identical reference space and thus enables a reference-based evaluation in the first place. Similar to Elhabian *et al.* [34], five datasets with randomly selected moving volumes at each corruption level were considered, yielding in total 20 evaluation datasets.

Data Description

The evaluation datasets (acquisition 1, acquisition 2) were acquired on a 3T whole-body clinical scanner (Siemens Prisma, Siemens Medical Solutions, Erlangen, Germany) with a gradient strength of 80 mT/m. A single shot EPI technique with a twice refocused spin echo diffusion preparation was employed [91] using the following parameters:

repetition time (TR)/echo time (TE) 6300/79 ms, field of view 250 mm, data matrix of 100×100 , yielding an in-plane resolution of 2.5 mm, 40 axial slices with a thickness of 2.5 mm and no gap. With these parameters, 5 unweighted ($b = 0$) and 30 weighted images at each $b = 1000$ and $b = 3000 \text{ s} \cdot \text{mm}^{-2}$ shells were acquired. The gradient directions were distributed uniformly on a half-sphere. The SNR in the in-vivo volumes was estimated from the signal in corpus callosum (using the PIESNO method for noise variance estimation [62]) to 38.5 in the unweighted dataset, 17.8 ± 4.17 in $b = 1000 \text{ s} \cdot \text{mm}^{-2}$ and 8.47 ± 3.68 in $b = 3000 \text{ s} \cdot \text{mm}^{-2}$. Using registration of the motion-corrupted volumes to their unmoved counterparts, the mean subject's movement during the second acquisition was estimated as 10.8 ± 4.82 deg for rotation and 6.94 ± 3.318 mm for translation.

4.3 Evaluation Metrics

The ground-truth available by the validation data construction described in the previous sections allows for quantitative evaluation of the precision of the proposed correction method. This is done either in a direct way by means of the target registration error on synthetic data or in an indirect, reference-based comparison on the in-vivo datasets. The TRE available in synthetic data experiments further allows to study the impact of different parameter choices for the PSO update and the initial position distribution on the method's precision. To analyze the multi-objective aspect of the approach, and especially the importance of each single objective function involved, a new measure is introduced in the next section.

4.3.1 Objective's Behavior

The role of each objective function throughout the multi-objective registration can be derived from its behavior in the global update phases. As shown in Fig. 4.3, the finally chosen particle (i.e. transform) for each registered image can be traced back through its history of PSO update steps, marking the particle's leader in each phase. Particles with good fitness values are most likely to get selected as leading particles and a particle selected as a leader indicates the capability of the corresponding local search associated with the particles' swarm. Summed up over all image registration tasks, the frequency of appearance as the leading particle reveals the importance of the linked objective function in the whole memetic setup. The aggregation of these traces can further reveal




Figure 4.3: Exemplary illustration of the trace observed during memetic registration of one diffusion-weighted image. The finally chosen particle (i.e. transform) started at its initial position (S) and then followed one leading particle in each PSO update step. The objective function (NW, MB or NN) of the respective leader is marked in solid magenta.

the importance of each objective function at different stages (pyramid levels) of the registration process.

4.3.2 Qualitative Evaluation

In addition to the quantitative evaluation, the impact of the correction can be evaluated qualitatively, by means of visual inspection of the corrected raw data, the results of fiber-tractography and the resulting signal maps. A detailed description of the performed experiments is provided in the last two sections of this chapter.

4.4 Evaluation on Synthetic Data

After describing the multi-objective correction method and the methods for construction of validation data, this section describes the individual experiments for evaluation of the proposed approach performed on synthetic data. In all experiments, the proposed multi-objective method is evaluated against the single-objective registration approaches for each objective function. Firstly, a quantitative evaluation of the registration precision for different configurations of the memetic search are performed, by measuring the TRE against the ground-truth transforms provided by the simulation. Secondly, the role of the individual objective functions are evaluated and lastly, a qualitative evaluation

of the impact of the correction precision is performed by fiber tracking. The detailed listing of the synthetic datasets used in this experiments is provided in Table 4.2.

4.4.1 Precision

In order to evaluate the alignment precision, the proposed multi-objective method for three different parameter settings

MULTIOBJ (6P): $C_1 = 0, C_2 = 1, N = 6$

MULTIOBJ (6P, RANDOM): $C_1 = 0.5, C_2 = 1, N = 6$

MULTIOBJ (12P, RANDOM): $C_1 = 0.5, C_2 = 1, N = 12$

and the single-objective correction methods (unweighted, model-based and nearest neighbor) were applied to each of the 12 simulated datasets (see Table 4.2). The precision of each experiment was calculated by evaluating the TRE of the estimated transforms against the ground-truth transforms (see Equation (4.12) in Section 4.2.1). For the TRE computation a set of $m = 27$ virtual landmarks placed on a uniform grid located within the brain was used.

4.4.2 Particle History

In an additional evaluation to the experiments from previous paragraph, the particle history as described in Section 4.3.1 was computed.

4.4.3 Initial Position Variance

The spread of initial positions can significantly influence the corrective function of the PSO settings by providing a certain degree of diversity already from the beginning. The particle spread radius can be controlled by the variance $\sigma(N)$ introduced in (4.6). The variance, or rather variance function, should increase with the number of particles N . To evaluate the impact of the initial spread, four different choices for $\sigma(N)$ were made

- | | |
|--|---------------------------------------|
| (i) $\sigma_1 = \frac{N}{\ \mathbf{D}\ },$ | (iii) $\sigma_3 = N \cdot \log_2(N),$ |
| (ii) $\sigma_2 = N,$ | (iv) $\sigma_4 = N^2.$ |

The TRE was evaluated for the simulated datasets with the highest amount of added motion at the lower SNR—datasets Ic₁₀ and IIc₁₀ for four different number of particles $N = 6, 12, 18$ and 24.

4.4.4 Fiber Tracking

The impact of the motion and artifact correction method on the fiber tracking results was evaluated on the higher-artifacted dataset without additional translation at SNR of 10 (dataset IIa₁₀). The most stable method according to the evaluation of tracking algorithm by Neher *et al.* [75]—the deterministic Constrained Spherical Deconvolution (CSD) fiber tracking method—was applied to the datasets. After the correction, the $b = 3000 \text{ s}\cdot\text{mm}^{-2}$ shell was extracted, the amount of 18 volumes allows to use the spherical harmonics base of order 4 in the reconstruction, and the CSD-tracking algorithm as provided within the *MRtrix* toolbox [109] was applied to the whole brain. From the resulting tractograms, the *fornix* tracts were extracted using a passing Region of Interest (ROI) defined in the ground-truth tractogram.

4.5 Evaluation on In-vivo Data

In the in-vivo experiments, the impact of the multi-objective method as well as the state-of-the-art single-objective methods on diffusion MRI analysis was evaluated by a comparison with a reference in-vivo acquisition without artifacts: firstly for the scalar maps derived from the higher-order NODDI signal model, secondly in evaluating the role of each objective and lastly by fiber tracking experiments. All experiments were performed on the dataset collection constructed with different levels of motion corruption as described in Section 4.2.2.

4.5.1 Tissue Parameter Estimation

For the evaluation, the single-objective corrections were represented by the unweighted reference method, which is a common correction method, and the extrapolation reference method using a DTI-based simulated datasets, which is best one in the correction of high b -value datasets. The tensors were fitted to the lower shell ($b = 1000 \text{ s}\cdot\text{mm}^{-2}$) pre-aligned with the unweighted reference method. The multi-objective method was considered in the *corrective* parametrization ($C_1 = 0$, $C_2 = 1$, $N = 6$). For each of the

constructed motion datasets, the NODDI-derived scalar indices Orientation Dispersion Index (ODI) and Intra-cellular Volume Fraction (ICVF) were computed after the correction. These values were evaluated against the corresponding indices computed from the unmoved reference dataset (acquisition 1).

The NODDI model further provides the main fiber orientation estimate in each voxel. This allows to evaluate the angular deviation of the estimated fiber orientation in the reference and in the corrected motion-corrupted datasets. To account for directional ambiguity, the direction vectors were corrected so that they point to the upper half sphere ($z \geq 0$) before evaluating the angular deviation.

The indices ODI and ICVF and also the mean fiber orientation provide meaningful values only in the anisotropic (white-matter) compartments of the brain. To account for the different compartments, the mean differences were computed over a tissue mask provided by the reference dataset. The tissue mask was constructed by applying an upper threshold of 0.85 to the NODDI-estimate of voxel isotropy (f_{iso}) to mask out the isotropic voxels, formed in the majority by the CSF-compartment.

4.5.2 Particle History

In an additional evaluation to the in-vivo experiments described in previous paragraph, the particle history (see Section 4.3.1) was computed for each motion corruption level. For each of the four added motion levels (25, 50, 75 and 90%), the particles' traces were aggregated across all bootstrap samples.

4.5.3 Fiber Tracking

The impact of the correction methods on fiber tracking was evaluated on two datasets at motion corruption level of 75% and 90% respectively. Similar to the tractography evaluation on the synthetic data, the deterministic CSD fiber tracking was applied to the $b = 3000 \text{ s} \cdot \text{mm}^{-2}$ shell of the corrected datasets and the unmoved reference dataset. For the reconstruction, spherical harmonics base of order 4 was used and the deterministic CSD-tracking algorithm from the MRtrix toolbox was used to obtain whole-brain fiber tracts. From the resulting tractograms, the *fornix* and the *cingulum* tracts were extracted using a passing ROI defined in the unmoved reference dataset.

Chapter 5

Results

In this chapter, the results of the experiments described in the previous chapter are presented. Firstly, the constructed synthetic data and the results of the synthetic data experiments, showing the precision and the outlier reduction achieved by the proposed multi-objective registration and its impact on fiber tractography, are presented in Section 5.1. Then, the in-vivo validation data and the results of the in-vivo experiments including the evaluation of NODDI model, fiber tractography and the roles of each objective function are described in Section 5.2.

5.1 Evaluation on Synthetic Data

The following sections summarize the results of the synthetic data experiments. First, the synthetic evaluation datasets together with the ground-truth transform estimation are shown (5.1.1), followed by the results of registration precision evaluation (5.1.2), role of each objective function (5.1.3), impact of the correction on fiber tracking performance (5.1.5) and the processing time (5.1.4). For a detailed description of the corresponding experiments please refer to Section 4.4 in the previous chapter where the notation for the evaluation datasets is listed as well in Table 4.2.

5.1.1 Synthetic Validation Data

The synthetic evaluation datasets were constructed for two different artifact levels (moderate, severe) at different b -shells (1000, 3000 $\text{s}\cdot\text{mm}^{-2}$) and different added noise levels (SNR of 10, 20). Exemplary slices of the datasets illustrating this variety of




Figure 5.1: Exemplary axial slices from the simulated synthetic datasets used for evaluation. (A–D) Images with moderate simulated eddy currents and motion artifacts, (E–H) images with more severe artifacts—doubled eddy currents, ghosting and ringing. (Upper row) Images simulated with Rician noise to reach SNR of 20. (Lower row) Images simulated with higher Rician noise to reach SNR of 10.

simulated effects are shown in Figure 5.1. The eddy-currents-caused artifacts were simulated for all (weighted) images and manifested themselves in form of *enlargement* effects on the signal (Figure 5.1(C,F)). An opposite manifestation in form of compression of the signal is shown in images from the higher artifacted dataset (G,H). Here, also the increase in the eddy currents effect size with increased b -value is reflected by the higher compression of the $b=3000 \text{ s}\cdot\text{mm}^{-2}$ image (H) compared to the lower shell image with similar gradient direction (G). The effect of doubled eddy current strength in the severely artifacted datasets is visible in image (F) where a higher shearing in the up-down direction can be observed compared to image (D) which was simulated with the same gradient direction but only moderate artifact influence. The higher graining of the signal in image (D) on one side and the smooth white-matter signal in image (B) on the other, illustrate the differences between images with lower and higher SNR levels. In the datasets with more severe artifacts, the simulated ghosting effect is clearly visible in the upper and lower parts of the unweighted image (Figure 5.1(E)) as well as in the lower b -shell image (G). In the latter one, also the ringing artifacts in form of a wave pattern are visible particularly in the lower part of the brain tissue signal.




Figure 5.2: Ground-truth estimation with reference registration. In each column from top to bottom: (1) the original moving image with distortions and noise (SNR=10), (2) the moving image with distortions but without noise used to estimate the reference registration, (3) the fixed image without noise and distortions, (4) the moving image after non-linear registration with an overlaid contour (green) computed from the white-matter and the outer border of the (noiseless) fixed image.

Ground-truth Estimation

The quality of the non-linear registration was assessed by visual inspection of the voxel-wise mean error between the transformed source and the target image. Exemplary axial slices from different b -shells are shown in Figure 5.2. The noiseless moving image is shown with the uncorrected rotational motion to better emphasize the difference to the noisy original image. These ground-truth rotations were corrected prior to the non-linear registration. The last row (Contour after registration) illustrates the registration result by augmenting the moving images after registration with the contour computed by applying a threshold filter that captures the white matter and the outer border in the respective fixed image.

5.1.2 Precision

To measure the registration precision, the proposed memetic multi-objective method was evaluated against the single-objective methods with respect to the Target Registration Error (TRE). For the datasets with moderate artifacts and lower portions of added motion (Ia_{20} , Ia_{10} and Ib_{20} , Ib_{10}) shown in Figure 5.3, box-plots (a,b), the average Target Registration Error achieved by the proposed method remained below the voxel size for both b -shells, whereas the reference single-objective methods managed to reach an error below the voxel size only for the lower ($b=1000 \text{ s} \cdot \text{mm}^{-2}$) shell. The overall best TRE in the higher shell ($b=3000 \text{ s} \cdot \text{mm}^{-2}$) is achieved by the multi-objective method with the parametrization ($C_2 = 0.5$ and $N = 12$). For this settings, the error remained below the single voxel size. From the single-objective methods, only the model-based reference method remained below the double voxel size (4 mm).

For the more severely artifacted data (shown in the lower part of Figure 5.3) at the lower b -shell, all methods achieved a similar average TRE below the double voxel size (4 mm) for all three added motion variants (a,b,c) and both SNR levels (10,20). At the higher b -shell ($3000 \text{ s} \cdot \text{mm}^{-2}$), the reference methods failed to register the data, yielding an average error of above triple voxel size, while the proposed multi-objective method succeeded in reducing the registration outlier of the single-objective methods. It maintained a precision below the double voxel size in the first two motion cases at both SNR levels (Ia_{20} , Ib_{20} and Ia_{10} , Ib_{10}). The number of outliers and the average TRE decreased with the PSO parametrization to $C_2 = 0.5$ and $N = 6$. Doubling the number of particles ($N = 12$) at this parametrization resulted in a further improvement of the average TRE, but rather a small one compared to the improvement by the multi-objective method against the best single-objective method and at the costs of increased processing time




Figure 5.3: Target Registration Error (TRE) for synthetic data with moderate (upper plots) and severe (lower plots) artifacts at different motion corruption levels (a,b,c), different SNR (10,20) and different b -values (1000, 3000 $\text{s}\cdot\text{mm}^{-2}$). Each plot shows the TRE for the three single objective methods (unweighted, nearest-neighbor and model-based) and three variants of the proposed multi-objective method. The dashed line marks the (isotropic) voxel size of 2 mm, the dotted line the two-voxel size of 4 mm. ©2016 IEEE, originally published in [49].

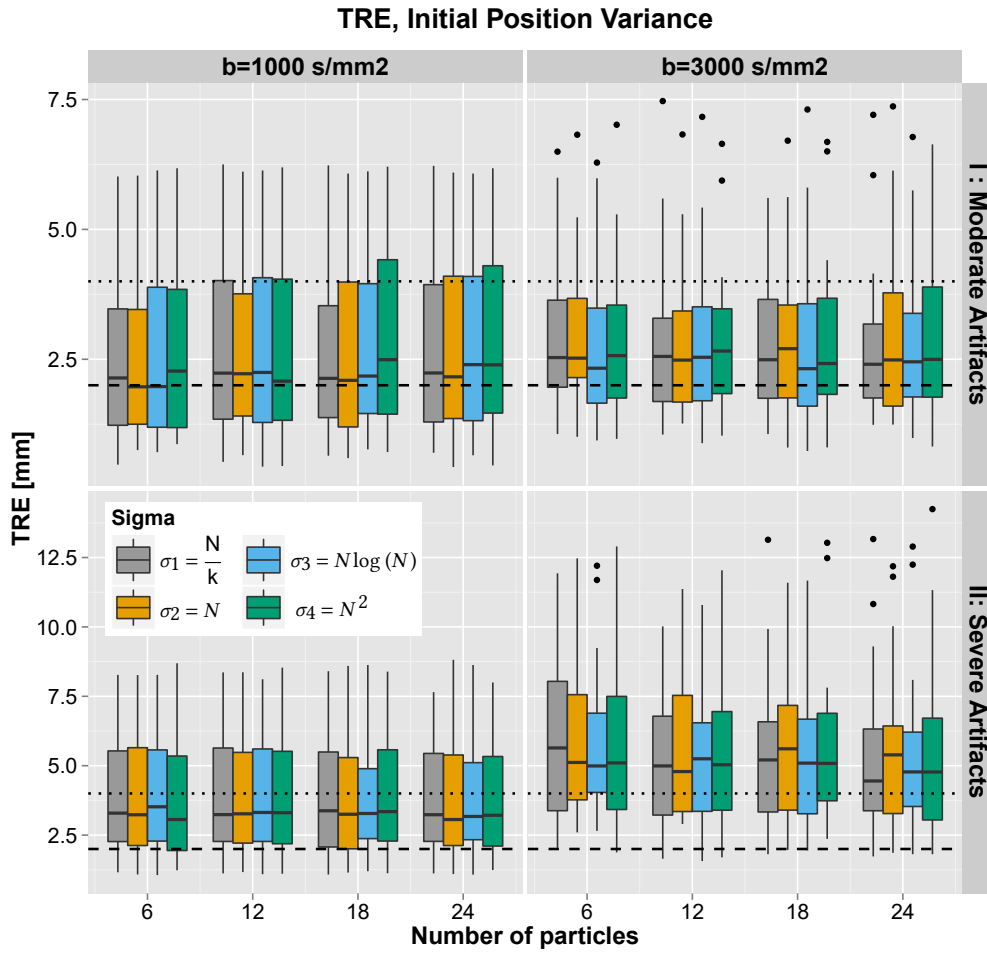


Figure 5.4: Target Registration Error (TRE) for different initial position variance choices for different numbers of particles ($N=6, 12, 18$ and 24) evaluated for the synthetic datasets Ic_{10} (upper row) and Iic_{10} (lower row). The dashed line marks the (isotropic) voxel size of 2 mm, the dotted line the two-voxel size of 4 mm.

as shown later on in Section 5.1.4. Datasets with the highest added level of motion, shown on the box-plots (c), presented a challenge for all correction methods at both moderate and severe artifact levels. The error of the unweighted reference method in the lower shell, and hence in the pre-alignment registration for the nearest-neighbor and model-based reference methods, almost reached the double voxel size already on the moderately artifacted datasets. With the nearest-neighbor method, these initialization errors amplified in the higher b -shell. Although the TRE dropped compared to the single-objective reference methods, the multi-objective method did not improve, in contrast to the other synthetic experiments, with the $C_2 = 0.5$ parametrization or with increasing number of particles neither at the lower nor at the higher b -shell. This behavior was also observed in the evaluation of the initial position spread and the number of particles used. For the moderately artifacted dataset (Ic_{10}), none of the tested choices of the initial position variance $\sigma(N)$ showed a predominant behavior

as depicted in Figure 5.4. For the datasets with the overall highest distortions—severe artifacts and the highest added motion (Πc_{20} , Πc_{10})—all methods, including the proposed method in all variants, yielded an average TRE above the double voxel size. For this dataset also the variation of the initial position spread and increasing the number of particles did not provide a clear improvement.

5.1.3 Particle History

The aggregated particle history graphs in Figure 5.5 show how the roles of each of the three reference objectives (Non-weighted reference objective (NW), Model-based reference objective (MB) and Nearest-neighbor reference objective (NN)) varied for different stages of the registration process and different SNR and artifact levels. The model-based objective had the most appearances among the objectives for the less artifacted data at higher SNR (Figure 5.5A). This effect corresponds with the TRE yielded

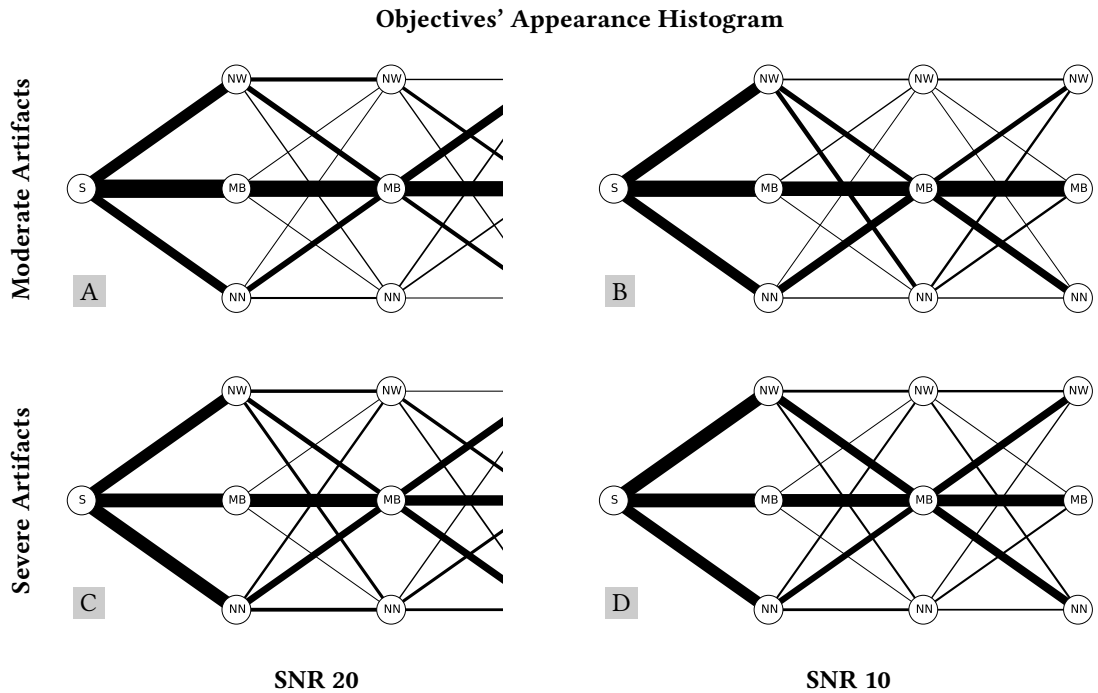


Figure 5.5: Aggregated history of particles after the optimization process for the synthetic datasets. (A) Simulation I at SNR of 20, (B) at SNR of 10, (C) Simulation II at SNR of 20 and (D) at SNR of 10. The node (S) represents the start position. During optimization, the particles move from left to right through the three different PSO updates, each associated with a leading particle. The edges in the graph represent the path of the particles with respect to the different leaders they followed. The edge thickness represents the number of appearances summed over all particles' history.

by the model-based reference method. It also dominated in the second local phase at all artifact and SNR levels. However, the model-based objective lost its importance at the last local search phase, i.e. the full-resolution pyramid level. This loss was more pronounced rather when increasing the artifact levels (see Figure 5.5A→C, B→D)) than for decreasing SNR level and lower SNR levels (see Figure 5.5A→B, C→D)).

5.1.4 Runtime

All experiments were performed on a desktop PC equipped with an Intel®-i7-Core (8 core) processor and 24 GB RAM. The mean processing times for the correction of synthetic datasets are provided in Table 5.1. The processing times of the multi-objective method scaled according to the number of particles. For the single-objective correction methods with the model-based and the nearest neighbor reference objectives, the processing time increased almost identically by the pre-alignment, which was computed on the lower b -shell prior to model fitting and estimation of similar gradient directions, respectively.

RUNTIME ANALYSIS

Approach	Runtime	StdDev
Unweighted reference	11 min 26 s	23 s
Nearest neighbor reference	17 min 01 s	31 s
Model-based reference	16 min 43 s	42 s
Proposed multi-objective (6p)	49 min 47 s	1 min 51 s
Proposed multi-objective (12p)	1 h 43 min 55 s	53 s

Table 5.1: Processing times for the synthetic datasets. In the column (Runtime) the mean processing time across the application to all evaluation datasets is listed, with the corresponding deviation (StdDev).

The method was implemented within the open-source C++ Medical Imaging and Interaction Toolkit (MITK) [80] and the source code is available from the MITK repository at www.mitk.org.

5.1.5 Fiber Tracking

The reconstructed *fornix* tracts (see Figure 5.6) revealed the impact of the correction misalignment. The initial registration error for the higher b -shell in the dataset before fiber reconstruction is listed for each correction method in Table 5.2. As a result of an

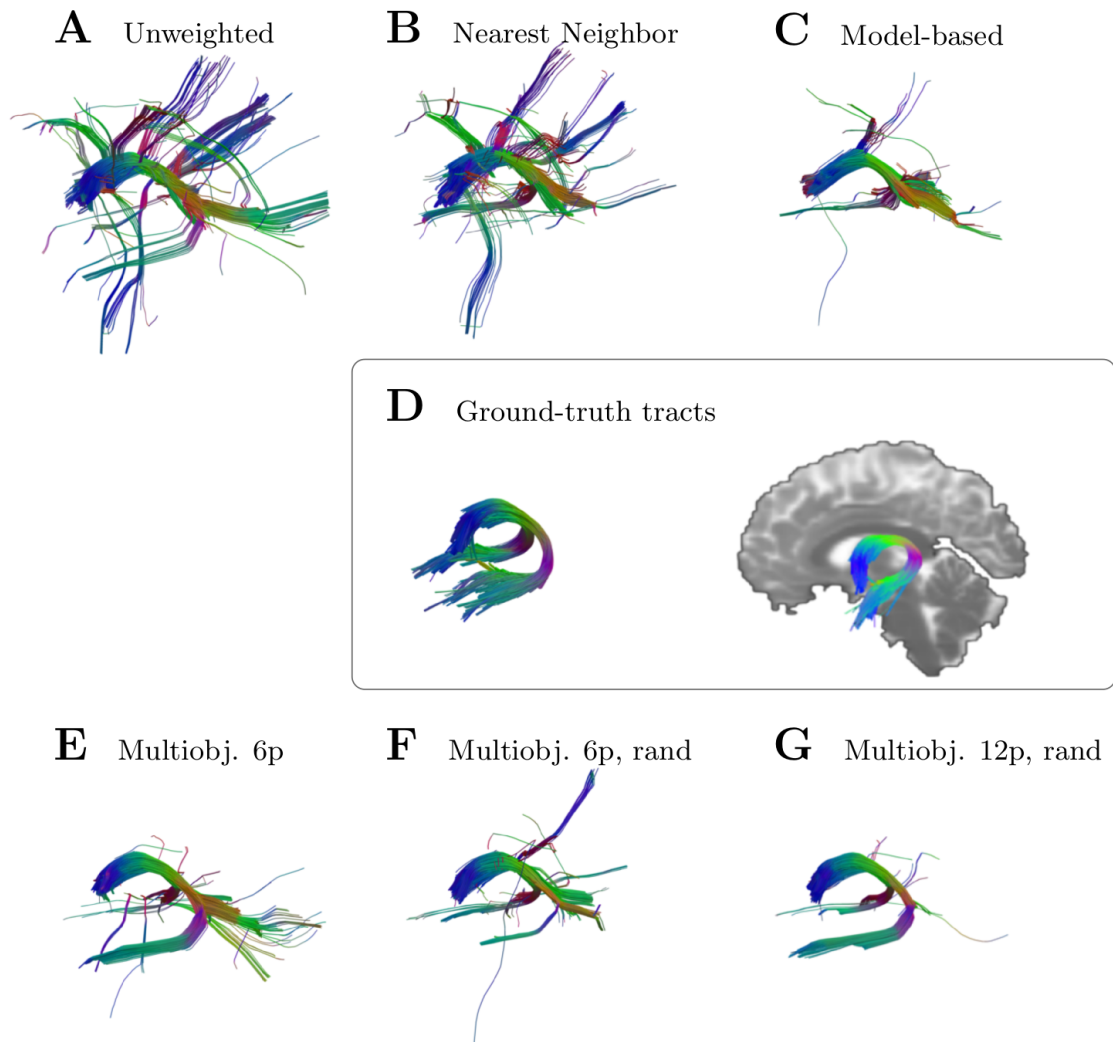


Figure 5.6: The *fornix* tracts reconstructed from the simulated dataset after motion and distortion correction. Each correction method was applied to the simulated dataset with higher artifacts and added rotation ($\Pi_{A_{10}}$) and the CSD tracking was performed on the $b=3000 \text{ s} \cdot \text{mm}^{-2}$ shell. The tracts were then extracted using a passing ROI defined in the ground-truth tractogram. ©2016 IEEE, originally published in [49].

TRE FOR DATASET $\Pi_{A_{10}}$						
Method	A	B	C	E	F	G
Mean TRE [mm]	6.32	6.63	7.43	4.29	3.39	3.07
StDev	± 2.27	± 2.85	± 3.83	± 2.70	± 2.07	± 1.96

Table 5.2: Initial alignment error in the higher b -shell of the evaluated correction methods on the dataset $\Pi_{A_{10}}$ used to perform the fiber tracking.

initial alignment error above three times the voxel size, the reference single-objective

methods produced a high amount of spurious fibers in the resulting tractogram. With decreasing initial TRE by the multi-objective approach with $N = 6$ particles, the amount of reconstructed spurious fibers reduced. The best results were provided on the basis of the proposed multi-objective approach with the parameter choice of $C_2 = 0.5$ and $N = 12$ particles with an initial error of 3.07 mm, i.e. 1.5-times the voxel size.

5.2 Evaluation on In-vivo Data

This section summarizes the results of the bootstrap experiments performed on the in-vivo datasets, which were constructed at four different motion corruption levels (25%, 50%, 75% and 90%) by replacing images in a reference acquisition by randomly selected gradient images from a distorted dataset (for detailed description please refer to Section 4.5). The evaluation datasets were corrected by the single-objective methods (unweighted and model-based reference) and the proposed multi-objective method, the ground-truth values for the evaluation originate from the unmoved reference acquisition.

5.2.1 In-vivo Validation Data

The in-vivo evaluation datasets were constructed in a bootstrap-like manner by interleaving the volumes from the unmoved acquisition of the subject with volumes from an intentionally distorted acquisition. To allow a visual inspection of the alignment across the gradient images, the data are shown in a specific ordering illustrated by Figure 5.7. To achieve this, a single line of voxels defined in the target reference space is shown for all gradient images next to each other. In this way, a perfect alignment is characterized by clean linear structures across the volumes as illustrated in Figure 5.7C.

Figure 5.8 uses the registration validation view to show the in-vivo datasets by displaying the data of one bootstrap experiment at each motion corruption level (25%, 50%, 75% and 90%). At the lowest level (25%) the disruption of the linear structures is visible at the few positions and the amount of artifacted slices grows according to the increasing level of motion corruption.

Raw Image Data differences

To illustrate the difference in correction performance of the evaluated methods, exemplary volumes from a 75% motion-corrupted dataset after the alignment are shown

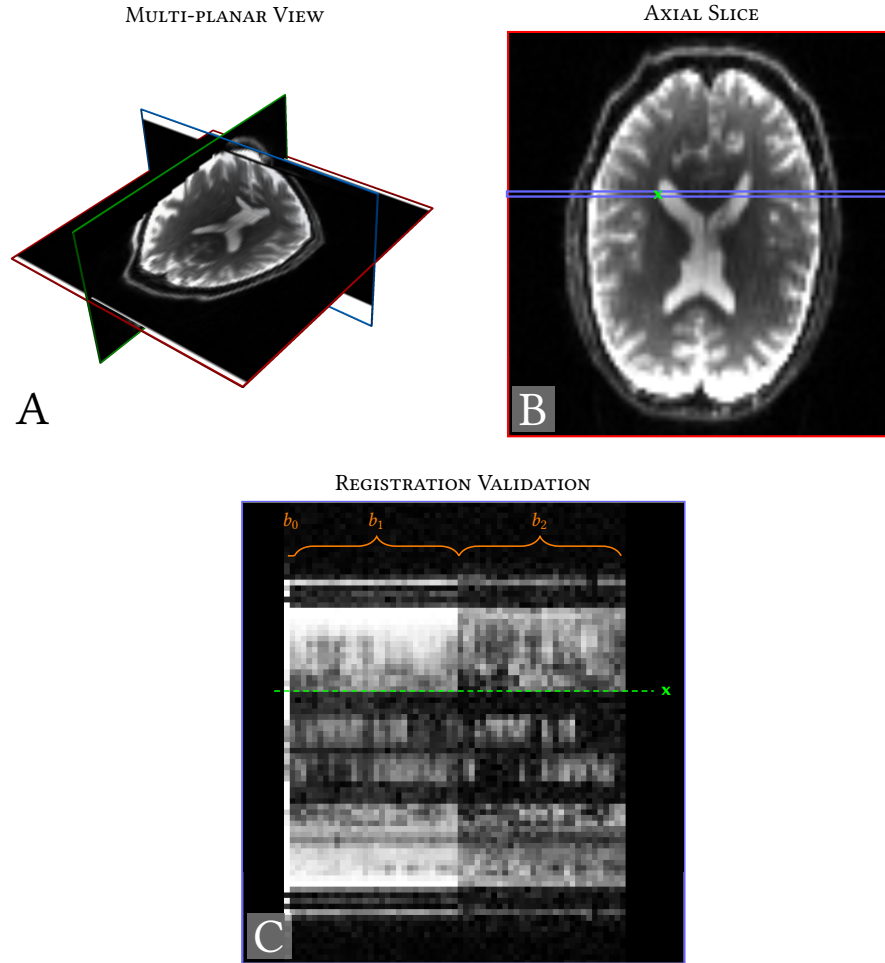


Figure 5.7: The principle of visual registration validation illustrated on an undistorted multi-shell DW image ($b_0 = 0, b_1 = 1000$ and $b_2 = 3000 \text{ s} \cdot \text{mm}^{-2}$). (A) The multi-planar reconstruction of the unweighted image showing the three anatomical orthogonal planes (*axial, sagittal, coronal*). (B) An axial slice from the unweighted volume with a line of voxels marked in blue and (C) the marked line of voxels shown for all gradient images. The different b -values are emphasized by the braces on the top of the image, the dashed green line shows the alignment of a single voxel (x), also labeled in (B). The contrast in the high b -value part of image (C) was enhanced for printing.

in the registration validation view in Figure 5.9(A-C). The proposed method showed the best motion correction and no outlier. The single-objective unweighted reference showed a high number of outliers in images from the higher b -shell, while it provided an outlier-free correction of the lower shell. For the single-objective model-based method, a single outlier occurred in the lower b -shell at this dataset, however, the number of outliers in the higher b -shell was lower than in unweighted reference method.

A further insight on the correction performance is provided by Figure 5.10, which shows the average error in the higher b -shell of the corrected data compared to the reference




Figure 5.8: A registration validation view of the in-vivo dataset showing one bootstrap sample at each motion corruption level. The corruption levels are marked in the bottom-right corner of each image. The contrast in the high b -value part was enhanced for printing.

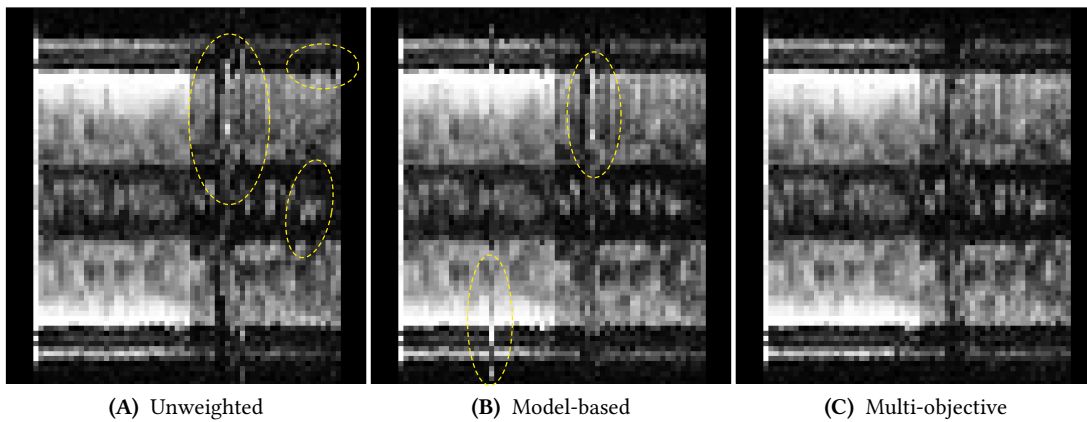


Figure 5.9: Illustration of the alignment of an in-vivo dataset using the registration validation view. The images show one dataset with 75% motion corruption after correction with (A) the unweighted reference objective, (B) the model-based reference objective and (C) the proposed multi-objective registration method. The yellow ellipses exemplify indicate some of the registration outliers. The contrast in the high b -value part was enhanced for printing.

acquisition, the mean computed across the bootstrap samples. At each corruption level, the proposed method showed the smallest relative difference to the unmoved acquisition, the single-objective unweighted reference method yielded the highest difference, already at the lowest corruption level. The high error rim (red color) in the frontal area shows the eddy-currents-caused signal enlargements.




Figure 5.10: Mean error of the corrected raw diffusion-weighted data at the higher b -shell for all motion corruption levels. In all images, the color indicates the mean error within the range of the colorbar given in the top-left corner ($[0.0, 60.0]$).

5.2.2 Tissue Parameter Estimation

The proposed method provided the smallest error when compared to the reference for almost all NODDI-based measures (see Figure 5.13(a,b)) including the angular deviation of the estimated principle fiber direction (Figure 5.13c) for all motion corruption levels (25, 50, 75 and 90%). At the lowest corruption level (25%), the corrupted volumes apparently did influence the fitting results only to a small extent and the average error in all measures of the *uncorrected* original data outperformed the proposed method. With increasing corruption levels, the error of the *uncorrected* data increased with the steepest slope. For the NODDI index ODI (Figure 5.13a) at the lowest corruption

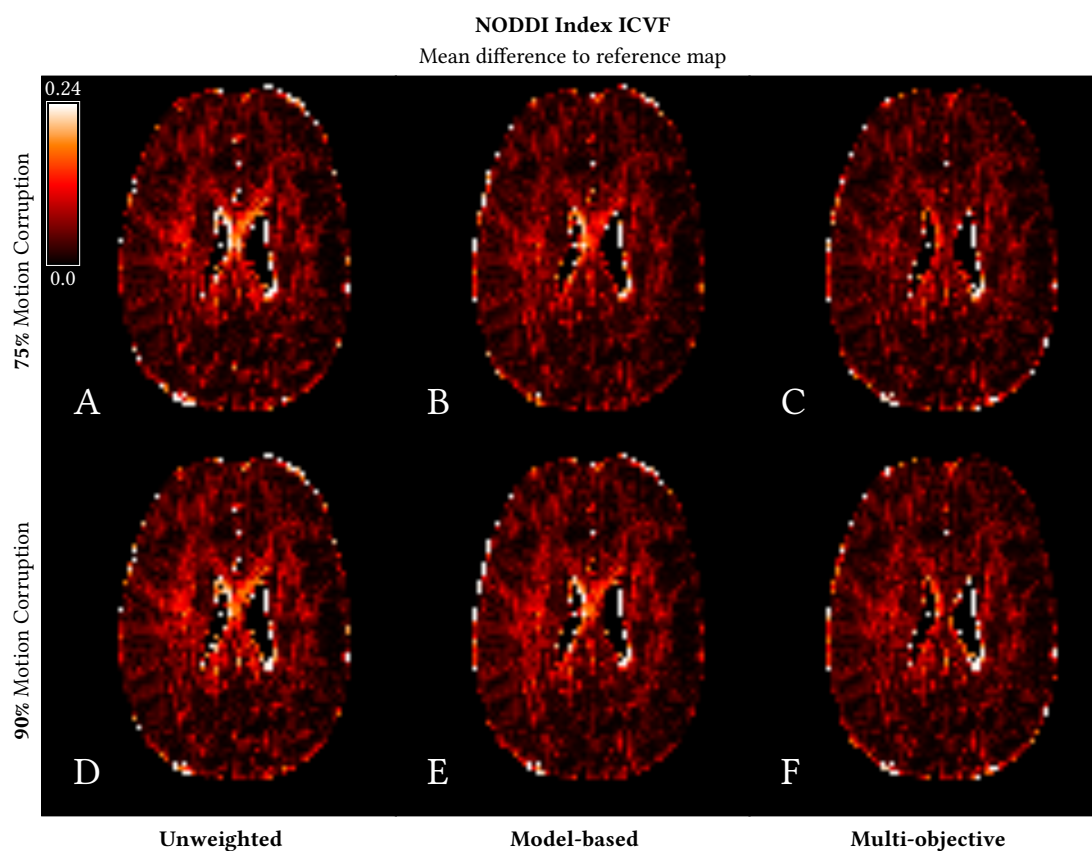


Figure 5.11: Visualization of the error in the NODDI Index ICVF for the two highest motion corruption levels. An axial slice from the in-vivo dataset, showing the mean difference between the ICVF values fitted to the corrected datasets and the reference values. (A,D) For the datasets corrected with the single-objective unweighted reference method, (B,E) with the single-objective model-based reference method and (C,F) with the multi-objective reference method. The mean value is computed across the bootstrap samples at the depicted corruption level—75% in the upper row, 90% in the lower row. The identical colormap with the range [0.0, 0.24] used for all images is shown in the top-left corner.

level, the model-based reference method yielded a slightly smaller error than the multi-objective correction method.

Visualizations of the mean differences of the estimated maps to the reference are shown in Figure 5.11 for the NODDI index ICVF and in Figure 5.12 for the index ODI. The high error areas for the ICVF index were located in the middle areas, the *corpus callosum* and at the borders of the lateral ventricles.

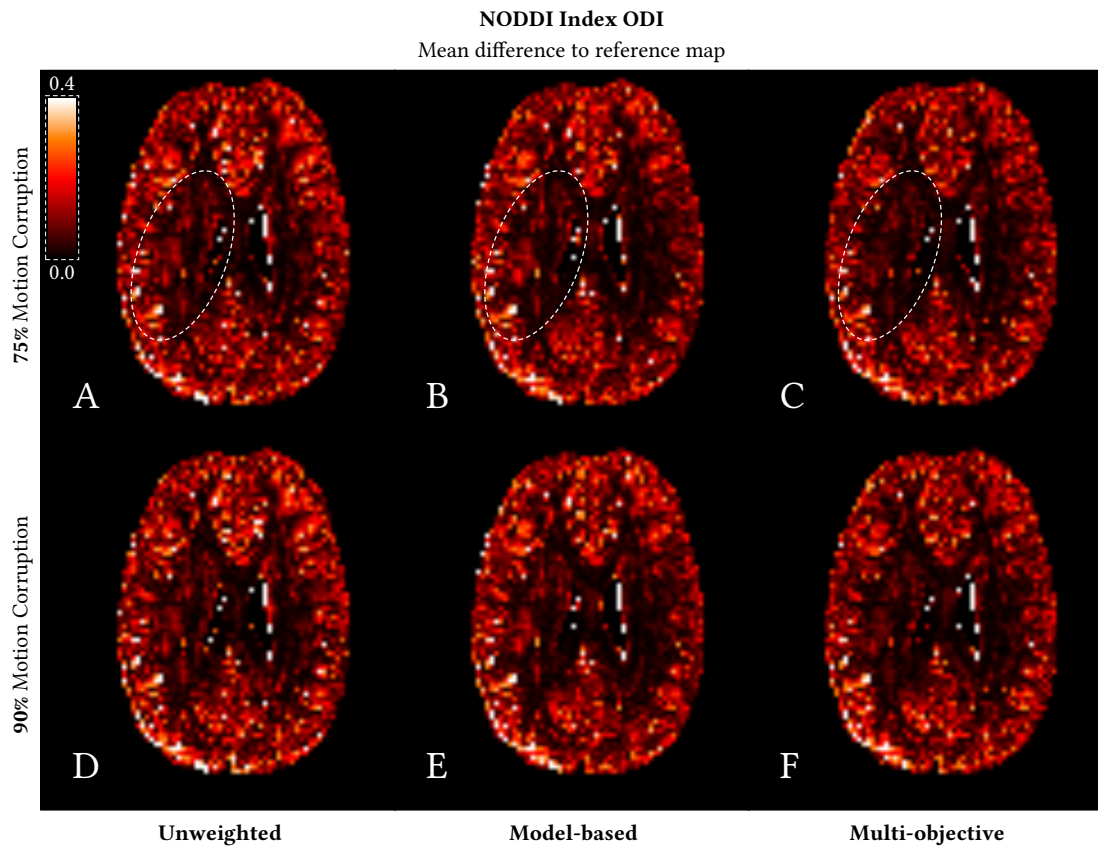


Figure 5.12: Visualization of the error in the NODDI Index ODI for the two highest motion corruption levels. An axial slice from the in-vivo dataset, showing the mean difference between the ODI values fitted to the corrected datasets and the reference values. (A,D) For the datasets corrected with the single-objective unweighted reference method, (B,E) with the single-objective model-based reference method and (C,F) with the multi-objective reference method. The mean value is computed across the bootstrap samples at the depicted corruption level—75% in the upper row, 90% in the lower row. The identical colormap with the range $[0.0, 0.4]$ used for all images is shown in the top-left corner.

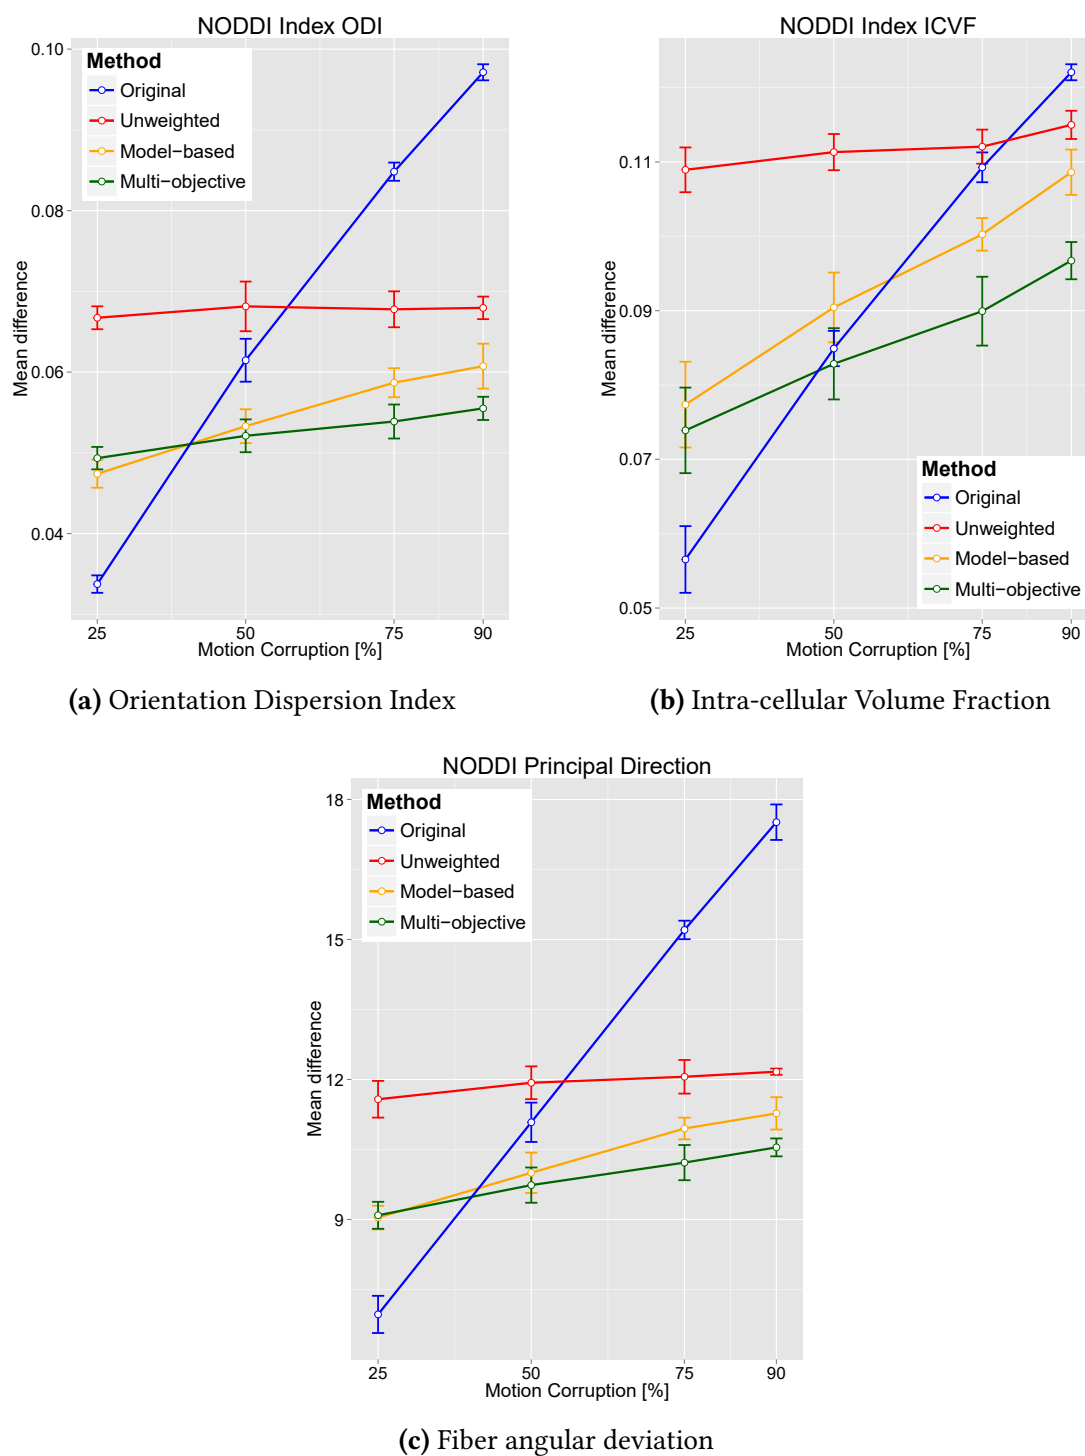


Figure 5.13: Reference-based evaluation of the correction schemes on basis of NODDI-derived measures in dependency to the level of motion corruption. (a) The mean error in the scalar index ODI, (b) the mean error in the scalar index ICVF and (c) fiber orientation deviation between the reference and the NODDI-fitted data. The estimated values are compared with the ground-truth maps without motion. The data points are plotted at the mean value, the error bars indicate the variance of the estimated parameters in the different bootstrap configurations.

5.2.3 Particle History

The aggregated particle history graph for the registration of the in-vivo data at the different motion corruption levels is shown in Figure 5.14. The model-based objective showed the most appearances at the two last final pyramid levels, whereas the unweighted objective provided the most leading particles at the first, the most coarse pyramid level. Similar to the behavior observed on the synthetic datasets with increasing artifacts and noise levels, the dominance of the model-based objective narrowed with increasing level of motion corruption. At the two highest levels, the importance of the individual objectives almost equalized at all levels.

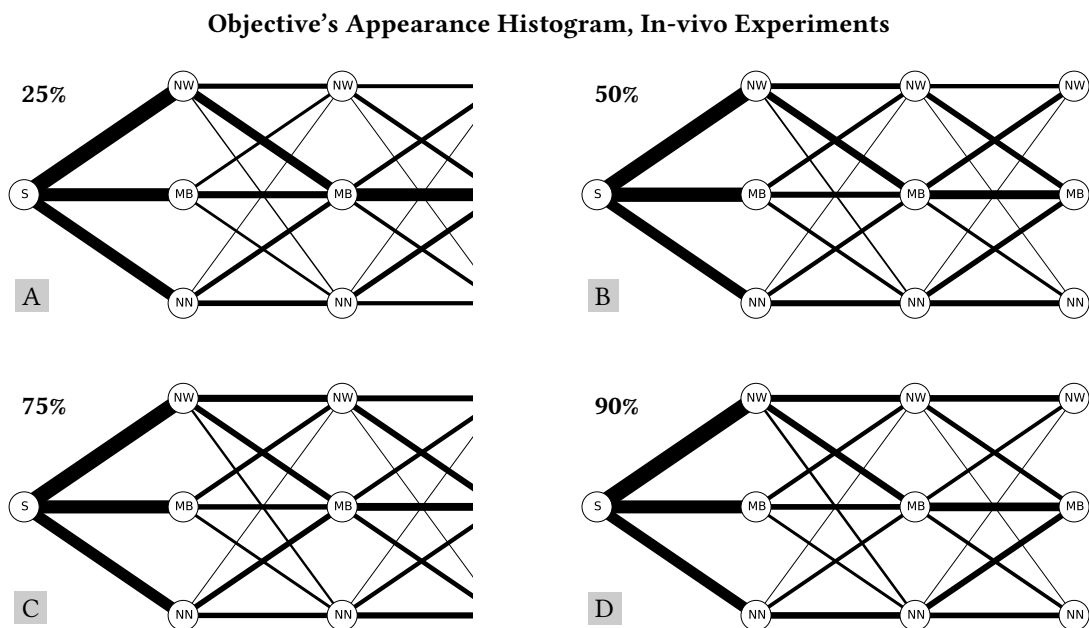


Figure 5.14: Aggregated history of particles after the optimization process for the in-vivo datasets at different motion corruption levels marked in the upper left corner of each subplot (A-D). The node (S) represents the start position. During optimization, the particles move from left to right through the three different PSO updates, each associated with a leading particle. The edges in the graph represent the path of the particles with respect to the different leaders they followed. The edge thickness represents the number of appearances summed over all particles' history.

5.2.4 Fiber Tracking

The results of the fiber reconstruction of the *fornix* tracts are shown in Figure 5.15, the reconstruction of the *cingulum* tracts are shown in Figure 5.16. For both tracts reconstructions, the proposed multi-objective method provided the most visually similar

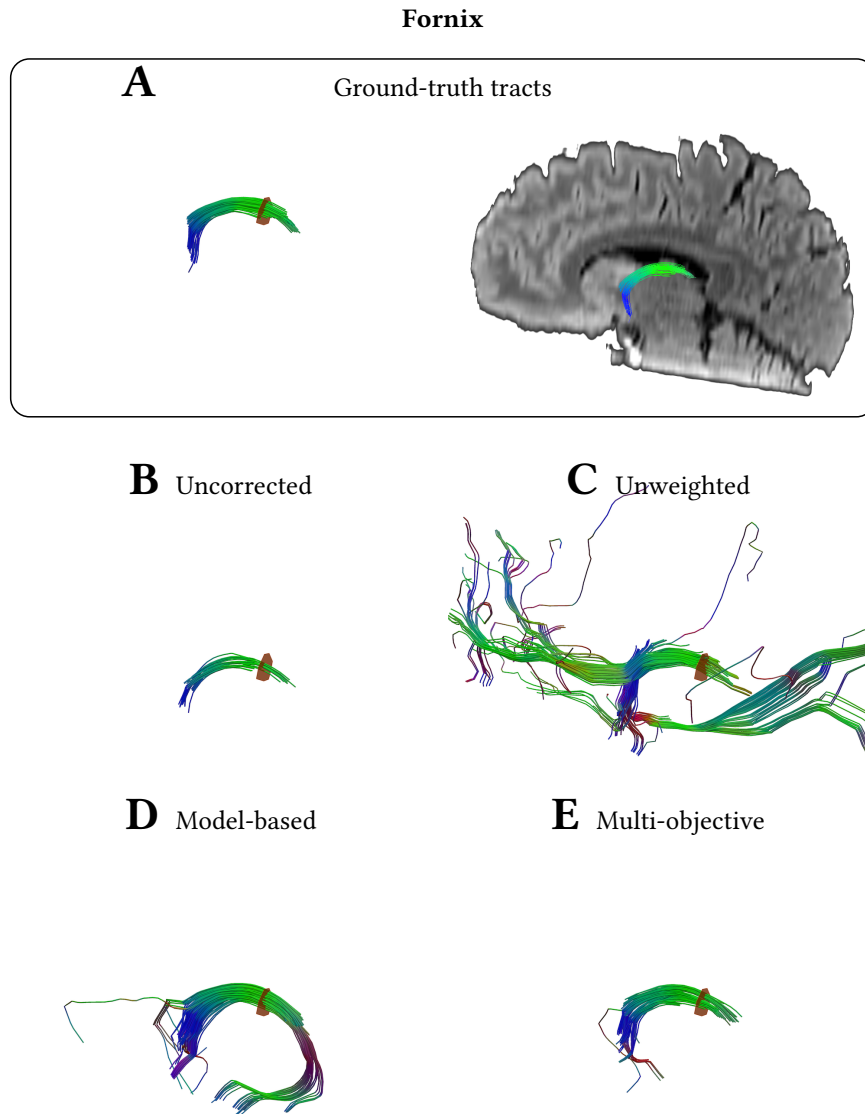


Figure 5.15: The *fornix* tracts reconstructed from an in-vivo dataset with 75% motion corruption after motion and distortion correction. Each correction method was applied to the dataset and the CSD tracking was performed on the $b=3000 \text{ s} \cdot \text{mm}^{-2}$ shell. The tracts were extracted using a passing ROI defined in the unmoved reference dataset (acquisition 1). (A) The reference tracts in a 3D-view and in a 2D-view overlaid over a sagittal slice of the unweighted image. (B) Tracts reconstructed from the uncorrected datasets, (C-E) tracts reconstructed from the datasets corrected with the respective methods.

result compared to the reference reconstruction. The single-objective methods produced a higher amount of spurious fibers, whereas the tracking on the uncorrected data resulted in an incomplete reconstruction with overall low number of tracts. In the *fornix* tracking experiment, at first sight, the reconstructed tracts from the uncorrected datasets resembled the reference tracts. However, it can be observed that the number of tracts are reduced and the reconstruction of the frontal arc is missing. Furthermore,

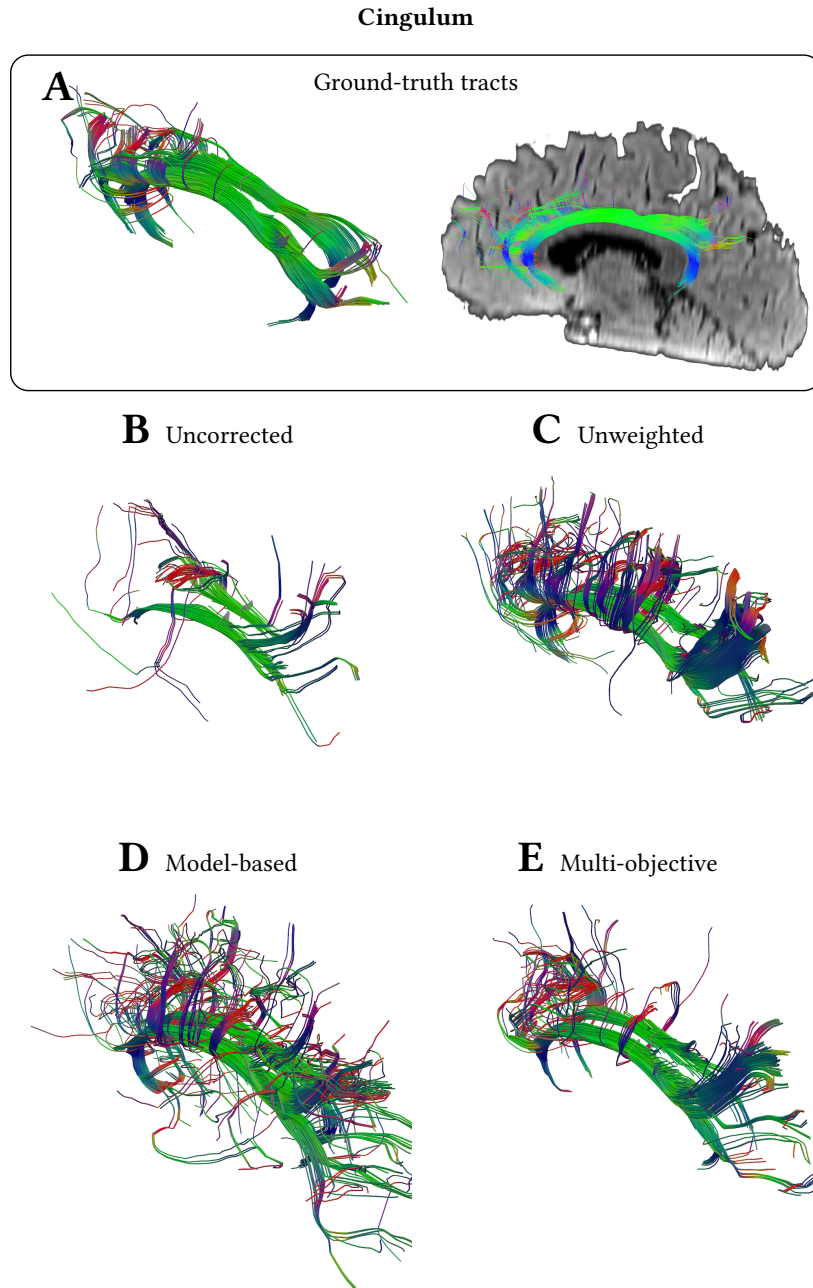


Figure 5.16: The *cingulum* tracts reconstructed from an in-vivo dataset with 90% motion corruption after motion and distortion correction. Each correction method was applied to the dataset and the CSD tracking was performed on the higher shell ($b=3000 \text{ s}\cdot\text{mm}^{-2}$). The tracts were extracted using a passing ROI defined in the unmoved reference dataset (acquisition 1). (A) The reference tracts in a 3D-view and in a 2D-view overlaid over a sagittal slice of the unweighted image. (B) Tracts reconstructed from the uncorrected datasets, (C-E) tracts reconstructed from the datasets corrected with the respective methods.

this reconstruction appears only for the *for**nix* tracts, which form a rather isolated bundle structure and thus not intermixed with signal from other white matter areas in case of imprecise correction. The tracking results on the more “centrally” located *cingulum* reflect once more the differences in the correction precision of the evaluated method. On the uncorrected dataset, the fiber tracking failed to reconstruct the major parts of the *cingulum* tracts.

Chapter 6

Discussion

The central objective of the thesis was to develop a robust retrospective method for correction of motion and acquisition artifacts in diffusion MRI data. The proposed, fundamentally new distortion correction approach follows a multi-objective memetic search optimization principle that allows to exploit the different intensity similarity relationships between gradient volumes in a diffusion-weighted image. The strength of the method was demonstrated quantitatively on realistic motion and eddy-currents-distorted synthetic datasets at different artifact and SNR levels. The advantages of the proposed method were further illustrated by a reference-based evaluation on in-vivo datasets with different levels of motion corruption. In the synthetic datasets, the TRE of the proposed method stayed below the voxel size for datasets with moderate artifacts and below double voxel size in the more severely artifacted data. In the in-vivo experiments, the proposed method yielded a mean error in the NODDI-derived measures well below the errors by the single objective correction approaches. The precision differences between the single-objective and the proposed multi-objective methods were further illustrated by the results of a fiber reconstruction performed on the corrected data. In conclusion, the main contributions of this thesis are:

- (A) A novel and robust multi-objective image registration method based on the memetic optimization principle, combining local search phases with global PSO phases and capable of exploiting the different intensity similarity relationships between the gradient images in a diffusion-weighted dataset.
- (B) A method for constructing realistic synthetic diffusion-weighted datasets with controlled addition of motion and acquisition artifacts, allowing a quantitative evaluation of retrospective correction approaches in diffusion MRI.

- (C) A quantitative evaluation of the proposed multi-objective method in terms of the registration error, as well as a qualitative evaluation of the impact on higher-order signal modeling (NODDI) as well as fiber tractography on both synthetic (B) and real in-vivo data.

The following sections discuss these contributions in further detail.

6.1 Evaluation Data

The realistic synthetic datasets which allowed the quantitative evaluation of the correction method in the first place were constructed by extending the simulation framework Fiberfox in two parts. Firstly, the bulk motion was simulated by applying the ground-truth rigid transforms to the tractogram used by Fiberfox for the signal simulation. Secondly, the eddy-currents-caused artifacts were simulated by adding an additional signal decay to the k -space simulation.

One limitation of the data simulation used for generating synthetic datasets is that it does not model all motion-induced effects that can occur in real acquisitions. The presented extension of the simulation framework Fiberfox allows to add motion only between single volume acquisitions, omitting artifacts caused by rapid motion during k -space readout such as signal drop-outs or even complete signal extinction in a single slice. However, such artifacts were present in the in-vivo datasets with similar results of the experiments regarding the improved correction precision achieved by the multi-objective correction method.

The missing ground-truth deformations for the simulated eddy currents effects were another limitation of the synthetic experiments. However, the deformation cannot be obtained in an analytical way since it originates from an exponential signal decay in the Fourier-space. Thus, using a non-linear deformable registration was a logical choice, since the conditions for a precise estimation were given. First, the moving image was available without any additional noise, second, the fixed image had almost identical image contrast since it was constructed by a simulation with identical settings but without artifacts, and third, the ground-truth rigid transform was known and thus the deformation described the pure effects by eddy currents. The non-linear registration uses a metric derived from the MI as well as a gradient descent optimizer, i.e. it is based on a similar registration procedure as the distortion correction methods. For this reason, a minor bias in this process cannot be ruled out completely. However, since a similar bias is applied to all correction methods, and since the reference estimation provided

a high-quality result by means of low voxel-wise error that was visually verified, a significant influence of such potential bias on the TRE evaluation is unlikely.

To enable the voxel-wise comparison to the reference images in the in-vivo experiments, all datasets must occupy an identical registration reference space that is defined by the unweighted ($b=0 \text{ s}\cdot\text{mm}^{-2}$) image. To meet this requirement all datasets were constructed with the undistorted unweighted volume from the reference acquisition. Even though a distortion may occur also in the unweighted volume, this choice is unlikely to bias the evaluation. The eddy-currents-caused effects scale proportionally with the b -value and thus the unweighted volume is predominantly prone to the rigid bulk motion which in turn only shifts and rotates the reference space.

6.2 Synthetic Experiments

The results of the single objectives on the moderately artifacted simulated datasets at the lower levels of added motion resembled the common experience in dMRI motion and distortion correction. While the average TRE in the lower b -shell ($b=1000 \text{ s}\cdot\text{mm}^{-2}$) stayed well below the single voxel size, the unweighted reference objective failed to precisely correct the higher b -shell. The model-based objective provided a good alignment of both shells with average TRE slightly above the single voxel size.

The severe artifacts in the simulated datasets with lower motion levels seem to force particularly the model-based objective into outliers in the $b=3000 \text{ s}\cdot\text{mm}^{-2}$ shell, yielding an average TRE even higher than the unweighted reference objective. However, for these datasets, the global update phase was capable to capture these, maintaining an average TRE below the double voxel size. At first, the severely artifacted simulated datasets with the highest added motion revealed the limits of the multi-objective approach. Although the average TRE was—compared to the single-objective methods—significantly reduced, it remained above the double voxel size. As demonstrated by the fiber tracking experiments, errors of such magnitudes affect the further processing steps to a great extent. The results of the fiber reconstruction on the synthetic datasets on the example of the *fornix* tract clearly reflected the average TRE in the higher b -shell after correction. The visual impression, particularly when focusing on the amount of spurious fibers leads to the conclusion that a registration error above the double voxel size has a strong negative influence on the resulting tracts. The tracts most similar to the ground-truth and with the least spurious fibers were reconstructed on the datasets corrected by the multi-objective method with 12 particles, with initial average TRE

above the single voxel size. However, this initial error still negatively influenced the tracking and the reconstructed tracts did not resemble the ground-truth to a full extent.

Experiments on the distribution of the initial particle positions showed no significant differences between the four evaluated choices for the variance. A possible reason for the weaker influence of the initial distribution is the overall attraction to strong solution candidates. Once a strong candidate is reached during local search, it then dominates the remaining candidates in the global PSO update phase. A further improvement in terms of the final TRE could be achieved by replacing the current “unconstrained” parameter initialization with one bounded by an estimate of the extent of motion and distortion. This could be for example provided by repeatedly acquiring multiple unweighted volumes at regular intervals through the complete acquisition session and estimating their alignment.

The frequency of occurrence of the objective functions in the leading role reflects the TRE achieved on the synthetic data. Among the reference, single-objective methods, the model-based approach yielded the lowest error and had a high number of appearances in the leading particle counts, particularly in the moderately artifacted dataset. In the more severely artifacted dataset and at higher b -values, however, the model-based approach lost its leading role because the model fitting was hampered due to misalignments in the lower b -shell after initial registration with the unweighted reference objective.

6.3 In-vivo Experiments

In the quantitative evaluation on synthetic datasets, the average TRE by the nearest-neighbor reference objective ranges between the average error of the unweighted and the model-based method, which are also the two mostly applied methods in published studies. Thus, in the in-vivo experiments, the multi-objective approach was evaluated only against the unweighted and the model-based single-objective correction methods.

At the lowest motion corruption level (25%) in the in-vivo experiments, both the model-based reference objective as well as the proposed method were outperformed by the uncorrected original data, which is a consequence of robust model fitting in the NODDI Toolkit. Also the resampling procedure applied during correction could be partially responsible for the increased error when compared to the uncorrected original dataset. Also the tensor fitting seemed to perform robustly at the lowest level, as the model-based objective slightly outperformed the multi-objective method in the measured ODI

index, which is also reflected by the particle history graphs showing a particularly high number of appearances of the model-based objective at the finest pyramid level.

The computed error in the raw data provides a further insight on the performance of the evaluated methods. The rotationally-dependent diffusion-weighted signal and the uncorrected eddy currents effects in the reference dataset limit the explanatory power of the direct signal comparison. However, the true eddy-currents-caused distortions affect only the front and the rear areas, as shown by the high error rim in the model-based and the multi-objective method at the 25% corruption level. The signal dependency to rotations, in contrast, affects the complete white-matter regions and thus parts of the overall increase in the signal differences can be accounted to this effect. The effect size is further reduced by considering the mean error across the randomly assembled bootstrap samples, though. For the unweighted reference objective, the measured high error already at the lowest corruption level reflects the known limitations of the method's precision in correcting higher b -shell images.

The areas of higher error in the NODDI index ICVF correspond to the differences in the raw signal in the data located in the mid-area of the displayed axial slice (in the *corpus callosum*). In this area, the eddy-currents-caused distortions, which diminish the significance of the reference-based evaluation, are less pronounced and the measured differences can thus be accounted to precision of the evaluated methods in correcting the motion-caused distortions. In accordance with the measured smaller mean differences in the NODDI index ODI, the visualized error is not so dominant. In general, the ODI values are low in homogeneous and dense white-matter structures like the *corpus callosum*. Though, the error was increased inside the area left to the *corpus callosum* which is occupied by multiple fibers with different orientations (the *superior longitudinal*, *superior fronto-occipital fasciculus* and the *superior corona radiata*).

6.4 Local Search Method

In the presented implementation, the local search phase is iterated with the PSO update steps at a few points naturally provided by the pyramidal registration scheme. With this choice, a relatively high importance is assigned to the local search compared to the global update phase. Both quantitative and qualitative experiments support the claim that the local search provides robust results in the most cases and thus the PSO update phase can be employed only at the few points to efficiently detect outliers.

As a consequence of the strong emphasis of the local search method, the overall performance of the memetic search relies on the ability of the local search to locate the true minimum, which is one of the necessary “boundary” conditions for a successful image registration method. This ability is of particular importance in the last, full-resolution pyramid level. If the local search violates this *global minimum localization* condition and is failing for multiple objectives at a high number of images, it first provides candidate solutions of poor quality and second, the registration outliers cannot be prevented in most cases, regardless of the starting position provided by the global *PSO* update. Such behavior was observed in the experiments on the synthetic datasets with severe artifacts and the highest added motion. To quantify the influence of the choice of local search method on the result, an additional experiment on a




Figure 6.1: Preliminary results illustrating capabilities of different local search methods on a more severely artifacted in-vivo dataset at 75% motion corruption level. (A,B) Higher amount of outliers by the single-objective methods with the Regular Step Gradient Descent (RSGD) optimizer. (C) Incomplete outlier reduction by the multi-objective method with RSGD-based local search. (D,E) Improved correction by the single-objective methods using the Adaptive Stochastic Gradient Descent (ASGD) optimizer with outliers only in the higher b -shell that are (F) corrected by the multi-objective method with ASGD-based local search. The contrast in the high b -value part was enhanced for printing.

in-vivo dataset with severe motion artifacts was performed. In this dataset, a weak performance was encountered for the local search using the RSGD optimizer. As shown in Figure 6.1(A,B), the registration using RSGD optimizer failed to align the data for both unweighted and model-based objectives, even for lower b -value images ($b=1000 \text{ s} \cdot \text{mm}^{-2}$). As a consequence, the multi-objective method (Figure 6.1(C)) could not correct all registration outliers, although it managed to remove some of them.

For the additional experiment, the RSGD optimizer in the local search was replaced by the ASGD optimizer. As shown for the unweighted reference objective on Figure 6.1(D), the ASGD succeeded in aligning the lower b -shell, while producing several outliers in the higher one. Thus, the ASGD-based local search behaved like the RSGD-based one in the evaluation experiments. This in turn provided a solid base for the model fitting and the extrapolation of the reference images for the model-based objective, resulting in a better alignment of the higher b -shell (Figure 6.1(E)). However, the model-based reference objective seemed to overestimate the scaling as shown by the almost complete rim in the mean difference image (see Figure 6.2(B,E)).

In contrast to the additional in-vivo experiment, the results on the highly artifacted simulated datasets show that the local search with RSGD optimizer is capable of locating the minimum, as the average TRE drops rapidly and almost reaches the single voxel

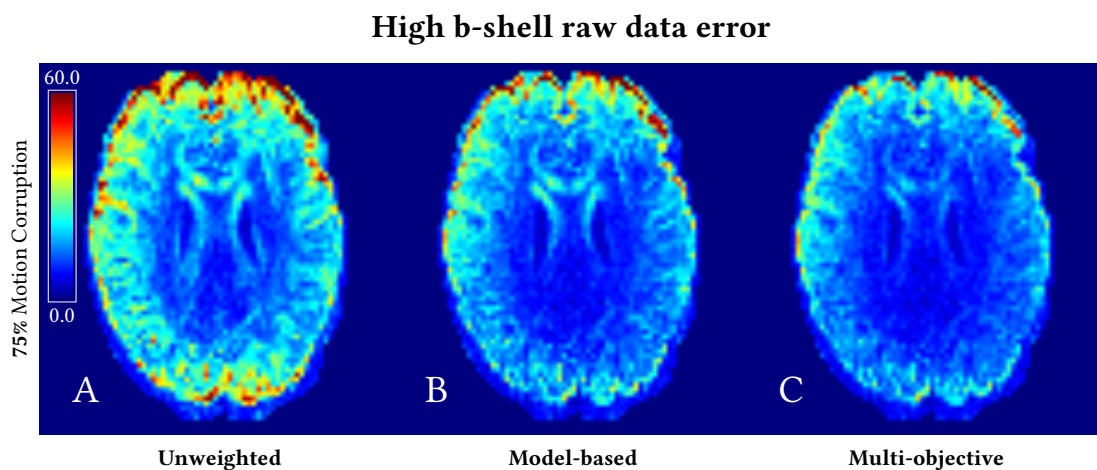


Figure 6.2: Preliminary results of ASGD-based local search methods evaluation on a more severely artifacted in-vivo datasets at 75% motion corruption level. Each figure shows the mean raw data error in the higher b -shell after correction with the ASGD-based local search. (A) Overall high error by the unweighted reference method. (B) High error rim by the model-based reference method indicates a imprecise scaling in addition to corrected eddy currents effects. (C) The error in the model-based objective resulted almost entirely from the corrected eddy currents. The identical colormap with the range $[0.0, 60.0]$ used for all images is shown in the top-left corner.

size, although the single objectives yielded an average registration error of above three times the voxel size (6 mm). This seems not to hold for the most complex case (Π_{C10}), where a reduction of the average TRE is achieved by the multi-objective method, this, however, still stayed above 5 mm.

In conclusion, the additional in-vivo experiment has revealed the necessity of using a stable local search method within the memetic search multi-objective approach. It also demonstrated the generality of the proposed approach by showing the improved performance with a different local search component.

6.5 Global Optimization Method

The choice of using only the few points provided by the pyramidal registration scheme was made under the assumption that the local search is able to provide robust results in most cases. The PSO update phase can then efficiently detect outliers as it was demonstrated in the synthetic data experiments and by the in-vivo experiments on subject 1. As discussed in the previous section on the example of the RSGD optimizer in the special case of the severely artifacted in-vivo dataset, a local search unable of locating the minimum precisely will have an impact on the performance of the complete memetic search, especially because it provides the final candidate solution set. Still, the multi-objective method achieved a partial correction of registration outliers, which leaves room for expectations in a stronger emphasis of global PSO update phase. This is also motivated by the improvement of the average TRE in the synthetic data experiments with increasing number of particles used, which is one of the central parameter choices improving the exploration capabilities provided by the global phase on one hand, but also the parameter with a strong influence on the processing time on the other. The preliminary results of the correction performance in the presence of a weak local search discussed in the previous section allowed also to inspect the effects of higher global search emphasis in such configurations. Increasing the number of particles resulted in a further outlier reduction especially in the higher b -shell, however, the resulting alignment is not satisfactory as documented in Figure 6.3.

Hence, the poor performance of a weak local search method cannot be fully compensated by the increase of particles used and this cannot be expected either from increasing the number of global phase steps which would more and more resemble the pure PSO approach, which is known to lack the precision in locating the global minimum [112]. However, with a suitable local search, the proposed multi-objective

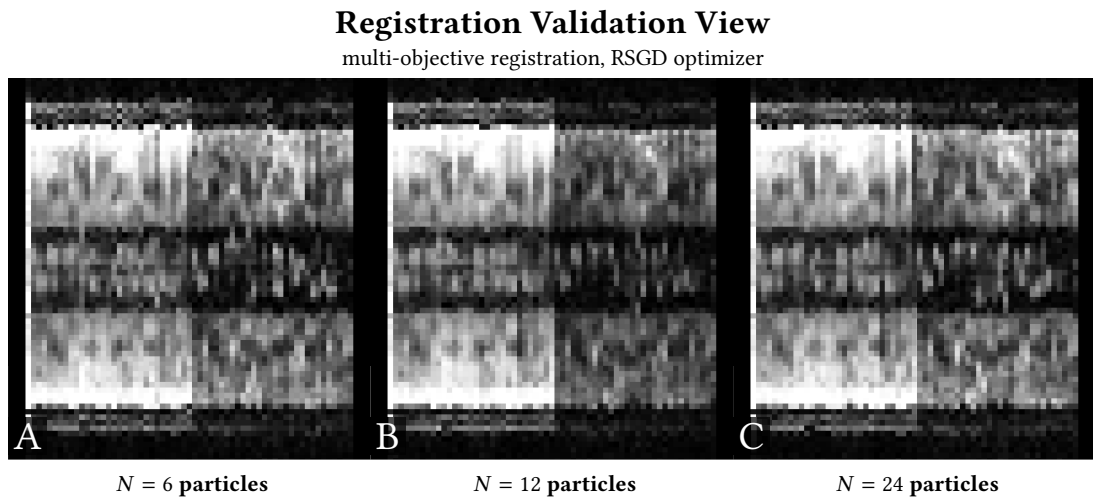


Figure 6.3: Preliminary results of the correction performance of a weak local search (RSGD optimizer) in dependency to different number of particles used in the global update phase. All figures show the additional, more severely artifacted in-vivo dataset at 75% motion corruption level after correction. (A) Visible multiple registration outliers in the both b -shells for $N = 6$ particles. (B) Improved correction in both shells with $N = 12$, but incomplete outlier reduction and (C) further outlier reduction in the lower shell with $N = 24$ particles, still uncorrected outliers in the higher shell. The contrast in the high b -value part was enhanced for printing.

method provides an improved precision already with low numbers of particles and with the global phase updates at the end of each pyramid level.

The maximal number of pyramid levels strongly depends on the matrix size and the number of acquired slices. A typical diffusion Magnetic Resonance (MR) head acquisition with 2-2.5 mm isotropic voxel spacing will comprise of at least 40 slices, which allows for a pyramid construction with 3 (or even 4 levels) as described. This condition is not necessarily fulfilled in other acquisitions like a thick-slice diffusion MR of the spinal cord, where the small number acquired axial slices would be strongly limiting the pyramid construction. However, in such cases an anisotropic pyramid schedule could be used to still allow for at least 3 pyramid levels.

Together with the number of particles, the number of pyramid levels figures as one of factors in the additional term of the global phase time complexity. However, for the presented configurations using only few particles, the prevalent influence on the overall runtime remains in the linear increase with the number of particles.

6.6 Relevance for Diffusion MRI Analysis

The NODDI-based indices evaluated in the in-vivo experiments were already used as biomarkers in group studies on neuro-degenerative diseases: Kamagata *et al.* [56] have considered the ODI and ICVF values in Parkinson’s diseased patients and the evaluated regions were significantly altered with mean group differences of 0.06 for the ICVF index and 0.03 for the ODI index. With this result, the discriminative values lie within the range of the voxel-wise errors produced by the correction schemes. Thus, the choice of motion and distortion correction method is of critical importance in this context.

Although the processing time of the multi-objective correction scheme increases with the number of particles, the applications in micro-structural tissue analysis or in connectomics are in the most cases not time-sensitive and rather valuing precise evaluation. Furthermore, the particles’ local search phases are independent of each other and thus allow for a reduction of the processing time through highly-parallel implementation.

The multi-objective method will provide only a little improvement for DTI acquisitions with single shell and moderate b -values. For such data, the unweighted reference objective is capable of providing a robust distortion correction, which is mirrored by the popularity of this correction approach in DTI studies.

For another popular imaging technique, the single-shell HARDI acquisition with dense q -space sampling, the nearest-neighbor objective in the presented definition would coincide with the unweighted reference objective. However, the objective is not restricted to gradient similarity across b -shells and the dense sampling of the gradient directions in the HARDI-case allows a different definition of the similarity operator selecting the nearest-neighbor on the same b -shell.

6.7 Outlook and Further Work

Very recently, the research in the motion and distortion correction have brought forward a novel extrapolation method [79], which could also improve the proposed multi-objective approach either as a further objective function or as a replacement of the currently employed model-based metric. In combination with the recently published result on simulation of diffusion MRI validation data with improved simulation of non-white-matter areas [43] (*note*: at the time of writing, the data was not available), this could provide a further insight on the presented memetic search based approach.

In the application to motion and distortion correction in diffusion MRI, the proposed method can fully utilize the different intensity relationships and thus the different registration objectives. However, the method itself is not only restricted to applications where different image metrics between the moving and the fixed image are available. For example in case of registering images with locally mis-matching contrast areas, like in the presence of lesions, the different objectives for the memetic settings could be formulated by restricting the metric computation only to a sub-region of the original domain. In this way, some of the sub-regions will not contain the critical area and likely provide a better transform estimate.

Chapter 7

Conclusion

Robust data preprocessing becomes increasingly important in diffusion-weighted MR imaging as a consequence of the growing number of acquired images per subject that need to be processed to provide a stable data basis for the higher-order signal models. Application of such models becomes increasingly common in dMRI analysis as they provide a better insight into micro-structural tissue properties. Beside the steadily growing number of acquired images, which directly increases the acquisition time and the proneness to subject's motion, these models also require acquisitions at higher b -values, which are first more susceptible to eddy-currents-caused artifacts and second come with lower signal-to-noise ratios. Altogether, a robust motion and artifact correction becomes a critical prerequisite for all subsequent dMRI analysis steps. In consequence of the challenging image contrasts at high b -values, current retrospective correction approaches, which are based on pairwise image registration, can easily produce a remarkable amount of registration outliers.

Central objective of this work was to address this behavior by proposing a novel multi-objective registration scheme based on the memetic optimization principle. The memetic formulation of the approach, that is based on alternating local search (pairwise image registration) and a global update (particle swarm optimization) phases, allowed to simultaneously follow all available image metrics without any *a priori* weighting. The design of the method followed the main hypothesis of this thesis, that a more consequent exploration of the different intensity relationships between the gradient volumes of a diffusion-weighted acquisition could enhance the capabilities of solving such challenging registration problems. An optimal method would be able to prevent registration outliers by choosing the right combination of the different objectives and efficiently lead the way towards the global minimum.

The precision and the significantly reduced amount of registration outliers on realistic synthetic data, simulated with different levels of motion and artifacts distortion, supported the main hypothesis. The proposed method can successfully correct for outliers even when the single registration objectives yielded an average registration error above the double voxel size. The dynamically changing role of the different objectives within the multi-objective settings, depending on the resolution and the level of corruption, could not be achieved when considering only an *a priori* weighting instead of the memetic search approach.

The improved precision in the in-vivo data correction, particularly at higher b -values, had an impact on scalar indices derived from the higher-order NODDI model as well as on fiber tracking. While the NODDI model is increasingly applied in the search for biomarkers for disease-induced tissue alterations and the correction is likely to affect the discriminative value of such, uncorrected artifacts in the fiber tracking can influence the network-graph construction and following evaluations as done in connectomics studies.

One of the most noteworthy issues to consider in the configuration of image registration based on memetic search is the choice of a suitable local search method. With a well-working local search phase, the memetic search allows a conservative parametrization of the global PSO phase, including the employment of only few particles, to provide a robust registration result. A poor performance of a weak local search method, as was observed on some of the severely artifacted datasets, can be corrected to a certain degree through a higher emphasis of the exploration capabilities provided by the global PSO update phase. However, this issue manifested itself only in the correction of highly artifacted synthetic and in-vivo datasets. Hence, this issue is important rather from a general methodical point of view than for practical considerations in dMRI-based evaluation studies. In these, such artifacts would lead either to exclusion of the subject or to a re-acquisition, since a subject's motion of such magnitude leads also to signal loss across the complete slices. For the moderately artifacted evaluation data, which are far more likely to appear in typical in-vivo acquisition, the proposed conservative parametrization successfully achieved the main goal of the thesis and provided a reduction of the registration outliers.

There are still possible intensity relationships between the gradient volumes of a diffusion-weighted image that were not yet exploited in the multi-objective settings. Taking advantage of the straightforward extensibility, inclusion of new objectives could further improve the robustness of the presented method.

The presented multi-objective approach to retrospective motion and artifact correction in dMRI datasets enables the application of recent higher-order diffusion models in a robust and stable fashion. This will play an important role not only for advanced characterization of micro-structural tissue properties but also for the subsequent diffusion MRI analysis approaches like fiber tractography and connectomics, including the clinical implication of such studies.

The robustness of the multi-objective approach could further prove valuable in the registration of medical image data with non-trivial intensity relationships, like in the presence/absence of lesions in pre- and post-operative acquisitions.

Thanks to the general formulation as a meta-heuristic approach, the presented method is not only restricted to retrospective correction of diffusion-weighted images. Hence, this research will hopefully inspire future work in meta-heuristic optimization in diffusion MRI as well as image registration in general.

Appendix A

Own Publications

The proposed multi-objective method and some of the results shown in this thesis were published in the journal *IEEE Transactions on Medical Imaging* [A]. Further, a preliminary approach to the construction of realistic validation data with motion artifacts was presented on a conference and published within the proceedings [B].

[A] **Jan Hering**, Ivo Wolf, and Klaus H. Maier-Hein, *Multi-objective Memetic Search for Robust Motion and Distortion Correction in Diffusion MRI*, *IEEE Transactions on Medical Imaging* (2016), [ePub ahead of print].

[B] **Jan Hering**, Peter F. Neher, Hans-Peter Meinzer, and Klaus H. Maier-Hein, *Erzeugung von Referenzdaten für Kopfbewegungskorrektur in Diffusion-MRI*, *Bildverarbeitung für die Medizin 2014* (Thomas Martin Deserno, Heinz Handels, Hans-Peter Meinzer, and Thomas Tolxdorff, eds.), *Informatik aktuell*, Springer Berlin Heidelberg, January 2014, pp. 342–347.

The following publications present preliminary results or related studies in the field of dMRI processing.

Peer-reviewed conferences

Jan Hering, Ivo Wolf, Tawfik Moher Alsady, Hans-Peter Meinzer, and Klaus H. Maier-Hein, *A Memetic Search Scheme for Robust Registration Of Diffusion-Weighted MR Images* in *Bildverarbeitung für die Medizin 2015* (Thomas Martin Deserno, Heinz Handels, Hans-Peter Meinzer, and Thomas Tolxdorff, eds.), *Informatik aktuell*, Springer Berlin Heidelberg, 2015, pp. 113–118.

Jan Hering, Ivo Wolf, Hans-Peter Meinzer, and Klaus H. Maier-Hein, *Model-based motion correction of reduced field of view diffusion MRI data* in SPIE Medical Imaging 2014 (Robert C. Molthen and John B. Weaver, eds.), International Society for Optics and Photonics, p. 90381L.

Jan Hering, Peter F. Neher, Bram Stieltjes, and Klaus H. Maier-Hein, *DTI Tractography Challenge 2014 - MITK Global Tractography* in Proceedings of 2014 MICCAI Tractography Challenge, Boston, USA.

Jan Hering, Ivo Wolf, Hans-Peter Meinzer, Bram Stieltjes, and Klaus H. Maier-Hein, *A Quantitative Evaluation of Errors Induced by Reduced Field-of-View in Diffusion Tensor Imaging*, Computational Diffusion MRI and Brain Connectivity, Springer, 2014, pp. 35–44.

Conference abstracts

Jan Hering, Peter F. Neher, Hans-Peter Meinzer, and Klaus H. Maier-Hein, *Construction of ground-truth data for head motion correction in diffusion MRI* in Proceedings of International Society of Magnetic Resonance in Medicine (ISMRM), 2014.

Co-authored publications

Caspar J. Goch, Bram Stieltjes, Romy Henze, **Jan Hering**, Luise Poustka, Hans-Peter Meinzer, and Klaus H. Maier-Hein, *Quantification of changes in language-related brain areas in autism spectrum disorders using large-scale network analysis*, International Journal of Computer Assisted Radiology and Surgery **9** (2014), no. 3, 357–365.

Acknowledgements

First of all, I want to express my gratitude to my supervisors Prof. Dr. Bernd Jähne and Prof. Dr. Hans-Peter Meinzer (in the beginning) and PD. Dr. Klaus H. Maier-Hein (later on) who always offered me advice and support. I further thank the Faculty of Mathematics and Computer Science for the opportunity to work on this thesis. This work was conducted at the *Division of Medical and Biological Informatics* (MBI) at the German Cancer Research Center (DKFZ) in Heidelberg and was carried out with the support of the German Research Foundation (DFG) as a part of the research project WO 1218/3-1.

A special thanks goes to my advisors PD Dr. Klaus H. Maier-Hein and Prof. Dr. Ivo Wolf for their support in realization of the project, for the thorough discussions of new research ideas as well as for entrusting me the freedom to explore and realize them, while keeping up my motivation. I also want to thank all the colleagues at the MBI department, and especially those from the H832 office and the *Medical Image Computing* group, for their effort in shaping such creative working environment with great atmosphere, for their selfless help and the readiness to discuss any possible topic. I also want to thank the colleagues from *Medical Physics* and *Radiology* Departments for the productive collaborations.

Finally, I want to thank my parents, my family and my closest friends for their unconditional support and motivation throughout the years of my whole studies.

Bibliography

- [1] Iman Aganj, Christophe Lenglet, Guillermo Sapiro, Essa Yacoub, Kamil Ugurbil, and Noam Harel, *Reconstruction of the orientation distribution function in single- and multiple-shell q -ball imaging within constant solid angle*, Magnetic Resonance in Medicine **64** (2010), no. 2, 554–566.
- [2] Murat Aksoy, Christoph Forman, Matus Straka, Stefan Skare, Samantha Holdsworth, Joachim Hornegger, and Roland Bammer, *Real-time optical motion correction for diffusion tensor imaging*, Magnetic Resonance in Medicine **66** (2011), no. 2, 366–378.
- [3] A. Alhamud, P.A. Taylor, der Kouwe van, and E.M. Meintjes, *Real-time measurement and correction of both B_0 changes and subject motion in diffusion tensor imaging using a double volumetric navigated (DvNav) sequence*, NeuroImage **126** (2016), 60–71.
- [4] A. Alhamud, M.D. Tisdall, A.T. Hess, K.M. Hasan, E.M. Meintjes, and A.J.W. van der Kouwe, *Volumetric navigators for real-time motion correction in diffusion tensor imaging*, Magnetic Resonance in Medicine **68** (2012), no. 4, 1097–1108.
- [5] Jesper L. R. Andersson, *Geometric distortions in diffusion MRI*, Diffusion MRI: from quantitative measurement to in-vivo neuroanatomy (Heidi Johansen-Berg and Timothy E. J. Behrens, eds.), Elsevier/Academic Press, London, UK ; Waltham, MA, 2nd ed ed., 2014, pp. 63–86.
- [6] Yaniv Assaf and Peter J. Basser, *Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain*, NeuroImage **27** (2005), no. 1, 48–58.
- [7] Yaniv Assaf and Ofer Pasternak, *Diffusion Tensor Imaging (DTI)-based White Matter Mapping in Brain Research: A Review*, Journal of Molecular Neuroscience **34** (2008), no. 1, 51–61.

- [8] 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** (2014), 44.
- [9] Brian B. Avants, Nicholas J. Tustison, Hongzhi Wang, Andrew J. Asman, and Bennett A. Landman, *Standardized Registration Methods for the SATA Challenge Datasets*, *MICCAI 2013 Segmentation Algorithms, Theory and Applications (SATA)* (Andrew J. Asman, Alireza Akhondi-Asl, Hongzhi Wang, Nicholas J. Tustison, Brian B. Avants, Simon K. Warfield, and Bennett A. Landman, eds.), 2013.
- [10] Yu Bai and Daniel C. Alexander, *Model-based registration to correct for motion between acquisitions in diffusion MR imaging*, 2008 5th IEEE International Symposium on Biomedical Imaging: From Nano To Macro, IEEE, May 2008, pp. 947–950.
- [11] Mark E. Bastin, *Correction of eddy current-induced artefacts in diffusion tensor imaging using iterative cross-correlation*, *Magnetic Resonance Imaging* **17** (1999), no. 7, 1011–1024.
- [12] Shani Ben-Amitay, Derek K. Jones, and Yaniv Assaf, *Motion correction and registration of high b-value diffusion weighted images*, *Magnetic Resonance in Medicine* **67** (2012), no. 6, 1694–1702.
- [13] Thomas Benner, André J. W. van der Kouwe, and A. Gregory Sorensen, *Diffusion imaging with prospective motion correction and reacquisition*, *Magnetic Resonance in Medicine* **66** (2011), no. 1, 154–167.
- [14] Jesse A. Brown, Jeffrey D. Rudie, Anita Bandrowski, John D. Van Horn, and Susan Y. Bookheimer, *The UCLA multimodal connectivity database: a web-based platform for brain connectivity matrix sharing and analysis*, *Frontiers in Neuroinformatics* **6** (2012).
- [15] Emmanuel Caruyer, Iman Aganj, Christophe Lenglet, Guillermo Sapiro, and Rachid Deriche, *Motion Detection in Diffusion MRI via Online ODF Estimation.*, *International Journal of Biomedical Imaging* **2013** (2013), 849363.
- [16] Emmanuel Caruyer, Alessandro Daducci, Maxime Descoteaux, Jean-Christophe Houde, Jean-Philippe Thiran, and Ragini Verma, *Phantomas: a flexible software library to simulate diffusion MR phantoms*, *ISMRM*, 2014.
- [17] Lin-Ching Chang, Derek K. Jones, and Carlo Pierpaoli, *RESTORE: robust estimation of tensors by outlier rejection.*, *Magnetic Resonance in Medicine* **53** (2005), no. 5, 1088–1095.

- [18] 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: Informed RESTORE*, Magnetic Resonance in Medicine **68** (2012), no. 5, 1654–1663.
- [19] Yi Shin Chang, Julia P. Owen, Nicholas J. Pojman, Tony Thieu, Polina Bukshpun, Mari L. J. Wakahiro, Jeffrey I. Berman, Timothy P. L. Roberts, Srikantan S. Nagarajan, Elliott H. Sherr, and Pratik Mukherjee, *White Matter Changes of Neurite Density and Fiber Orientation Dispersion during Human Brain Maturation*, PLoS ONE **10** (2015), no. 6, e0123656.
- [20] Daan Christiaens, Frederik Maes, Stefan Sunaert, and Paul Suetens, *Global tractography of multi-shell HARDI data*, Joint annual meeting ISMRM-ESMRMB, 2014.
- [21] Thomas G. Close, Jacques-Donald Tournier, Fernando Calamante, Leigh A. Johnston, Iven Mareels, and Alan Connelly, *A software tool to generate simulated white matter structures for the assessment of fibre-tracking algorithms*, NeuroImage **47** (2009), no. 4, 1288–1300.
- [22] Claire J. Cochrane and Klaus P. Ebmeier, *Diffusion tensor imaging in parkinsonian syndromes A systematic review and meta-analysis*, Neurology **80** (2013), no. 9, 857–864.
- [23] Carlos A. Coello Coello, Gary B. Lamont, and David A. Van Veldhuizen, *Evolutionary algorithms for solving multi-objective problems*, 2. ed ed., Genetic and evolutionary computation series, Springer, New York, NY, 2007.
- [24] N. Colgan, B. Siow, J. M. O’Callaghan, I. F. Harrison, J. A. Wells, H. E. Holmes, O. Ismail, S. Richardson, D. C. Alexander, E. C. Collins, E. M. Fisher, R. Johnson, A. J. Schwarz, Z. Ahmed, M. J. O’Neill, T. K. Murray, H. Zhang, and M. F. Lythgoe, *Application of neurite orientation dispersion and density imaging (NODDI) to a tau pathology model of Alzheimer’s disease*, NeuroImage **125** (2016), 739–744.
- [25] M. A. Cote, G. Girard, A. Bore, E. Garyfallidis, J. C. Houde, and M. Descoteaux, *Tractometer: Towards validation of tractography pipelines*, Medical Image Analysis **17** (2013), 844–857.
- [26] A. Danilchenko and J.M. Fitzpatrick, *General Approach to First-Order Error Prediction in Rigid Point Registration*, IEEE Transactions on Medical Imaging **30** (2011), no. 3, 679–693.

- [27] Arpita Das and Mahua Bhattacharya, *Affine-based registration of CT and MR modality images of human brain using multiresolution approaches: comparative study on genetic algorithm and particle swarm optimization*, Neural Computing and Applications **20** (2010), no. 2, 223–237.
- [28] Swagatam Das and Ponnuthurai Nagaratnam Suganthan, *Differential Evolution: A Survey of the State-of-the-Art*, IEEE Transactions on Evolutionary Computation **15** (2011), no. 1, 4–31.
- [29] Richard Dawkins, *The selfish gene*, new ed ed., Oxford University Press, Oxford; New York, 1989.
- [30] Juan Du, Songyuan Tang, Tianzi Jiang, and Zhensu Lu, *Intensity-based robust similarity for multimodal image registration*, International Journal of Computer Mathematics **83** (2006), no. 1, 49–57.
- [31] J. Dubois, C. Poupon, F. Lethimonnier, and D. Le Bihan, *Optimized diffusion gradient orientation schemes for corrupted clinical DTI data sets.*, MAGMA **19** (2006), no. 3, 134–143.
- [32] S. Elhabian, Y. Gur, J. Piven, M. Styner, I. Leppert, G. Bruce Pike, and G. Gerig, *Motion is inevitable: The impact of motion correction schemes on HARDI reconstructions*, Proceedings of the MICCAI 2014 Workshop on Computational Diffusion MRI, September 2014.
- [33] Shireen Elhabian, Yaniv Gur, Clement Vachet, Joseph Piven, Martin Styner, Ilana Leppert, G. Bruce Pike, and Guido Gerig, *A preliminary study on the effect of motion correction on HARDI reconstruction*, 2014 IEEE 11th International Symposium on Biomedical Imaging (ISBI), Institute of Electrical & Electronics Engineers (IEEE), April 2014.
- [34] Shireen Elhabian, Yaniv Gur, Clement Vachet, Joseph Piven, Martin Styner, Ilana R. Leppert, G. Bruce Pike, and Guido Gerig, *Subject–motion correction in HARDI acquisitions: choices and consequences*, Brain Imaging Methods **5** (2014), 240.
- [35] Shawna Farquharson, J.-Donald Tournier, Fernando Calamante, Gavin Fabinyi, Michal Schneider-Kolsky, Graeme D. Jackson, and Alan Connelly, *White matter fiber tractography: why we need to move beyond DTI*, Journal of Neurosurgery **118** (2013), no. 6, 1367–1377.

- [36] Uran Ferizi, Torben Schneider, Eleftheria Panagiotaki, Gemma Nedjati-Gilani, Hui Zhang, Claudia A. M. Wheeler-Kingshott, and Daniel C. Alexander, *A ranking of diffusion MRI compartment models with in vivo human brain data*, Magnetic Resonance in Medicine **72** (2014), no. 6, 1785–1792.
- [37] Els Fieremans, Jens H. Jensen, and Joseph A. Helpert, *White matter characterization with diffusional kurtosis imaging*, NeuroImage **58** (2011), no. 1, 177–188.
- [38] J. M. Fitzpatrick and J. B. West, *The distribution of target registration error in rigid-body point-based registration*, IEEE Transactions on Medical Imaging **20** (2001), no. 9, 917–927.
- [39] Alex Fornito, Andrew Zalesky, and Michael Breakspear, *The connectomics of brain disorders*, Nature Reviews Neuroscience **16** (2015), no. 3, 159–172.
- [40] Ezequiel Gleichgerricht, Madison Kocher, and Leonardo Bonilha, *Connectomics and graph theory analyses: Novel insights into network abnormalities in epilepsy*, Epilepsia **56** (2015), no. 11, 1660–1668.
- [41] Caspar J. Goch, Bram Stieltjes, Romy Henze, Jan Hering, Luise Poustka, Hans-Peter Meinzer, and Klaus H. Maier-Hein, *Quantification of changes in language-related brain areas in autism spectrum disorders using large-scale network analysis*, International Journal of Computer Assisted Radiology and Surgery **9** (2014), no. 3, 357–365.
- [42] Qiyong Gong and Yong He, *Depression, Neuroimaging and Connectomics: A Selective Overview*, Biological Psychiatry **77** (2015), no. 3, 223–235.
- [43] Mark S. Graham, Ivana Drobnjak, and Hui Zhang, *Realistic simulation of artefacts in diffusion MRI for validating post-processing correction techniques*, NeuroImage **125** (2016), 1079–1094.
- [44] Patric Hagmann, Leila Cammoun, Xavier Gigandet, Stephan Gerhard, P. Ellen Grant, Van Wedeen, Reto Meuli, Jean-Philippe Thiran, Christopher J. Honey, and Olaf Sporns, *MR connectomics: Principles and challenges*, Journal of Neuroscience Methods **194** (2010), no. 1, 34–45.
- [45] Joseph Hajnal, David Hawkes, and Derek Hill (eds.), *Medical Image Registration*, Biomedical Engineering, vol. 5, CRC Press, June 2001.
- [46] John C. Haselgrove and James R. Moore, *Correction for distortion of echo-planar images used to calculate the apparent diffusion coefficient*, Magnetic Resonance in Medicine **36** (1996), no. 6, 960–964.

- [47] Renjie He and Ponnada A. Narayana, *Global optimization of mutual information: application to three-dimensional retrospective registration of magnetic resonance images.*, Computerized Medical Imaging and Graphics **26** (2002), no. 4, 277–292.
- [48] Jan Hering, Peter F. Neher, Hans-Peter Meinzer, and Klaus H. Maier-Hein, *Erzeugung von Referenzdaten für Kopfbewegungskorrektur in Diffusion-MRI*, Bildverarbeitung für die Medizin 2014 (Thomas Martin Deserno, Heinz Handels, Hans-Peter Meinzer, and Thomas Tolxdorff, eds.), Informatik aktuell, Springer Berlin Heidelberg, January 2014, pp. 342–347.
- [49] Jan Hering, Ivo Wolf, and Maier-Hein, Klaus H., *Multi-objective Memetic Search for Robust Motion and Distortion Correction in Diffusion MRI*, IEEE Transactions on Medical Imaging (2016), [ePub ahead of print].
- [50] Jan Hering, Ivo Wolf, Hans-Peter Meinzer, Bram Stieltjes, and Klaus H Maier-Hein, *A Quantitative Evaluation of Errors Induced by Reduced Field-of-View in Diffusion Tensor Imaging*, Computational Diffusion MRI and Brain Connectivity, Springer, 2014, pp. 35–44.
- [51] Ben Jeurissen, Jacques-Donald Tournier, Thijs Dhollander, Alan Connelly, and Jan Sijbers, *Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data*, NeuroImage **103** (2014), 411–426.
- [52] Y. Jin and H. Branke, *Evolutionary optimization in uncertain environments - A survey*, IEEE Transactions on Evolutionary Computation **9** (2005), no. 3, 303–317.
- [53] Heidi Johansen-Berg, *Human connectomics — What will the future demand?*, NeuroImage **80** (2013), 541–544.
- [54] Hans J. Johnson, M. McCormick, L. Ibanez, and The Insight Software Consortium, *The ITK Software Guide*, third ed., "Kitware, Inc.", 2013.
- [55] George C. Kagadis, Konstantinos K. Delibasis, George K. Matsopoulos, Nikolaos A. Mouravliansky, Pantelis A. Asvestas, and George C. Nikiforidis, *A comparative study of surface- and volume-based techniques for the automatic registration between CT and SPECT brain images.*, Medical Physics **29** (2002), no. 2, 201–213.
- [56] Koji Kamagata, Taku Hatano, Ayami Okuzumi, Yumiko Motoi, Osamu Abe, Keigo Shimoji, Kouhei Kamiya, Michimasa Suzuki, Masaaki Hori, Kanako K. Kumamaru, and et al., *Neurite orientation dispersion and density imaging in the substantia nigra in idiopathic parkinson disease*, European Radiology (2015).

- [57] J. Kennedy and R. Eberhart, *Particle swarm optimization*, Proceedings of ICNN'95 - International Conference on Neural Networks, vol. 4, IEEE, 1995, pp. 1942–1948.
- [58] 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** (2010), no. 1, 196–205.
- [59] Stefan Klein, Josien P. W. Pluim, Marius Staring, and Max A. Viergever, *Adaptive Stochastic Gradient Descent Optimisation for Image Registration*, International Journal of Computer Vision **81** (2009), no. 3, 227–239.
- [60] Stefan Klein, Marius Staring, and Josien P. W. Pluim, *Evaluation of Optimization Methods for Nonrigid Medical Image Registration Using Mutual Information and B-Splines*, IEEE Transactions on Image Processing **16** (2007), no. 12, 2879–2890.
- [61] Cheng Guan Koay, Lin-Ching Chang, John D. Carew, Carlo Pierpaoli, and Peter J. Basser, *A unifying theoretical and algorithmic framework for least squares methods of estimation in diffusion tensor imaging*, Journal of Magnetic Resonance **182** (2006), no. 1, 115–125.
- [62] Cheng Guan Koay, Evren Özarslan, and Carlo Pierpaoli, *Probabilistic Identification and Estimation of Noise (PIESNO): A self-consistent approach and its applications in MRI*, Journal of Magnetic Resonance **199** (2009), no. 1, 94–103.
- [63] Tobias Kober, Rolf Gruetter, and Gunnar Krueger, *Prospective and retrospective motion correction in diffusion magnetic resonance imaging of the human brain*, NeuroImage **59** (2012), no. 1, 389–398.
- [64] Denis Le Bihan and Mami Iima, *Diffusion Magnetic Resonance Imaging: What Water Tells Us about Biological Tissues*, PLoS Biol **13** (2015), no. 7, e1002203.
- [65] Alexander Leemans and Derek K. Jones, *The B-Matrix Must Be Rotated When Correcting for Subject Motion in DTI Data*, Magnetic Resonance in Medicine **61** (2009), no. 6, 1336–1349.
- [66] Wen Li, Nancy C. Andreasen, Peg Nopoulos, and Vincent A. Magnotta, *Automated parcellation of the brain surface generated from magnetic resonance images*, Frontiers in Neuroinformatics **7** (2013).
- [67] Chen-Lun Lin, Aya Mimori, Yen-Wei Chen, Chen-Lun Lin, Aya Mimori, and Yen-Wei Chen, *Hybrid Particle Swarm Optimization and Its Application to Multimodal 3D Medical Image Registration*, Hybrid Particle Swarm Optimization and Its Application to Multimodal 3D Medical Image Registration, Computational Intelligence

- and Neuroscience, Computational Intelligence and Neuroscience **2012**, **2012** (2012), e561406.
- [68] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, *Multimodality image registration by maximization of mutual information*, IEEE Transactions on Medical Imaging **16** (1997), no. 2, 187–198.
 - [69] Klaus H. Maier-Hein, Romuald Brunner, Kira Lutz, Romy Henze, Peter Parzer, Nina Feigl, Jasmin Kramer, Hans-Peter Meinzer, Franz Resch, and Bram Stieltjes, *Disorder-Specific White Matter Alterations in Adolescent Borderline Personality Disorder*, Biological Psychiatry **75** (2014), no. 1, 81–88.
 - [70] Klaus H. Maier-Hein, Carl-Fredrik Westin, Martha E. Shenton, Michael W. Weiner, Ashish Raj, Philipp Thomann, Ron Kikinis, Bram Stieltjes, and Ofer Pasternak, *Widespread white matter degeneration preceding the onset of dementia*, Alzheimers & Dementia **11** (2015), no. 5, 485–493.
 - [71] Jan Modersitzki, *Numerical Methods for Image Registration*, Oxford University Press, Oxford ; New York, 2003.
 - [72] Siawoosh Mohammadi, Harald E. Möller, Harald Kugel, Dirk K. Müller, and Michael Deppe, *Correcting eddy current and motion effects by affine whole-brain registrations: evaluation of three-dimensional distortions and comparison with slice-wise correction*, Magnetic Resonance in Medicine **64** (2010), no. 4, 1047–1056.
 - [73] Pablo Moscato and others, *On evolution, search, optimization, genetic algorithms and martial arts: Towards memetic algorithms*, Caltech concurrent computation program, C3P Report **826** (1989).
 - [74] Amir Moussavi-Biugui, Bram Stieltjes, Klaus Fritzsche, Wolfhard Semmler, and Frederik B. Laun, *Novel Spherical Phantoms for Q-Ball Imaging Under In Vivo Conditions*, Magnetic Resonance in Medicine **65** (2011), no. 1, 190–194.
 - [75] Peter F. Neher, Maxime Descoteaux, Jean-Christophe Houde, Bram Stieltjes, and Klaus H. Maier-Hein, *Strengths and weaknesses of state of the art fiber tractography pipelines – A comprehensive in-vivo and phantom evaluation study using Tractometer*, Medical Image Analysis **26** (2015), no. 1, 287–305.
 - [76] Peter F. Neher, Frederik B. Laun, Bram Stieltjes, and Klaus H. Maier-Hein, *Fiberfox: facilitating the creation of realistic white matter software phantoms.*, Magnetic Resonance in Medicine **72** (2014), no. 5, 1460–1470.

- [77] Peter F Neher and Klaus H Maier-Hein, *A machine learning based approach to fiber tractography using classifier voting*, Lecture Notes in Computer Science, Springer Science + Business Media, 2015, pp. 45–52.
- [78] Peter F Neher, Bram Stieltjes, Marco Reisert, Ignaz Reicht, Hans-Peter Meinzer, and Klaus H Maier-Hein*, *MITK global tractography*, SPIE medical imaging, International Society for Optics and Photonics, 2012, pp. 83144D–83144D.
- [79] Markus Nilsson, Filip Szczepankiewicz, Danielle van Westen, and Oskar Hansson, *Extrapolation-Based References Improve Motion and Eddy-Current Correction of High B-Value DWI Data: Application in Parkinson’s Disease Dementia*, PLoS ONE **10** (2015), no. 11, e0141825.
- [80] Marco Nolden, Sascha Zelzer, Alexander Seitel, Diana Wald, Michael Müller, Alfred M. Franz, Daniel Maleike, Markus Fangerau, Matthias Baumhauer, Lena Maier-Hein, Klaus H. Maier-Hein, Hans Peter Meinzer, and Ivo Wolf, *The medical imaging interaction toolkit: challenges and advances, 10 years of open-source development*, International Journal of Computer Assisted Radiology and Surgery **8** (2013), no. 4, 607–620.
- [81] Francisco P. M. Oliveira and Joao Manuel R. S. Tavares, *Medical image registration: A review*, Computer methods in biomechanics and biomedical engineering **17** (2014), no. 2, 73–93.
- [82] Eleftheria Panagiotaki, Torben Schneider, Bernard Siow, Matt G. Hall, Mark F. Lythgoe, and Daniel C. Alexander, *Compartment models of the diffusion MR signal in brain white matter: A taxonomy and comparison*, NeuroImage **59** (2012), no. 3, 2241–2254.
- [83] Kerstin Pannek, David Raffelt, Christopher Bell, Jane L. Mathias, and Stephen E. Rose, *HOMOR: Higher Order Model Outlier Rejection for high b-value MR diffusion data*, NeuroImage **63** (2012), no. 2, 835–842.
- [84] Konstantinos E. Parsopoulos and Michael N. Vrahatis, *Multi-Objective Particles Swarm Optimization Approaches*, Multi-Objective Optimization in Computational Intelligence (Lam Thu Bui and Semeer Alam, eds.), IGI Global, 2008, pp. 20–42.
- [85] Ofer Pasternak, Nir Sochen, Yaniv Gur, Nathan Intrator, and Yaniv Assaf, *Free water elimination and mapping from diffusion MRI*, Magnetic Resonance in Medicine **62** (2009), no. 3, 717–730.
- [86] Y. G. Petalas, K. E. Parsopoulos, and M. N. Vrahatis, *Memetic particle swarm optimization*, Annals of Operations Research **156** (2007), no. 1, 99–127.

- [87] Bart D. Peters, J. Blaas, and Lieuwe de Haan, *Diffusion tensor imaging in the early phase of schizophrenia What have we learned?*, Journal of Psychiatric Research **44** (2010), no. 15, 993–1004.
- [88] Josien P. W. Pluim, J. B. Antoine Maintz, and Max A. Viergever, *Mutual-information-based registration of medical images: a survey*, IEEE Transactions on Medical Imaging (2003), 986–1004.
- [89] Sonia Pujol, William Wells, Carlo Pierpaoli, Caroline Brun, James Gee, Guang Cheng, Baba Vemuri, Olivier Commowick, Sylvain Prima, Aymeric Stamm, Maged Goubran, Ali Khan, Terry Peters, Peter Neher, Klaus H. Maier-Hein, Yundi Shi, Antonio Tristan-Vega, Gopalkrishna Veni, Ross Whitaker, Martin Styner, Carl-Fredrik Westin, Sylvain Gouttard, Isaiah Norton, Laurent Chauvin, Hatsuho Mamata, Guido Gerig, Arya Nabavi, Alexandra Golby, and Ron Kikinis, *The DTI Challenge: Toward Standardized Evaluation of Diffusion Tensor Imaging Tractography for Neurosurgery: The DTI Challenge on Tractography for Neurosurgery*, Journal of Neuroimaging **25** (2015), no. 6, 875–882.
- [90] Pim Pullens, Alard Roebroek, and Rainer Goebel, *Ground truth hardware phantoms for validation of diffusion-weighted MRI applications*, Journal of Magnetic Resonance Imaging **32** (2010), no. 2, 482–488.
- [91] T.G. Reese, O. Heid, R.M. Weisskoff, and V.J. Wedeen, *Reduction of eddy-current-induced distortion in diffusion MRI using a twice-refocused spin echo*, Magnetic Resonance in Medicine **49** (2003), no. 1, 177–182.
- [92] Margarita Reyes-Sierra and CA Coello Coello, *Multi-objective particle swarm optimizers: A survey of the state-of-the-art*, International journal of computational intelligence research **2** (2006), no. 3, 287–308.
- [93] Jeffrey D. Riley, David L. Franklin, Vicky Choi, Ronald C. Kim, Devin K. Binder, Steven C. Cramer, and Jack J. Lin, *Altered white matter integrity in temporal lobe epilepsy: Association with cognitive and clinical profiles*, Epilepsia **51** (2010), no. 4, 536–545.
- [94] G. K. Rohde, A. S. Barnett, P. J. Basser, S. Marengo, and C. Pierpaoli, *Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI*, Magnetic Resonance in Medicine **51** (2004), no. 1, 103–114.
- [95] Nicholas Said, W. Jeff Elias, Prashant Raghavan, Alan Cupino, Nicholas Tustison, Robert Frysinger, James Patrie, Wenjun Xin, and Max Wintermark, *Correlation of diffusion tensor tractography and intraoperative macrostimulation during deep*

- brain stimulation for Parkinson disease: Clinical article*, Journal of Neurosurgery **121** (2014), no. 4, 929–935.
- [96] J. Santamaría, O. Cordon, and S. Damas, *A comparative study of state-of-the-art evolutionary image registration methods for 3D modeling*, Computer Vision and Image Understanding **115** (2011), no. 9, 1340–1354.
- [97] J. Santamaría, O. Cordon, S. Damas, I. Aleman, and M. Botella, *A scatter search-based technique for pair-wise 3D range image registration in forensic anthropology*, Soft Computing **11** (2007), no. 9, 819–828.
- [98] J. Santamaría, O. Cordon, S. Damas, J. M. García-Torres, and A. Quirin, *Performance evaluation of memetic approaches in 3D reconstruction of forensic objects*, Soft Computing **13** (2008), no. 8-9, 883–904.
- [99] J. David Schaffer, *Multiple Objective Optimization with Vector Evaluated Genetic Algorithms*, Proceedings of the 1st International Conference on Genetic Algorithms (Hillsdale, NJ, USA), L. Erlbaum Associates Inc., 1985, pp. 93–100.
- [100] J.A. Schnabel, C. Tanner, A.D. Castellano-Smith, A. Degenhard, M.O. Leach, D.R. Hose, D.L.G. Hill, and D.J. Hawkes, *Validation of nonrigid image registration using finite-element methods: application to breast MR images*, IEEE Transactions on Medical Imaging **22** (2003), no. 2, 238–247.
- [101] Farshid Sepehrband, Kristi A. Clark, Jeremy F. P. Ullmann, Nyoman D. Kurniawan, Gayeshika Leanage, David C. Reutens, and Zhengyi Yang, *Brain tissue compartment density estimated using diffusion-weighted MRI yields tissue parameters consistent with histology*, Human Brain Mapping **36** (2015), no. 9, 3687–3702.
- [102] Kiran K. Seunarine and Daniel C. Alexander, *Multiple Fibres: Beyond the diffusion tensor*, Diffusion MRI: From quantitative Measurements to In-vivo Neuroanatomy (Heidi Johansen-Berg and Timothy E.J. Behrens, eds.), Academic Press, 2012, pp. 105–123.
- [103] A. Sotiras, C. Davatzikos, and N. Paragios, *Deformable Medical Image Registration: A Survey*, IEEE Transactions on Medical Imaging **32** (2013), no. 7, 1153–1190.
- [104] E. O. Stejskal and J. E. Tanner, *Spin Diffusion Measurements: Spin Echoes in the Presence of a Time-Dependent Field Gradient*, The Journal of Chemical Physics **42** (1965), no. 1, 288.
- [105] Aline M. Studerus-Germann, Jean-Philippe Thiran, Alessandro Daducci, and Oliver P. Gautschi, *Diagnostic approaches to predict persistent post-traumatic*

- symptoms after mild traumatic brain injury - a literature review*, International Journal of Neuroscience **126** (2016), no. 4, 289–298.
- [106] C. Studholme, D. L. G. Hill, and D. J. Hawkes, *An overlap invariant entropy measure of 3D medical image alignment*, Pattern Recognition **32** (1999), no. 1, 71–86.
- [107] Yu Sun, Yu Chen, Renick Lee, Anastasios Bezerianos, Simon L. Collinson, and Kang Sim, *Disruption of brain anatomical networks in schizophrenia: A longitudinal, diffusion tensor imaging based study*, Schizophrenia Research **171** (2016), no. 1-3, 149–157.
- [108] J.-Donald Tournier, Fernando Calamante, and Alan Connelly, *Determination of the appropriate b value and number of gradient directions for high-angular-resolution diffusion-weighted imaging*, NMR in biomedicine **26** (2013), no. 12, 1775–1786.
- [109] J.-Donald Tournier, Chun-Hung Yeh, Fernando Calamante, Kuan-Hung Cho, Alan Connelly, and Ching-Po Lin, *Resolving crossing fibres using constrained spherical deconvolution: Validation using diffusion-weighted imaging phantom data*, NeuroImage **42** (2008), no. 2, 617–625.
- [110] Andrea Valsecchi, Sergio Damas, and Jose Santamaria, *Evolutionary Intensity-based Medical Image Registration: A Review*, Current Medical Imaging Reviews **9** (2013), no. 4, 283–297.
- [111] Petra E. Vértes and Edward T. Bullmore, *Annual Research Review: Growth connectomics - the organization and reorganization of brain networks during normal and abnormal development*, Journal of Child Psychology and Psychiatry **56** (2015), no. 3, 299–320.
- [112] M.P. Wachowiak, R. Smolikova, Yufeng Zheng, J.M. Zurada, and A.S. Elmaghraby, *An approach to multimodal biomedical image registration utilizing particle swarm optimization*, IEEE Transactions on Evolutionary Computation **8** (2004), no. 3, 289–301.
- [113] Jinhui Wang, Xindi Wang, Mingrui Xia, Xuhong Liao, Alan Evans, and Yong He, *GRETN: a graph theoretical network analysis toolbox for imaging connectomics*, Frontiers in Human Neuroscience **9** (2015).
- [114] Dominik Weishaupt, Victor D. Köchli, and Boruk Marincek, *How does MRI work? An Introduction to the Physics and Function of Magnetic Resonance Imaging*, Springer, 2008.

- [115] Gavin P. Winston, Pankaj Daga, Mark J. White, Caroline Micallef, Anna Misorocchi, Laura Mancini, Marc Modat, Jason Stretton, Meneka K. Sidhu, Mark R. Symms, David J. Lythgoe, John Thornton, Tarek A. Yousry, Sebastien Ourselin, John S. Duncan, and Andrew W. McEvoy, *Preventing visual field deficits from neurosurgery*, *Neurology* **83** (2014), no. 7, 604–611.
- [116] Xiaoyan Xu and R.D. Dony, *Differential evolution with Powell's direction set method in medical image registration*, *IEEE International Symposium on Biomedical Imaging: Nano to Macro*, 2004, April 2004, pp. 732–735 Vol. 1.
- [117] Anastasia Yendiki, Kami Koldewyn, Sita Kakunoori, Nancy Kanwisher, and Bruce Fischl, *Spurious group differences due to head motion in a diffusion MRI study*, *NeuroImage* **88** (2014), 79–90.
- [118] Hui Zhang, Penny L. Hubbard, Geoff J.M. Parker, and Daniel C. Alexander, *Axon diameter mapping in the presence of orientation dispersion with diffusion MRI*, *NeuroImage* **56** (2011), no. 3, 1301–1315.
- [119] Hui Zhang, Torben Schneider, Claudia A. Wheeler-Kingshott, and Daniel C. Alexander, *NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain*, *NeuroImage* **61** (2012), no. 4, 1000–1016.
- [120] Yudong Zhang, Shuihua Wang, and Genlin Ji, *A Comprehensive Survey on Particle Swarm Optimization Algorithm and Its Applications*, *Mathematical Problems in Engineering* (2015), 931256.
- [121] Aimin Zhou, Bo-Yang Qu, Hui Li, Shi-Zheng Zhao, Ponnuthurai Nagarathnam Suganthan, and Qingfu Zhang, *Multiobjective evolutionary algorithms: A survey of the state of the art*, *Swarm and Evolutionary Computation* **1** (2011), no. 1, 32–49.
- [122] Jiancheng Zhuang, Zhong-Lin Lu, Christine Bouteiller Vidal, and Hanna Damasio, *Correction of eddy current distortions in high angular resolution diffusion imaging*, *Journal of Magnetic Resonance Imaging* **37** (2013), no. 6, 1460–1467.
- [123] Barbara Zitová and Jan Flusser, *Image registration methods: a survey*, *Image and Vision Computing* **21** (2003), no. 11, 977–1000.