

DISS. ETH NO. 30087

ADVANCED SIMULATION, ENCODING AND
INFERENCE FOR CARDIAC DIFFUSION TENSOR
IMAGING

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES
(Dr. sc. ETH Zurich)

presented by

JONATHAN LEANDER WEINE
M.Sc. in Medizinphysik,
Technische Universität Dortmund & Ruhr-Universität Bochum

24. February 1992

accepted on the recommendation of

Prof. Dr. Sebastian Kozerke
Prof. Dr. Maxim Zaitsev

2024

Abstract

Cardiovascular diseases are a significant global public health challenge, necessitating early detection of cardiac pathophysiological changes. Cardiovascular Magnetic Resonance (CMR) has emerged as a versatile diagnostic tool, providing clinical assessments of anatomy, function, and perfusion. Standard CMR protocols, including cine imaging and contrast-enhanced techniques, offer a comprehensive evaluation of global cardiac parameters and anatomy. In addition, advanced CMR methods, such as relaxometry, strain imaging, and diffusion tensor imaging (DTI), provide information for detailed tissue characterization and local function assessment. Despite the recent increase of interest in cardiac DTI (cDTI), it primarily remains a research tool, hindered by a lack of consensus on optimal acquisition parameters for clinical translation.

The challenges in cDTI acquisition and processing include high demands on hardware, sensitivity to artifacts and the necessity to address respiratory and contractile motion. A comprehensive comparison of proposed techniques based on simulations is envisioned to support clinical translation. To this end, an open-source MR simulation framework incorporating complex motion and flow is presented. In addition, an MR sequence definition software package was implemented, contributing to reproducibility and supporting comparison in the research community.

Deep learning (DL) algorithms present a promising avenue to accelerate and improve cDTI processing. However, due to the inherent complexity of cDTI, high-quality training data is scarce. Furthermore, careful consideration is required to address potential biases in the training datasets to facilitate generalization for varying acquisition parameters and to minimize risk of systematic errors. Here, a novel approach exclusively using in-silico data to train a DL model for diffusion tensor inference is proposed, demonstrating improved performance on in-vivo data while allowing for control over the training data distribution. Thereby, a method for creating benchmark datasets to test and train reliable algorithms is proposed.

The macroscopic assessment of contractile cardiac motion can be achieved by strain imaging. As cardiac contraction is directly determined by the state of tissue microstructure, combining strain imaging and cDTI is a promising approach. To avoid the scan time penalty of acquiring separate scans, a joint encoding strategy for cDTI and tissue velocity mapping (TVM) is proposed, showing good agreement of tissue velocities with established reference measurements and negligible effect on the estimated diffusion tensors.

In conclusion, this dissertation contributes methodological advances in simulation, encoding and inference in cDTI, supporting a potential clinical translation.

Kurzfassung

Kardiovaskuläre Erkrankungen stellen eine erhebliche Belastung für das Gesundheitswesen dar. Deshalb sind diagnostische Verfahren für das frühzeitige Erkennen von pathophysiologischen Veränderungen des Herzens erforderlich. Kardiovaskuläre Magnetresonanzbildgebung (CMR) ist ein etabliertes Verfahren, das die klinische Beurteilung von Anatomie, Funktion und Perfusion des Herzens ermöglicht. Standard CMR-Protokolle beinhalten Cine-Bildgebung und kontrastmittelbasierte Methoden, welche die Bewertung globaler kardialer Parameter zulassen. Darüber hinaus erlauben komplexere CMR-Methoden wie Relaxometrie, Strain-Bildgebung und Diffusionstensorbildgebung (DTI) eine detaillierte Charakterisierung des Gewebes. Trotz des wachsenden Interesses an DTI im Herzen (cDTI) wird es derzeit hauptsächlich in der Forschung genutzt. Das Fehlen eines Konsens über die optimale Methodik und ihrer Parameter erschwert dabei die klinische Anwendung.

Herausforderungen bei der Aufnahme und Verarbeitung von cDTI-Daten sind die hohen Hardwareanforderungen, die Anfälligkeit für Bildartefakte und das Berücksichtigen von Atem- und Kontraktionsbewegungen. Ein detaillierter Vergleich der bestehenden Methoden, basierend auf Simulationen, könnte die Entwicklung der klinischen Anwendung von cDTI unterstützen. Hierzu wird ein Open-Source MR-Simulationsframework vorgestellt, das komplexe Bewegungen und Blutfluss berücksichtigt. Zusätzlich wurde ein Softwarepaket zur Sequenzdefinition implementiert, um den Vergleich verschiedener Pulssequenzen zu verbessern.

Deep Learning (DL)-Algorithmen stellen einen vielversprechenden Ansatz zur Beschleunigung und Verbesserung der Verarbeitung von cDTI-Daten dar. Aufgrund der inhärenten Komplexität von cDTI sind hochwertige Trainingsdaten jedoch rar. Darüber hinaus sollten öffentliche Datensätze mit Vorsicht genutzt werden, sodass statistische Verzerrungen erkannt, die Generalisierbarkeit ermöglicht und das Risiko systematischer Fehler minimiert wird. In dieser Arbeit wird eine Methode vorgestellt DL-Modelle ausschließlich mit In-Silico-Daten zur Inferenz von Diffusionstensoren zu trainieren. Diese Methode liefert verbesserte Resultate auf In-Vivo-Daten und ermöglicht gleichzeitig die Kontrolle über die Trainingsdaten. Damit können Benchmark-Datensätze erstellt werden, um Algorithmen zuverlässig zu trainieren und zu testen.

Kardiale Kontraktion kann durch Strain-Bildgebung beurteilt werden. Da die Kontraktion durch den Zustand der Gewebemikrostruktur bestimmt wird, ist die Kombination von Strain-Bildgebung und cDTI ein vielversprechender Ansatz. Um eine höhere Scanzeit durch separate Aufnahmen zu vermeiden, wird in dieser Arbeit eine Strategie zur simultanen Kodierung von cDTI und Gewebe-Geschwindigkeitsquantifizierung (TVM) präsentiert. Die Ergebnisse der Gewebe-Geschwindigkeiten

stimmen mit etablierten Referenzmessungen überein und der Einfluss auf die Diffusionstensoren ist vernachlässigbar.

Zusammenfassend trägt diese Dissertation zu methodischen Fortschritten in der Simulation, Kodierung und Inferenz in cDTI bei und unterstützt damit potenziell die Entwicklung klinischer Anwendungen.

Contents

Abstract	i
Kurzfassung	iii
Contents	v
Introduction	1
Outline	5
Contributions of this Thesis	7
1 Physiology of the Heart	9
1.1 Anatomy	9
1.2 Microstructure	10
1.3 Contraction and Strain	12
2 Motion Encoding in Cardiac MRI	15
2.1 Physiological Motion - Compensation Strategies	15
2.2 Encoding Displacements	17
2.3 Encoding Tissue Velocity	20
2.4 Encoding Diffusional Motion	23
2.5 Related Cardiac Pathologies	29
3 Spin Echo cDTI	31
3.1 Diffusion Encoding Waveforms	31
3.2 Echo Planar Imaging	32
3.3 Fat Suppression	33
3.4 Reduced Field of View	34
3.5 Off-Resonances	35
3.6 Eddy Currents and Concomitant Fields	36
4 Simulation and In-silico Ground Truth	39
4.1 Lagrangian vs Eulerian Description	39
4.2 MR Signal Simulation	40
4.3 Digital Cardiac Phantoms	43
4.4 Sequence Definition	44
5 CMRsim - A Python Package for Cardiovascular MR Simulations Incorporating Complex Motion and Flow	47

5.1	Introduction	47
5.2	Methods	48
5.3	Results	58
5.4	Discussion and Conclusion	63
5.5	Appendix A	65
5.6	Appendix B - Comparison to External Implementations	65
5.7	Appendix C - Performance and Scaling	66
6	Parameter Estimation / Inverse Problems	67
6.1	Linear Models	67
6.2	Bayesian Modelling	68
6.3	Learning Based Methods	69
7	Synthetically Trained Convolutional Neural Networks for Improved Tensor Estimation from Free-Breathing Cardiac DTI	71
7.1	Introduction	71
7.2	Methods	72
7.3	Results	78
7.4	Discussion	83
7.5	Appendix	86
8	Joint Encoding of Diffusion Tensors and Tissue Velocities in the Heart	89
8.1	Introduction	89
8.2	Methods	90
8.3	Results	96
8.4	Discussion	101
8.5	Conclusion	106
	Discussion and Outlook	107
	References	109
	List of Publications	131
	Acknowledgements	133

Introduction

As cardiovascular diseases are among the leading global public health issues, early diagnosis of pathophysiological changes in the heart is urgently required¹. Cardiac Magnetic Resonance (CMR) has emerged as a versatile, clinically feasible tool to assess parameters such as anatomy, function and perfusion and thereby providing diagnostic value^{1,2}. Standard CMR protocols include balanced Steady-State Free Precession (bSSFP) cine imaging in a variety of cardiac views to measure global parameters such as ventricular volumes and ejection fraction. Furthermore, contrast agent enhanced imaging such first pass or stress perfusion and late gadolinium enhancement as well as blood-flow quantitation with phase-contrast (PC) MRI^{3,4} are routinely performed.

Additionally, tissue characterization techniques and local function assessment of the heart is possible using advanced CMR methods^{5,6}, such as relaxometry⁷, strain imaging⁸⁻¹⁰ as well as Diffusion Tensor Imaging (DTI)^{11,12}. While some of the advanced methods recently started gaining popularity in clinical settings³, cDTI, with its capability to non-invasively assess microstructural information of the cardiac muscle^{11,12}, is still mostly used in research⁵. The association of diffusion tensors with the cardiac tissue microstructure has been validated in multiple studies; ex-vivo including histology¹³⁻¹⁵ as well as in-vivo¹⁶. Moreover, values for diffusion metrics have been reported in several pathological cardiac conditions¹⁷⁻²⁵. However, no consensus regarding the optimal parameters such as resolution, type of selective excitation, readout strategies or diffusion weightings has been reached yet, which is a crucial step for clinical translation²⁶.

The acquisition and processing of cDTI data is challenging due to high demands on hardware^{11,12}, the necessity for a theoretically sound tissue model to interpret the signal²⁷, as well as the sensitivity to a variety of artifacts²⁸. In addition, respiratory motion as well as the inevitable, complex contractile motion of the heart introduces technical challenges, hampering the clinical translation of cDTI⁵. As free-breathing acquisitions are generally preferred in clinical CMR¹, Spin Echo (SE) with single shot echo-planar imaging (EPI), allowing the acquisition within a single cardiac cycle, is advantageous. However, this requires the diffusion encoding gradient waveforms to be motion compensated^{11,12,29,30}, which increases their duration. Moreover, the gradient-induced eddy currents alter the the EPI readout, causing image distortions³¹. Furthermore, off-resonance fields e.g., due to blood flow in cardiac vasculature can create image distortions³².

The macroscopic assessment of contractile cardiac motion can be achieved with strain imaging⁸⁻¹⁰. Although ultrasound-based strain imaging is wide-spread due to its low cost, CMR strain imaging with its 3D capability is the gold standard due to its high accuracy and reproducibility³³. On the one hand, certain global

strain metrics are routinely used in diagnostics³³, on the other hand regional strain estimation can provide valuable information in a variety of conditions^{8–10}. One method, called feature tracking, uses post-processing algorithms on bSSFP cine images to obtain strain values. All other strain imaging techniques require separate scans, while also providing superior spatial resolution^{8–10}.

Cardiac function is directly determined by the state of the cardiac microstructure³⁴. Therefore, combining microstructural and strain information promises to yield a more sensitive and accurate biomarker for early pathological changes. To this end, recent work has used strain and cDTI data of separate scans to project the macroscopic deformation onto the dominant microstructural orientations^{35–38}. The strain along the myocyte orientation is a promising biomarker for local function alterations³⁴. However, the long protocol duration due to many different measurements as well as low scan-efficiency due to prospective respiratory motion gating are key barriers to clinical usability, hence acquiring the required separate scans is not ideal³.

In all areas of MRI, there is an increasing use of machine learning algorithms (often referred to as Artificial Intelligence (AI))³⁹, ranging from reconstruction⁴⁰, image processing steps like registration⁴¹, segmentation^{42–45}, denoising⁴⁶ or parameter estimation⁴⁷, to even optimization of the measurement sequence along with the reconstruction^{48,49}. The application of AI algorithms in cDTI is a promising approach to reduce the required time for scanning and processing as well as to increase the robustness of the inferred metrics^{50–54}. However, biases in the training dataset affect the generalization capabilities of AI algorithms, requiring special care to be taken when using public data for training⁵⁵. This is especially relevant in medical image processing, where the behaviour for out-of-sample data and worst-case scenarios is critical⁵⁶.

Due to the complex nature of cDTI, the diffusion signal is inherently coupled to the acquisition technique⁵⁷. With respect to training AI algorithms for processing and inference in cDTI, this poses the problem of obtaining sufficient training data with matching imaging parameters. Since ground-truth is inherently not available and high-resolution and -quality scans are practically infeasible to acquire with cDTI, it is fundamentally difficult to create a high-quality dataset for training and testing the critical cases⁵⁸.

Simulations are an important and established tool in MR research to evaluate the performance of reconstruction algorithms, estimate effects of sequences on the signal or investigate the validity of post-processing algorithm in quantitative mapping⁵⁹. Publicly available, open-source software to implement such simulations supports the reproducibility of the results. However, currently available frameworks have not specifically addressed the inclusion of complex motion and flow as present in CMR^{60–69}. In addition to the simulation software, there is an ongoing effort in creating digital phantoms capturing healthy and diseased population statistics of

all aspects of the heart⁷⁰, potentially serving as input to the simulation. Generating images from personalized digital twin models can improve early diagnosis⁷¹, and facilitate a systematic comparison of methodology in acquisition and post-processing, including the training and testing of AI algorithms.

Outline

Chapter 1 provides a general review of the physiology of the heart, its microstructure and its connection to macroscopic function.

The role of motion in CMR, as a technical challenge as well as a biomarker is summarized in chapter 2. Firstly, the commonly applied strategies to deal with respiratory and contractile motion in CMR are presented. Secondly, the principles of encoding bulk tissue motion, i.e. the macroscopic contraction, are highlighted. Thirdly, an overview of concepts for inferring micro-structural information with cDTI is provided. Finally, cardiac pathologies studied with the presented methods are briefly reviewed. The basics of SE-cDTI with single shot EPI are discussed in chapter 3.

In chapter 4, the foundations of simulating cDTI data are summarized, which include the representation of digital phantoms as well as the generation of contrasts and signal formation based on solving the Bloch equations. Furthermore, available digital phantoms and sequence definition frameworks are reviewed. The publication in chapter 5 provides a detailed description of the open-source simulation framework CMRsim, specifically targeting cardiovascular MR simulations.

A general introduction to inverse problems, and the application of learning-based methods to solve them, is provided in chapter 6. This is followed by the published paper in chapter 7, in which a pipeline to train deep neural network on synthetic cDTI data is proposed.

Finally, the paper in chapter 8 presents a method to jointly encode cardiac tissue velocity and diffusion tensor maps without additional scan time.

Contributions of this Thesis

The work presented in this thesis explores the possibility of improving the value and usability of cDTI in three aspects.

Firstly, an open-source framework for MR simulations, including complex motion and flow, is presented. The aim of this project was to increase reproducibility, robustness and to facilitate comparison of CMR methods. Extensive documentation, including examples, was created and is publicly available. Along with the simulation framework, an open-source package for sequence definition was implemented. It follows the concepts of the increasingly popular PulSeq framework, while the scope was extended to build a library of sequences, supporting vendor-independent methodological comparisons.

Secondly, the feasibility of training a neural network on synthetic cDTI data is demonstrated. The parameterized training data generation incorporated a sampling process to obtain random eigenvalue triplets with uniform marginal distributions on MD and FA. Thereby, it reflects the expected statistics for healthy and infarcted tissue. Furthermore, the spatial structure of the random samples was imposed by tensor-valued interpolation in radial and circumferential coordinates in left-ventricular (LV) segmentations obtained from in-vivo data. The network inference on in-vivo data showed improved diffusion metrics while avoiding over-smoothing and allowing for a reduction of required averages. Thereby, a pipeline for training algorithms for new acquisition methods, as well as rigorous testing of AI algorithms is presented.

Finally, strain-rate and diffusion tensor encoding were combined into a single scan without scan time penalty, potentially increasing the value cDTI. This was achieved by encoding tissue velocity into the image phase, while using the signal magnitude for diffusion estimation. Simulations showed negligible influence of the addition strain-rate encoding gradients on diffusion metrics and in-vivo measurements showed good agreement with reference methods.

1. Physiology of the Heart

1.1 Anatomy

The heart is the central organ of the cardiovascular system and consists of four chambers, left and right atria and ventricles as illustrated in Figure 1.1a. It is responsible for pumping blood throughout the body by periodically contracting and relaxing. The deoxygenated blood from the entire body arrives in the right atrium via the vena cavae, from which it flows through the tricuspid valve into the right ventricle during diastole. In systole, the right ventricle ejects the blood into the pulmonary arteries to reach the lungs, where it is oxygenated. Oxygenated blood flows back to the left atrium via the pulmonary veins. During diastole, blood from the left atrium flows into the left ventricle through the mitral valve. Finally, during systole, the LV ejects the oxygenated blood into the aorta from which it is distributed through the entire body. The mitral and tricuspid valves are supported by the papillary muscles (Figure 1.1).

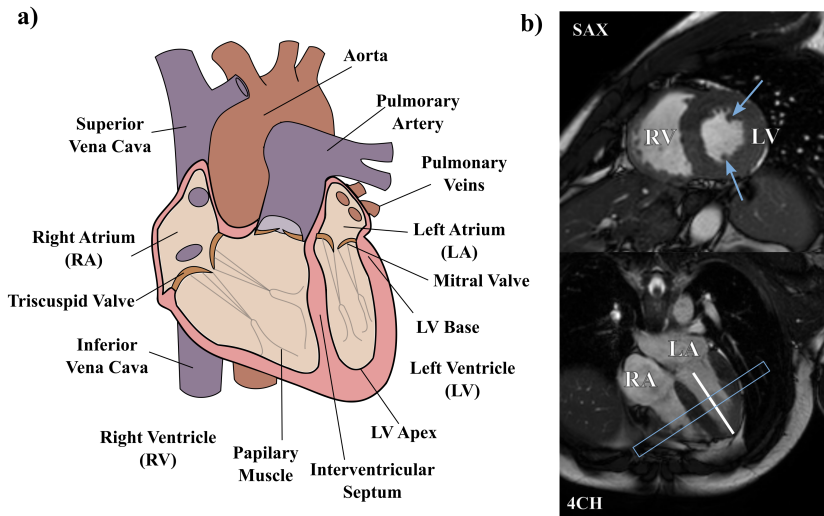


FIGURE 1.1: a) Illustration of the cardiac anatomy from an anterior view, based on [72]), b) bSSFP cine image in short-axis view (SAX) and four chamber view (4CH). In the SAX view (top), ventricles are annotated with RV (right ventricle) and LV (left ventricle). Papillary muscles are indicated by the blue arrows. In the 4CH view (bottom), the white line indicates the LV long axis, and the blue box indicates the positioning of a mid-ventricular SAX slice.

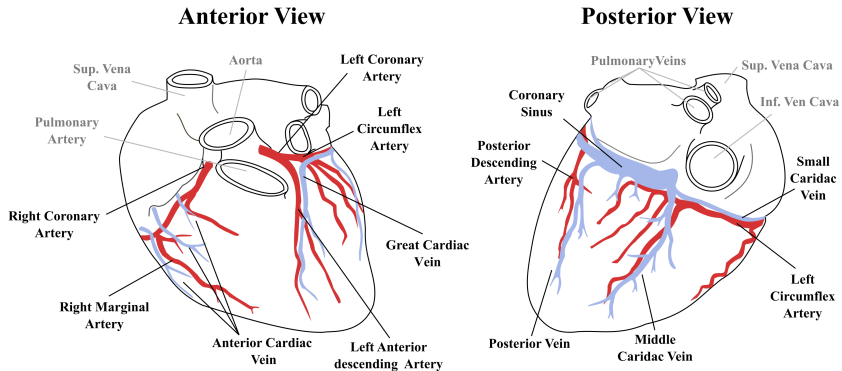


FIGURE 1.2: Schematic illustration of cardiac vasculature (based on [74]) from an anterior view (left) and a posterior view (right). Coronary arteries are drawn in red and coronary veins are drawn in blue. The anterior view corresponds to the orientation in Figure 1.1a.

To reproducibly capture the anatomy of the heart in 2D images, a set of standardized views are defined³. Two of them, namely short-axis (SAX) and the four-chamber (4CH) view are presented in Figure 1.1b. The short axis is perpendicular to the long axis, which is defined by the line between the apex and the mitral valve⁷³.

The heart itself is perfused with nutrients and oxygen via the coronary arteries. The corresponding draining of deoxygenated blood takes place via the cardiac veins. Figure 1.2 illustrates the vasculature, where the two cardiac arteries emerge from the aorta and the cardiac veins connect to the right atrium. Most of the myocardium is drained by the three large veins (great, middle and small cardiac vein) which run parallel to the coronary arteries⁷⁴. As the contraction of the ventricles results in significant tissue pressure, the flow in the coronary arteries and veins follow a complex pattern. The majority of the antegrade flow in the coronary arteries happens during diastole, while the peak flow in the cardiac veins is reached in late systole⁷⁵. The flow of oxygen-rich and deoxygenated blood induces a dynamic pattern in tissue susceptibility, and hence a varying off-resonance field for the tissue in proximity of the coronary vasculature, directly impacting CMR applications⁷⁶.

1.2 Microstructure

The myocardial microstructure is complex but highly organized. It consists of myocytes and fibroblasts embedded in an extra-cellular matrix predominantly composed of collagen⁷⁷. Cardiomyocytes are approximately 100 μm to 150 μm long and 10 μm to 30 μm wide^{78,79} and are the fundamental contractile elements. They




FIGURE 1.3: a) Illustration of the transmural variation of myocyte orientation characterized by the Helix Angle. Helical orientation at endo- and epicardium are depicted as red and purple lines on the corresponding surface. Orientations between epi and endo-cardium are only depicted in a small area to improve visibility. b) Schematic illustration of sheetlets sliced perpendicular to the myocyte orientation. The orange and green ellipses highlight the local sheetlet orientation. Based on [12].

are arranged in continuously branching, axially interconnecting sequences. The local predominant structure, defined as axial myocyte-aggregate orientation, is referred to as fiber⁸⁰. The fiberorientation forms a left-handed helix at the epicardium and a right-handed helix at the endocardium with a gradual transition through the circumferential configuration in the mesocardium. The organizational structure can be characterized by the so-called Helix Angle (HA), which transmurally progresses from -60° to 60° relative to the short-axis plane⁸¹. The transmural variation of the helix angle is illustrated in Figure 1.3a.

As a secondary organizational structure, groups of 5 to 10 myocytes are organized in a lamina structure, separated by collagen layers filled with fluid. These secondary structures are defined as sheetlets, extending in the direction of fiber orientations and parallel to the shear layers⁸¹. The sheet orientation can be characterized by the angle between the sheet plane and the cross-myocyte plane. While the predominant

pattern of sheet orientation is the subject of ongoing research, areas of sheets with similar orientation can be observed. The sheet orientation varies smoothly within the same area and shows abrupt changes in the boundary zones⁸⁰. Figure 1.3b illustrates the sheetlet structure sliced in a cross-myocyte direction (perpendicular to the local myocyte orientation).

1.3 Contraction and Strain

The deformation pattern occurring during contraction of the myocardium involves longitudinal shortening, radial thickening and circumferential shortening. Additionally, a twist of the LV is induced by the counter-rotation of apex and base. This includes a brief initial clockwise rotation of the apex at the onset of contraction, followed by a global counter-clockwise LV rotation near the apex and a clockwise rotation near the LV base during ejection⁸².

The changes in the macroscopic configuration of the heart are caused by the shortening and subsequent relaxation of myocytes combined with reorientation at tissue scale. During contraction the myocytes shorten by $\sim 15\%$ and thicken by $\sim 8\%$ ⁸³. While longitudinal ventricular shortening is directly connected to the myocyte, further mechanics are required to explain the observed wall thickening^{12,34}. During contraction sheets re-orient, such that they are more perpendicular to the tangential epicardial plane during systole^{84,85}. This reorientation in combination with concomitant relative shearing displacement between layers is hypothesized to be the main contribution to LV wall thickening, with minor contribution from cardiomyocyte thickening¹².

On a macroscopic scale LV contraction can be characterized by local tissue strains, defined as the relative material deformation in response to the applied load. The formal definition of strain depends on the frame of reference (see 4.1). In the Lagrangian formulation, following material points through space, strain is defined as the change in length dL of the myocardial tissue compared to a reference state L_0 . Considering all directions in 3D, the strain definition becomes a symmetric tensor¹⁰:

$$\boldsymbol{\epsilon} = \begin{pmatrix} \frac{du_x}{dx} & \frac{1}{2} \left(\frac{du_y}{dx} + \frac{du_x}{dy} \right) & \frac{1}{2} \left(\frac{du_z}{dx} + \frac{du_x}{dz} \right) \\ \frac{1}{2} \left(\frac{du_y}{dx} + \frac{du_x}{dy} \right) & \frac{du_y}{dy} & \frac{1}{2} \left(\frac{du_z}{dy} + \frac{du_y}{dz} \right) \\ \frac{1}{2} \left(\frac{du_z}{dx} + \frac{du_x}{dz} \right) & \frac{1}{2} \left(\frac{du_z}{dy} + \frac{du_y}{dz} \right) & \frac{du_z}{dz} \end{pmatrix}, \quad (1.1)$$

where $\mathbf{u} = (u_x, u_y, u_z)^T$ are the displacements of material point along the three orthogonal directions x, y, z . The analysis of myocardial strains is commonly performed using the local coordinate system of radial, circumferential and longitudinal directions as shown in Figure 2.7a).

The temporal change of the strain defined in equation 1.1, is called the strain-rate tensor⁸⁶:

$$\dot{\epsilon} = \frac{1}{2} \left((\nabla \mathbf{v}) + (\nabla \mathbf{v})^T \right), \quad (1.2)$$

where $\mathbf{v} = \dot{\mathbf{u}}$ is the velocity of the material point and $\nabla \mathbf{v}$ denotes the velocity Jacobian. As strain-rate theoretically encodes the same information of wall motion, it can also serve as an equivalent biomarker. The choice between the two definitions often depends on the imaging technique³³, assessing either displacements or velocities (see 2.2 and 2.3). Clinically, the most relevant parameter is the global longitudinal strain (corresponding to the total longitudinal shortening with respect to the end-diastolic shape) with physiological healthy values between -18.9% to -20.4% ³³. Normal peak systolic strain rates in longitudinal direction are reported from 1.0 s^{-1} to 1.4 s^{-1} . Typically, the radial strain component, mainly related to wall thickening, is the largest and is larger in the free wall than in the septum¹⁰. Healthy global radial strains were reported within 35% to 59% ³³.

While used as clinical indicators, the macroscopic projections in radial, circumferential and longitudinal direction do not coincide with the fiber orientation of the myocardium. When strains are projected onto the fiber orientation, the resulting *myofiber strain* is more spatially homogeneous (-0.14 ± 0.2)³⁵ and strain-line patterns emerge^{34,87–89}.

2. Motion Encoding in Cardiac MRI

2.1 Physiological Motion - Compensation Strategies

The most prominent effects of motion in CMR are potential spatial mismatch between image acquisitions, as well as ghosting and blurring caused by motion during sampling. Moreover, motion occurring while gradients are played out creates additional signal phase. On one hand, this phase can be used to encode motion-related information as described in 2.3 and 2.4, on the other hand, it can also lead to unintended signal modulation or even signal loss. Aside from involuntary motion, two types of physiological motion need to be addressed: respiratory motion and contractile cardiac motion.

Motion compensation strategies can be categorized into prospective and retrospective. Prospective strategies use physiological real-time information to synchronize and modify parameters of the acquisition⁹⁰. Retrospective methods correct motion during reconstruction or post processing.

2.1.1 Contractile Motion

The Electrocardiogram (ECG) is the physiological signal most commonly used for scan synchronization in CMR. The R-peak of the ECG signal is used as a reference point of the cardiac cycle. The duration between the R-peak and the acquisition is defined as trigger delay t_{trig} . The concept of an ECG-triggered acquisition is illustrated in Figure 2.1. For steady-state imaging, k-space segments are acquired continuously and are retrospectively sorted into bins according to their temporal distance to the last R-peak, allowing the reconstruction of one image per cardiac phase⁹⁰. The ECG signal can be significantly degraded by strong MR gradients and RF pulses, which is especially relevant in diffusion-weighted imaging, hampering the accuracy for ECG triggering. Peripheral pulse units (PPU) detect the pulse wave of systolic ejection in peripheral vasculature. The intrinsic long delay, however, makes the PPU unsuitable for prospective triggering.

2.1.2 Respiratory Motion

The conceptually easiest strategy to mitigate motion effects is breath-held acquisition. This requires splitting the entire measurement into parts, which are sufficiently short. Furthermore, the breath-holds must be consistent with respect to the respiratory state, such as full expiration. As the strategy relies on patient compliance, which can be problematic in clinical settings, free-breathing scans are preferable. To this end, navigator echoes are commonly interleaved with image acquisition to assess




FIGURE 2.1: a) Transversal plane of a 3D survey scan with a 1D navigator placement on right hemidiaphragm, indicated by the white rectangle b) Sagittal plane corresponding to the navigator placement shown in a). c) A schematic breathing curve (blue) with a navigator acceptance window indicated by horizontal green lines. Above the breathing curve, an ECG signal is plotted. A prospectively ECG-triggered and respiratory-gated acquisition is indicated by the boxes in the ECG signal. Data is not acquired if the navigator yields a breathing state outside of the acceptance window (indicated by the red cross).

the current breathing state. The basic version of such navigators, is a 1-dimensional column placed on the right hemi-diaphragm as depicted in Figure 2.1a,b^{90,91}. The lung-liver interface shows as a transition from dark to bright in the navigator signal, providing a traceable feature. It has been shown that the respiratory motion of the heart can be mainly described by a large superior-inferior and a smaller anterior-posterior translation, which is proportional to the diaphragm position⁹². In prospective navigator gating, which is particularly useful for single shot imaging, the actual image acquisition is accepted/rejected based on the position of the interface as illustrated in Figure 2.1. In retrospective strategies, the acquired k-space points are binned according to the navigator position to reconstruct images per breathing state. Depending on the sequence, more sophisticated strategies using 2D navigators, the fat-signal⁹³ or self-navigation⁹⁴ have been proposed to overcome specific limitations of the 1D diaphragm navigator.

In cardiac diffusion imaging most methods use single shot acquisitions per diffusion weighting¹² (see 3.2). Typical protocols contain on the order of 100 images per dataset, which are assumed to be in the same motion state, i.e. at the same cardiac and respiratory phase. A single-breath-hold acquisition ensuring spatial consistency is infeasible due to the long scan duration. Multiple breath holds with navigator gating potentially suffer from longer total scan times, depending on patient compliance. For small gating windows, the preferable free-breathing acquisition becomes relatively inefficient, while larger gating windows provide less spatial consistency.

To facilitate the use of larger gating windows, image registration is commonly applied as retrospective motion compensation⁹⁵. The registration methods differ by the degrees of freedom in the applied transformation, such as rigid motion, affine

transformation and deformable registration^{96,97}. As the contrast varies significantly between different diffusion weightings, registration can be challenging. To address this, a variety of approaches including low-rank tensor decomposition prior to registration⁹⁸, respiratory binning based on spectral embedding⁹⁹, deep learning algorithms⁵⁰ and the inclusion of the full tensor information¹⁰⁰ have been proposed.

2.2 Encoding Displacements

2.2.1 DENSE

In displacement encoding with stimulated echoes (DENSE), the displacement with respect to a reference (usually the end-diastolic motion state) is encoded in the image phase for multiple points within the RR interval¹⁰¹. To this end, a mono-polar displacement encoding gradient is added between the initial excitation and the second RF pulse in a Stimulated Echo Acquisition Mode (STEAM) sequence (see Figure 2.2). After the third RF pulse, which generates the echo used for imaging, the displacement encoding gradient is played out again, resulting in zero net phase for static tissue. For a gradient pulse G_E of duration t_{enc} , resulting in the encoding strength $k_{\text{enc}} = \gamma G_E t_{\text{enc}}$ along e.g. the x-axis, a material point that has moved by Δx during the mixing time T_M (defined by the temporal gap between the second and third RF pulse), accumulates the phase $\phi = k_{\text{enc}} \Delta x$. While the Lagrangian strain (equation 1.1) can theoretically be computed directly from the measured phase¹⁰², phase wraps occurring in practice need to be addressed in data processing.

To obtain strain over the cardiac cycle, the sequence is repeated with increasing mixing time T_M . For 3-dimensional displacement information, each mixing time has to be acquired with displacement encoding along each of the spatial axes separately^{102,103}.

2.2.2 Tagging

There is a variety of tagging methods, all using sequences consisting of a preparation stage and an imaging stage. In the preparation stage, magnetization is spatially modulated, thereby creating patterns (referred to as "tags") in the signal magnitude, phase, or both. The deformation of the tags caused by tissue motion between tagging and image acquisition is derived during image processing. In the imaging stage, usually a spoiled gradient-recalled echo (GRE), EPI or a bSSFP sequence is used¹⁰⁴.

The most widely used tagging strategy is SPatial Modulation of Magnetization (SPAMM)¹⁰⁵ (Figure 2.3) and its improved versions. In SPAMM, tagging is performed by applying two RF pulses and a gradient pulse between them. The first pulse flips the magnetization into the transverse plane, where the gradient dephases the magnetization. The second RF pulse flips the magnetization back to the longitu-

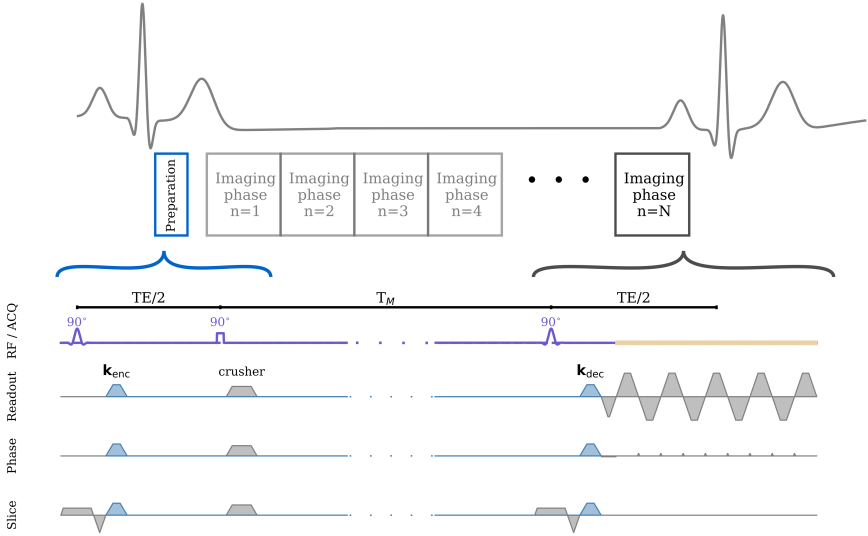


FIGURE 2.2: Schematic sequence diagram (bottom) of displacement encoding with stimulated echoes (DENSE) with reference to the ECG signal (top). The preparation or encoding module (blue box), consisting of slice-selective excitation, an encoding gradient pulse k_{enc} (blue trapezoidal gradients) and 90° hard pulse followed by a crusher is played out at the reference time. Imaging phases containing the third 90° pulse, the decoding gradient k_{dec} rewinding the k_{enc} -phase for static tissue and a segmented EPI readout, each correspond to one frame. The mixing time T_M defines which imaging phase is imaged. Each imaging phase corresponds to one frame of the cine acquisition.

dinal direction, followed by a crusher. In the subsequent imaging stage, the image magnitude is modulated sinusoidally due to the spatially modulated saturation. The encoding moment k_{enc} defines the distance as well as the direction of the tag lines. The highest possible spatial resolution of the obtained deformation field is defined by the minimum detectable distance of the tag-lines.

To create sharper edges of the tags, higher-order polynomial tagging pulses¹⁰⁶ or the Delays Alternating with Nutation for Tailored Excitation (DANTE) (Figure 2.3) method can be used. In DANTE, a series of block pulses is applied during a tagging gradient, resulting in a Dirac-comb like response function¹⁰⁷ (Figure 2.3). To address the effect of tag fading due to T_1 recovery, Complementary-SPAMM (CSPAMM)^{108,109} was proposed, where the tagging is acquired a second time with inverted magnetization. By subtracting the two tagged images the tag fading is compensated. To achieve 2D deformation fields, the tagging stage can be repeated in a second direction resulting in a grid of taglines. Alternatively, radial, ring or variable density tagging may be applied¹⁰⁴.

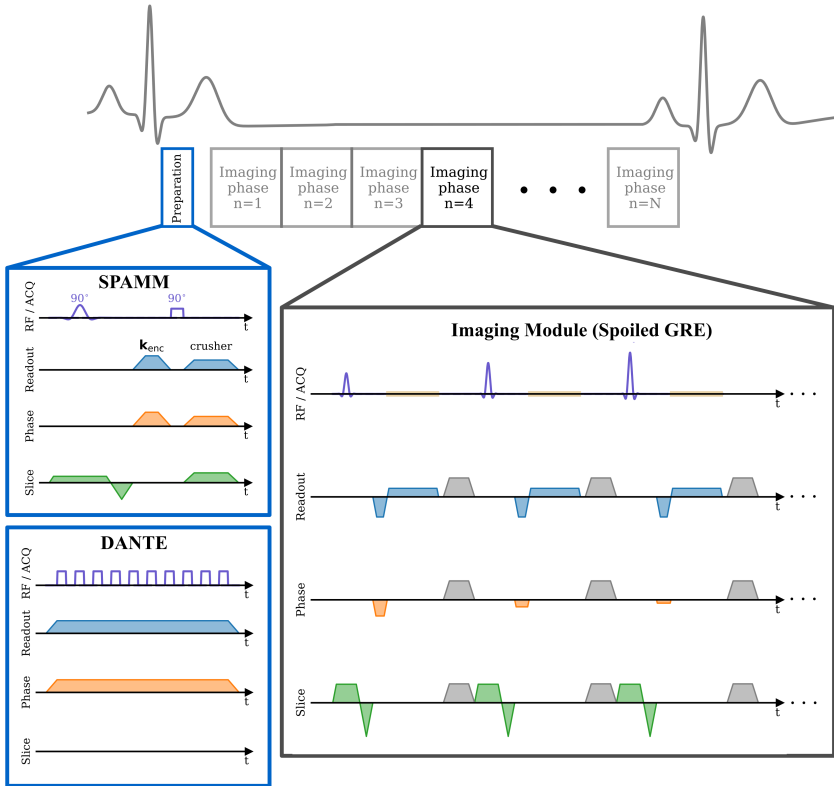


FIGURE 2.3: Schematic sequence diagram of a tagging sequence relative to the ECG signal. The preparation or tagging module (blue box) is played out at the reference contractile state. Two types of tagging modules are illustrated: SPAMM and DANTE. As imaging module, a spoiled gradient echo sequence with increasing flip-angle is depicted. Each imaging phase corresponds to one frame of the cine acquisition

2.2.3 Feature Tracking

In contrast to the methods described above, feature tracking (FT) does not require a separate scan, as it uses bSSFP cine images routinely acquired in cardiac protocols^{8,10}. The tracked features are anatomic structures rather than magnetization tags, which are detected by image-processing algorithms. Features are extracted from the neighbourhood of initially defined reference points. The trajectory of the points is tracked based on optical flow algorithms, using the extracted features¹¹⁰. As the signal in the myocardium is homogeneous, this approach mostly works based on points along the myocardial borders with high contrast. Strain estimation based on FT is therefore subject to the accuracy of segmentation and positioning

of reference points. Furthermore, since feature tracking is typically performed on 2D images, features moving out of the imaging plane potentially hamper the FT performance¹¹⁰.

2.3 Encoding Tissue Velocity

2.3.1 Phase Contrast MRI

In an MR measurement with static and homogeneous magnetic field B_0 and an applied, arbitrary gradient waveform $\mathbf{G}(t)$, assuming the Lagrangian framework of moving material points (see 4.1) with trajectories $\mathbf{r}(t)$, each material point accumulates a phase $\phi(t)$ according to

$$\phi(t) = \gamma \int_{t_0}^t \mathbf{G}(t') \cdot \mathbf{r}(t') dt' \quad (2.1)$$

$$= \gamma \sum_{n=1}^{\infty} \frac{1}{n!} \left. \frac{d^n \mathbf{r}(t)}{dt^n} \right|_{t_0} \cdot \int_{t_0}^t \mathbf{G}(t') (t' - t_0)^n dt'. \quad (2.2)$$

Here, equation 2.2 is the Taylor expansion of the phase defined in equation 2.1¹¹¹. The integral in the n^{th} -order expansion term, including the gradient waveform $\mathbf{G}(t')$, is called the n^{th} -order gradient moment $\mathbf{m}_n(t)$. By defining the gradient waveform, such that all moments up to a given order evaluate to zero, the accumulated phase $\phi(T)$ directly corresponds to the n^{th} -order term of the particle-trajectories' Taylor expansion.

In PC MRI, this principle is used to encode the velocity of flowing blood or moving tissue¹¹². To this end, a bipolar gradient with a non-vanishing first order moment $\mathbf{m}_1(T)$ is added between excitation and readout of a given MR sequence. For a sufficiently short encoding duration, the phase corresponds to the velocity by $\mathbf{v} \cdot \mathbf{g} = \phi_v(T)/\gamma |\mathbf{m}_1(T)|$ along the encoding direction $\mathbf{g} = \mathbf{m}_1(T)/|\mathbf{m}_1(T)|$. This approximation assumes, that the blood flow is accurately represented by the first-order term over the duration T of the applied gradient waveform $\mathbf{G}(t)$ ¹¹¹.

As the phase is cyclic, high velocities can result in phase wraps. To avoid wraps and utilize the full dynamic range of velocity encoding, $\mathbf{m}_1(T)$ is chosen such that the maximum expected velocity results in a phase of $\pi = \gamma \mathbf{m}_1(T) \cdot \mathbf{v}_{\text{max}}$. Accordingly, the encoding strength $\text{VENC} = \pi/\gamma |\mathbf{m}_1(T)|$ can be defined, yielding the following expression

$$\mathbf{v} = \mathbf{g} \frac{\phi_v}{\pi} \text{VENC}. \quad (2.3)$$

In addition to the motion-induced phase accumulation, other effects such as off-resonance, eddy currents, and concomitant fields result in phase contributions, contributing to a background phase term ϕ_0 . To separate the phase contribution of the three-dimensional velocity ϕ_v from ϕ_0 , at least four measurements with varying velocity encoding gradients are necessary. A set of N measurements, where each measurement n uses the encoding strength $VENC_n$, can be formulated as encoding matrix $\mathbf{E}_v \in \mathcal{R}^{N \times 4}$

$$\Phi = \mathbf{E}_v \cdot \left(\phi_0, \frac{\pi}{VENC_n} v_x, \frac{\pi}{VENC_n} v_y, \frac{\pi}{VENC_n} v_z \right)^T, \quad (2.4)$$

where $\Phi \in \mathcal{R}^N$ contains the image phase of a single image pixel. The simplest encoding matrix corresponds to a reference measurement and three velocity encodings along the Cartesian axes resulting in the matrix¹¹³:

$$\mathbf{E}_{v,\text{Simple}} = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix}. \quad (2.5)$$

Another commonly used 4-point encoding scheme, referred to as Hadamard encoding^{113,114}, is given by the matrix:

$$\mathbf{E}_{v,\text{Hadamard}} = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix}. \quad (2.6)$$

The simultaneous encoding of all directions in each acquisition results in directional independence of the estimation variance, which is determined by the measurement noise. Furthermore, for a given encoding moment, the estimation variance is reduced by a factor of 2 using the Hadamard scheme compared to the simple 4-point scheme¹¹³. On the other hand, the dynamic range, defined as the maximum velocity that can be encoded in all directions without a phase wrap, also is reduced by a factor of $\sqrt{2}$ for a fixed encoding moment¹¹³. Hadamard encoding can lead to non-integer phase wraps¹¹⁵ that are directionally coupled¹¹³ and thus more difficult to remove in post-processing than phase-wraps in simple 4-point encoding.

2.3.2 Practical acquisition

The most basic MR sequence used in Tissue Velocity Mapping (TVM) is a spoiled GRE sequence¹¹⁶. To achieve velocity sensitization, bipolar gradients are inserted between excitation and readout^{117,118} as illustrated in Figure 2.4, showing the sequence kernel for a single repetition. One repetition corresponds to a single k-space line, and a single velocity encoding. The sequence kernel is repeated continuously during each RR interval, which is subdivided into a given number of bins. The number of bins determines the temporal resolution of the cine acquisition, as each bin corresponds to one image. Kernel repetitions within a bin can be used to acquire multiple k-space lines as illustrated in 2.4b. Thus, the number of required RR intervals to acquire all k-space lines for each of the bins depends on the number of TRs per bin and the image resolution.

Most of the published TVM work used breath-held acquisitions to reduce the phase contribution due to breathing motion^{119–124}. However, depending on the temporal and spatial resolution, the required breath-hold duration becomes infeasible. To this end, view-sharing of bins has been proposed^{119,120}, where the acquisition order of k-space lines allows to reconstruct intermediate frames. The increased scan efficiency allows to acquire the reference and two velocity encoding directions, for a single slice, with a temporal resolution of 68 ms in 16–19 RR-intervals. Furthermore, free-breathing approaches using interleaved respiratory navigators¹²⁵ or self-navigated golden angle spiral acquisition¹²⁶ have been proposed.

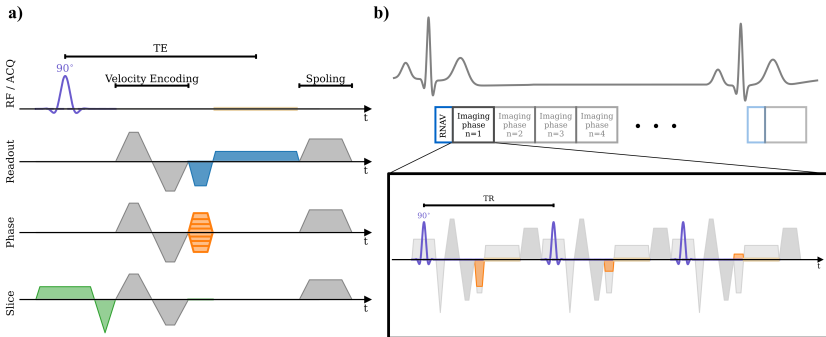


FIGURE 2.4: a) Sequence diagram of showing a single repetition of a spoiled gradient echo with a velocity-sensitizing bipolar gradient and Cartesian readout. Phase encoding is indicated for multiple lines b) Sequence with three TR in a single imaging phase, with increasing phase encoding relative to an ECG signal.

2.4 Encoding Diffusional Motion

As described in 1.2, the relevant length scales of the myocardial microstructure are on the order of $10\ \mu\text{m}$ to $100\ \mu\text{m}$, which is well beyond the achievable image resolution of clinical CMR. To assess information at this scale, diffusion weighting with bipolar diffusion encoding gradients has been used¹²⁷. The diffusion gradients impose a phase on randomly moving spins. The random motion creates a phase distribution in the spin population within a voxel. This results in a magnitude modulation of the ensemble magnetization⁵⁷.

2.4.1 Signal Model & Encoding

An effective description of the ensemble magnetization, including the effect of diffusional spin motion, is provided by the Bloch-Torrey equation^{128,129}:

$$\frac{\partial}{\partial t} m(t, \mathbf{r}) = \frac{\partial}{\partial \mathbf{r}} \left[D(\mathbf{r}) \frac{\partial}{\partial \mathbf{r}} m(t, \mathbf{r}) \right] - [R_2(\mathbf{r}) + i\Omega(t, \mathbf{r})] m(t, \mathbf{r}), \quad (2.7)$$

where $D(\mathbf{r})$ and $R_2(\mathbf{r})$ are the diffusion coefficient and transverse relaxation rate at location \mathbf{r} . The Larmor frequency $\Omega(t, \mathbf{r}) = \Omega(\mathbf{r}) + \mathbf{G}(t) \cdot \mathbf{r}$ consists of the sample-intrinsic term $\Omega(\mathbf{r})$ and the field variation due to the applied gradient $\mathbf{G}(t)$. This description can only be fully determined for a specified spatio-temporal scale, as its parameters are scale dependent¹²⁹. The scale is implicitly defined by the measurement duration and the instantaneous diffusion constants within the compartments. One fundamental challenge of obtaining information about the microscopic scale is to identify which feature of the effective parameters remain observable after the intrinsic spatial averaging of the MR acquisition¹²⁹.

Assuming a spatially homogeneous relaxation rate and Larmor frequency on a length scale that is larger than the diffusion length, equation 2.7 can be rewritten:

$$\frac{\partial}{\partial t} \psi(t, \mathbf{r}) = \frac{\partial}{\partial \mathbf{r}} \left[D(\mathbf{r}) \frac{\partial}{\partial \mathbf{r}} \psi(t, \mathbf{r}) \right] - \mathbf{G}(t) \psi(t, \mathbf{r}), \quad (2.8)$$

with

$$\psi(t, \mathbf{r}) = e^{-R_2 t - i(\Omega) t}. \quad (2.9)$$

There is no universal signal expression when applying arbitrary diffusion gradients. The signal has a functional dependency on the gradient applied during the measurements $S[\mathbf{g}(t)]$. To obtain an explicit expression of this dependency, equation 2.8 can be treated for small diffusion weightings. This yields the general expression^{57,129}:

$$S[\mathbf{g}(t)] = S_0 \langle e^{i\phi(t)} \rangle \quad (2.10)$$

with

$$\phi(t) = \int_0^t dt' \mathbf{g}(t') \cdot \mathbf{r}(t'), \quad (2.11)$$

where $\phi(t)$ is the phase a single spin with position $\mathbf{r}(t)$ accumulates due to the applied gradients $\mathbf{g}(t)$. The angle brackets represent the averaging over all spins in one voxel. Under the condition of balanced gradients $\int_0^T dt' (\mathbf{g}(t')) = 0$ over the duration T , integration by parts in equation 2.11 yields^{57,130}:

$$\phi(t) = - \int_0^t dt' \mathbf{q}(t') \cdot \mathbf{v}(t') \quad \text{with} \quad \mathbf{q}(t') = \int_0^{t'} dt'' \mathbf{G}(t''). \quad (2.12)$$

To associate the general signal representation in equation 2.10 to the diffusional motion of the spins, the Taylor expansion of the logarithmic signal can be written as

$$\begin{aligned} \ln S[\mathbf{q}(t)] &= i \int u_a^{(1)}(t_1) q_a(t_1) dt_1 \\ &- \frac{1}{2!} \int u_{ab}^{(2)}(t_1, t_2) q_a(t_1) q_b(t_2) dt_1 dt_2 \\ &- \frac{i}{3!} \int u_{abc}^{(3)}(t_1, t_2, t_3) q_a(t_1) q_b(t_2) q_c(t_3) dt_1 dt_2 dt_3 \\ &+ \frac{1}{4!} \int u_{abcd}^{(4)}(t_1, t_2, t_3, t_4) q_a(t_1) q_b(t_2) q_c(t_3) q_d(t_4) dt_1 dt_2 dt_3 dt_4 \\ &+ \dots \end{aligned} \quad (2.13)$$

where the rank n tensor $u^{(n)}$ is the n -th cumulant of the molecular velocity distribution and assuming normalization by S_0 . Here $abcd\dots$ indexes the three spatial dimensions and repeating indices in the integral imply summation according to Einstein's convention. As an example, the 3×3 rank 2 tensor constituting the second-order cumulant (also known as variance) can be written as¹³⁰:

$$u_{ab}^{(2)} = \langle v_a(t_1) v_b(t_2) \rangle - \langle v_a(t_1) \rangle \langle v_b(t_2) \rangle. \quad (2.14)$$

The first term of equation 2.14 is the averaged product of molecular velocity at $n = 2$ time points, called the n -point correlation function. This correlation function contains information about the microstructure. As previously stated, a fundamental challenge in probing the tissue micro-structure with MRI is to develop a model of the tissue capable of explaining the velocity correlations on the scales of the measurement duration^{27,57}. An exemplary explanation for the two-point correlation of velocities is collisions with barriers introducing negative correlations, where the correlation time depends again on the presence of other restrictions.

Equation 2.13 is a general signal representation, where its applicability is limited by the encoding strength of the applied gradients. Commonly, the diffusion weighting strength is stated as b -value for a measurement with echo time TE :

$$b = \int_0^{TE} \mathbf{q}^2(t) dt. \quad (2.15)$$

Achievable b -values on clinical scanners with high-performance gradient systems (80 mT/m and 100 mT/(m ms)) are on the order of 500 s/mm². The achievable b -value is relatively low when compared with diffusion experiments in other organs, because of the additional constraints arising from cardiac motion (see 3.1). The regime of b -values is well within the convergence radius of the cumulant expansion (equation 2.13)¹³⁰. Furthermore, the signal is well represented by the first two terms of equation 2.13 for such low b -values¹³¹, where the first term evaluates to zero in case of zero net bulk motion (corresponding to 2.3.1). Using only cumulants up to the second order is referred to as the Gaussian phase approximation providing the framework for DTI.

In the most basic form of DTI, the signal is fitted to a simplified version of the Gaussian phase approximation:

$$S(b, n_i) = S_0 e^{-bn_i D_{ij}^{app} n_j}, \quad (2.16)$$

where b is the b -value defined in equation 2.15, g_i is the direction of the applied linear diffusion encoding and D_{ij}^{app} is an entry of the symmetric 3×3 apparent diffusion tensor. Again, repeating indices imply summation.

The integrals in equation 2.13 contain the gradient waveforms (inside the definition of $\mathbf{q}(t)$) which implies that the measured diffusion tensor in equation 2.16 depends on the gradient duration and shape. To formalize this dependency the velocity correlation $u_{ab}^{(2)}$ is expressed using its retarded form, reflecting the causality of the correlation:

$$\mathcal{D}_{ab}(t) = u_{ab}^{(2)} \theta(t_2 - t_1), \quad (2.17)$$

where θ is the step function. To evaluate the actual signal response, it is more convenient to rewrite the signal equation in the Fourier domain by using $\mathcal{D}(\omega) = \mathcal{F}\{\mathcal{D}(t)\}$ and $\mathbf{q}(\omega) = \mathcal{F}\{\mathbf{q}(t)\}$ ¹³²:

$$\ln(S(\omega, \mathbf{q}(\omega))/S_0) = - \int \frac{d\omega}{2\pi} q_i(\omega) \mathcal{D}_{ij}(\omega) q_j(\omega). \quad (2.18)$$

Figure 2.5 illustrates how the gradient spectrum corresponds to the diffusion sampling process. For free diffusion, the velocities at different times are completely uncorrelated, which translates to a second-order cumulant represented by a delta function, hence a flat spectrum. If the correlation time is on the same order as the diffusion gradient waveform duration, the dependence of the measured apparent diffusion tensor

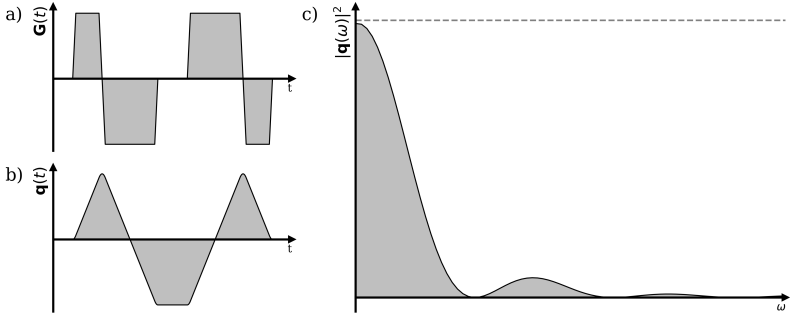


FIGURE 2.5: a) Second-order motion-compensated gradient waveform $G(t)$. b) $q(t)$ (equation 2.12) corresponding to the waveform in a). c) Power spectrum of the function depicted in b). The dashed line corresponds to a the spectrum of free diffusion.

$$D^{app} = \frac{\int \frac{d\omega}{2\pi} q_i(\omega) \mathcal{D}_{ij}(\omega) q_j(\omega)}{\int \frac{d\omega}{2\pi} |q(\omega)|^2} \quad (2.19)$$

on the actual gradient shape becomes obvious. The denominator in equation 2.19 is equivalent to the time-domain definition of the b-value in equation 2.15 and connected via Parseval's theorem. As a consequence of the time dependency, directly comparing measured diffusion values only makes sense if identical sequences are applied.

2.4.2 Tensor Inference

For diffusion encoding schemes with a given set $N \geq 6$ different encodings with b-value b and direction \mathbf{g} yielding $\mathbf{b}_n = b\mathbf{g}\mathbf{g}^T$, the simplest, pixel-wise estimation of the diffusion tensors D_{ij} and unweighted signal S_0 from the measurements S_n is given by the direct matrix inversion. Therefore, let

$$\tilde{\mathbf{d}} = (\log(S_0), D_{xx}, D_{yy}, D_{zz}, D_{xy}, D_{xz}, D_{yz})^T, \quad (2.20)$$

$$\tilde{\mathbf{b}}_n = (-1, b_{xx}, b_{yy}, b_{zz}, b_{xy}, b_{xz}, b_{yz})_i^T \rightarrow \mathbf{B} = (\tilde{\mathbf{b}}_1, \tilde{\mathbf{b}}_2, \dots), \quad (2.21)$$

and the signal model for Gaussian noise be

$$\mathbf{x} = (\log(S_1), \log(S_2), \dots)^T = \mathbf{B}\tilde{\mathbf{d}} + \eta, \quad (2.22)$$

then the solution to the optimization problem

$$\arg \min_{\tilde{\mathbf{d}}} \|\mathbf{x} - \mathbf{B}\tilde{\mathbf{d}}\|_2^2 \quad (2.23)$$

is the pseudo inverse of \mathbf{B}^{133} . More theoretical background on inverse problems is provided in chapter 6.

The SNR in strongly weighted images is lower due to the intended magnitude modulation, which potentially introduces a bias during fitting if the signal is close to the noise floor¹³⁴. To address this effect, advanced fitting algorithms such as weighted linear least squares¹³⁵ or iterative linear and nonlinear inversion^{136–138} have been proposed.

2.4.3 Interpretation

To infer information about myocardial microstructure, the apparent diffusion tensor is decomposed into its three eigenvalue-eigenvector pairs $\lambda_{i=1,2,3}$ and $\mathbf{e}_{\lambda,i=1,2,3}$. Assuming descending order of the eigenvalues $\lambda_1 > \lambda_2 > \lambda_3$, the first eigenvector $\mathbf{e}_{\lambda 1}$ has been shown to align with the direction of the cardiomyocytes^{13,15,139,140}. The combination of $\mathbf{e}_{\lambda 1}$ and $\mathbf{e}_{\lambda 2}$ spans a plane, with $\mathbf{e}_{\lambda 3}$ being the normal, that is associated with the sheetlet orientation as described in 1.2^{141,142}. Figure 2.6a illustrates the correspondence between eigenvectors and myocyte as well as sheetlet orientation.




FIGURE 2.6: a) Illustration of correspondence between diffusion tensor eigenvectors $\mathbf{e}_{\lambda 1}$ (red), $\mathbf{e}_{\lambda 2}$ (blue), $\mathbf{e}_{\lambda 3}$ (green) and myocyte (red cylinders) / sheetlet (opaque white box) orientations. b) Illustration of diffusion tensor shape resulting from myocyte-orientation (dis-)order; overall alignment results in high anisotropy (left) while disarray results in a more spherical tensor. Figure is based on [12].

To obtain scalar measures, the first eigenvector is described by the helix angle (HA) on the one hand and the transverse angle (TA) on the other. HA and TA can be computed relative to the local coordinate system of radial, circumferential and longitudinal direction, as shown in Figure 2.7a. The angle between the radial-circumferential plane and $\mathbf{e}_{\lambda 1}$ defines HA and the angle between the projection of $\mathbf{e}_{\lambda 1}$ onto that plane and the circumferential basis vector defines TA¹⁴³ (Figure 2.7a). The sheetlet angle can be computed as the angle between the cross-myocyte orientation \mathbf{e}_{CF} and the projection of $\mathbf{e}_{\lambda 3}$ onto the cross-fiber plane¹⁴³ as illustrated in Figure 2.7c .

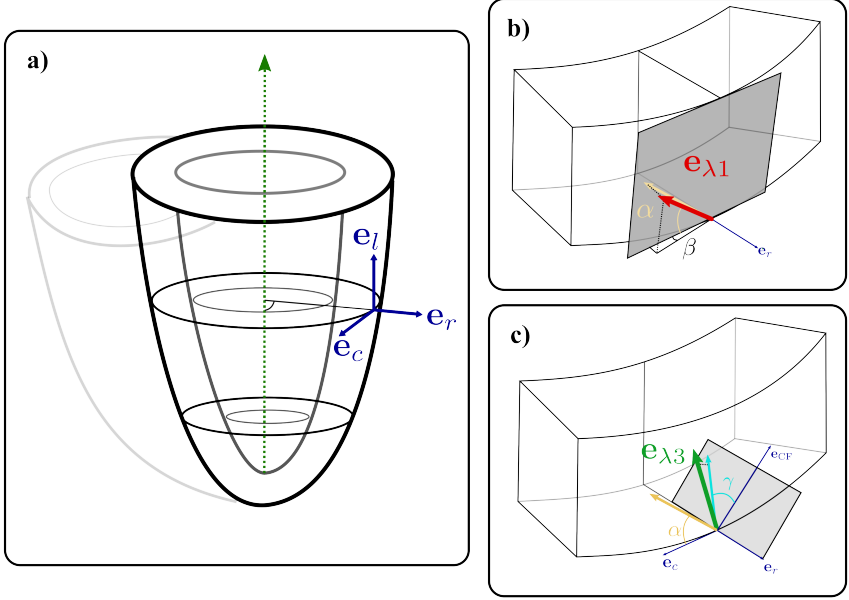


FIGURE 2.7: a) Definition of the local coordinate system with radial e_r , circumferential e_c and longitudinal e_l basis vectors. The long axis of the ventricle is depicted as dotted green arrow, from apex to base. b) Graphical definition of the helix angle (α) and transverse angle (β). The local epicardial plane (perpendicular to the local radial direction e_r) is depicted as gray rectangle. The first diffusion-eigenvector $e_{\lambda 1}$ is depicted as red arrow and its projection onto the epicardial plane is depicted as beige arrow. c) Graphical definition of the sheetlet angle (γ). The cross-myocyte plane is depicted as gray rectangle, the beige arrow is the $e_{\lambda 1}$ -projection from b), e_r and e_c are local radial and circumferential directions and e_{CF} defines the cross-myocyte direction. The third eigenvector $e_{\lambda 3}$ and its projection onto the cross-myocyte plane are depicted as green and turquoise arrows. Definitions are based on [143].

In addition to the fiber and sheetlet orientation, two more tensor metrics are defined to characterize the tissue. The mean of eigenvalues or the mean diffusivity (MD) is calculated as:

$$MD = \text{Tr}(\mathbf{D}^{app})/3, \quad (2.24)$$

and captures the isotropic component of the diffusional motion. A high MD generally hints at low hindrance for water molecules, associated with damaged microstructure (e.g. due to infarction)^{144,145}. Fractional anisotropy is defined as:

$$FA = \sqrt{\frac{3 \sum_i (\lambda_i + MD)^2}{2 \sum_i \lambda_i^2}}, \quad (2.25)$$

and is the most common measure of anisotropy, which is associated with the degree of fiber order within the voxel^{11,12} as depicted in Figure 2.6b.

2.5 Related Cardiac Pathologies

A number of pathologies are intrinsically coupled to the state and contractility of the myocardial microstructure. Additionally, the global function of the heart is influenced by pre- and afterload as well as its LV geometry³³. While strain in general is a widely used marker for cardiac diseases³³, the underlying factors influencing strain vary between diseases⁸⁹. In the clinics, mainly global strains at certain time points in the cardiac cycle are regarded^{8-10,33,146}. However, local variations may provide additional information relevant for risk estimation and treatment planning¹⁴⁷.

One condition, cDTI was shown to provide detailed information about is myocardial infarction and subsequent scar formation¹⁴⁸⁻¹⁵⁰. As myocytes in the infarcted region are damaged, diffusion becomes less restricted, especially in cross-myocyte orientation. This results in decreased FA and increased MD values^{17,22}. Furthermore, tissue strain is regionally impaired in ischemic heart disease and subsequent infarction³³, directly correlating with infarct size, distinguishing transmural as well as viability of tissue^{147,151}. As tissue viability is connected to the integrity of myocytes, combining strain and cDTI information could improve treatment planning and outcome prediction^{17,152,153}.

The heart undergoes a process of remodelling in response to various pathophysiological alterations of the body¹⁵⁴. For example, as a result of infarction, the rest of the myocardium remodels to compensate for the loss of contracting tissue^{14,21,149,153,155}. In hypertrophic cardiomyopathy, the heart muscle wall thickness and stiffness increase, which is connected to reduced sheetlet mobility. The reduced sheetlet mobility can be measured by comparing diastolic and systolic sheetlet angles from cDTI data^{18,20} as well as reduced FA¹⁵⁶. The reduced macroscopic contractility was shown with TVM¹⁵⁷. The interconnection of strain and microstructure in dilated cardiomyopathy was investigated using tagging and cDTI¹⁹, revealing decreased sheetlet reorientation and systolic strain, as well as steeper HAs. Remodelling can also occur due to LV overload¹⁵⁸ as in aortic stenosis, showing increased MD and decreased FA as well as elevated HA slope²⁵. Higher cardiac load in obesity was shown to impair peak diastolic tissue velocity¹⁵⁹ as well as increased MD with decreased FA and HA transmural⁵³.

cDTI was also used to investigate the microstructural changes in cardiac amyloidosis^{23,24} and congenital heart disease^{160,161}. Moreover, early cardiac toxicity of certain treatments such as radiotherapy can be monitored by strain imaging¹⁶².

3. Spin Echo cDTI

There are two sequences commonly used for cDTI, namely STEAM and SE. Both sequences rely on consistent ECG triggering (see section 2.1.1). STEAM-cDTI looks similar to the DENSE acquisition but the excitation and readout are spread over two RR-intervals, thus encoding and decoding of displacement occurs at the same trigger delay with matching contraction state (see Figure 3.1a). Therefore, instead of bulk motion, diffusion displacement of spins is encoded¹². In SE diffusion sequences excitation and readout are performed within the same RR interval as illustrated in Figure 3.1b. This allows acquisition during free-breathing, hence it is more suitable for clinical applications. The following sections provide a basic overview for all blocks of SE-cDTI.

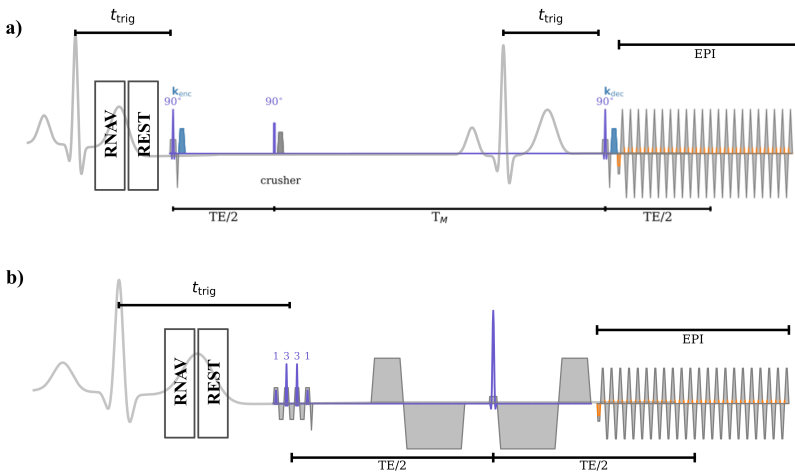


FIGURE 3.1: Comparison of a diffusion-weighted a) STEAM and b) SE sequence with reference to an ECG-signal. The trigger-delay is marked as t_{trig} . Both sequences use a single-shot EPI readout and REgional Saturation Technique (REST) slabs as well as a leading navigator. STEAM is acquired over two RR intervals, while the SE acquisition is performed within a single heart beat. The diffusion encoding k_{enc} and decoding k_{dec} gradients in STEAM are applied with the same delay to the previous RF-pulse. The SE sequence uses a binomial spectral-spatial pulse and a second-order motion-compensated diffusion gradient waveform. Based on [163].

3.1 Diffusion Encoding Waveforms

Compared to diffusion MRI applied in other anatomical regions (e.g. neurology), the contractile motion of the heart creates restrictions for certain concepts. Foremost, tissue strain results in intra-voxel phase gradients if diffusion encoding waveforms

are applied. These intra-voxel gradients lead to signal magnitude modulation, potentially even leading to signal dropout¹⁶⁴. This signal modulation alters the estimation of diffusion tensors, such that the diffusivity is artificially increased^{164–166}.

For STEAM, one strategy of mitigating those strain effects is the so-called *sweet spot* imaging¹⁶⁷. This strategy, the trigger delay is chosen as one of two time points (one systolic and one diastolic) in which the integral of the strain evaluates to zero, therefore minimizing the effect on the measured diffusion tensor¹⁴³.

For SE acquisitions the preferred strategy to mitigate these signal losses is to construct the diffusion encoding waveforms such that the moments, as defined in 2.1, are zero up to a sufficient order^{29,30,166,168}. This strategy is only feasible if the diffusion encoding and acquisition are within intervals of the cardiac cycle in which the material point trajectories are described by a Taylor polynomial of low order. In early systole and around peak contraction this condition is violated, resulting in signal loss despite motion compensation¹⁶⁹. While higher-order motion compensation promises robustness against strain effects, the minimum achievable duration for the waveform increases¹¹¹. Second-order motion compensation has been shown to be robustly applicable while yielding waveforms short enough^{12,29,30,170}.

More recently, an advanced diffusion encoding strategy called q-space trajectory encoding (QTE)^{171–173} was applied to cardiac imaging^{170,174}. In QTE, instead of using unidirectional gradients, the applied encoding waveforms describe a three-dimensional trajectory in q-space (as defined in eq. 2.12). Therefore, instead of a scalar b-value, eq. 2.15 generalizes to the so-called B-tensor B_{ij} of rank 2. The exponent in the signal model (equation 2.16) then evaluates as $-\Sigma_{ij}B_{ij}D_{ij}$. This allows e.g. to measure the diffusion tensor trace with a single waveform¹⁷⁰. Furthermore, if signal contributions associated with higher than second order cumulants are considered, these waveforms allow access to information on distributions of tensors within voxels^{171,174}. Motion compensation is also necessary in these advanced encoding waveforms¹⁷².

3.2 Echo Planar Imaging

This section is based on chapter 16 of the textbook [111]

EPI is one of the fastest encoding strategies in MRI, consisting of a train of gradients in readout direction with alternating amplitude, combined with intermittent phase encoding (referred to as blips). Each lobe of the readout gradient train generates a gradient-recalled echo, corresponding to one k-space line. The phase blips shift the trajectory to the next phase encoding (PE) step, which results in a meandering k-space trajectory as illustrated in Figure 3.2b. Due to the fast gradient switching and long acquisition window, EPI is prone to effects altering the effective trajectory, often resulting in geometric distortions and so-called EPI ghosts. The ghost artifacts

in PE direction stem from a mismatch of lines depending on the polarity of the readout lobe. A calibration scan without blips allows to correct the phases per line.

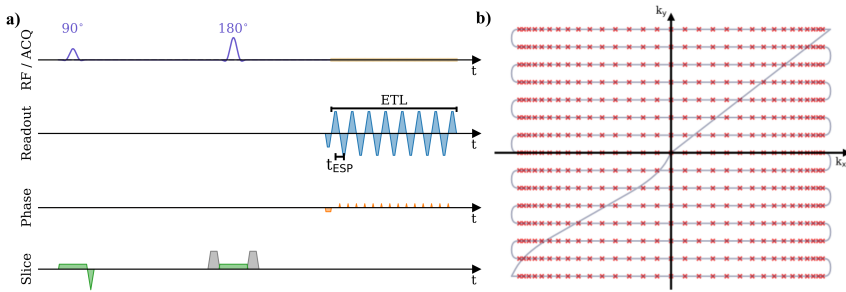


FIGURE 3.2: a) Sequence diagram of a spin-echo acquisition with single-shot EPI readout. Echo Train Length (ETL) and echo-spacing t_{ESP} b) k-space trajectory (blue line) and sampling points (red markers) of the EPI readout.

The number of echos in the train is denoted as echo train length (ETL), where the echos are commonly uniformly spaced by t_{ESP} . The bandwidth in PE direction is defined by $1/t_{ESP}$, which determines the degree of geometric distortions and displacements due to local off-resonance fields as described below. Therefore, the minimal achievable echo spacing (ESP) (corresponding to highest bandwidth) is desirable for cDTI. In most cases the readout gradient train is gradient slew-limited, hence, a short t_{ESP} requires sampling during the gradient ramps resulting in non-equidistant k-space points.

With diffusion encoding gradients played out in the SE sequence, the achievable minimum echo time (TE) is mainly determined by the ETL. One approach to reduce the ETL is to skip the first couple of outer k-space lines at the beginning of the echo train (referred to as partial Fourier acquisition).

3.3 Fat Suppression

This section is based on chapters 4 and 5 of the textbook [111]

The myocardium is surrounded by tissues containing fat, whose signal can be spatially shifted depending on the acquisition bandwidth. Therefore, fat suppression should be applied during cDTI acquisition. The two most common techniques to achieve this are the use of SPectral-SPatial (SPSP) pulses for excitation and Spectral Pre-saturation with Inversion Recovery (SPIR).

All variants of SPSP pulses are composed of multiple RF subpulses¹¹¹. Concurrently, an oscillating bipolar slice-selection gradient is applied, determining the spatial selectivity, while the RF-envelope, modulating the subpulse flip angles, defines

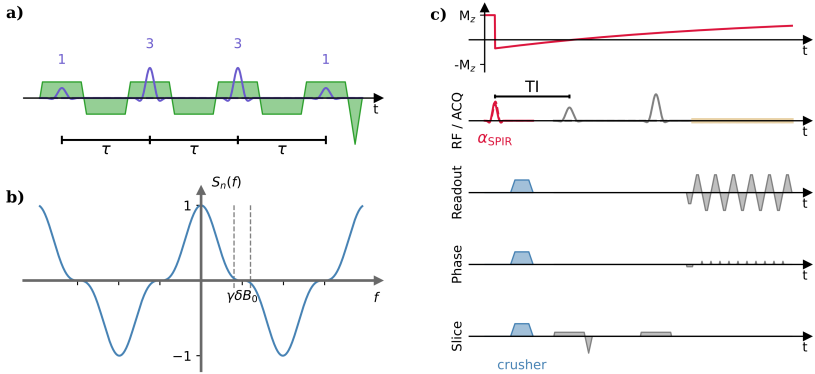


FIGURE 3.3: a) Diagram of a slice-selective binomial 1-3-3-1 spatial-spectral excitation pulse and b) the corresponding frequency response. The suppressed frequency range is marked by vertical dashed lines around the target frequency $\delta B_0 \gamma$. c) Sequence diagram of a SPIR fat-saturation pre-pulse, where the inversion is applied such that the zero crossing occurs at SE excitation. The spatially selective inversion pulse α_{SPIR} has a frequency offset matching the chemical shift of fat.

spectral selectivity¹¹¹. Binomial SPSP pulses are the most commonly used class, where the subpulse flip angles follow the binomial coefficients as illustrated in Figure 3.3a. The frequency response of these pulses is given as $S_n(f) = \cos^n(\pi f \tau)$, with n being the order of the binomial coefficient and τ is the spacing between the subpulses. Thus, choosing $\tau = 1/2f$, with f determined by the chemical shift frequency of fat, results in zero-crossings of the frequency response at the fat peak, thus suppressing fat (see Figure 3.3b).

In SPIR, a spectrally selective pulse is played out before the actual image acquisition, as illustrated in Figure 3.3c. This pulse excites the fat selectively followed by a crusher gradient. The excitation of the image acquisition is timed such that the recovery of the fat magnetization is at the zero crossing (see 3.3c, top row) suppressing the fat signal contributions to the image.

3.4 Reduced Field of View

The resolution (field-of-view (FOV) divided by the matrix size) in the PE direction mainly determines the echo time of the SE cDTI scan. Therefore, it is desirable to reduce the number of acquired k-space lines corresponding to the number of blips in the EPI train. To avoid fold-over artifacts when reducing the FOV in PE direction below the anatomic extent, the magnetization outside of the FOV must be suppressed. Commonly used techniques to achieve this are REST (also referred to as pre-saturation pulses), inner-volume selection by rotated slice selective excitation, and two-dimensional excitation pulses.

Two types of multi-dimensional excitation pulses have been used in SE-cDTI: 2D echo planar pulses⁹⁸ and pencil-beam excitation with spiral readouts³¹. For an extensive review of the general principle of multi-dimensional excitation pulses the reader is referred to [111, pp. 125].

REST

In REST a series of slice-selective 90° pulses are applied followed by a crusher, prior to the actual acquisition. The slice position and thickness are chosen such that the magnetization in the slabs around the volume of interest is saturated. A sequence diagram for a simple SE scan and an illustration of the regional saturation slab placement is shown in Figure 3.4.

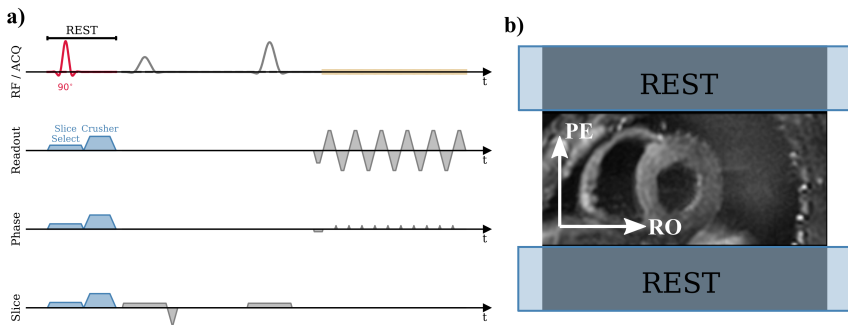


FIGURE 3.4: a) Sequence diagram of a regional saturation (REST) pre-pulse followed by a spin-echo acquisition with single-shot EPI readout. b) Illustration of the saturation slab placement relative to the heart in a short-axis view. Phase encoding (PE) and readout (RO) directions are indicated by arrows.

Inner Volume Excitation

When two slice-selective RF-pulses are used in a SE sequence, one way to suppress signal in the PE direction is to rotate the slice selection gradient of the refocusing pulse as illustrated in Figure 3.5. While a perpendicular configuration produces the sharpest border in the PE direction and works well in single slice acquisitions, the saturation of neighbouring slices and the resulting signal reduction is detrimental for the acquisition of multiple slices. An adapted strategy of slice orientations in combination with REST, as illustrated in 3.5b, is beneficial in this case¹⁷⁵.

3.5 Off-Resonances

Tissue interfaces can introduce susceptibility gradients and hence local field homogeneity^{176,177}. Furthermore, the oxygenation of blood within coronary vasculature can introduce significant differences in susceptibility compared to the surrounding tissue⁷⁶. These

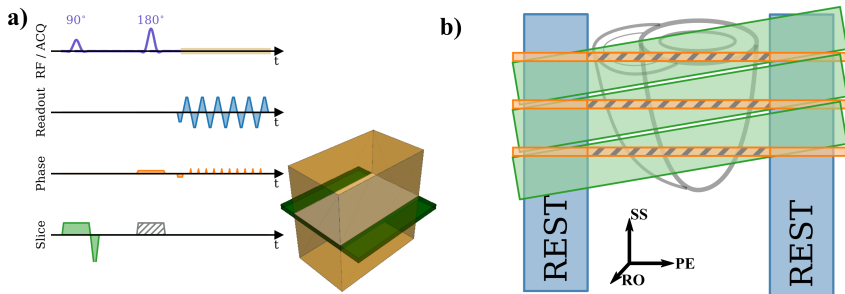


FIGURE 3.5: a) Sequence diagram of a spin-echo acquisition with single-shot EPI readout with orthogonal inner volume excitation. The hatched gradient pulse corresponds to a simple slice-selective refocusing. Instead of this pulse, the gradient is played out in PE direction, where the lower amplitude results in a thicker slab to be refocused. Next to the sequence diagram is a 3D rendering, which illustrates the resulting image (white) where the colored slices correspond to the slice-selective excitations. b) Illustration of a SAX angulated inner volume excitation combined with REST slabs. Green boxes depict slice-selective excitation and orange boxes depict the slice-selective SE refocusing, resulting in the imaging volume marked by the hatches. REST slabs are placed in PE direction as in Figure 3.4.

local off-resonances mainly affect the SE EPI acquisition in two ways; namely background phase contributions and geometric distortions.

For static off-resonance fields SE sequences refocus the resulting phase dispersion at echo time, thereby not affecting the signal. If the off-resonance fields are, however, dynamically changing over the sequence duration, refocusing is not achieved, hence resulting in residual background phase. This can be caused by e.g. blood inflow in the coronary arteries or motion of the myocardium-liver or heart-lung interfaces^{177,178}. If the phase contribution is very strong, thus resulting in considerable intra-voxel phase gradients, the magnitude can be modulated³².

As EPI has a finite duration, samples are acquired before and after the exact echo time, and as such even static off-resonance fields can result in geometric distortion. Depending on the sign of the local off-resonance and the orientation of phase encoding blips, these field variations result in localized stretching or compression of the image³². This most prominently happens close to the posterior vein. These distortions are possible to correct for during reconstruction using a B_0 map.

3.6 Eddy Currents and Concomitant Fields

Motion compensated diffusion encoding waveforms, in SE-cDTI induce considerable eddy currents in conductive parts of the MR scanner. This results in temporally and spatially varying field alterations, on the one hand contributing to the image phase and on the other hand adversely affecting the actual k-space vectors at sampling events. In EPI readouts this can result in ghosting artifacts, apparent translation

of objects in the image or even signal loss³¹. If available, a scanner-specific eddy current response function can be used to correct for image distortions by adapting the k-space locations during reconstruction.

Concomitant fields are a direct consequence of Maxwell's equations, resulting in a magnetic field contribution according to equation 4.14. The resulting phase contribution due to diffusion gradients is zero at the echo time for SE-cDTI and symmetric waveforms. Therefore, constructing diffusion encoding waveforms often includes this symmetry around the refocusing pulse. One advantage of asymmetric waveforms^{179,180}, is the echo time reduction with unchanged EPI readout.

Multiple frameworks have been proposed to optimize diffusion encoding waveforms with system and sequence timings constraints, while minimizing the effects of eddy currents¹⁸¹ and concomitant fields^{172,182}.

4. Simulation and In-silico Ground Truth

Simulations require a definition of a computational domain, that captures the physiology of the heart as well as a signal model for MRI. The computational domain, also referred to as in-silico phantom, contains the physical properties of the tissue, that are relevant for the simulation.

4.1 Lagrangian vs Eulerian Description

The representation of phantoms can be based on the Eulerian or the Lagrangian framework. In the Eulerian framework, the phantom geometry is discretized, and the resulting mesh is used to compute spatial derivatives of fields. Using these derivatives, the governing differential equations of the domain can be numerically solved for given boundary conditions. The evolution of the fields is observed at prescribed nodal locations in the mesh. Examples for this approach are Computational Fluid Dynamics (CFD) simulations solving the Navier-Stokes equations to obtain the blood flow in, for example, the aorta¹⁸³, simulating the diffusion-weighted MR-signal by solving the Bloch-Torrey equations (equation 2.7) on meshed virtual tissue models¹⁸⁴ or simulating the mechanical contraction of the LV of personalized shape models¹⁸⁵.




FIGURE 4.1: a) Illustration of a meshed (tetrahedral) left ventricle used for biophysical simulation in the Eulerian description. The cell edges are plotted as black lines. b) Particles representing a slice of the left ventricle moving in space (Lagrangian). The trajectories are drawn as lines and the particles are depicted as spheres c) Particles representing flowing blood in a meshed aorta. The particle trajectories are obtained by solving the kinematics of particles in the flow field inside the aorta (Euler-Lagrangian method).

In the Lagrangian framework, the domain is populated by particles which are associated with a set of properties. They can move in the domain and interact. The particle trajectories are obtained e.g. by iteratively solving the kinematics inside the domain fields. Applications of this framework are e.g. simulations of MR-signals of flowing blood by seeding particles in the velocity fields obtained from CFD simulations or modelling diffusion processes as random walkers within restricted compartments, where collisions with boundaries are checked based on a meshed domain¹⁸⁶.

4.2 MR Signal Simulation

This section is based on the appendix of the paper in chapter 4, referencing specific implementations in CMRsim (see chapter 5)

Although nuclear magnetic resonance is a quantum mechanical phenomenon, the emergent macroscopic behaviour of spin $\frac{1}{2}$ particles (such as hydrogen) in the classical limit of quantum statistics can be described by the Bloch-equations:

$$\frac{d\mathbf{M}}{dt} = \gamma\mathbf{M} \times \mathbf{B} + \frac{M_0 - M_z}{T_1}\mathbf{e}_z - \frac{M_x\mathbf{e}_x + M_y\mathbf{e}_y}{T_2}. \quad (4.1)$$

These differential equations describe the evolution of the macroscopic magnetization vector \mathbf{M} in a prescribed volume V in a magnetic field \mathbf{B} . The change of magnetization is characterized by the longitudinal relaxation time T_1 and transversal relaxation time T_2 , γ denotes the gyromagnetic ratio, and M_0 is the equilibrium magnetization. This formulation assumes mass-less and non-interacting particles, which makes the Lagrangian description of the digital phantom preferable. On the contrary, the Eulerian description is capable of incorporating interactions captured by differential equations, such as the Bloch-Torrey equation (equation 2.7).

As the signal picked up by the receive coils in MR measurements is the voltage induced by the precession of all magnetization vectors, it is necessary to integrate $\mathbf{M}(t)$ over the entire digital phantom. In the discrete case, this is achieved by summing all magnetization vectors, weighted by the volume each of them represents. As for all discrete representations of a continuous reality, the discretization must be fine enough to avoid simulation artifacts. In the Lagrangian particle-based description for MR simulations, this requires using sufficiently many particles per image-pixel as input, to avoid e.g. spurious echoes or residual magnetization after spoiling. Furthermore, to prevent inverse crimes when using the simulation as ground truth in reconstruction algorithms, the phantom must be defined on a multiple times finer scale than the image, which is especially relevant when using voxelized or regularly gridded phantoms¹⁸⁷.

Since the gradients and RF pulses alter the magnetic field \mathbf{B} , finding a solution to equation 4.1 is specific to the actual sequence definition. Theoretically, it is

possible to find analytical solutions to the Bloch equations for sequences, however, these solutions often assume specific simplifications to subdivide the sequence into manageable sections. Nevertheless, analytical signal equations are a useful tool in many cases, as highlighted in section 5.2.4. The most general approach to find a solution is to numerically integrate the evolution of magnetization. While adaptive step-width methods like Runge-Kutta accelerate the numerical simulation for static phantoms, their application with complex motions is not straight forward. The following paragraph describes the application of simple forward integration with fixed step size. An overview of available simulation frameworks is provided in section 5.1.

The evolution of magnetization is described by the Bloch equations, with the magnetization vector $\mathbf{M}_n(t)$ of a single particle n represented as

$$\begin{aligned} \mathbf{M}_n(t) &= [M_{+n}(t), M_{-n}(t), M_{zn}(t)]^T, \\ \text{with } M_+ &= M_x + iM_y, M_- = M_x - iM_y, \end{aligned} \quad (4.2)$$

where M_x , M_y and M_z are the spatial components of the magnetization. Solving the Bloch equations in the rotating frame of reference can be achieved by discretizing in time using a temporal step size of δt and iteratively updating the magnetization corresponding to asymmetric operator splitting¹⁸⁸

$$\mathbf{M}_n(t + \delta t) = T_n(\delta t) \circ \Phi_n(t, \delta t, \mathbf{r}(t)) \circ R(t, \delta t, \mathbf{r}(t)) \circ \mathbf{M}_n(t). \quad (4.3)$$

Here, $R(t, \delta t, \mathbf{r}(t))$ is the rotation operator corresponding to RF application:

$$R(t, \delta t) = \begin{pmatrix} \cos^2 \frac{\alpha(t, \delta t)}{2} & e^{i2\theta} \sin^2 \frac{\alpha(t, \delta t)}{2} & -ie^{i\theta(t, \delta t)} \sin \alpha(t, \delta t) \\ e^{i2\theta(t, \delta t)} \sin^2 \frac{\alpha(t, \delta t)}{2} & \cos^2 \frac{\alpha(t, \delta t)}{2} & ie^{-i\theta(t, \delta t)} \sin \alpha(t, \delta t) \\ -\frac{i}{2} e^{-i\theta(t, \delta t)} \sin \alpha(t, \delta t) & \frac{i}{2} e^{i\theta(t, \delta t)} \sin \alpha(t, \delta t) & \cos \alpha(t, \delta t) \end{pmatrix}. \quad (4.4)$$

Using a trapezoidal discrete integration step, the flip angle α and the phase θ of the applied B_1 (complex) RF field is given as:

$$\alpha(t, \delta t) = \gamma |B_1(t + \delta t) + B_1(t)| \frac{\delta t}{2} \quad (4.5)$$

$$\theta(t, \delta t) = \angle \left(\frac{B_1(t + \delta t) + B_1(t)}{2} \right). \quad (4.6)$$

The subsequent precession of the n -th particle is defined as

$$\Phi_n(t, \delta t, \mathbf{r}_n(t)) \circ \mathbf{M}_n = \text{diag} \left(e^{i\phi_n(t, \delta t, \mathbf{r}_n)}, e^{-i\phi_n(t, \delta t, \mathbf{r}_n)}, 1 \right) \cdot \mathbf{M}_n, \quad (4.7)$$

where the precession phase $\phi(t, \delta t, \mathbf{r}_n)$ depends on the phase-generating effects captured by dedicated sub-modules φ_s . These sub-modules are e.g. off-resonance or concomitant fields, and depend on the dynamic properties of the particle or of the fields that the particle is moving through, thus requiring an update of these

quantities $\Gamma(\mathbf{r}_n(t), t)$. The precession phase for a single integration step is thus given by

$$\begin{aligned} \phi(t, \delta t, \mathbf{r}_n) = & \gamma (\mathbf{G}(t + \delta t) \cdot \mathbf{r}_n(t + \delta t) + \mathbf{G}(t) \cdot \mathbf{r}_n(t)) \frac{\delta t}{2} \\ & \Sigma_s \varphi_s (\mathbf{G}(t), \mathbf{r}_n(t), \Gamma(\mathbf{r}_n(t), t)), \end{aligned} \quad (4.8)$$

where $\mathbf{G}(t)$ is the gradient waveform defined by the MR-sequence. Lastly, the relaxation operator $T_n(\delta t)$ is defined as

$$T_n(\delta t) \circ \mathbf{M}_n = \text{diag} \left(e^{-\delta t/T_{2n}}, e^{-\delta t/T_{2n}}, e^{-\delta t/T_{1n}} \right) \cdot \mathbf{M}_n + \left[0, 0, (1 - e^{-\delta t/T_{1n}}) \right]^T, \quad (4.9)$$

where T_{1n}, T_{2n} are relaxation times assigned to a single particle n . The acquisition of signal $s_m(t)$ at sampling times defined in the MR-sequence, weighted by spatially varying coil sensitivities $C_m(\mathbf{r}(t))$ can be evaluated as

$$s_m(t) = \sum_{n=1}^N M_{+n} C_m(\mathbf{r}_n(t)) e^{-i\vartheta_{ACQ}(t)}, \quad (4.10)$$

where ϑ_{ACQ} denotes the receiver phase.

Off-resonance induced phase accumulation can be incorporated into the simulation by using the corresponding sub-module computing the phase contribution for a single particle n as $\varphi_{\Delta B_0}(\omega_n(t), \delta t) = \gamma \omega_n(t) / 2\pi \delta t$. Examples for frequency difference with respect to the rotating frame $\omega_n(t)$ are the particles representing myocardial tissue close to a cardiac vein or a blood particle flowing through an off-resonance field defined by the surrounding tissue. A chemical shift translates to a constant off-resonance frequency in this description.

Modelling T_2^* in the Lagrangian framework, can be achieved by assigning each particle a random phase rate $\phi_{T_2^*n}$, which is sampled from a zero mean Lorentzian distribution

$$\phi_{T_2^*n} \sim P(x, T_2') = \frac{T_2'}{\pi (1 + (x/\gamma)^2)} \text{ with } T_2' = \frac{1}{\frac{1}{T_2} - \frac{1}{T_2}}, \quad (4.11)$$

assuming sufficient particles within each voxel. Relaxation times T_2 and T_2^* correspond to a tissue type as defined in the digital phantom. To avoid extreme values of $\phi_{T_2^*n}$ in smaller particle populations, the sampling process of $P(x, T_2')$ involves a cut-off in the intermediate uniform sample $u \sim \mathcal{U}(0.01, 0.99)$, which is used to calculate the phase rate as

$$\phi_{T_2^*n} = \frac{1}{T_2'} \tan \left(\pi \left(u - \frac{1}{2} \right) \right). \quad (4.12)$$

This phase rate is added as phase contribution for each simulation step and particle within the corresponding sub-module

$$\varphi_{T_2^*}(\phi_{T_2^* n}, \delta t) = \phi_{T_2^* n} \delta t. \quad (4.13)$$

Phase contributions due to concomitant fields B_c for any given particle n are calculated based on the applied gradients $\mathbf{G}(t)$ and particle positions $\mathbf{r}_n(t)$ as

$$\begin{aligned} \varphi_{B_c}(\mathbf{G}(t), \mathbf{r}_n(t), \delta t) = & \frac{\gamma}{2B_0} \left[\left(G_x(t)z_n(t) + \frac{1}{2}G_z(t)x_n(t) \right)^2 \right. \\ & \left. + \left(G_y(t)z_n(t) + \frac{1}{2}G_z(t)y_n(t) \right)^2 \right] \delta t \end{aligned} \quad (4.14)$$

4.3 Digital Cardiac Phantoms

Simulating the signals obtained from encoding motion in CMR at the macroscopic (see 2.2 and 2.3) and microscopic scales (see 2.4) and their interaction requires cardiac phantoms that incorporate macroscopic deformation as well as information about tissue microstructure. Therefore, the following sections review the available approaches to represent the digital phantom in MR simulations at both scales, as well as one approach to combine them.

4.3.1 Macroscopic Scale

The approaches to represent cardiac digital phantoms can be categorized into: voxel-based, analytical, hybrid and shape models¹⁸⁹. Voxel-based phantoms are obtained from real data by labeling the tissues of interest from images^{190–192}. The less realistic, analytical phantoms are based on mathematical descriptions, of organ shapes, tissue structure and contraction, where variation of anatomical features or population statistics can be applied by modifying the model parameters. In hybrid models the anatomy and its temporal configuration is represented e.g. by 4-dimensional B-splines¹⁹³, allowing some variation with respect to the reference phantom. The coupling of the phantom with a biophysical model of the heart allows to include pathological changes^{194,195}. More recently, cardiac shape models have been proposed, yielding an expressive description of the dominant anatomic features^{185,196–199}. When coupled to a biophysical model even healthy and pathological population variability can be represented¹⁸⁹. The nodes of these cardiac phantoms and their trajectories can subsequently serve as magnetization bearing particles (Lagrangian description) and used in Bloch simulations.

4.3.2 Microscopic Scale

Simulating diffusion-weighted MR signals for a single voxel is used to investigate the influence of microstructural alteration on the signal. One simulation strategy is to consider randomly moving particles, restricted within a synthetic microstructure. The particles carry magnetization, which evolves according to the Bloch equations²⁰⁰. Another approach solves the Bloch-Torrey equation on meshed microstructure¹⁸⁴.

Synthetic microstructure can be obtained by building virtual tissue models as a composition of simple geometric objects mimicking myocytes, such as cylinders or cuboids with varying size and orientation^{201–206}. A more complex and parameterized strategy to create in-silico tissue models was proposed²⁰⁷. Here bundling into sheetlets, with their respective orientation as well as a packing of sheetlets in the myocardial wall, is included into the algorithm of composing cylinders. The tissue parameters are chosen according to the statistics and disorder class of healthy and diseased tissue. While the most realistic and detailed tissue phantom are obtained from three-dimensional segmentation of histology²⁰⁸, these models cannot be easily generalized to healthy and pathological statistics.

4.3.3 Model-Based Cardiac Diffusion

Performing detailed diffusion Monte-Carlo simulations in realistic tissue involving scales of μm , embedded in the anatomical structure and its motion on the macroscopic scale (mm to cm), quickly becomes computationally infeasible. However, given the practically feasible duration and diffusion weighting of the encoding waveforms in cDTI, it is justifiable to separate large and small scale motions in the simulation. Following the approach of distributions of diffusion tensors within a voxel as signal model¹⁷⁴, the combination of macroscopic motion and diffusion contrast can be achieved by assigning diffusion tensors to the nodes of a mesh. The diffusion tensor orientations can be constructed from statistical description of myocyte aggregates as in²⁰⁹. The scaling of eigenvalues depends on the applied sequence, which can be obtained from Monte-Carlo simulations in tissue models. In chapter 5 such a phantom is used to simulate cDTI images. A 3D generalization for constructing random diffusion tensors in the LV based on marginal distributions of MD and FA and spatial correlation, as in chapter 7, is publicly available at <https://gitlab.ethz.ch/ibt-cmr/modeling/cmr-random-diffmaps>.

4.4 Sequence Definition

This section is based on a conference contribution [210]

Defining MR sequences without vendor-specific nomenclature plays a central role in simulation, as well as in the efforts to promote open and reproducible science. In recent years, multiple software packages have been proposed aiming to provide

a vendor-agnostic framework for sequence definitions^{211–214}. Accompanying the works presented in this thesis, the CMRseq Python package was implemented (available at https://gitlab.ethz.ch/ibt-cmr/mri_simulation/cmrseq), building on concepts from the increasingly popular Pulseseq framework²¹¹. One core functionality of sequence definition frameworks is the representation and composition of MR-sequence atoms, such as RF pulses, gradient waveforms, and sampling or acquisition events. Furthermore, to obtain realistic sequence timings, MR-systems must be incorporated. The hierarchical composition of sequence atoms into recurring blocks, up to the level of full MR-measurements, is part of CMRseq. The package structure supports building a comprehensive library of MR sequences in a maintainable community effort.

5. CMRsim - A Python Package for Cardiovascular MR Simulations Incorporating Complex Motion and Flow

Jonathan Weine¹, Charles McGrath¹, Pietro Dirix¹, Stefano Buoso¹, Sebastian Kozerke¹

¹Institute for Biomedical Engineering, University and ETH Zurich, Zurich, Switzerland

Jonathan Weine and Charles McGrath contributed equally to this work.

Published as: Weine J, McGrath C, Dirix P, Buoso S, Kozerke S. CMRsim - A python package for cardiovascular MR simulations incorporating complex motion and flow. *Magn Reson Med.* 2024 ; 1-17. doi: 10.1002/mrm.30010

5.1 Introduction

Simulations are an important tool to evaluate the performance of MRI acquisition, reconstruction, and postprocessing algorithms⁵⁹. Simulation frameworks may target specific applications with assumptions on physiology, anatomy, and MR physics^{31,215–220}, or may take a more general approach, seeking to accommodate a range of possible applications^{60–69}. Additionally, the widespread deployment of machine learning algorithms makes simulations increasingly relevant, such as to optimize sampling strategies and/or sequences^{48,49,221}, to facilitate access to simulated ground-truth information for postprocessing algorithms, and to train and test models on synthetic data^{183,222}.

Requirements such as simulation time and availability of computational resources play a major role in determining what level of simulation detail can be considered. The most general approach, which has previously been implemented in multiple available software projects^{60–69}, is based on the numerical solution of the Bloch equations for a selected sequence (defining gradient and RF waveforms as well as sampling events). However, when generating large synthetic data sets, it can be preferable to evaluate signal equations obtained from analytic rather than numeric solutions to the Bloch equations, trading off generality for lower computational cost. Although analytic simulation paradigms have been implemented for specific signal models^{218,223–226}, they are based on strong physical assumptions, which usually limit extensibility to complex motion patterns.

In cardiovascular MR (CMR), motion plays a central role, not only as a challenge to be addressed in acquisition and reconstruction, but also as a biomarker for health and disease^{8,227,228}. For example, CMR is uniquely suited to assess time-resolved blood flow in three dimensions. Accordingly, the development of dedicated simulations for phase-contrast (PC) CMR has been of significant interest. These works can be approximately divided into two categories: (1) Euler-Lagrangian frameworks, in which the Bloch equations are solved for magnetization moving according to velocity-vector fields derived from computational fluid dynamics

(CFD)^{186,215,220,229-234} ; and (2) Eulerian approaches, in which the Bloch equations are solved in a fixed reference frame and motion is incorporated through transport equations^{216,235-239}.

Discretizing the digital anatomic structure into noninteracting particles is a common strategy in MR simulations, which can also be deployed for incorporating complex organ motion such as cardiac contraction. Although there are many openly available simulation frameworks⁶⁰⁻⁶⁹, to the best of our knowledge, none have focused specifically on the complex motion requirements of CMR. While some frameworks allow for some form of motion to be incorporated, they do not permit efficient embedding of complex cardiovascular motion, such as nonrigid deformation and flow, due to limited analytic descriptions of motion⁶³, computationally infeasible preloading of trajectories²³² or limited compatibility to specific phantoms²⁴⁰.

To this end, the objective of the work at hand is to describe an open-source MR simulation framework, referred to as CMRsim, which facilitates the incorporation of complex motion and flow on arbitrary digital phantoms for studying CMR acquisition and reconstruction. Furthermore, special emphasis is put on modular architecture with extensive documentation.

5.2 Methods

This section provides an overview of the package architecture and a formal definition of the simulation, followed by an introductory simulation incorporating cardiac and breathing motion, describing the basic workflow and functionality of CMRsim. Next are two advanced exemplary simulation experiments highlighting CMRsim's capability and extensibility. A discussion on performance and scaling of the simulation as well as a conceptual comparison to external simulations is provided Appendices B and C. A complete overview of CMRsim functionality, including additional examples, is provided as Jupyter notebooks. All source code, notebooks, and API documentation are publicly available at <https://cmr.ethz.ch/research/software.html>.

5.2.1 Package Overview

CMRsim requires as inputs a target object (dynamic digital phantom) and a specific MR sequence. Two simulation paradigms are supported: (1) numerical Bloch Simulations and (2) analytic signal models (hereby referred to as Bloch simulation and analytic simulation, respectively). The package is subdivided into the following components, as schematically shown in Figure 5.1. Italicization indicates class or objects implementations.

- **Trajectory representation:** Collection of *Trajectory Modules* creating a representation of motion and time-varying physical properties defined by a dynamic digital phantom.
- **Bloch Simulation:** Numerically solving the Bloch equations for arbitrary waveforms (RF, gradients) and motion. Additional submodules can be specified to incorporate effects such as off-resonance, T_2^* and concomitant fields. Implemented as *Bloch Simulation* module.
- **Analytic Simulation:** Evaluating signal models obtained from analytical solutions to the Bloch equations, composed of a *Contrast Model* and an *Encoding Module*.
- **Dataset handling:** Handling of dynamic digital phantoms and implementation of utilities such as mesh refinement, generation of in- and outflow boundary conditions and computation of motion fields on regular grids, implemented as *Dataset* classes.

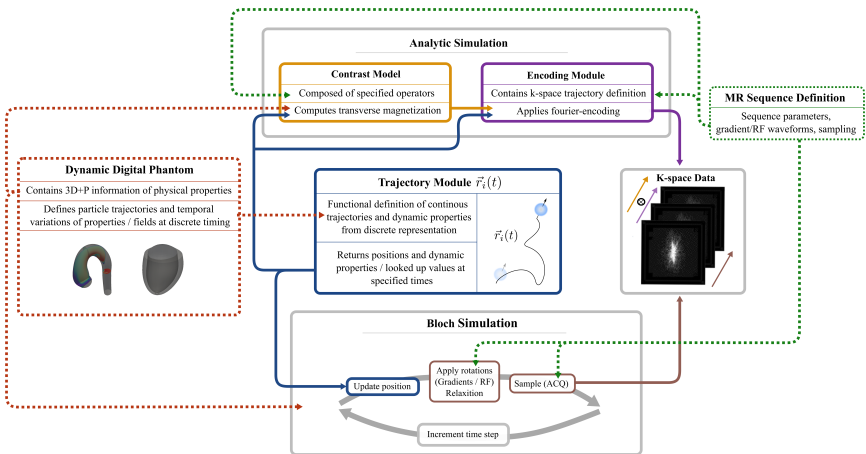


FIGURE 5.1: Overview of the CMRsim framework. Dashed-line blocks indicate required external inputs for the definition of the dynamic digital phantom and the MR sequence. Solid-line blocks define the constitutive components of CMRsim. The dynamic digital phantom (red box) contains P fields, corresponding to properties such as velocity or off-resonance, in addition to the 3D mesh (a meshed aorta containing velocity fields and a left-ventricle mesh model are shown as example). The core part of the package is the definition of the *Trajectory Modules* (blue box), providing a continuous representation of particle positions and dynamic particle properties. *Trajectory Modules* are used for analytic and numeric *Bloch Simulations*, both yielding k-space signals as output. Analytic simulations evaluate signal models based on analytical solutions of the Bloch equations, where the signal model is composed of a *Contrast Model* and an *Encoding Module*. The required parameters for the *Contrast Model* and the *Encoding Module* as well as the sequence definition (green box) for *Bloch Simulations* must be provided by the user. ACQ, acquisition.

The description of dynamic digital phantoms in CMRsim is based on the Lagrangian framework with noninteracting "massless" particles. Each particle n with $\iota_n = \{\mathbf{r}_n(t), \Gamma(\mathbf{r}_n(t), t)\}$ is defined by a time-dependent trajectory $\mathbf{r}_n(t)$ and a set of properties $\Gamma(\mathbf{r}_n(t), t) = \{M_0(\mathbf{r}_n(t), t), T_1(\mathbf{r}_n(t), t), T_2(\mathbf{r}_n(t), t), \dots\}$ including available magnetization $M_0\delta V\rho$, where δV is the volume represented by the particle and ρ is its proton density, and relaxation times T_1 and T_2 . Although the phantom resources necessary to run the exemplary notebooks are provided in the online repository, the CMRsim package does not include any digital phantoms. Therefore, phantoms containing all required material properties must be obtained externally. To avoid restricting the use of any external phantoms, a generic representation as vtk files or NumPy arrays ensures compatibility with CMRsim. The simplest representation usable in CMRsim is a dictionary of arrays, each containing one property of all particles. Common phantoms defined on 3D uniform grids (e.g., Shepp-Logan) can be easily transformed into the required format.

5.2.2 Trajectory Representation

Particle motion is described using *Trajectory Modules*, providing an interface to query positions $\mathbf{r}(t)$ at any given time point t . This abstraction decouples motion implementation from MR signal simulation and therefore ensures modularity and extensibility. As such, it is possible to extend motion implementation (e.g., fit a Taylor expansion for each particle, using a neural network for motion state retrieval, or even iteratively solving the kinematics in a turbulent flow field) without changing the simulation logic. All information required for instantiating a specific *Trajectory Module* must be provided by the dynamic digital phantom. A complete list of available modules is provided in the API reference.

In addition to particle position information, the *Trajectory Module* also returns a dictionary containing user defined field-lookups or dynamic particle properties $\Gamma(\mathbf{r}(t), t)$ (e.g., off-resonance, coil sensitivity, hereby referred to as additional fields) at given positions. All valid Trajectory Modules must implement at least two query methods using the signatures defined as follows:

1. The call function

$$f(\mathbf{r}_0, t, \dots) \rightarrow \{\mathbf{r}(t), \Gamma(\mathbf{r}(t), t)\}, \quad (5.1)$$

with $\mathbf{r}_0 \in \mathbb{R}^{N \times 3}$, $t \in \mathbb{R}^T$ and $\mathbf{r}(t) \in \mathbb{R}^{N \times T \times 3}$.

This function provides positions $\mathbf{r}(t)$ as well as properties $\Gamma(\mathbf{r}(t), t)$ for all given times contained in vector t , for a current batch of N particles. In addition to time points, the function takes the initial position of the particles as input arguments and can be extended with keyword arguments.

2. Iteratively advancing particle positions for an interval of δt

$$f(\mathbf{r}(t), \delta t) \rightarrow \{\mathbf{r}(t + \delta t), \Gamma(\mathbf{r}(t + \delta t), t + \delta t)\} \text{ with } \mathbf{r}(t) \in \mathbb{R}^{N \times 3}. \quad (5.2)$$

This method takes the current positions $\mathbf{r}(t)$ as well as the temporal step size δt and returns the updated position $\mathbf{r}(t + \delta t)$ and additional fields $\Gamma(\mathbf{r}(t + \delta t), t + \delta t)$. The method is used inside the *Bloch Simulation* loop in every integration step.

Examples of *Trajectory Modules* are *PODTrajectory*, *TaylorTrajectory*, and *TurbulentTrajectory* which are used in the demonstrations of this work as described further down.

5.2.3 Bloch Simulation

The evolution of magnetization is described by the Bloch equations, with the magnetization vector $\mathbf{M}_n(t)$ of a single particle n represented as

$$\mathbf{M}_n(t) = [M_+^n(t), M_-^n(t), M_z^n(t)]^T, \quad (5.3)$$

with $M_+ = M_x + iM_y, M_- = M_x - iM_y$

where M_x, M_y and M_z are the spatial components of the magnetization. Solving the Bloch equations in the rotating frame of reference is achieved by discretizing in time and iteratively updating the magnetization using a temporal step size of δt according to

$$\mathbf{M}_n(t + \delta t) = T_n(\delta t) \circ \Phi_n(t, \delta t, \mathbf{r}(t)) \circ R(t, \delta t, \mathbf{r}(t)) \circ \mathbf{M}_n(t), \quad (5.4)$$

where $R(t, \delta t, \mathbf{r}(t))$ is the rotation operator corresponding to RF application, $\Phi_n(t, \delta t, \mathbf{r}(t))$ is the rotation operator corresponding to precession and $T_n(\delta t)$ is the relaxation operator. During every iterative step, the particle positions are updated using the *Trajectory Module*. Additional fields $\Gamma(\mathbf{r}(t), t)$, which are required for precession-related effects, are looked up in the same call. Detailed descriptions of all operations involved in solving the Bloch equations, including the application of concomitant fields, T_2^* and off-resonances are provided in 4.2.

The MR sequence definition including RF and gradient waveform input in CMRsim (cf. Figure 5.1) is vendor-agnostic as the Bloch Simulations are instantiated using waveforms in the form of NumPy arrays. Thereby, gradient and complex-valued RF waveforms as well as sampling (ACQ) events with their corresponding receiver phases are provided on the (not necessarily uniform) simulation time grid.

5.2.4 Analytic Simulation

Evaluating signal models derived from analytic solutions to the Bloch equations is commonly used to avoid the computational cost of numeric Bloch simulations^{31,218,223-225}. Although signal models are typically only valid for a specific sequence and set of assumptions, certain terms are transferrable among models. To enable reuse of such terms for simulation experiments, CMRsim defines operators ψ_i that can be assembled into a *contrast model* Ψ , which maps each particle t_n to transverse magnetization $\tilde{m}_{xy,n}$

$$\tilde{m}_{xy,n} = \underbrace{(\psi_L \circ \dots \circ \psi_0)}_{\Psi}(t_n) \quad (5.5)$$

This composition corresponds to a pipe-and-filter approach, in which the piped data are a set of arrays containing the properties of each particle and all L return the apparent magnetization capturing the MR contrast defined by the operators. To facilitate parallel computation, such as when sweeping flip angles in a gradient-echo sequence, the array containing the complex magnetization $\tilde{m}_{xy,n}$ has the shape (N, R, K) . N is the (batched) number of particles, R s the number of scan repetitions (corresponding to different image contrasts, e.g. defined by a set of flip angles) computed in parallel, and K is the number of k-space vectors used for Fourier encoding. Operators can expand axis R to efficiently handle the evaluation of multiple parameters in each of the operators (e.g., using multiple coils for each scan repetition). The resulting shape of the complex apparent magnetization before and after the application of an operator ψ_i is (N, R, K) and $(N, e_i * R, K)$, where e_i is the operator's expansion factor (e.g., number of coils or flip angles). Every operator defines which properties must be provided in the input, and all properties are automatically available to all operators through keyword arguments.

The apparent magnetization returned by the model is passed to the *Encoding module* which calculates a discretized Fourier integral to obtain the k-space signal $s(t_s)$ at sampling times t_s , as follows

$$s(t_s) = \sum_{n=0}^{N-1} \tilde{m}_{xy,n}(t_s) e^{i\mathbf{k}(t_s) \cdot \mathbf{r}_n(t_s)} + \eta. \quad (5.6)$$

The shape of $s(t_s)$ is $(I I_i e_i, R, N_\eta, K)$, where N_η is the number of independent noise instantiations η with specified standard deviation. The volume δV_n that particle n represents, which is necessary to evaluate the discrete Fourier integration, is contained in the apparent magnetization $\tilde{m}_{xy,n}$ via the particle property $M_{0,n} = \delta V_n \rho_n$ referred to as available magnetization, as described above. The k-space vectors and sampling times used for evaluating eq. 5.6 can be explicitly specified on module instantiation.

5.2.5 Dataset Handling

Dataset classes bundle recurring functionality to transform and handle dynamic digital phantoms, such that corresponding *Trajectory Modules* can be instantiated. One example illustrating the use for flow simulations is the *RefillingFlowDataset*. This class requires a meshed domain containing at least a mean velocity vector field. Since *Trajectory Modules* for flow require a projection of the meshed domain onto a regular grid, this is implemented as a method in *RefillingFlowDataset*. Furthermore, the class implements density estimation for a given set of particles within the mesh, as well as a random uniform seeding and pruning of particles that have left the meshed domain (outflow). To address inflow, a region can be specified where the dataset seeds new particles based on the density estimation to achieve uniform filling. Another example used in this work is the *CardiacMeshDataset*, implementing a contracting left ventricle obtained from biomechanical simulations¹⁸⁵. A complete list of available classes can be found in the API documentation.

5.2.6 Implementation Notes

CMRsim is compatible with Python versions later than 3.6. Pre-built docker images that contain all dependencies to run simulations are provided to decrease the burden of installation. As speed is often a focus in MR simulations, core functions are implemented using the Python API of TensorFlow2²⁴¹ (version > 2.11) enabling GPU acceleration, while also maintaining compatibility for non-GPU workflows. Furthermore, handling and visualizing 3D phantoms can be conveniently achieved using PyVista²⁴². The project including a package and container registry is hosted at https://gitlab.ethz.ch/ibt-cmr/mri_simulation/cmrsim, where the link to the full API documentation as well as a collection of examples and tutorials can be found. Additionally, an entry on the Python packaging index <https://pypi.org/project/cmrsim/> is available, containing all necessary links.

5.2.7 Introductory Simulation Example

To demonstrate the necessary steps to run a CMRsim simulation, an introductory example is provided in the following, including pseudo code in Figure 5.2. The example pertains to free-breathing cine balanced steady-state free precession (bSSFP) imaging incorporating both cardiac and respiratory motion. All details including sequence parameters can be found in the corresponding Jupyter notebook.

A dynamic digital phantom of a contracting left ventricle was obtained from biomechanical simulations¹⁸⁵. Time-resolved deformation fields were saved for every mesh node every 4 ms. Contractile motion was assumed to be periodic and uniform and tissue specific T_1 and T_2 values were assigned to all nodes. Mesh nodes were then resampled to a higher density to alleviate discretization artifacts¹⁸⁷.

Preparation / Load resources

```
READ dynamic digital phantom snapshots
INIT CardiacMeshDataset FROM snapshots
CALL CardiacMeshDataset refinement to add particles
CALL CardiacMeshDataset cell size computation
GET trajectories, M0, and static properties per particle for all
    particles FROM CardiacMeshDataset
READ breathing curve
```

Initialize Modules

```
INIT PODTrajectory FROM particle trajectories
INIT PODTrajectory FROM particle M0
INIT SimpleBreathingMotion FROM breathing curve and PODTrajectory
INIT Bloch Simulation FROM RF, gradients and sampling events arrays
```

Run simulation

```
INIT magnetization and positions for all particles
FOR batch of particles in all particles
    FOR repetition in number of TRs
        UPDATE M0 of particles using the M0 PODTrajectory
        CALL Bloch Simulation with TR index RETURN magnetization,
        position
        WRITE magnetization and position for next TR
    END
END
```

Retrieve results

```
FOR TR in repetitions
    GET k-space samples from accumulator of Bloch Simulation
    WRITE samples to result array
END
```

FIGURE 5.2: Pseudocode for the introductory example. The program is divided into sections: first loading and preparing the dynamic digital phantom, secondly initializing the required CMRsim objects, thirdly calling the CMRsim objects to perform the simulation and finally saving the k-space data.

To achieve consistent proton density over time, the available equilibrium magnetization M_0 per particle (defining a weight for spatial integration) was computed from mesh-cell volumes and evaluated for each repetition time (TR) of the sequence. Breathing motion was assumed to be translational along the foot-head direction and implemented using the *SimpleBreathingMotion* module, which uses time-dependent breathing amplitudes and direction as input. Cardiac motion was parameterized using Proper Orthogonal Decomposition (POD), implemented in the *PODTrajectory* module, allowing to arbitrarily define the time resolution for sampling the deformation fields^{243,244}. The *Bloch Simulation* module was instantiated containing the RF and gradient waveforms as well as ACQ events for all sequence TR s. Although the temporal reference of these arrays is not required to be uniform, all given arrays must be defined using the same temporal grid. The *Bloch Simulation* module was called once for each TR , using the particle definitions obtained from the *CardiacMesh-Dataset* module as input. For each TR , the *Bloch Simulation* module accumulates the MR signal at timings according to the defined temporal grid. The complete k-space data can be retrieved by reading the samples from all signal accumulators contained in the *Bloch Simulation* module. Supporting information S1 contains a Jupyter notebook showing the corresponding Python code.

5.2.8 Advanced Simulation Examples

Turbulent Flow Imaging Downstream of a Stenosis

This example illustrates a spoiled gradient echo (GRE)^{113,116} simulation of turbulent flow downstream of a stenotic section. Figure 5.3 shows a flowchart of the simulation experiment. The NumPy arrays defining gradient and RF waveforms, as well as sampling events were generated using the *CMRseq*²¹⁰ package. In addition to the unweighted reference, six velocity encodings (VENCs) were used: $(550, 50, 250, 50, 100, 50)\text{cm s}^{-1}$ in directions $[(0, 0, 1), (0, 0, 1), (0, 1, 0), (0, 1, 0), (1, 0, 0), (1, 0, 0)]$. The mean velocity within the meshed domain was approximately 80 cm s^{-1} . The sequence parameters were $TR = 10\text{ ms}$, $TE = 5\text{ ms}$ and a flip angle of 15° . A schematic illustration is included in Figure 5.3c. Imaging resolution was set to $2\text{ mm} \times 2\text{ mm}$ with a FOV $22.2\text{ cm} \times 14.2\text{ cm}$. The temporal grid size was set to $10\text{ }\mu\text{s}$. The target density for initial seeding as well as for reseeded was 5 mm^{-3} resulting in a total of 1.5 million particles. Reconstruction was performed by inverse Fourier transformation.

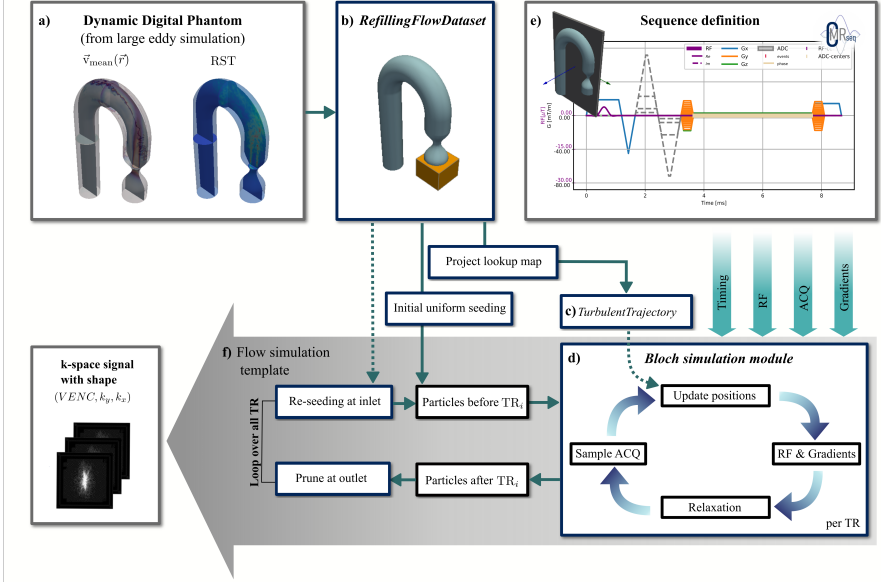


FIGURE 5.3: Simulation of turbulent flow imaging downstream of a stenosis. a) A *Dataset* class is instantiated using the mesh, mean velocity, and Reynolds stress tensors (RST) obtained from a prior computational fluid dynamics (CFD) simulation. b) The orange box shows the inflow region where density-based reseeded is performed. c) The instantiation of the *Trajectory Module* requires the CFD mesh to be projected onto a uniform grid serving as a lookup map for the velocity and stress tensors. d) When calling the *Bloch Simulation* module, the particle positions are updated by the *Trajectory Module*. e) The MR sequence definition is used to get an array representation of timing, RF pulses, gradients, and acquisition (ACQ) events per sequence TR . f) Looping over TR s and interleaving the reseeded is captured by a *flow simulation* template. Calling the simulation yields k -space data for each velocity encoding (VENC). The functionality implemented by CMRsim is marked by dark blue boxes.

The dynamic digital phantom, consisting of the meshed domain containing mean velocity and Reynolds stress tensor fields (c.f. Figure 3a), was generated using a Large Eddy Simulation (LES) on a previously published geometry¹⁸³. The mesh was used to create a *RefillingFlowDataset* instance. The inflow region was placed upstream of the stenosis as shown in Figure 5.3b. Reseeding and pruning of outflow were performed after every TR .

The positional update for all particles is performed by a *TurbulentTrajectory* instance. For initial position $\mathbf{r}_n(t) = [x_n(t), y_n(t), z_n(t)]^T$, the updated position at time $t + \delta t$ is calculated as

$$\mathbf{r}_n(t + \delta t) = \mathbf{r}_n(t) + \delta t \mathbf{v}_n(t), \quad (5.7)$$

where the velocity vector $\mathbf{v}_n(t)$ is obtained by a lookup of the velocity field at location $\mathbf{r}_n(t)$ containing the mean velocity $\mathbf{U}(\mathbf{r}_n)$, Lund transformation²⁴⁵ based on the Cholesky decomposition of the Reynolds stress tensor $a(\mathbf{r}_n) \in \mathbb{R}^{3 \times 3}$ and Lan-

grangian integral time scale $\tau(\mathbf{r}_n)$. For each particle, a turbulent velocity component $\mathbf{u}_n(t)$ is computed by solving the Markov chain integration of the modified classical Langevin equation²⁴⁶

$$\mathbf{u}_n(t + \delta t) = \mathbf{u}_n(t)e^{-\delta t/\tau(\mathbf{r}_n(t))} + \zeta(\mathbf{r}_n(t))a(\mathbf{r}_n(t))\sqrt{1 - e^{-2\delta t/\tau(\mathbf{r}_n(t))}}, \quad (5.8)$$

where $\zeta(\mathbf{r}_n(t)) = \mathcal{N}^3(0, 1)$. Accordingly, the total velocity vector update per step follows as

$$\mathbf{v}_n(t) = \underbrace{\mathbf{U}(\mathbf{r}_n(t))}_{\text{mean}} + \underbrace{\mathbf{u}_n(t)}_{\text{turbulent}}. \quad (5.9)$$

To evaluate the functionality of the *TurbulentTrajectory* module, trajectories were tracked in a separate experiment only involving motion simulation. This included re-seeding assuming 50 TRs each 5 ms long and a target density of 1 mm^{-3} . To evaluate density conservation, the final particle density was calculated. A detailed description of all steps as well as additional illustrations are available in the corresponding Jupyter notebook.

Cardiac Diffusion Tensor Imaging

The experiment described in this section aims at studying the effect of cardiac motion-induced phase on cardiac diffusion tensor estimation when using acceleration-compensated diffusion encoding with spin-echoes (SE-Mo12)²⁹. The SE-Mo12 single-shot EPI sequence including binomial excitation pulses was defined using CMRseq²¹⁰. The temporal grid for RF and gradient waveforms was set to $100 \mu\text{s}$. The EPI readout was defined using a 91×49 matrix and an in-plane resolution of $2 \text{ mm} \times 2 \text{ mm}$. Diffusion encoding included 3 low (100 s/mm^2) and 9 high (450 s/mm^2) b-values²⁴⁷. All timings and waveforms required in the analytic signal model were obtained from this sequence definition as shown in Figure 5.4. The model used for computing the coil- and diffusion weighted apparent transverse magnetization is given by

$$\tilde{m}_{xy,n}(t_s) = \left[\psi_C \circ \psi_{Diff} \circ \psi_{T_2^*} \circ \psi_{\Delta B_0} \circ \psi_\phi \circ \psi_{SE} \circ \iota_n \right], \quad (5.10)$$

with

$$\psi_{SE} = M_{0,n} e^{-\frac{TE}{T_{2,n}}} \left(1 - e^{-\frac{TR}{T_{1,n}}} \right), \quad (5.11)$$

$$\psi_\phi e^{i\gamma \int_{t_a}^{t_b} \mathbf{G}(t) \cdot \mathbf{r}_n(t) dt}, \quad (5.12)$$

$$\psi_{\Delta B_0} = e^{-i\gamma \Delta B_{0,n} |t_s - TE|}, \quad (5.13)$$

$$\psi_{T_2^*} = e^{-\frac{|t_s - TE|}{T_{2,n}^*}}, \quad (5.14)$$

$$\psi_{Diff} = e^{-b\mathbf{g}^T\mathbf{D}_n\mathbf{g}}, \quad (5.15)$$

$$\psi_C = C(\mathbf{r}_n), \quad (5.16)$$

$$(5.17)$$

where $T_{2,n}^*$, $T_{2,n}$, $T_{1,n}$, $\Delta B_{0,n}$, \mathbf{D}_n , and $M_{0,n}$ refer to relaxation times, off-resonance, diffusion tensor and available magnetization. Parameters are coil sensitivities $C(\mathbf{r}_n)$ at particle location \mathbf{r}_n , sampling times t_s , linear diffusion weightings (b, \mathbf{g}), gradient waveforms $\mathbf{G}(t)$ defined from t_a to t_b , echo time TE and repetition time TR . k-space vectors for sampling times are denoted $\mathbf{k}(t_s)$, complex Gaussian noise is denoted as η .

The MRXCAT 2.0 workflow¹⁸⁹ was used to generate a dynamic digital phantom, yielding a co-registered, texturized background and a mesh model of a contracting left ventricle (LV) with 4801 mesh nodes. The off-resonance frequency for each particle for one static configuration was calculated as described in¹⁷⁶ starting from typical susceptibilities as implemented in the *RegularGridDataset* module. The off-resonance per particle was assumed to be approximately constant over the duration of the EPI readout.

To reduce the effect of discretization artifacts, the *CardiacMeshDataset* implements functionality to refine the mesh as well as to render and inspect the motion trajectories of all mesh nodes. The trajectory of the refined mesh nodes for all snapshots between 150 ms and 350 ms is captured by the *PODTrajectory* module. To evaluate the diffusion weighting operator, spatially coherent random diffusion tensors were assigned to all LV mesh nodes according to the sampling procedure described in²²². A slab of 30 mm thickness was extracted from the LV as well as the background phantom surrounding the target field of view with a margin of 20 mm. Breathing motion was incorporated by assuming a periodic global translation of the slab and the LV using the *SimpleBreathingMotion* module. Figure 5.4 provides an overview of the process.

To obtain an estimate of the diffusion tensors without the necessity of registering the data, a second simulation run without breathing motion was conducted. All reconstructions were performed using the NUFFT implementation of the BART Toolbox²⁴⁸. A detailed description of all for all steps as well as additional illustrations are available in the corresponding notebook.

5.3 Results

5.3.1 Introductory Simulation Example

The range of motion of the contracting left ventricle over the duration of image acquisition, including the constant position of the excitation slice is illustrated in

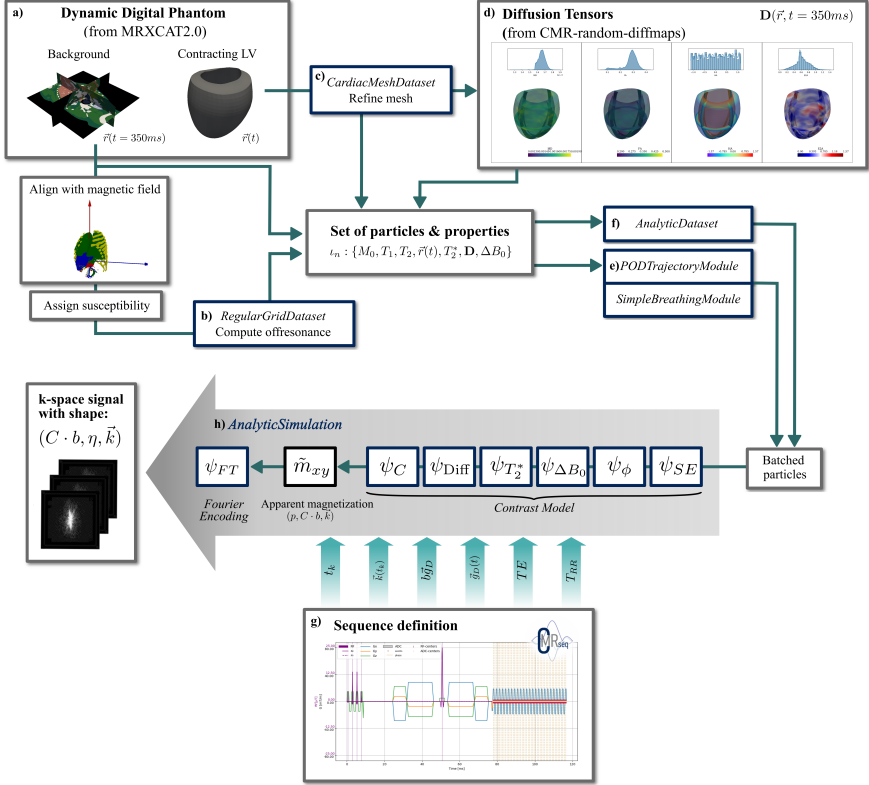


FIGURE 5.4: Simulation of cardiac diffusion tensor imaging. a) The dynamic digital phantom consists of a contracting left-ventricular (LV) mesh and background derived from MRXCAT2.o. b) Off-resonances for all particle locations are computed using the *RegularGridDataset* module. c,d) The LV mesh is refined using the *CardiacMeshDataset* module c) and used as input to generate random diffusion tensors per mesh node using the *cmr-random-diffmaps* package d). e,f) The *Trajectory Module* e) and *AnalyticDataset* f) (responsible for streaming of data to GPU) are instantiated from all particles and their properties contained in the refined dynamic digital phantom. The properties $M_0, T_1, T_2, T_2^*, \mathbf{r}(t), \mathbf{D}, \Delta B_0$ refer to available magnetization, relaxation times, the discretized particle trajectory, diffusion tensor, and off-resonance. g, h) From the sequence parameters and waveforms g), the *AnalyticSimulation* module h) including the *Contrast Model* and *Encoding module* ψ_{FT} is constructed. Parameters derived from the sequence definition are sampling times t_k and k-space vectors $\mathbf{k}(t_k)$, linear diffusion weighting with b-value b and direction \mathbf{g} , discretized diffusion gradient waveform $\mathbf{g}_D(t)$, echo time TE , and RR-interval time T_{RR} . The Contrast Model consists of the operators ψ_C (coil sensitivity weighting), ψ_{Diff} (model-based diffusion weighting), $\psi_{T_2^*}$ (T_2^* -weighting), $\psi_{\Delta B_0}$ (off-resonance weighting during readout), ψ_ϕ (phase accumulation during diffusion encoding) and ψ_{SE} (spin-echo contrast). The apparent magnetization \tilde{m}_{xy} computed by the *Contrast Model* has the shape (#particles, #coils · #diffusion-weightings, #k-space samples). Calling the *AnalyticSimulation* instance yields k-space data of shape (#coils · #diffusion-weightings, #noise-instantiations, #k-space samples).

Figure 5.5a). Figure 5.5b) shows the image resulting from the simulation as described in the methods section. For comparison the simulation was repeated using only contractile motion (Figure 5.5c) and no motion (Figure 5.5d). The number of dummy shots was set to 81, which results in steady state magnetization within the imaging slice for the static phantom. The ghosts appearing in Figures 5.5b) and c) are caused by non-steady state magnetization entering the excitation slice due to through-plane motion. Motion blurring corresponding to breathing and cardiac motion can be seen in Figures 5.5b) and c) but not in d) as expected.




FIGURE 5.5: a) Illustration of the range of motion including breathing and contraction over 750 ms (duration of the simulated sequence). The red arrow indicates the translation of the center of mass from the initial position (light blue mesh) to the final position (light green mesh). The constant slice for excitation is indicated by the black opaque box. The images correspond to simulation runs with b) combined breathing and contractile motion, c) contractile motion only and d) static phantom.

The total simulation time was 128s when motion of the approximately 385.000 particles was included, while for the static case the simulation took 62s on a NVIDIA TITAN RTX. For all three cases the sequence duration was about 750 ms using a maximal temporal step size of 10 μ s.

5.3.2 Turbulent Flow Imaging Downstream of a Stenosis

Figure 5.6a) and 5.6b) contain the density estimation at the end of particle tracking as two-dimensional histograms in y - z and x - y plane. The histograms show homogeneity in the outflow and inflow region. Figure 5.6c) shows the trajectories for a subset of particles. Velocity fluctuations in the post-stenotic region as well as the laminar flow at the inlet of the U-bend can be seen.

The images resulting from the simulation are shown in Figure 7. The third and fourth rows show the phase difference from which the mean-velocity can be evaluated as well as Turbulent Kinetic Energy (TKE). As expected, multiple phase wraps are present when the VENC is too low, specifically along the flow direction (phase difference in Figure 5.7c). With a larger VENC phase wrapping is not present and a jet can be seen (phase difference in Figure 5.7b). Additionally, as the Cartesian readout is parallel to the jet, the zero-VENC case shows a jet structure in the




FIGURE 5.6: a) 2D histogram of particles at their final location within a central 10mm thick slice in the y-z plane after repeated re seeding and trajectory integration according to eq. (17) b) 2D histogram in the x-y-plane of the same particle configuration as in a); c) shows the integrated trajectories for a set of particles initially placed in a line right before the stenosis.

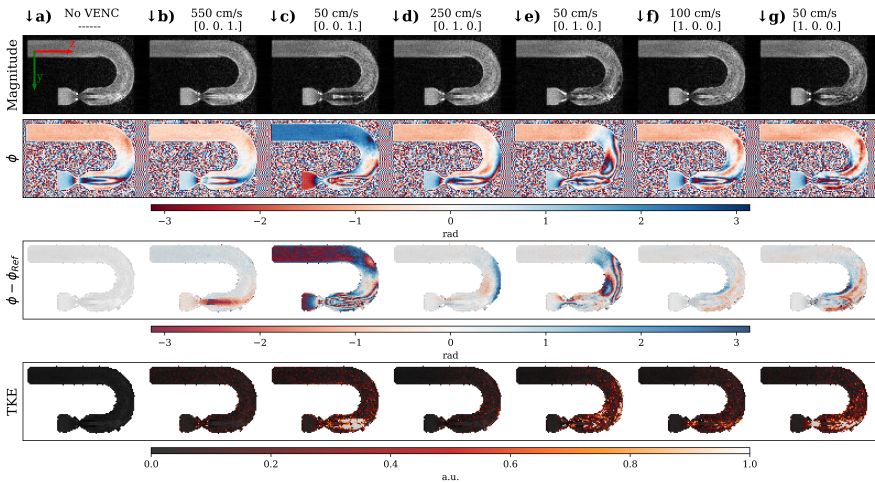


FIGURE 5.7: Simulation for the phase contrast imaging of turbulent flow downstream of a stenosis in a UBend. Columns a) to g) correspond to the applied velocity encoding. Rows one and two show the simulated magnitude, and phase images. Row three shows the phase difference relative to the non-encoded reference image a). Bottom row shows the estimated turbulent kinetic energy. Color bars correspond to all images contained in the row above it. TKE, turbulent kinetic energy; VENC, velocity encoding.

phase image (phase in Figure 5.7a). At larger VENC, TKE estimates show minimal structure, however at lower VENC turbulence around the jet is seen (Figures 5.7c,e,g). All magnitude images show some degree of signal loss due to intra-voxel dephasing, as expected.

The total duration of the simulation for all seven VENCs was 29 minutes on a single NVIDIA TITAN RTX (24 Gb) including re-seeding. The number of particles was approximately 1.5 million. Based on the sequence definition (71 k-space lines with

a TR of 10 ms) approximately 71.000 temporal integration steps were performed. Re-seeding per TR took approximately 20% of the simulation time.

5.3.3 Cardiac Diffusion Tensor Imaging

Figure 5.8 shows the magnitude and phase of a subset of simulated images. Rows correspond to magnitude and phase, with each column corresponding to one diffusion weighting. The overlay of the LV mask obtained from the motion state of the unweighted image emphasizes the displacements due to breathing motion. Furthermore, magnitude decay due to diffusion weighing can be seen. Phase images show both phase variation due to the coil sensitivities as well as residual motion sensitization of the diffusion waveform in the left ventricle.




FIGURE 5.8: Magnitude (upper row) and phase images (lower row) obtained from analytic simulation. Each column corresponds to a single diffusion weighting. Residual motion sensitivity can be seen by the diffusion-weighting dependent phase inside the LV. The translational displacement between TRs is due to breathing motion.

The estimation of tensor metrics shown in Figure 5.9 was performed on images simulated without breathing motion to mitigate the necessity of registration. The upper row of maps / blue curves correspond to a simulation without motion-induced phase, while orange curves and the second row of maps show results with motion phase. As expected, the phase-gradient introduced by strain within the LV results in increased mean diffusivity estimation (MD). The distributions of and fractional anisotropy (FA) correspond to the randomly sampled diffusion tensors (compare Figure 5.4).

Evaluation of the signal model took 6 s, using approximately 200 000 particles. As the motion-induced phase due to the diffusion gradient waveforms was only evaluated for the LV, the simulation time per image was approximately 10 s including the computation of trajectories on a $100 \mu\text{s}$ grid for approximately 50 000 particles on a single NVIDIA TITAN RTX (24 Gb).

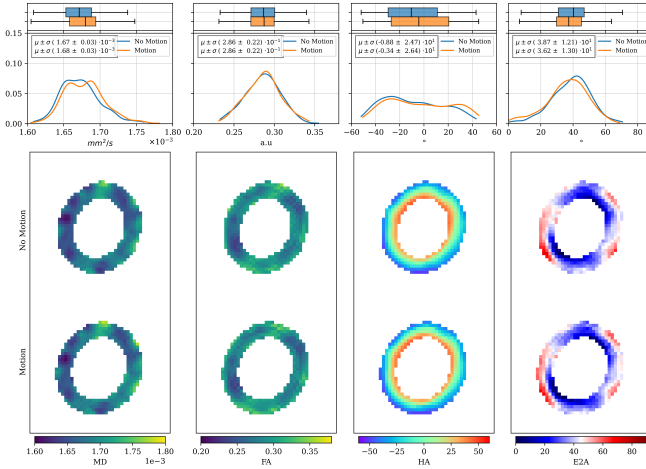


FIGURE 5.9: Diffusion tensor estimation on the magnitude data obtained from simulation without breathing motion. Columns from left to right show mean diffusivity (MD), fractional anisotropy (FA), helix angle (HA) and sheetlet angle (E2A). The histograms in the top row show the distribution with the masked maps in the bottom row. The first row of images corresponds to the blue lines and markers in the histogram plots, therefore including no motion induced phase in the tensor estimation. Lower row images and orange plots correspond to simulation incorporating motion phase respectively.

5.4 Discussion and Conclusion

In this work we have presented an open-source Python CMR simulator facilitating the incorporation of complex organ motion and flow. Our numerical experiments show that CMRsim results in expected motion effects such as motion-induced phase accumulation and magnitude modulation due to intra-voxel phase gradients and incoherent motion.

The most significant and relevant feature of CMRsim is the decoupling of motion implementation from the MR signal simulation by introducing *Trajectory Modules*. This allows the incorporation of virtually any Lagrangian motion description and reduces the requirements for adding custom modules. The choice of not relying on a rasterized phantom comes at the cost of computing magnetization density. This limitation also holds for calculating fields that usually depend on the phantom state such as off-resonance or B₁ maps. The definition of CMRsim *Trajectory Modules* easily allows the incorporation of these properties as additional fields, if they are computed prior to MR simulation and therefore contained in the dynamic digital phantom. Accordingly, in contrast to frameworks that include pre-defined digital phantoms or require a regular spatial grid, some aspects are shifted to the choice of the correct modeling approach for each dynamic digital phantom. While this is

not fundamentally included in CMRsim, the package does include functionality to handle specific phantoms as presented in this work.

The limitations of the modeling approach can be discussed based on the presented diffusion simulation example. Here a separation of scales between diffusional and bulk motion is assumed. Furthermore, the diffusion encoding is assumed to only capture the Gaussian component, which holds true only for certain b-value ranges^{27,57}. If detailed microstructure is of interest, Monte-Carlo simulations in realistic microstructures^{206,249,250} are a suitable method, which in principle could also be implemented within a *Trajectory Module*.

Furthermore, for flow in deformable meshes incorporating fluid–structure interactions, feasibility depends on whether the projected flow fields (as currently required by CMRsim) are sufficient for particle tracking methods. Although technically this is possible to implement in CMRsim (similar to the readily available pulsatile flow simulation), it is untested and requires future development. Another aspect that is easily implemented in CMRsim but needs practical consideration is the interleaving of particle interactions or projections of the particle distribution to a mesh or grid, such as used for the density estimation in refilling. If the required interval for these interleaved calculations is about 100 times the temporal step of the Bloch solver, the Lagrangian description may become computationally inefficient, and combined MR–computational fluid dynamics simulations can be preferable.

A key objective of CMRsim was to develop a package in Python focusing on maintainability, usability and extensibility. Firstly, Python is widespread in science, is seen as comparably beginner friendly, has a large community with many scientific packages and its GPL compatible license supports open-source availability. Secondly, due to the absence of a build process and the dynamic nature of Python, code changes are easier than in a compiled language. Finally, using TensorFlow for the computationally heavy tasks reduces the maintenance component of checking compatibility for system-specific drivers. Leveraging the provided containers and python packages provided by TensorFlow ensures long-term maintenance and, as such, transferability to modern systems is anticipated to be more easily feasible than reported for i.e. JEMRIS and MRILab⁶¹.

To achieve good extensibility, CMRsim abstracts recurring computations into modules with minimal coupling. Therefore, changing functionality inside single operators does not require changes in other code locations. This allows, for example, to implement new contrast models without cloning the CMRsim repository. Aside from code quality, providing development environments is crucial to reduce the threshold to use and extend the functionality of a package. To this end, several docker containers are provided in the project container registry. To make the package usable, documentation including examples is crucial. Furthermore, catching breaking changes increases robustness to feature extensions. To this end, automatic

unit-test execution, and package documentation as well as docker image builds using GitLab's CI/CD functionality are set up for CMRsim.

While other frameworks require sequence definition files either internally defined or obtained from other packages, CMRsim uses NumPy arrays. CMRsim deliberately does not contain functionality to define MR sequences as this functionality is not inherent to the simulation logic and therefore should be implemented separately, for example using (Py)Pulseq^{211,212}. Sequence definitions for the experiments described in this work were based on CMRseq²¹⁰.

In conclusion, CMRsim is a simulation framework that allows the incorporation of complex motion to systematically study advanced CMR acquisition and reconstruction approaches. The open-source package features modularity and transparency facilitating maintainability and extensibility in support of reproducible research.

5.5 Appendix A

This section was moved to 4.2.

5.6 Appendix B - Comparison to External Implementations

This section is meant to provide a perspective on how incorporating motion into MR simulations potentially can be achieved without using CMRsim. JEMRIS⁶⁰ is chosen as reference, as it is the only available framework allowing to specify particle trajectories in Bloch simulations. In JEMRIS, these trajectories are read from a file; hence, the user is required to calculate and correctly write particle trajectories using the same temporal grid used in the MR sequence definition before MR simulation.

For flow simulations, this specifically requires the user to perform particle tracking, reseeding, and trimming themselves, resulting in particle trajectories that need to be written/read for each TR in the sequence. Furthermore, as often more than 4 million particles are used for flow simulation, these files can easily exceed 10 GB, hence adding significant read/write overhead.

Finally, a rough comparison can be made for simulation times based on the results of²³². They presented an example of a one-dimensional PC simulation in JEMRIS, with a simulation time of 330 min when using 600 CPU cores on the ROMEO HPC center. In comparison, a simulation with matching particle count, TR (in ms), and number of TRs in CMRsim takes about 180 min on a Titan RTX. Furthermore, it is unclear whether the JEMRIS simulation time includes particle tracking and reseeding, which is included in the CMRsim simulation time.

5.7 Appendix C - Performance and Scaling

Memory use in CMRsim is driven primarily by the size of the digital phantom mesh and the number of particles, with memory scaling linearly with both. Both simulation paradigms are implemented such that the particles are processed in batches, assuming the trivial decomposition of noninteractive particles during signal evaluation. Therefore, memory requirements can generally be alleviated by choosing smaller batch sizes.

An upper estimate for a large mesh ($250^3 \times 25$ timesteps for pulsatile flow \times 4 fields) and 100 million particles requires about 12 GB of memory (6 GB for each); however, this level of detail is often unnecessary and typical simulations require less than 5 GB. With modern GPUs commonly exceeding 20 GB of memory, these simulations are easily attainable. Time complexity can be divided into reseeding and core simulation, with both scaling linearly with particle count, while the core simulation additionally scales linearly with number of timesteps.

Both analytic and Bloch simulations obey the same scaling laws; however, generally, analytic simulations can allow lower particle densities and reduced number of timesteps per TR, decreasing the simulation time.

Finally, simulation time can be decreased by reducing particle count or number of timesteps, although this may lead to discretization errors or particle density errors. Otherwise, methods of reducing real-world scan time work as well, such as reduced matrix size, increased gradient performance, or undersampling.

6. Parameter Estimation / Inverse Problems

As shown in the previous sections, diffusional and tissue motion are encoded into the complex-valued MR signal. Aside from spin motion, the signal includes statistic effects such as sample noise as well as deterministic effects, i.e. the point-spread-function (PSF), partial-volume effects, local field homogeneity and eddy currents. A forward model capturing all relevant physical processes can be represented as a mapping $\mathcal{F}(\rho|\Omega)$ with parameters Ω , which maps some ground-truth parameters ρ to a set of signals \mathbf{s} .

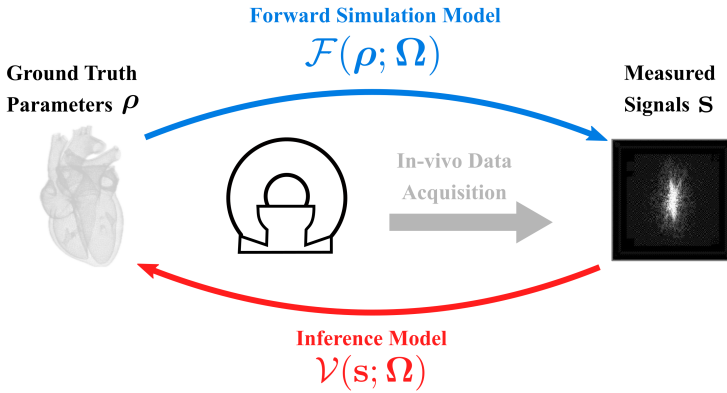


FIGURE 6.1: Illustration of the inverse problem of estimating tissue parameters ρ from a set of measured signal \mathbf{s} , assumed to result from the forward model $\mathcal{F}(\rho|\Omega)$ containing the model parameters Ω , using an inference model $\mathcal{V}(\mathbf{s}; \Omega)$.

Since the acquisition is partially stochastic, this mapping corresponds to the conditional probability $p(\mathbf{s}|\rho)$. For in-vivo measurements, the prior distribution of true tissue parameters $p(\rho)$ is not known. Estimating the tissue parameters from given \mathbf{s} , hence the inversion of mapping $\mathcal{F}(\rho|\Omega)$ usually corresponds to finding a point estimate based on the posterior distribution. To this end, the following sections introduce three commonly used point estimates and their corresponding background.

6.1 Linear Models

This section is based on [251]

Linear models are used as a simple tool to describe processes in quantitative fields of study. Their advantage is the simplicity and favourable analytic properties, while

providing sufficiently accurate results in many cases. The description of a linear model yielding the observation y as linear combination of a set of L predictors x_l

$$y = \sum_{l=0}^L x_l \beta_l + \epsilon, \quad (6.1)$$

where the constant term $x_0 = 1$ is introduced to incorporate the intercept value β_0 , and ϵ is a random error, e.g. due to measurement noise. The model is linear in the coefficients β_i , such that non-linear transformations of the predictors (such as the exponential magnitude decay in diffusion) can be included without losing the property of linearity. The random error is typically assumed to be a zero-mean and uncorrelated between measurements, while having the same variance in each measurement. When processing magnitude data as in cDTI it should be noted, that the noise is Rician distributed with non-zero mean²⁵², which becomes especially relevant for low signal-to-noise ratio (SNR) data. For a set of n measurements contained in the vector \mathbf{y} , Equation 6.1 can be written in matrix form

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}, \quad (6.2)$$

where $\mathbf{X} \in \mathbb{R}^{n \times (L+1)}$ is the so-called design matrix, defined by the forward model and its parameters $\mathcal{F}(\boldsymbol{\rho}|\Omega)$.

Minimizing the residual sum of square (or least squares approach) under the assumption of zero-mean random error, is guaranteed to result the unbiased estimator $\hat{\boldsymbol{\beta}}$ with minimal variance

$$\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta}} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}. \quad (6.3)$$

This also holds in case of non-spherical errors, which means the variance varies for different predictors x_i and are allowed to be correlated. The generalized least squares solution in this case is given by the unbiased, minimum variance estimator

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}^T \mathbf{V}^{-1} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{V}^{-1} \mathbf{y}, \quad (6.4)$$

where the positive definite matrix \mathbf{V} is given by the covariance matrix of the error $\text{cov}(\boldsymbol{\epsilon}) = \sigma^2 \mathbf{V}$.

Assuming, the error terms to be a (multivariate) Gaussian random variable, the least squares solution (eq. 6.3/6.4) is equal to the maximum likelihood estimator. Referencing Figure 6.1, this means, the estimated tissue parameters are the ones maximizing the likelihood of obtaining the given signal given by $p(\mathbf{s}|\boldsymbol{\rho})$.

6.2 Bayesian Modelling

This section is based on [251]

In Bayesian methods, the full posterior distribution of the tissue parameter given the signal $p(\rho|\mathbf{s})$ is modelled by as composition of parameterized, closed-form distributions. Using Bayes' rule the posterior evaluates to

$$p(\rho|\mathbf{s}) = \frac{p(\mathbf{s}|\rho)p(\rho)}{\int p(\mathbf{s}|\rho)p(\rho)d\rho}, \quad (6.5)$$

where the denominator is independent of the parameters and can be calculated by sampling the learned priors. The learning process consists of choosing a statistical model for $p(\mathbf{s}|\rho)$ and updating the priors according to the observed data points. While Bayesian hierarchical modelling has been used to improve Intra-Voxel Incoherent Motion (IVIM) inference^{253–255}, explicitly modelling the distribution for tensor-valued inference quickly becomes too complex to be efficiently handled. A commonly used point estimate is the maximum posterior (MAP)

$$\hat{\rho}_{\text{MAP}} = \arg \max_{\rho} \{p(\rho)p(\mathbf{s}|\rho)\}. \quad (6.6)$$

When the full posterior is not available, the MAP can be obtained via regularized least square estimates, where the L_2 -regularization corresponds to a Gaussian prior on the parameters $p(\rho) \sim \mathcal{N}(\mu, \sigma)$ and the L_1 -regularization corresponds to a Laplacian prior.

6.3 Learning Based Methods

In recent years, the success of neural networks in a variety of fields and their basic principles have been discussed. For a detailed introduction into theoretical background, implementation aspects, parameter optimization and application see [256]. This section is restricted to stating a few implications, relevant for using neural networks as inference model in diffusion tensor imaging.

All neural networks are fundamentally a parameterized function estimator $\mathcal{V}_{\theta}(\mathbf{s}) \rightarrow \rho$ that inverts the forward model $\mathcal{F}(\rho|\Omega)$, where θ contains all parameters of the network. The process of minimizing a loss function, e.g. the squared difference between network prediction and ground truth with respect to the network parameters θ over a joint distribution $p(\mathbf{s}, \rho)$

$$\min_{\theta} \mathbb{E}_{p(\mathbf{s}, \rho)} |\rho - \mathcal{V}_{\theta}(\mathbf{s})|_2^2 \quad (6.7)$$

is called supervised training. To this end, a training data set \mathcal{T} containing data and the corresponding ground truth (referred to as labels) is required. The optimality condition for the functional \mathcal{V}_{θ} in eq. 6.7, obtained using the corresponding Euler-Lagrange equation shows that the point estimate from a neural network trained this way, yields the mean of the posterior. In case of a multi-modal posterior, e.g. due to mixed signal of healthy and diseased tissue, this point estimate is not particularly

good as the modes are averaged. One strategy to address this, is to introduce a variational approximation of the posterior distribution as output of the network and using the Kullback-Leibler divergence

$$\mathcal{D}_{KL}(p||q) = \sum_{X \in \mathcal{T}} p(X) \cdot \log \left(\frac{p(X)}{q(X)} \right), \quad (6.8)$$

as loss (measuring the similarity of two distributions p and q), such that the training becomes

$$\min_{\theta} \mathcal{D}_{KL}(p(\rho|\mathbf{s})||\mathcal{V}_{\theta}(\mathbf{s})), \quad (6.9)$$

where X are the examples in the training set \mathcal{T} . Another approach is the field of variational networks, where the prediction of the network becomes non-deterministic and by sampling outputs an estimate of the learned posterior distribution can be obtained²⁵⁷.

As real ground truth is intrinsically not available in CMR the construction of training datasets for tissue parameters requires special attention. Biases in the training dataset were shown to transfer to the predictive performance of networks⁵⁵. To address the problem of missing ground truth data, the approach of training networks on synthetic data obtained from simulation has been proposed^{58,258}. Using parameterized anatomic models and tailored priors of tissue parameters allows to generate training sets with control over statistical properties (see chapter 4). Furthermore, the out-of-sample behaviour, i.e. how the network predicts when being presented a sample, that lies outside of the ranges of the training data, can be accessed this way. This is especially relevant in the context of medical imaging to estimate the robustness of Deep Learning (DL) assisted classifiers⁵⁶.

A type of neural network architectures particularly successful in the field of image processing are Convolutional Neural Networks (CNN). In CNNs, convolutions with learnable filters with non-linear activation functions as well as down- and up-sampling operations are concatenated. The inductive bias of picking up local spatial information as intrinsic feature of convolutions makes these networks very efficient in many computer vision tasks. The concatenation of convolutions and down sampling operations increase the area from which information is processed in the subsequent step, thereby implicitly performing a multi-scale analysis of the image. To prevent the network from learning spurious correlations, e.g. based on image position, organ shape or similar, the architecture should be chosen according to the expected features in the data.

7. Synthetically Trained Convolutional Neural Networks for Improved Tensor Estimation from Free-Breathing Cardiac DTI

Jonathan Weine¹, Robbert J H van Gorkum¹, Christian T Stoeck^{1, 2}, Valery Vishnevskiy¹, Sebastian Kozerke¹

¹Institute for Biomedical Engineering, University and ETH Zurich, Zurich, Switzerland

²Division of Surgical Research, University Hospital Zurich, University of Zurich, Zurich, Switzerland

Published as: Weine J, van Gorkum RJH, Stoeck CT, Vishnevskiy V, Kozerke S. Synthetically trained convolutional neural networks for improved tensor estimation from free-breathing cardiac DTI. *Comput Med Imaging Graph.* 2022; 99:102075. doi:10.1016/J.COMPMEIMAG.2022.102075

7.1 Introduction

Cardiac Diffusion Tensor Imaging (cDTI) provides invaluable information about the state of the myocardial microstructure^{20,23,24,163}. Although it is a promising technique, challenges with respect to motion sensitivity remain to be addressed^{11,12}. Cardiac motion-induced signal loss during data acquisition in spin-echo sequences can be alleviated by using second-order motion-compensated diffusion gradient waveforms^{29,30,145,163,170}. Despite the use of respiratory navigators, patient-friendly free-breathing acquisition strategies often result in spatial misalignment of the data. Therefore, the use of non-rigid registration prior to parameter inference is required. However, this registration step is a challenging task because of the varying contrast, the intrinsically low signal-to-noise ratio (SNR), geometrical distortions due to local off-resonances in the heart and partial volume effects, which are especially relevant at the myocardial borders. The standard data processing approach includes the use of non-rigid intensity-based registration algorithms as e.g. the elastix toolbox⁹⁶ or parameterized total variation (pTV) registration⁹⁷. Motion registration can result in significant differences of estimated diffusion tensors when compared to a breath-hold acquisition⁹⁸. Part of this issue is associated with the ill-conditioning of the registration problem, which can cause larger variations in the subsequent diffusion tensor estimation. As these variations are especially prominent at myocardial borders, the evaluation of tensor metrics is commonly performed within a mask excluding the left-ventricular (LV) borders¹².

To reduce the variations in estimated diffusion metrics, regularized fitting algorithms can be applied. For example, a hierarchical Bayesian framework has been used to model the expected marginal distribution of the estimated parameters as well as their spatial correlations²⁵³⁻²⁵⁵. While these works directly infer diffusion tensor metrics, such as mean diffusivity (MD) and fractional anisotropy (FA), their application to obtain full diffusion tensors from motion-affected data is not trivial, as explicit modeling the spatial correlations of tensors and motion can quickly become too

complex to be handled efficiently. Furthermore, taking both healthy and lesion tissue into account requires Gaussian mixture models, which increase complexity and compromise stability of Bayesian inference. Another approach to incorporate the data's spatial structure for tensor inference is the application of convolutional neural networks (CNNs)⁵³. However, when training such networks on in-vivo data, parameter distributions of the training datasets can introduce potential biases into tensor estimation²⁵⁹. Furthermore, ground truth diffusion tensor information of in-vivo cDTI data is inherently not available. Synthetic data, on the other hand, can provide ground truth, facilitating quantitative analysis of inference accuracy and precision. Additionally, when exclusively using synthetically generated data to train neural networks, all distributions of parameters are controllable. This allows to test the algorithm's performance also on out-of-distribution samples and increases the interpretability of results.

In the work at hand we propose a parameterized pipeline to synthesize free-breathing cDTI data with physiologically plausible statistics and subsequently train a CNN to infer full diffusion tensors directly from unregistered in-silico and in-vivo cDTI data, exclusively using synthetic data for training. We hypothesize that integrating the explicit, intensity-based registration step into the synthetically CNN yields higher precision and accuracy of the estimated diffusion tensors. The quantitative analysis on synthetic test data shows reduced errors with the CNN compared to the reference method, which uses conventional intensity-based registration followed by linear least-square (LLSQ) tensor estimation. Furthermore, we show that training on synthetic data transfers well to in-vivo data. Moreover, when only subsets of data averages are used, the network still yields high quality predictions. By augmenting the in-vivo data with artificial lesions we demonstrate that our approach is able to discriminate healthy from diseased tissue.

7.2 Methods

7.2.1 In-vivo Acquisition Protocol

Data of 5 healthy volunteers (2 male) was acquired on a 1.5 T clinical MR system (Philips Healthcare, Best, Netherlands) upon written informed consent according to ethics and institutional guidelines. The gradient system of the scanner delivers 80 mT m^{-1} at $100 \text{ mT m}^{-1} \text{ s}^{-1}$ slew-rate. A 32-channel cardiac coil was used as receive array. Imaging was performed using a spin-echo sequence with single-shot echo-planar imaging (EPI) readout and acceleration-compensated diffusion gradient waveforms (Mo12)²⁹. The protocol consisted of navigator-gated (5 mm - window) free-breathing acquisitions of 12 diffusion directions with 10 averages each, including 3 directions with 100 s/mm^2 and 9 directions with 450 s/mm^2 ¹⁶³. Three slices were obtained (apical, mid-ventricular, basal). The data was reconstructed using ReconFrame (GyroTools LLC, Winterthur, Switzerland). For reference, all images

were also registered to the first diffusion-weighted image using the pTV-registration toolbox⁹⁷. The left ventricle was manually segmented based on the first average of the first diffusion direction.

7.2.2 Synthetic Data Generation

The data synthesis pipeline is illustrated in Figure 7.1. All parts of the pipeline were implemented in Python (3.8) and TensorFlow (2.3)²⁴¹. To obtain realistic binary LV masks, varying in shape and size, we used myocardial segmentations from systolic short-axis views provided in the ACDC dataset²⁶⁰. No other information than the binary masks were used from the ACDC data. The set of masks was then split into two subsets of 491 training and 69 test examples, while no subject case of the ACDC dataset was included in both datasets. The masks were resized, shifted and embedded into a reference short-axis slice of the XCAT phantom^{193,218} such that the center of mass and size aligned with the reference LV. The spatial dimensions of the phantom, the field of view and the physical properties (proton density, T₁, T₂) per tissue type were taken from the MRXCAT framework²¹⁸.

For each dataset, one LV mask was randomly chosen, and a helix angle (HA) map as well as an absolute sheetlet angle (E2A) map were generated. Given that more training datasets than unique LV mask were generated, the random choice of LV masks was performed with replacement. The helix angle linearly varied from -60° to 60° over the transmural direction, while the E2A map was constructed randomly to contain blob-like areas of high and low values with smooth transitions (see 7.5.1). A subset of LV mask pixels was randomly sampled as reference positions. For each reference position, an eigenvalue triple $\Lambda = (\lambda_1, \lambda_2, \lambda_3)$ was sampled to generate diffusion tensors according to $\mathbf{D} = \hat{\mathbf{U}}\text{diag}(\Lambda)\hat{\mathbf{U}}^T$, where the eigenbasis *hatU* was defined by HA and E2A. The rejection-acceptance sampling process yielding the eigenvalue triples, resulted in uniform distributions of MD and FA over specified ranges (see 7.5.2). The ranges were set to match the values reported in a patient study by Das et al.²⁰. The spatial structure of the tensor maps was introduced by interpolating the reference tensors in log-Euclidian space using a kernel of radial basis functions (RBF) $\kappa_\sigma(d)$ with kernel-width σ ²⁶¹:

$$\log \mathbf{D}(\mathbf{r}_n) = \frac{\sum_m \kappa_\sigma(d(\mathbf{r}_n, \mathbf{r}_m)) \log \mathbf{D}(\mathbf{r}_m)}{\sum_m \kappa_\sigma(d(\mathbf{r}_n, \mathbf{r}_m))}, \quad (7.1)$$

where \log denotes the matrix logarithm and $\mathbf{D}(\mathbf{r}_n)$ denotes a diffusion tensor assigned to location \mathbf{r}_n . The distance $d(\mathbf{r}_n, \mathbf{r}_m)$ was defined as the geodesic in the coordinate system of transmural depth and polar angle (see 7.5.3).

In half of the training datasets a lesion was included. Healthy and lesioned tissue was defined by using different sampling intervals of MD and FA for the eigenvalue triples based on data reported in patient studies²⁰. For each example two tensor-maps were

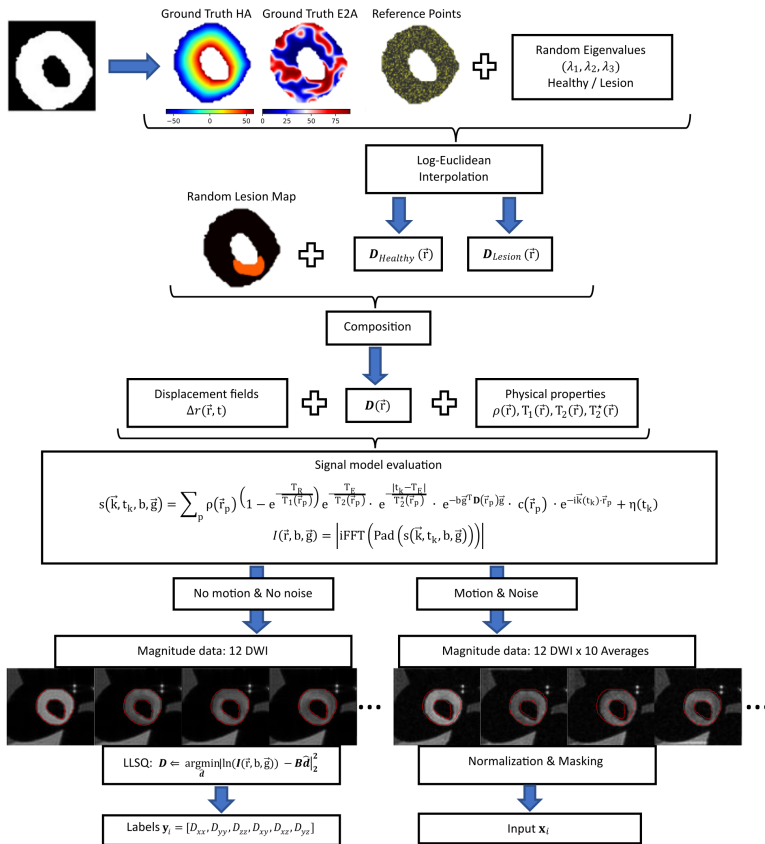


FIGURE 7.1: Flow chart of synthetic data generation. Starting from top left: based on a left-ventricular (LV) mask, helix angle (HA) and sheet angle (E2A) maps are generated. Combined with random eigenvalue triples, reference tensors are constructed at random locations within the mask, which are subsequently interpolated to all LV points. Random lesion maps are used to compose tensor maps of healthy and lesioned tissue. Displacement fields obtained from in-vivo data, the generated tensor map and physical property maps are used as input to the signal model. The signal model is evaluated once without motion and noise (left) and once with motion and noise (right). The former is used to generate the training labels, while the latter serves as input to the network.

generated, using healthy and lesioned tissue eigenvalue triples correspondingly. A randomly generated lesion mask (see Appendix A) was then used to compose them. The tensor maps for healthy and lesioned tissue were composed according to the mask defining the lesion location, while HA and E2A for both tensor maps were identical.

To introduce respiratory motion, deformation fields obtained from previous in-vivo data²⁵ were randomly selected and applied in simulations. To this end, the deformation fields were aligned and scaled to match location and size of the LV mask, and split across the training and test sets. For each synthetic data collection, a subset of 119 deformation fields were used and further perturbed to increase variability. As the first of the 120 images served as the reference motion configuration, no deformation field was applied in this case.

Image generation was achieved by evaluating the signal model using the digital phantom as input with its material point properties $\{\rho, T_1, T_2, T_2^*, \mathbf{D}\}$. The resulting k-space signal $s(\mathbf{k}, t_k, b, \mathbf{g})$ for a given diffusion weighting with b-value b and the unit orientation vector of the diffusion encoding gradient \mathbf{g} is calculated as:

$$s(\mathbf{k}, t_k, b, \mathbf{g}) = \sum_p c(\mathbf{r}_p) \rho(\mathbf{r}_p) \left(1 - e^{-T_R/T_1(\mathbf{r}_p)}\right) e^{-T_E/T_2(\mathbf{r}_p)} \dots \quad (7.2)$$

$$\dots e^{-|t_k - T_E|/T_2^*(\mathbf{r}_p)} e^{-b\mathbf{g}^T \mathbf{D}(\mathbf{r}_p) \mathbf{g}} e^{-i\mathbf{k}(t_k) \cdot \mathbf{r}_p} + \eta.$$

Index p runs over all material points in the digital phantom, which was defined on a 5-fold finer grid as compared to the target image resolution. Matrix c denotes coil sensitivity weighting and η denotes complex zero-mean Gaussian noise. The k-space vectors \mathbf{k} and the sampling times t_k were set according to the EPI trajectory used in the in-vivo measurement protocol²⁹. The noise standard deviation was varied to obtain a mean SNR between 15 and 25 inside the left ventricle. From k-space data $s(\mathbf{k}, t_k, b, \mathbf{g})$, sets of images \mathbf{x} for diffusion weightings $\mathbf{B} = (b_1 \mathbf{g}_1^T \mathbf{g}_1, b_2 \mathbf{g}_2^T \mathbf{g}_2, \dots)$ were reconstructed with standard inverse Fourier transform with zero padding and subsequent coil combination.

Since $\mathbf{D}(\mathbf{r}_p)$ maps of the digital phantom were generated on 5-fold higher resolution, they could not be used as training labels directly. Therefore, for each dataset, the simulation was repeated with all displacements set to zero and no noise added (referenced as ground truth hereafter). The diffusion tensor training labels at image resolution were calculated by solving the pixel-wise LLSQ problem $\arg \min_{\hat{\mathbf{a}}} \|\mathbf{I}(\mathbf{r}, b, \mathbf{g}) - \mathbf{B} \hat{\mathbf{a}}\|_2^2$ for the static data $\mathbf{I}(\mathbf{r}, b, \mathbf{g})$ ¹³³. The simulated image datasets were normalized by the LV mean intensity in the b_1 -image. Furthermore, the background was masked, by multiplying a morphologically dilated and smoothed version of the LV mask.

7.2.3 Network Training

As illustrated in Figure 7.2, an Inception-Resnet style²⁶² network was implemented. To monitor the training process and perform the hyperparameter search, a validation subset of size 100 was split off the training dataset of 5000 examples. Hyperparameters as well as learning rate were tuned by a grid search for the number of residual blocks (1, 2, 4, 6), layers per residual block combined with filters per layer [64, 128, (64, 64), (32, 64, 128), (64, 64, 64)], dropout rate [0.05, 0.1, 0.15, 0.2, 0.25] and number of filters in the initial convolutional layers. The metrics to determine the best hyperparameters were mean absolute error (MAE) as well as explained variance score (R2) for mean diffusivity (MD) and fractional anisotropy (FA).

To account for the different orders of magnitude of the tensor entries, the outputs of the last convolutional layer of the network \hat{d}_i were scaled with the standard deviation and shifted by the mean, which were calculated from the marginal distributions per tensor entry over the entire training dataset. The loss function used for training was a combination of the squared error of tensor values d_i , FA and MD:

$$L = \sum_i \left(d_i^{label} - d_i^{pred} \right)^2 + \alpha_1 \left(FA(\mathbf{d}^{label}) - FA(\mathbf{d}^{pred}) \right)^2 + \alpha_2 \left(MD(\mathbf{d}^{label}) - MD(\mathbf{d}^{pred}) \right)^2, \quad (7.3)$$

where the superscripts stand for label and prediction and α_1, α_2 are hyperparameters to weigh the iso-/anisotropic parts of the loss. As optimizer, Adam with a initial learning rate of 2.5×10^{-3} and exponential decay rates for first/second moments of $\beta_1 = 0.9, \beta_2 = 0.99$ was used. The network was trained for 100 epochs with a batch size of 3. As reference a simple multilayer CNN as well as a U-Net²⁶³ were trained in the same fashion, however, were discarded due to insufficient predictions.

7.2.4 Data Analysis

The network performance was evaluated on synthetic test data. Alongside the network prediction (Network), three different LLSQ estimations were calculated: (i) LLSQ inference on static i.e. motion-free data ($LLSQ_{Stat}$) serving as lower error bound, (ii) LLSQ inference on unregistered motion-corrupted data ($LLSQ_{Unreg}$) serving as baseline reference, and, (iii) LLSQ inference on registered motion-corrupted data ($LLSQ_{Reg}$). The noise level for all test examples was set to yield a mean SNR of 18 inside the LV in the b_1 -image, approximately matching the average SNR of the available in-vivo data. For $LLSQ_{Reg}$, images were registered to the b_1 -image, which matched the motion configuration of simulated ground truth. The errors for

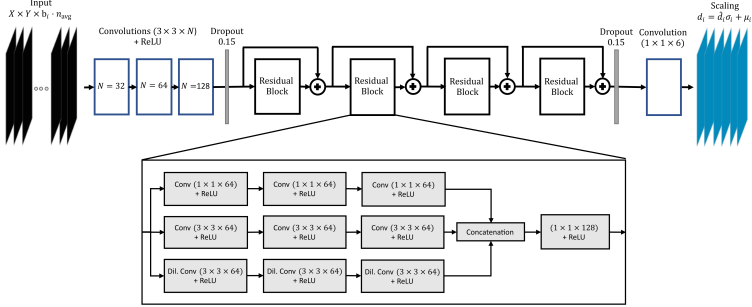


FIGURE 7.2: Network architecture. The lower part of the illustration provides detailed information on the residual blocks. Stated hyperparameters such as number of convolutional channels and dropout rates are identical to the configuration used to generate the results. ReLU refers to rectified linear unit while Dil. Conv denotes convolution with dilation rate 2. For convolutional layers, $(N \times M \times C)$ denotes the kernel size and number of channels. Padding for all convolutions was set to “same”. The last convolutional layer is not followed by an activation function and the output is scaled and shifted by standard deviation and mean of the training label entries.

all tensor estimates were calculated with respect to the ground truth obtained as described in 7.2.2.

The Network was compared to results obtained with $LLSQ_{Stat}$, $LLSQ_{Unreg}$ and $LLSQ_{Reg}$ based on MD, FA, eigenvalues, HA and E2A. Mean and standard deviation of the marginal distributions as well as MAE, root mean squared error (RMSE) and R2 were calculated. Furthermore, the potential discrimination of tissue types based on MD and FA was investigated. To this end, we calculated the Jensen-Shannon-distance between lesion and healthy distributions for ground truth as well as for Network versus $LLSQ_{Reg}$. The SNR dependency of the prediction quality was investigated by evaluating MAE and R2 of MD, FA, HA and E2A for SNRs between 10 and 50.

As no ground truth is available in-vivo, a qualitative comparison relative to the results obtained with in-silico data was performed. To this end, the spatial characteristics and marginal distributions of MD, FA, HA and E2A for the Network inference as well as $LLSQ_{Unreg}$ and $LLSQ_{Reg}$ were compared. Furthermore, the SNR dependency of the estimation quality was investigated by retrospectively reducing signal averages prior to tensor estimation. Finally, the Network’s sensitivity to detect potential lesions was studied by augmenting the in-vivo data with synthetic lesions. The Jensen-Shannon distance between distributions of healthy and lesioned tissue was calculated to assess whether their discriminability was preserved. Before passing the in-vivo data to the Network, the data was normalized by the mean LV signal of the b_1 -image and subsequently masked using a morphologically dilated

version of the LV mask. This pre-processing step corresponded to the procedure applied on the synthetic data.

7.3 Results

7.3.1 Synthetic Test Data

The marginal distributions of MD, FA and the eigenvalues are shown in Figure 7.3. The ground truth (blue bars) shows the metric prior implicitly modeled by data synthesis. Comparing mean, standard deviation and RMSE as stated in the plots, the Network inference is found to result in distributions closer to ground truth than $LLSQ_{Reg}$ for all metrics. The mean MD and FA of $LLSQ_{Reg}$ are significantly lower than ground truth.

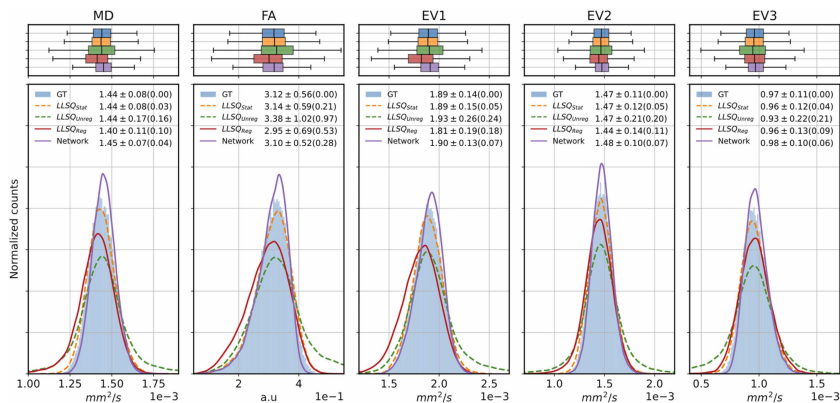


FIGURE 7.3: Marginal distributions of MD, FA and the three eigenvalues (EV1-3) for the synthetic test data. Blue bars show ground truth, dashed lines the optimal reference $LLSQ_{Stat}$ (orange) and least optimal reference $LLSQ_{Ureg}$ (green); solid lines correspond to $LLSQ_{Reg}$ (red) and Network inference (purple). The histogram bins per metric are identical for all algorithms. Numbers stated next to the legends are mean \pm standard-deviation of the distribution and (RMSE) with respect to ground truth. Top row shows boxplots of the distributions shown in the lower row with corresponding colors.

Figure 7.4 compares the marginal distributions of MD, FA and eigenvalues separated into lesion (orange) and healthy (blue) tissue. The normalization of counts was performed per class. The corresponding Jensen-Shannon distances between lesion and healthy distributions are stated in the subplots for MD and FA. The distances for MD are in better agreement with the ground truth for both $LLSQ_{Reg}$ and Network inference than for FA. The mean absolute errors for MD, FA, E2A, HA and all eigenvalues over the test set are reported in Tab. 1. The references $LLSQ_{Stat}$ and $LLSQ_{Reg}$ yielded the lowest and highest mean absolute errors correspondingly. Furthermore, Network inference resulted in lower errors than $LLSQ_{Reg}$ for all

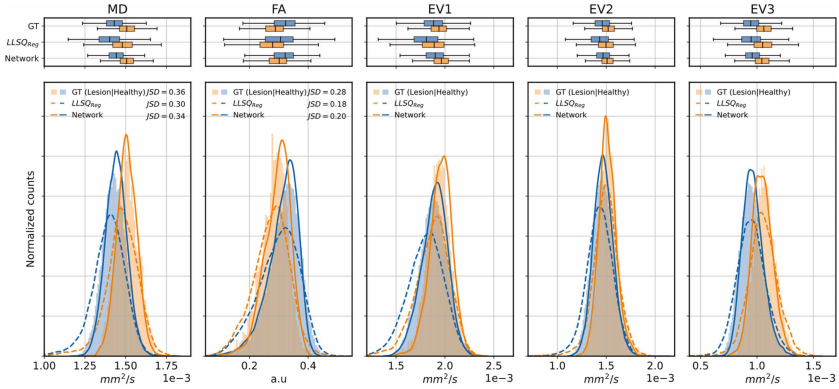


FIGURE 7.4: Marginal distributions for MD, FA and eigenvalues divided into lesion (orange) and healthy (blue) tissue for ground truth (bars), $LLSQ_{reg}$ (dashed lines), and Network inference (solid lines). The vertical axis shows normalized counts, where the normalization is performed per class (healthy/class). Top row shows the boxplots of the distributions shown in the lower row, where each pair of blue/orange box corresponds to one algorithm as indicated in the first column. For MD and FA, the Jensen-Shannon distance between healthy and lesion distributions are stated in correspondence to the legend.

metrics. Scatter plots comparing $LLSQ_{Reg}$ and Network versus ground truth (GT) for MD, FA, HA and E2A are given in Figure 7.5. The R2 scores are stated in the plot. The Network inference shows lower errors at high and low values of HA which correspond to locations at the myocardial borders. For the other metrics the Network inference is closer to the line of identity, resulting in higher R2 scores. The boxplots in Figure 7.5 illustrate the distribution of signed errors, demonstrating the lower number of outliers with Network inference.

TABLE 1: Mean absolute errors over test data for MD, FA, E2A, HA and eigenvalues of $LLSQ_{Stat}$, $LLSQ_{Unreg}$, $LLSQ_{Reg}$ and Network. Lowest error, other than the optimal reference $LLSQ_{Stat}$, is highlighted

	MD	FA	E2A	EV1	EV2	EV3	HA
	$10^{-4} \text{ mm}^2/\text{s}$	a.u.	$^{\circ}$	$10^{-4} \text{ mm}^2/\text{s}$	$10^{-4} \text{ mm}^2/\text{s}$	$10^{-4} \text{ mm}^2/\text{s}$	$^{\circ}$
$LLSQ_{Stat}$	0.246	0.016	3.2	0.414	0.381	0.345	3.5
$LLSQ_{Unreg}$	1.029	0.063	11.8	1.541	1.335	1.366	16.0
$LLSQ_{Reg}$	0.626	0.037	7.2	1.155	7.642	0.674	10.9
Network	0.337	0.022	5.6	0.571	0.537	0.484	5.3

The SNR dependency of the estimation quality is shown in Figure 7.6. MAE and R2 scores for MD, FA, HA and E2A are plotted against the simulated SNRs of (10, 14, 18, 22, 26, 30, 50). For higher SNR the errors of $LLSQ_{Stat}$ and $LLSQ_{Reg}$ decrease, while the Network inference is relatively insensitive to SNR changes. For SNR below 14, the Network yields even better estimates than $LLSQ_{Stat}$.

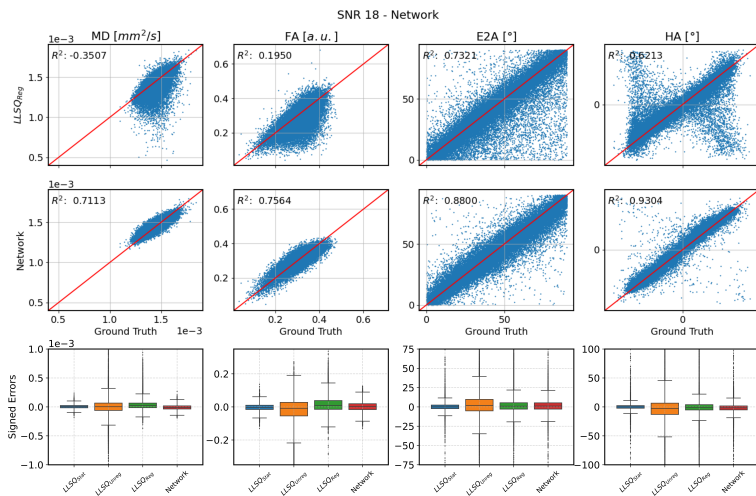


FIGURE 7.5: Overview of test errors (MD, FA, HA and E2A). First and second row show scatter plots of $LLSQ_{Reg}$ and Network inference versus ground truth. Line of identity is plotted in red, and the explained variance score is stated as text in each subplot. The bottom row shows boxplots of the signed error distributions.

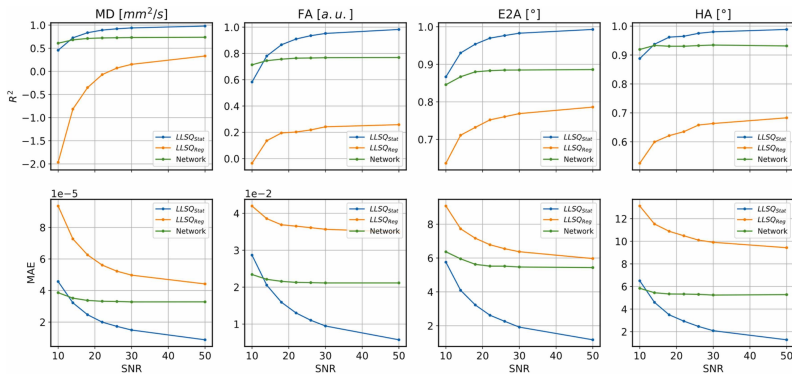


FIGURE 7.6: SNR dependency of R^2 score (top row) and mean absolute error (bottom row) of tensor metrics MD, FA, E2A and HA (from left to right) for Network inference (green) versus $LLSQ_{Stat}$ (blue) and $LLSQ_{Rep}$ (orange).

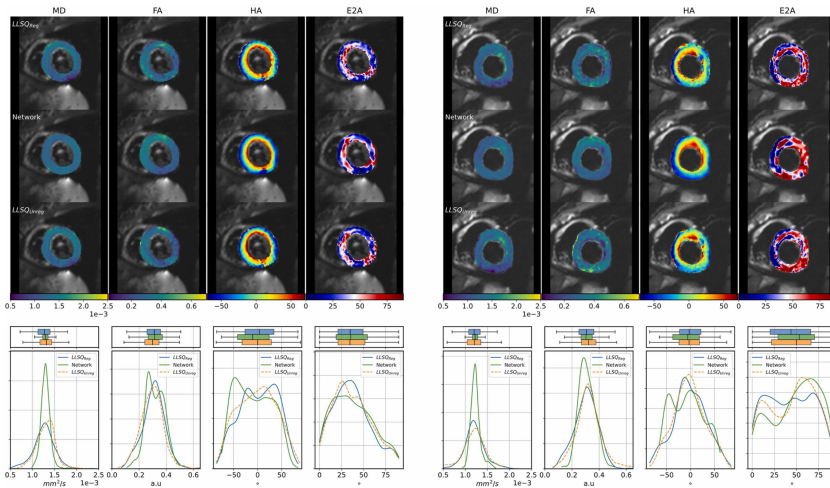


FIGURE 7.7: Exemplary in-vivo results of the mid-ventricular (left) and basal (right) slice of the same volunteer. Images show MD, FA, HA, and E2A maps (from left to right) of $LLSQ_{Reg}$ (top), Network (mid), and $LLSQ_{Unreg}$ (bottom). Corresponding histograms and boxplots are shown below the maps

7.3.2 In-vivo Data

Plots of MD, FA, HA and E2A maps for $LLSQ_{Reg}$, $LLSQ_{Unreg}$ and Network inference were created for all volunteer data and are provided as Supplementary Material. Alongside the maps, the corresponding marginal distributions are plotted. Figure 7.7 shows exemplary results for the mid-ventricular and basal slice of the same volunteer. The metric maps of the Network inference appear less grainy than with $LLSQ_{Reg}$. This is most prominent at the LV borders, where $LLSQ_{Reg}$ produces more outliers. Furthermore, the low MD values close to the liver interface are not present in the Network inference. The histograms in Figure 7.8 show that the Network yields narrower distributions for MD and FA without outliers, while the median values agree well. E2A maps are consistent between Network and $LLSQ_{Reg}$. The HA maps of the mid-ventricular slices contain a distorted region close to the posterior vein for both algorithms. In Figure 7.8 the basal slices of two different volunteers are shown, in which the data reveals signal dephasing close to the posterior vein and the lung interface, respectively. These areas show lower MD and higher FA in $LLSQ_{Reg}$ and $LLSQ_{Unreg}$, while the Network inference does not show these variations. The HA of all inferences on the first dataset show distortions at the location corresponding to the lowered signal.

The transmural helix angle variation averaged over all volunteers per slice location is shown in Figure 7.9. For all slice locations, both algorithms result in a linear transmural HA variation. The standard deviation of HA at the myocardial borders

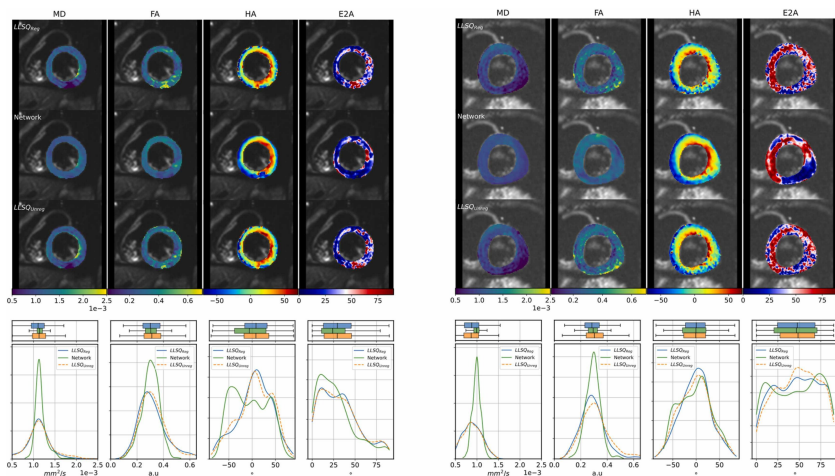


FIGURE 7.8: Exemplary in-vivo results of two different volunteers displaying under-estimation of MD and over-estimation of FA in areas of low SNR (close to the posterior vein and the lung interface). Images show MD, FA, HA, and E2A maps (from left to right) of $LLSQ_{Reg}$ (top), Network (mid), and $LLSQ_{Unreg}$ (bottom). Corresponding histograms and boxplots are shown below the maps.

shows the reduced number of outliers with Network inference, especially for basal and mid-ventricular slices.

The SNR dependency of inference quality for $LLSQ_{Reg}$, $LLSQ_{Unreg}$ and Network using the volunteer data shown in Figure 7.7, is illustrated in Figure 7.10. The curves show the MAE as a function of reduced signal averages with respect to inference using the full number of 10 signal averages per diffusion direction. In the mid-ventricular and basal slice, the Network inference (green) changes less than $LLSQ_{Reg}$ (blue). Especially for MD, FA, and HA the Network inference is less sensitive to SNR changes.

Figure 7.11 displays the marginal distributions of MD and FA of the augmented in vivo data as well as the Network inference. The distributions in healthy tissue in-vivo differ from the synthetic ones shown in Figure 7.4, as they are obtained from the grainy $LLSQ_{Reg}$ in-vivo inference. The parameter distribution for the lesioned tissue was sampled with the same procedure as used for the synthetic data.




FIGURE 7.9: Transmural variation of helix angles per slice over all in-vivo datasets. Area around the line corresponds to the standard deviation. Blue lines show data for $LLSQ_{Reg}$ and green lines show Network inference.

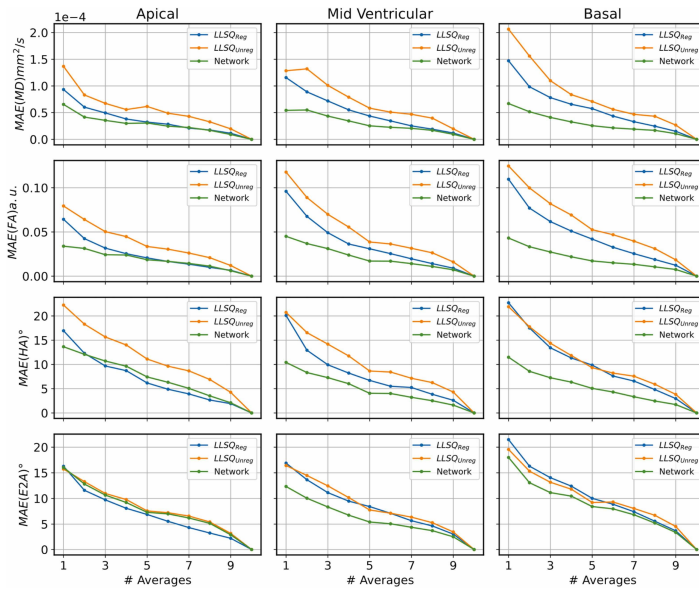


FIGURE 7.10: MAE of MD, FA, HA and E2A as a function of reduced signal averages for the volunteer data shown in Figure 7.9. Reference value for MAE is the inference using 10 averages. Columns correspond to slice location apical (left), mid-ventricular (middle) and basal (right). Rows correspond to metric. Colors correspond to algorithm: $LLSQ_{Reg}$ (blue), $LLSQ_{Unreg}$ (orange) and Network (green).

7.4 Discussion

In this work we have implemented a framework to generate synthetic data that captures the statistics of in-vivo cDTI data. Physiologically plausible priors on diffusion tensors, SNR variations, as well as respiratory-induced motion were

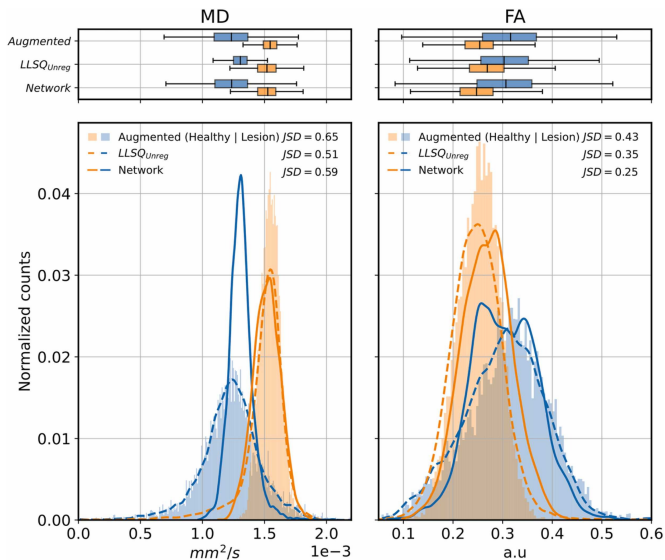


FIGURE 7.11: Marginal distribution of MD and FA using Network inference on lesion-augmented in vivo data. Blue and orange refer to healthy and lesioned tissue, respectively. Normalization of counts was performed per class. Blue bars illustrate the $LLSQ_{Reg}$ inference which was augmented with lesions (orange bars). Solid lines show the Network inference. Dashed lines show the $LLSQ_{Unreg}$ inference on the augmented data. The Jensen-Shannon distances (JSD) between healthy and lesion distribution are stated inside the subplots corresponding to the legend.

included into the training data. The residual neural network trained exclusively on synthetic data yielded consistent and robust tensor estimates from in vivo data, even in low SNR conditions, significantly outperforming current reference methods.

The results on the synthetic test data demonstrate the feasibility to train a CNN for tensor estimation without requiring a prior registration step. The marginal distributions of metrics as well as error distributions show the Network’s superior inference quality. Relative to the $LLSQ_{Reg}$ reference, noise-induced granularity was reduced with Network inference. Furthermore, for pixels close to the LV borders, the Network inference resulted in more consistent tensors as spatial correlations helped to address signal reduction caused by partial volume effects and susceptibility gradients.

The in-vivo results suggest a good transfer from synthetic to real data. The Network inference using in-vivo data as input showed similar features compared to using synthetic data as input, namely, the reduction of noise-induced granularity and a reduction of outliers especially at the LV borders. In two volunteers, the reduced SNR at the lung-myocardial interface and next to the posterior vein led to underestimation of MD and over estimation of FA by $LLSQ_{Reg}$, an issue also discussed in¹³⁴. The

Network's robustness against SNR changes, in contrast, prevented this effect in both cases.

In general, the MD and FA distributions of the Network inference were found to be narrower compared to the $LLSQ_{Reg}$ results. Comparing this finding with the results on synthetic test data revealed the same characteristic, although the differences were smaller. A possible explanation may relate to signal dephasing and deformation caused by field gradients³² as well as residual signal from the blood pool, which are not present in the synthetic test data, causing the error with $LLSQ_{Reg}$ inference to be smaller compared to the in-vivo situation.

The slice orientation for some in-vivo measurements was not exactly perpendicular to the long axis of the LV, which can be seen by the non-circular isolines in the HA maps. Since the network yielded the same pattern combined with good agreement of the E2A maps suggests that the network does not hallucinate unreasonable angulation maps. This is also supported by the non-regular HA map with $LLSQ_{Reg}$ as well as with Network inference on the apical slice in volunteer 5 which is positioned too far apical.

Our results from retrospective reduction of signal averages clearly demonstrate the increased SNR robustness of Network inference. Accordingly, the approach offers potential clinical value in cases, where patients are not able to complete the full examination. In general, our method allows to decrease the number of acquired averages, while also providing control over the training distributions and thereby it adds to studies which have trained neural networks on subsets of acquired in-vivo DTI datasets to reduce the number of acquisitions^{52,53,264}.

By augmenting the in-vivo data with artificial lesions it was demonstrated that discrimination between healthy and lesion tissue based on MD and FA is preserved. Residual overestimation of FA and under-estimation of MD in lesions can be explained by the sharp edges introduced into the metric maps by the mix-and-match strategy. Preserving the bimodality of MD, FA for lesion versus healthy tissue and being able to verify it is a key aspect of the presented work given that training data distributions have been shown to have an effect on potential inference biases²⁵⁹. Especially in medical applications, the interpretability of deep learning algorithms and knowledge about potential biases is very important.

As the current network implementation requires a LV mask to be applied before inference, the presented pipeline is not fully automatic yet. Various deep learning approaches have, however, shown the feasibility of obtaining an automatic LV segmentation^{50,265}. Combining such a segmentation algorithm with our approach could offer direct feedback for the MR system operator while scanning.

Improving the alignment of the LV in free breathing cDTI data is an active area of research^{50,98,99}. As no ground truth information about diffusion tensors is available in-vivo, these methods only take image intensities into account. Using full tensor

information for determining the registration transformation has been shown to improve tensor estimation¹⁰⁰. To this end, by adapting our network approach to estimate the full registration transformation, while having the tensor information available during training could improve the algorithm’s robustness even further. Additionally, the inclusion of further information into the forward model of our simulation pipeline is straight-forward. As the tensor-sampling strategy is parameterized, it is also transferable to other sequences. The corresponding diffusion tensor statistics could be estimated by Monte-Carlo simulations^{206,250}. This would enable the use of CNN based tensor fitting including its benefits, without the necessity to newly acquire large amounts of in-vivo data.

In conclusion, our synthetically trained CNN offers increased tensor estimation accuracy and precision of free-breathing in-vivo cDTI data when compared to conventional registration followed by least squares fitting. The approach is particularly beneficial in low SNR regimes, thereby enabling potential scan time reduction by reduced-average acquisition in-vivo.

7.5 Appendix

7.5.1 A: E2A and Lesion Map Generation

The algorithm to generate random E2A maps works on any ring-like 2D LV mask. Generating the maps involves two interpolation steps, first on a coarse resolution and subsequently on the original scale. To this end, the original 2D mask M_H of size (X, Y) is resized by the parameter f_r , resulting the coarser scale mask M_L of size $(X//f_r, Y//f_r)$, where $//$ denotes integer division. Therefore, the pixel size of M_L after resizing follows as $(f_r\Delta x_H, f_r\Delta y_H) \rightarrow (\Delta x_L, \Delta y_L)$. For each point in M_H and M_L , polar coordinates $\mathbf{p} = (r, \phi)$ are calculated, using the mask’s center of mass as coordinate origin. the positions belonging to the LV in M_H and M_L are denoted as $\mathbf{p}_{H,i}$ with $i \in [0, N_H]$ and $\vec{\mathbf{p}}_{L,j}$ with $i \in [0, N_L]$ respectively. Given $\mathbf{p}_{L,j}$, a subset of points is sampled randomly with uniform probability to serve as reference points $\mathbf{p}_{Ref,k}$ with $k \in [0, N_p]$. The number of reference points $N_p = \text{floor}(N_L f_s)$ is determined by the *seed-fraction* parameter f_s with $0 < f_s < 1$.

For all reference points $\mathbf{p}_{Ref,k}$ a sheetlet angles is assigned according to a Bernoulli distribution, where the probability for a high value *phigh* is specified as parameter. The resulting values are then interpolated to all $\mathbf{p}_{L,j}$ using a nearest-neighbour kernel, where the distances $d(\mathbf{p}_{Ref,k}, \mathbf{p}_{L,j})$ are evaluated as described in 7.5.3. This is followed by the second interpolation to $\mathbf{p}_{H,i}$, using $\mathbf{p}_{L,j}$ as reference points. the second interpolation uses a RBF kernel with width σ_{RBF} , again using the distance definition $d(\mathbf{p}_{L,j}, \mathbf{p}_{H,i})$ as described in 7.5.3. The weights for radial and circumferential direction in the distance definition $(\alpha_{rad}, \alpha_{circ})$ are specified as parameters.

The effect of each parameter on the map appearance can be investigated using the Jupyter-notebook provided at https://gitlab.ethz.ch/ibt-cmr-public/synthetic-dt-cnn/notebooks/Demo_E2A_Maps.ipynb. The main effects can be summarized as follows:

- f_R : Determines the map granularity. A higher f_R results in a coarser scale in the first interpolation step. The coarser this scale, the coarser is the structure of the resulting E2A map.
- $(\alpha_{rad}, \alpha_{circ})$: Determines the preferred direction of connected E2A areas. For a low α_{circ} blobs of high/low E2A are more likely to be elongated in circumferential direction.
- σ_{RBF} : Determines the smoothness of spatial transition from high to low sheetlet angles.
- p_{High} : Controls the ratio pixels with high/low sheetlet angles.

Generating random lesion maps, uses the same seeding and interpolation scheme as described for the sheetlet angle maps. However, a rounding step is included to obtain binary masks, which are subsequently modified by morphological opening to avoid too small or fractioned lesion areas. Furthermore, a rejection loop is used to ensure the number of connected lesion areas as well as the ratio of lesion vs LV area matches the specified values. To promote lesion positions close to the endocardial border, the probability is weighted by the transmural position $p_{High} \rightarrow p_{High}(r)$. A demonstration is provided in https://gitlab.ethz.ch/ibt-cmr-public/synthetic-dt-cnn/notebooks/Demo_Lesion_Maps.ipynb

7.5.2 B: Rejection-Acceptance Sampling of Eigenvalues

The aim of the implemented rejection sampling algorithm is the generation of eigenvalue-triples such that joint distribution $p(MD, FA)$ is uniform within specified box-constraints. The rejection is performed based on the FA value. MD uniformity is ensured by first sampling a uniform MD value followed by three uniform eigenvalues X_i from the interval $(0, 3)$ which are subsequently scaled according to $EV_i = MDX_i/\Sigma_i X_i$.

To obtain an approximation of the FA distribution for the specified MD interval, a burn-in phase of drawing 100 000 non-rejected eigenvalue-triples is used. The kernel density estimation of the burned in samples is used and stored as representation of the unrejected distribution. The rejection step follows the general rejection-acceptance sampling scheme²⁶⁶:

1. Let f_{FA} be the uniform target distribution on interval $[a, b]$
2. Generate Y from p_{FA}
3. Generate U from the uniform distribution $\mathcal{U}(0, cp_{FA}(Y))$

4. If $U < f_{FA}(Y)$ accept Y as sample

A demonstration including additional illustrations and implementation details can be found at https://gitlab.ethz.ch/ibt-cmr-public/synthetic-dti-cnn/notebooks/Demo_Tensor

7.5.3 C: Geodesic in Relative Polar Coordinates

Every position inside the LV mask is described by a transmural depth r and an angle ϕ . The geodesic between two points $\mathbf{p}_1 = (r_1, \phi_1)^T$ and $\mathbf{p}_2 = (r_2, \phi_2)^T$ is the parameterized path described as $\mathbf{fl}(t) = \mathbf{p}_1 + (\mathbf{p}_2 - \mathbf{p}_1)t = (r_1 + \delta r t, \phi_1 + \delta \phi t)^T$. Choosing the shortest path implies that $\delta \phi \leq 180^\circ$ for all points $\mathbf{p}_1, \mathbf{p}_2$. To calculate the length of the geodesic, the path is integrated as follows

$$d(\mathbf{p}_1, \mathbf{p}_2) = \int_0^1 |\dot{\gamma}|_2 dt. \quad (7.4)$$

Using the polar coordinate form $\gamma(t) = r(t)\mathbf{e}_r(t)$ and the basis vector definitions $\mathbf{e}_r(t) = (\cos \phi(t), \sin \phi(t))$, $\mathbf{e}_\phi(t) = (-\sin \phi(t), \cos \phi(t))$, the path derivative evaluates to $\dot{\gamma}(t) = \dot{r}(t)\mathbf{e}_r(t) + r(t)\dot{\phi}(t)\mathbf{e}_\phi$. From that follows:

$$d(\mathbf{p}_1, \mathbf{p}_2) = \int_0^1 \left| \begin{pmatrix} \Delta r \cos \phi(t) + r(t)\Delta \phi \sin \phi(t) \\ -\Delta r \sin \phi(t) + r(t)\Delta \phi \cos \phi(t) \end{pmatrix} \right|_2 dt \quad (7.5)$$

$$= \int_0^1 \sqrt{(\Delta r)^2 + (\Delta \phi)^2 (r_1 + \Delta r t)^2} dt. \quad (7.6)$$

For the special cases of $\Delta r = 0$ and $\Delta \phi = 0$ the integral becomes trivial, while for all other cases the solution is given as:

$$d(\mathbf{p}_1, \mathbf{p}_2) = \left[\frac{\Delta r + r_1}{2\Delta r} A + \frac{\Delta r}{2\Delta \phi} \ln(A + \Delta r \Delta \phi + \Delta r_1) \right] - \left[\frac{r_1}{2\Delta r} B + \frac{\Delta r}{2\Delta \phi} \ln(B + \Delta \phi r_1) \right] \quad (7.7)$$

$$A = \sqrt{(\Delta r)^2 + (\Delta \phi)^2 (r_1^2 + \Delta r)^2} \quad (7.8)$$

$$B = \sqrt{(\Delta r)^2 + (\Delta \phi)^2 r_1^2} \quad (7.9)$$

Weighting the coordinate differences $\tilde{\Delta r} = \alpha_{rad}\Delta r$ and $\tilde{\Delta \phi} = \alpha_{circ}\Delta \phi$ allows to stretch and compress the distance in radial/circumferential direction separately as described in 7.5.1.

8. Joint Encoding of Diffusion Tensors and Tissue Velocities in the Heart

Jonathan Weine¹, Christian T Stoeckl^{1, 2}, Stefano Buoso¹, Sebastian Kozerke¹

¹Institute for Biomedical Engineering, University and ETH Zurich, Zurich, Switzerland

²Division of Surgical Research, University Hospital Zurich, University of Zurich, Zurich, Switzerland

Submitted for peer Review

8.1 Introduction

Cardiovascular Magnetic Resonance (CMR) allows for the assessment of local ventricular tissue microstructure^{12,20,23,24,29,53,168,169,267} and tissue strain^{8–10} in various conditions. Tissue microstructural changes have been studied using cardiac diffusion tensor imaging (cDTI) in cases of hypertrophic and dilated cardiomyopathies^{12,18,163}, aortic stenosis²⁵, obesity⁵³ and infarction^{21,149,150,154,155,268,269}. The estimation of diffusion tensors is affected by myocardial motion and strain^{143,164,166,167,169}. To alleviate these effects second-order motion-compensated diffusion encoding waveforms in combination with spin echo sequences (Mo12-SE) have been shown to be widely applicable^{12,29,30,98,170}. Since the macroscopic contraction pattern of the heart is connected to its microstructure, the helical orientation of myocytes aggregates, their interconnections, and their organization in sheetlets in relation to macroscopic strain or its temporal derivative i.e. strain rate have been of particular interest^{12,153}.

The most widely used strain measurement approach utilizes feature tracking (FT) on standard cine image series^{9,10}. Alternative techniques such as myocardial tagging^{105,108,109}, Displacement ENcoding with Stimulated Echoes (DENSE)^{101,102}, or Tissue Velocity Mapping (TVM)^{117,121} require additional measurements.

TVM utilizes the phase-contrast principle¹¹² to encode three-directional tissue velocities into the image phase. From the estimated velocities, strain rates are obtained by computing spatial derivatives^{9,10}. Different sequences have been used for TVM, such as spoiled gradient echo (GRE)^{118,120,121,123,124,126,270} and spin echo^{271–273}. The value of TVM has been shown for myocardial infarction¹⁴⁷, obesity¹⁵⁹, and atrial fibrillation²⁷⁴. The systolic strain rate has been proposed as an imaging marker in diabetes²⁷⁵, ischemic diseases²⁷⁶, pulmonary hypertension²⁷⁷, hypertrophic cardiomyopathy¹⁵⁷ and myocardial dysfunction after radiation therapy¹⁶².

Myofiber strain connects to macroscopic function and it is considered to be spatially homogenous in healthy hearts³⁵. To infer upon myofiber aggregate strain, estimates of both macroscopic strain and local fiber orientation are required. Since projections of measured strains onto radial, circumferential and longitudinal directions do not

generally coincide with the microstructure, transmural variation of circumferential strain has been reported³⁵. Several works have proposed to combine cDTI data with DENSE³⁵⁻³⁸, or TVM^{272,273} all acquiring separate scans. Co-registering the data is challenging and a potential source of error³⁵.

Since the sequence timing of Mo12-SE cDTI with Cartesian readout includes a gap between excitation and the start of the diffusion encoding gradients, it is proposed to include additional bipolar gradients to encode tissue velocities without changing cDTI sequence timings. Biophysical simulations of a contracting left ventricle were used to assess accuracy and precision of the proposed encoding scheme. In-vivo feasibility of obtaining both cDTI and tissue velocity maps in a single scan was demonstrated on healthy volunteers.

8.2 Methods

8.2.1 Joint Encoding Sequence

Bipolar velocity encoding gradients were inserted into the Mo12-SE cDTI²⁹ sequence between the spectral-spatial excitation and the start of the diffusion encoding waveform. The velocity encoding strength (VENC) was set to 20 cm s^{-1} ¹²². To lower the required amplitude of the bipolar gradients per direction, a Hadamard encoding scheme¹¹⁴ was chosen.

The cDTI acquisition included three slices of the left ventricle (apical, mid-ventricular and basal) with a slice thickness of 10 mm and a gap of 10 mm. For each slice, three low b-values (100 s mm^{-2}) and nine high b-values (450 s mm^{-2}) were acquired plus an additional unweighted reference. For each diffusion weighing, eight averages were acquired, with half of them using inverted diffusion gradients to subtract eddy current phase contributions up to the linear terms. The velocity encoding direction was varied per average. Hence, each velocity encoding was acquired for both the original and inverted diffusion waveforms. Figure 8.1 illustrates the encoding scheme and sequence. The sequence timings correspond to a gradient system with 80 mT m^{-1} maximum gradient strength and $100 \text{ T m}^{-1} \text{ s}^{-1}$ slew rate per physical gradient axis.

8.2.2 Diffusion Tensor and Velocity Estimation

The individual single-shot images were registered using the parameterized Total Variation toolbox⁹⁷. Subsequently, left-ventricular (LV) masks were obtained using manual annotations and LV local coordinates and directions (radial, circumferential, and longitudinal) were derived¹⁶⁸.

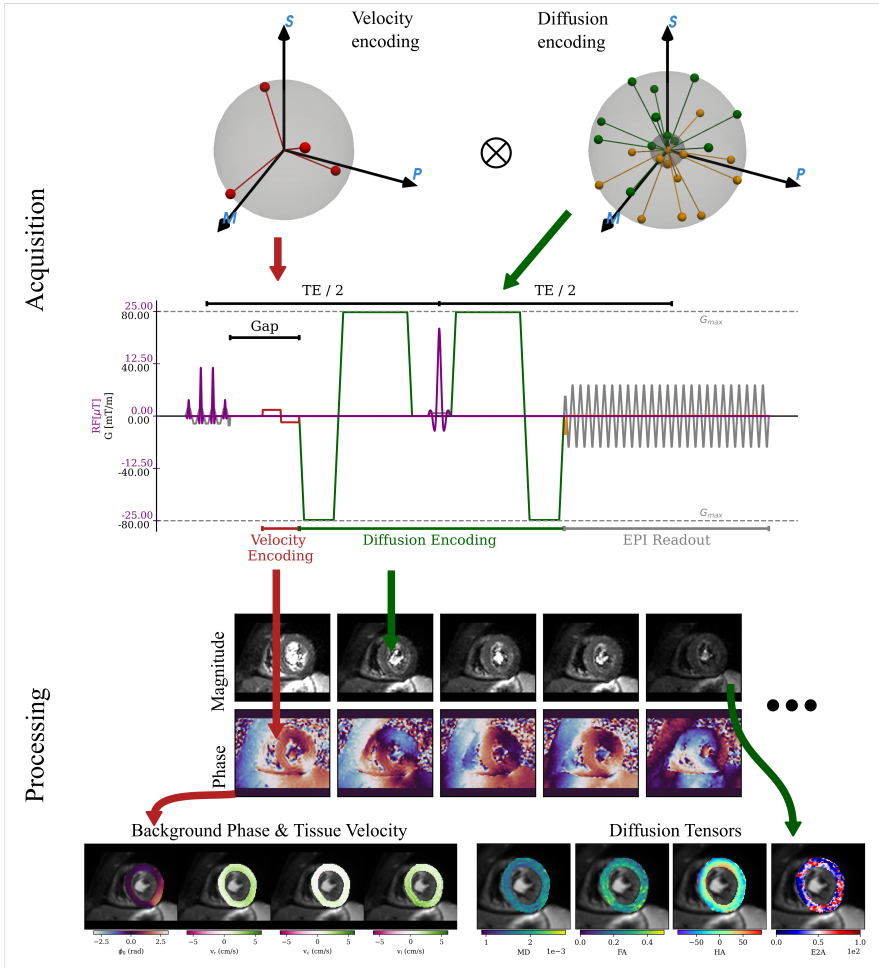


FIGURE 8.1: Illustration of the proposed sequence. The velocity encoding scheme (left) and the diffusion encoding scheme (right) are illustrated in the top row, where the green and orange spherical markers correspond to original and inverted diffusion gradients. Every combination of one velocity and one diffusion encoding waveforms yields a single image. The gradients corresponding to either velocity or diffusion encoding waveforms are indicated by red and green arrows/lines, respectively. The image phase is used to infer tissue velocity vectors (v_x, v_y, v_z) and background phase ϕ_0 (bottom left), while mean diffusivity (MD), fractional anisotropy (FA), helix angle (HA) and sheet angle (E2A) are derived from the magnitude images (bottom right)

Registered magnitude images were used to estimate the diffusion tensors $\hat{\mathbf{d}}$, using the `cdtipy` python package, solving the pixel wise linear least-squares minimization problem defined as

$$\mathbf{d} = [\log(s_0), D_{xx}, D_{yy}, D_{zz}, D_{xy}, D_{xz}, D_{yz}]^T, \quad (8.1)$$

$$\mathbf{x} = [l \log(s_1), \log s_2, \cdot]^T = \mathbf{B}\mathbf{d} + \boldsymbol{\eta} \quad (8.2)$$

$$\hat{\mathbf{d}} = \arg \min \|\mathbf{x} - \mathbf{B}\mathbf{d}\|_2^2, \quad (8.3)$$

where s_i are the pixel-wise magnitude values of the diffusion encoding i , \mathbf{B} is the diffusion encoding matrix, \mathbf{d} is the vector containing the unweighted signal intensity as well as the six unique entries of the symmetric diffusion tensor per pixel and $\boldsymbol{\eta}$ denotes random noise. The diffusion tensors were subsequently used to compute mean diffusivity (MD), fractional anisotropy (FA), helix angle (HA) and sheetlet angles (E2A).

Registered phase images were first unwrapped using the scikit-image python package (v0.21) and subsequently a model containing tissue- and respiratory-related velocities was fitted. In the following notation, subscripts q, n, h indicate an acquisition with diffusion weighting $q \in [0, 12]$, gradient inversion $n \in [0, 1]$, and velocity encoding $h \in [0, 3]$, and the subscript p indicates the pixel index. In the signal model, the pixel-wise phase for a selected acquisition $\hat{\phi}_{p,q,n,h}$, is defined by the contribution of the background phase Φ_p^0 , the cardiac-related velocity vector \mathbf{v}_p^c , the translational breathing-related velocity $\mathbf{v}_{q,n,h}^b$ and a phase term, $\phi_{q,p}^{\text{eddy}}$, accounting for the first-order eddy currents due to the diffusion gradients³¹. Breathing motion was assumed to be uniform for all LV pixels for a single-shot image and along a constant direction \mathbf{n}^b for all acquisitions, such that $\mathbf{v}_{q,n,h}^b = v_{q,n,h}^b \mathbf{n}^b$. The pixel-wise cardiac velocity \mathbf{v}_p^c and background phase Φ_p^0 were assumed to be identical for all acquisitions q, n, h .

Let $\boldsymbol{\varphi} = \frac{\pi}{\text{VENC}} \left(\frac{\text{VENC}\Phi_p^0}{\pi}, v_p^{c,x}, v_p^{c,y}, v_p^{c,z} \right) \in \mathbb{R}^{1 \times 4}$ be the phase contribution from background phase Φ_p^0 and cardiac velocities \mathbf{v}_p^c for a single pixel p . Furthermore, let $\boldsymbol{\vartheta}_{q,n,h} = \frac{\pi}{\text{VENC}} \left(0, v_{q,n,h}^b n_x^b, v_{q,n,h}^b n_y^b, v_{q,n,h}^b n_z^b \right) \in \mathbb{R}^{1 \times 4}$ be the phase contribution due to breathing motion. Using the Hadamard encoding matrix defined as

$$\mathbf{E}_v = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix}, \quad (8.4)$$

and further assuming the eddy-current effects to be captured by a phase ramp parameterized as

$$\phi_{p,q,n,h}^{\text{eddy}} = \mathbf{a}_{q,n,h} \cdot \mathbf{r}_p, \quad (8.5)$$

with the augmented in-plane pixel coordinate vector $\mathbf{r}_p' = (1, x_p, y_p)^T$ and the plane parameterization $\mathbf{a}_{q,n,h} = (\frac{w_0}{w_x}, \frac{w_0}{w_y}, \frac{w_0}{w_z})_{q,n,h}$, the total phase evaluates to

$$\hat{\phi}_{p,q,n,h} = \mathbf{E}_v \cdot \boldsymbol{\varphi}_p + \mathbf{E}_v \cdot \boldsymbol{\vartheta}_{q,n,h} + \phi_{p,q,n,h}^{\text{eddy}}. \quad (8.6)$$

Inverting the diffusion gradients is assumed to invert the eddy-current phase contributions according to $\phi_{p,q,0,h}^{\text{eddy}} = -\phi_{p,q,1,h}^{\text{eddy}}$. Estimation of all parameters Φ_p^0 , \mathbf{v}_p^c , $v_{q,n,h}^b$, \mathbf{n}^b , and $\mathbf{a}_{q,n,h}$ was achieved by minimizing the mean squared error of the phase-signal model and the unwrapped measured phase $I_{p,q,n,h}$, using the Adam optimizer implemented in TensorFlow (v2.12). The fitting process was divided into two steps, because fitting the full model to all data where in a systematic underestimation of through-plane contraction velocity. In the first step model eq 8.6 was fit according to

$$\arg \min_{\{\boldsymbol{\varphi}_p, \boldsymbol{\vartheta}_{q,n,h}, \mathbf{a}_{q,n,h}\}} \sum_{p,q,n,h} |\hat{\phi}_{p,q,n,h} - I_{p,q,n,h}|^2 + \alpha_1 \sum_p (|v_p^{c,x}| + |v_p^{c,y}|) + \alpha_2 \sqrt{\sum_p (v_p^{c,z})^2}, \quad (8.7)$$

using all data. The regularization weights α_1, α_2 were chosen based on simulation data. In the second pass, the fit of the through-plane contractile velocity was performed while fixing in-plane contractile velocities, background phase as well as the the breathing motion direction to the results of the first step. The fit model of the second step was identical to eq. 8.6, although only the non-diffusion weighted images were used as input and the longitudinal velocity was initialized with the direct Hadamard inversion. Furthermore, the breathing motion was regularized using the L2-norm of all velocity factors $v_{q,n,h}^b$ and a mean through-plane velocity offset $\langle v_p^{c,z} \rangle_p$, that was scaled with the expected motion amplitude was added. The calibration of the respective factor α_4 was also based on simulation data, resulting in following optimization problem

$$\arg \min_{v_p^{c,z}, v_{0,n,h}^b} \sum_{p,n,h} |\hat{\phi}_{p,0,n,h} - I_{p,0,n,h}|^2 + \alpha_3 \sqrt{\sum_{n,h} (v_{0,n,h}^b)^2} - \alpha_4 \langle v_p^{c,z} \rangle_p. \quad (8.8)$$

8.2.3 Biophysical Simulation

Diffusion tensor and velocity estimation was first tested in-silico on a numerical phantom generated based on a biophysical model of the left ventricle¹⁸⁵. A cardiac cycle of 950 ms was simulated, with peak systole occurring at 340 ms^{185,189}. The MR signal was simulated using the CMRsim package²⁷⁸. Breathing motion was incorporated as a sinusoidal breathing pattern with a period of 7 s, using CMRsim's breathing motion module. The simulation was repeated with varying breathing

amplitudes $A_b \in \{0, 2.5, 6.25, 12.5, 18.75, 25\}$ mm to investigate the impact of the breathing on the velocity estimation. Random diffusion tensors, based on a rule-based helix-angle model and a random sheetlet distribution were assigned at peak systole²²². The orientation of the tensors was warped during the simulation using the biophysical deformation field. CMRsim's proper orthogonal decomposition module was used to retrieve a consistent tensor map along with isochromat position and proton density. The corresponding single-shot k-space data were obtained using the analytical signal model

$$s(\mathbf{k}(t_s)) = \sum_p^N \delta V \tilde{m}_p e^{\mathbf{k}(t_s) \cdot \mathbf{r}_p(t_s)} + \eta, \quad (8.9)$$

where the particle apparent complex-valued magnetization \tilde{m}_p , representing the volume element δV , results from

$$\tilde{m}_p = \psi_C \circ \psi_{\text{Diff}} \circ \psi_{T_2^*} \circ \psi_\phi \circ \psi_{\text{SE}} \circ o_p, \quad (8.10)$$

with material points denoted as o_p and operator definitions for spin-echo contrast ψ_{SE} , phase due to motion during application of diffusion and velocity encoding (from t_a to t_b) ψ_ϕ , T_2^* -decay and B_0 inhomogeneity during readout, model-based diffusion weighting ψ_{Diff} , and coil-sensitivity weighting ψ_C given as:

$$\psi_{\text{SE}} = \rho_p e^{-\frac{TE}{T_{2,p}}} \left(1 - e^{-\frac{TR}{T_{1,p}}} \right), \quad (8.11)$$

$$\psi_\phi = e^{j\gamma \int_{t_a}^{t_b} \mathbf{G}(t) \cdot \mathbf{r}_p(t) dt}, \quad (8.12)$$

$$\psi_{T_2^*} = e^{-\frac{|t_s - TE|}{T_{2,p}^*}} e^{j\gamma \delta B_{0,p}(t_s - TE)}, \quad (8.13)$$

$$\psi_{\text{Diff}} = e^{-b\mathbf{g}^T \mathbf{D}_p \mathbf{g}}, \quad (8.14)$$

$$\psi_C = C(\mathbf{r}_p). \quad (8.15)$$

Tissue properties $\rho_p, T_{2,p}^*, T_{2,p}, T_{1,p}, \delta B_{0,p}, \mathbf{D}_p$ refer to proton density, relaxation times, off-resonance and diffusion tensor. Parameters are echo time TE and repetition time TR, diffusion and velocity encoding gradient waveforms $\mathbf{G}(t)$, sampling time t_s , diffusion weighting (b, \mathbf{g}) and coil sensitivities $C(\mathbf{r}_p)$ at particle location \mathbf{r}_p . k-space vectors at sampling times are denoted $\mathbf{k}(t_s)$ and the complex-valued Gaussian noise is denoted as η .

The phase accrual due to motion was numerically integrated using the gradient waveforms for velocity and diffusion encoding on a 10 μ s time grid. To investigate the effect of the velocity encoding gradients on diffusion estimates, the simulation was performed by varying $\mathbf{G}(t)$ in eq. 8.12, for the following four simulation cases: (i) no phase accrual; (ii) phase accrual due to velocity encoding; (iii) diffusion encoding only; and (iv) phase accrual due to both waveforms. The simulation was run

for multiple trigger delays $t_{\text{trig}} \in \{120, 129, 139, 148, 157, 166, 176, 185, 194, 204, 213, 222, 231, 241, 250\}$ ms.

To compare the simulation results, the ground truth LV mask for each slice was subdivided into 60 segments, where the transmural depth was divided into three bins and the circumference was divided into 20 segments. As ground truth for tissue velocity, the direct Hadamard inversion of the simulation results without breathing motion ($A_b = 0$), no noise and only velocity encoding (ii) was used. The calibration of α_4 (see eq. 8.8) was achieved by comparing the velocities obtained from applying the fitting routine to the data of simulation case (ii) to the tissue velocity ground truth. The effect of phase accrual during diffusion encoding, as well as the impact of respiratory motion on velocity estimates were investigated by computing the signed error of velocity estimation for all breathing amplitudes A_b comparing simulation cases (ii) and (iv). To investigate a potential effect of the additional velocity encoding on the diffusion metrics, the signed error of the diffusion metrics MD, FA, HA and E2A estimated for all breathing amplitudes A_b was calculated for simulation case (iv) when using the same breathing motion A_b as with simulation case (iii). Furthermore, the diffusion metrics for simulation cases (i) and (iii) were compared to evaluate whether the simulation captures the increase of MD caused by intravoxel phase gradients resulting from residual motion sensitivity of the diffusion waveforms. Finally, the effect of registration on the diffusion metric estimation was investigated by comparing results for non-zero breathing amplitudes A_b to $A_b = 0$ with simulation case (i).

8.2.4 In-vivo Data

The proposed sequence was implemented on a 1.5T clinical MRI scanner (Achieva, Philips Healthcare, Best, the Netherlands) with a 80 mT m^{-1} max strength and $100 \text{ T m}^{-1} \text{ s}^{-1}$ slew rate gradient system. In-vivo data was acquired in 5 healthy volunteers (3 males, 2 females, 28 ± 2.3 years) upon written informed consent according to ethics and institutional guidelines. The imaging protocol included the acquisition of a cine stack of short-axis slices covering the whole LV and a four-chamber cine dataset with 1.3 mm in-plane resolution and 25 heart phases. The cDTI protocol was acquired during a) breath hold (cDTI-TVM-BH) and b) during free breathing (cDTI-TVM-FB). During a single breath hold, four averages per diffusion encoding direction of the three slices were acquired, thus minimizing the breathing contribution per Hadamard encoding tuple. The acquisition resolution was set to $2.5 \text{ mm} \times 2.5 \text{ mm}$ with a slice thickness of 10 mm. The echo and repetition times were 92 ms and 3 RR intervals, respectively. The ECG trigger delay (t_{trig}) was set to 60% peak systole. All images were reconstructed using ReconFrame (GyroTools LLC, Winterthur, Switzerland) to a resolution of $1.25 \text{ mm} \times 1.25 \text{ mm}$ using zero filling. After reconstruction, images were registered as described above. For each of the three slices, LV masks were manually drawn. Based on the non-

diffusion weighted images, a phase ramp was fitted using points in the chest wall and liver. This phase ramp was subtracted from all images to reduce the impact of phase wraps on velocity estimation. The estimation of diffusion tensors and tissue velocity was performed as described above.

As reference, a velocity-encoded (3 directions + reference) phase-contrast GRE sequence was acquired with respiratory gating (7 mm gating window) for the mid-ventricular slice and with 50 heart phases. Velocity maps were obtained offline using ReconFrame (GyroTools LLC, Winterthur, Switzerland). LV masks for each frame were obtained by mapping the automatic segmentation of the cine data¹⁸⁹ to the GRE data using the imaging slice orientation matrices. Mean and standard deviation of the velocity projections were computed.

For each volunteer, DTI metrics (MD, FA, HA, E2A) as well as background phase and tissue velocities maps were computed. A sector-wise comparison of estimated velocities from the reference method (TVM), free-breathing cDTI-TVM-FB and breath-held cDTI-TVM-BH was performed. The sectors used correspond to the AHA definition for short-axis views⁷³. To evaluate the agreement of velocity estimation between TVM, cDTI-TVM-FB and cDTI-TVM-BH, Bland-Altman plots were calculated.

8.3 Results

8.3.1 Biophysical Simulation

Figure 8.2 compares velocity estimations for three of the six different breathing amplitudes at all trigger delays with the proposed fitting method as well as with the direct Hadamard inversion for breath-held data. For the simulation without any breathing motion (cf. Figure 8.2a) the fitting algorithm converged to the direct Hadamard inversion. Also, no bias is observed for the estimation of radial and circumferential velocities for shallow and deep breathing for all trigger delays (first two rows in Figure 8.2b,c). The interquartile ranges of in-plane velocity estimation errors are on the order of 2 mm s^{-1} . In the non-regularized case, longitudinal velocity estimation suffers from a bias dependent on the breathing amplitude, which is consistent for all trigger delays. Using regularization, the bias is removed (Figure 8.2, third row). The interquartile ranges of the longitudinal velocity estimation errors are below 1 mm s^{-1} .

The differences of the estimated velocities when comparing simulation (ii) to (iv) (isolating the effect of residual motion sensitization of the diffusion waveforms) show no bias and interquartile ranges on the order of 1 mm s^{-1} in-plane and negligible differences in longitudinal direction. Box plots as in Figure 8.2 for this comparison are provided in the supplementary material (Figure S2).




FIGURE 8.2: Velocity estimation errors using the proposed fitting algorithm (cDTI-TVM) when compared against direct Hadamard inversion (reference) for trigger delays from 120 ms to 250 ms. Configurations for reference and fit are illustrated in the top row. Data for both scenarios were simulated using simulation case (ii) i.e. only velocity encoding gradients in ψ_ϕ (eq. 8.12). For the reference breath-hold was assumed ($A_b = 0$ mm), while the cDTI-TVM data, processed with the two-step fit, incorporated a) no breathing ($A_b = 0$ mm), b) shallow breathing ($A_b = 6.25$ mm) and c) deep breathing ($A_b = 18.75$ mm). Rows of the boxplots correspond to (from top to bottom): radial velocity v_r , circumferential velocity v_c , longitudinal velocity v_l with calibrated second fitting step, and longitudinal velocity obtained without adjustment $v_{l,pre}$ ($\alpha_4 = 0$ in equation 8.8). Each box plot illustrates the distribution of signed errors over all mask segments per trigger delay.

Figure 8.3 compares the estimation of diffusion tensor metrics MD, FA, HA and E2A between simulation (i) to (iii). Breath-held (left column), shallow and deep breathing show an increased MD for simulation case (iii) with a negative bias of approximately 10^{-5} mm² s⁻¹ in the box plots of the first row. MD estimation without breathing motion (Figure 8.3a top row) shows a dependency on the trigger delay, where the variance of the signed error is the lowest at around 185 ms corresponding to approximately 60% systole. While the FA estimation is unbiased for the breath-held simulation, the registered breathing data results in an underestimation of FA of approximately 0.01. The FA estimation error for the breath-held data is lowest at 60% systole. The errors for both angles HA and E2A show no bias or dependency on trigger delay.

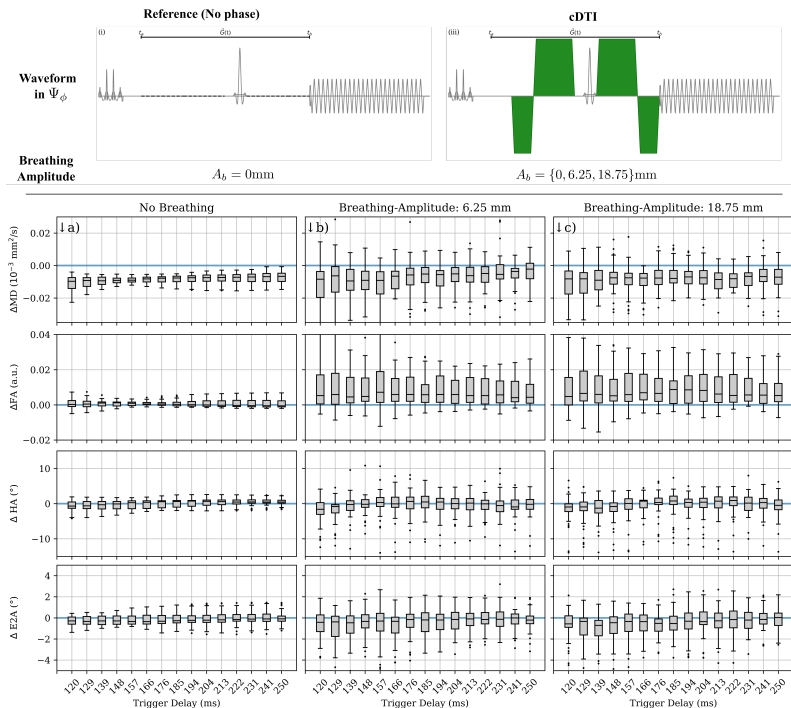


FIGURE 8.3: Box plots showing the differences in diffusion metrics when comparing simulation case (i) i.e. no phase accrual and (iii) i.e. with cDTI waveform as illustrated in the top row for trigger delays from 120 ms to 250 ms. The reference only used breath-hold ($A_b = 0$ mm), while the cDTI data incorporated a) no breathing ($A_b = 0$ mm) b) shallow breathing ($A_b = 6.25$ mm) and c) deep breathing ($A_b = 18.75$ mm). Each box plot illustrates the distribution of signed errors over all mask segments per trigger delay. Rows correspond to the diffusion metrics (from top to bottom): mean diffusivity (MD), fractional anisotropy (FA), helix angle (HA) and sheetlet angle (E2A).

Differences for diffusion metrics estimated for simulation case (iii) versus (iv), isolating the effect of the phase gradient induced by adding velocity encoding to the diffusion measurements, are provided as supplementary material (Figure S2). Interquartile ranges per metric are $IQR_{\Delta MD} \approx 3 \cdot 10^{-6} \text{mm}^2 \text{s}^{-1}$, $IQR_{\Delta FA} \approx 0.003$, $IQR_{\Delta HA} \approx 1^\circ$, and $IQR_{\Delta E2A} \approx 0.5^\circ$.

The results comparing diffusion metrics obtained from the simulation with breathing motion against no breathing motion both using simulation case (iii) are shown in detail in the supplementary material. While no bias in MD is present, FA is underestimated if the data incorporates breathing motion similar to the bias shown in Figure 8.3.

8.3.2 In-vivo Data

Figures 8.4 to 8.7 show the results for a single volunteer. The diffusion metrics estimated from cDTI-TVM-FB as well as cDTI-TVM-BH acquisition are presented in Figure 8.4, where except for an area in the anterior sector in the apical slice, all maps are consistent. The lateral part of the ventricle shows elevated MD and reduced FA in both cases, while the helix angle displays approximately linear transmural progression. Mean and standard deviation for MD and FA over all volunteers and slices were $(1.50 \pm 0.02) \cdot 10^{-3} \text{mm}^2 \text{s}^{-1}$ and 0.32 ± 0.08 for cDTI-TVM-BH and $(1.47 \pm 0.02) \cdot 10^{-3} \text{mm}^2 \text{s}^{-1}$ and 0.32 ± 0.08 for cDTI-TVM-FB.

Figure 8.5 shows the background phase as well as the velocity vectors obtained from cDTI-TVM-FB and cDTI-TVM-BH data. Mean and standard deviation of tissue velocity vectors over all LV pixels per slice are stated in Figure 8.5. The global velocity averages over all volunteers at 60% systole are given in Table 2. In all slices an off-resonance background phase induced by the lung-myocardium interface of the lateral LV is visible. Furthermore, the varying off-resonance induced by the posterior cardiac vein is seen in all slices. In the apical slice in the anterior part close to the great cardiac vein another area of elevated off resonance is apparent.

The mean and standard deviation of the projected velocities for the mid-ventricular slice over the cardiac cycle, obtained from the reference method (TVM), are plotted in the left column of Figure 8.6. The mean and standard deviation for the velocity vectors obtained from cDTI-TVM-FB and cDTI-TVM-BH at a single trigger delay (marked by the dashed gray line in Figure 8.6) are shown in the same plots as orange and green error bars. In the right column of Figure 8.6, the velocity maps of cDTI-TVM-FB and cDTI-TVM-BH as well as conventional TVM are shown.

TABLE 2: Global averages and standard deviations of tissue velocity projections in cm/s for radial (v_r), circumferential (v_c), and longitudinal (v_l) direction over all volunteers, for all three slice positions (apical, mid-ventricular, basal) and for all methods (TVM, cDTI-TVM-FB, and cDTI-TVM-BH).

	Apical		Mid-ventricular		Basal		
	cDTI-TVM-BH	cDTI-TVM-FB	cDTI-TVM-BH	cDTI-TVM-FB	TVM	cDTI-TVM-BH	cDTI-TVM-FB
$v_r (\text{cm s}^{-1})$	1.7 ± 0.4	1.8 ± 0.5	2.2 ± 0.5	2.1 ± 0.5	2.1 ± 1.3	2.2 ± 0.6	2.2 ± 0.9
$v_c (\text{cm s}^{-1})$	-0.8 ± 0.5	-0.7 ± 0.5	0.1 ± 0.6	0.2 ± 0.5	0.0 ± 1.5	1.1 ± 0.7	1.2 ± 0.6
$v_l (\text{cm s}^{-1})$	0.3 ± 1.1	0.3 ± 0.8	2.3 ± 0.9	2.4 ± 0.8	1.5 ± 1.5	3.8 ± 1.5	3.9 ± 1.1

A spatially resolved comparison of the velocity projections for the three methods is shown in Figure 8.7. As no reference acquisitions TVM (blue markers) for the basal and apical slices are available, only the cDTI-TVM-FB (orange) and cDTI-TVM-BH (green) data is shown for sectors 0 to 8 and 21 to 32. Each box plot illustrates the distribution of velocities inside the corresponding sector. The lower radial velocity in the septum for all slices, as well as the increasing (from apex to base) longitudinal

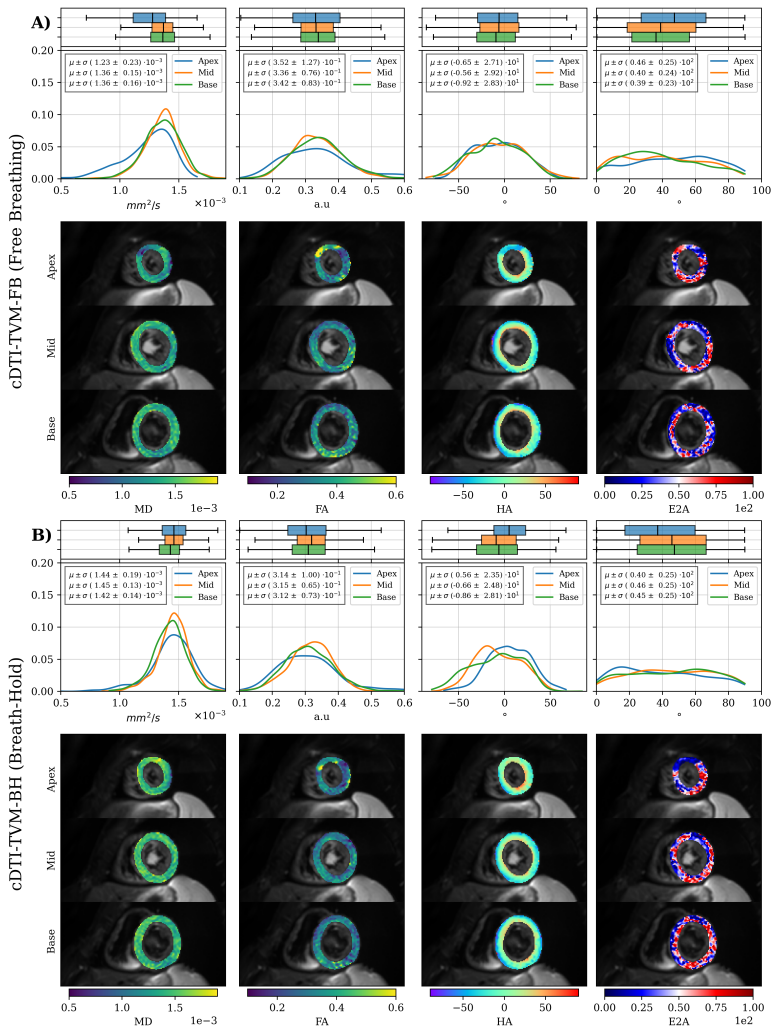


FIGURE 8.4: DTI metrics derived from A) cDTI-TVM-FB (free breathing) and the B) cDTI-TVM-BH (breath-hold) acquisition. From left to right: mean diffusivity (MD), fractional anisotropy (FA), helix angle (HA) and absolute sheetlet angle (E2A). Top rows show histogram plots for all three slices corresponding to the maps shown below. The boxplots correspond to the histograms. The mean and standard deviation for each histogram is given in the histogram axes, where the order corresponds to the legend right next to it.

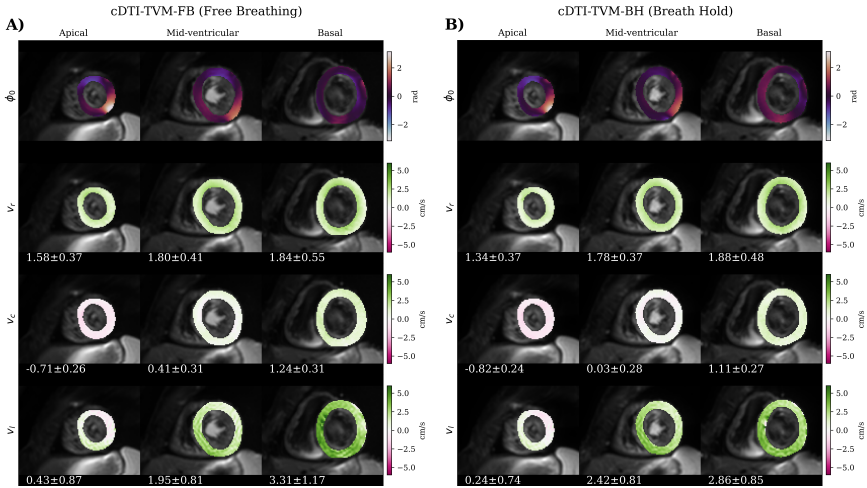


FIGURE 8.5: Background phase (top row) and velocity projection maps obtained from A) cDTI-TVM-FB (free breathing) and the B) cDTI-TVM-BH (breath-hold) acquisition. From second to last row: radial, circumferential, and longitudinal velocities. For both A) and B), columns correspond to the results for the apical (left), mid-ventricular (mid) and basal (right) slice. The numbers in the velocity plots state the mean and standard deviation of the corresponding velocity over all LV-pixels in cm s^{-1} .

velocity can be seen in the plot. TVM has higher intra-sector variability than cDTI-TVM-FB and cDTI-TVM-BH for all velocities, indicated by the larger interquartile range and fliers of the blue box plots.

Figure 8.8 shows the Bland-Altman plots comparing mean velocities for the sectors as defined in Figure 8.7 for conventional TVM versus cDTI-TVM-FB (top row), conventional TVM versus cDTI-TVM-BH (middle), and cDTI-TVM-FB versus cDTI-TVM-BH (bottom row). The points corresponding to the data from Figure 8.7 are represented by star-shaped markers. All comparisons show no systematic bias, with mean differences on the order of 1 mm s^{-1} . The limits of agreement for the comparisons against the reference method are only slightly better for the breath-hold data. Furthermore, the agreement between cDTI-TVM-BH and cDTI-TVM-FB for longitudinal velocities is better than the agreement in their respective comparison to the reference, while it is worse for the in-plane velocity components. Both comparisons of cDTI-TVM-BH and cDTI-TVM-FB relative to the conventional TVM show similar outliers, corresponding to sectors with high variability.

8.4 Discussion

Jointly estimating diffusion tensors and tissue velocities using a single sequence and without scan time penalty has been demonstrated to be feasible. The addition of the

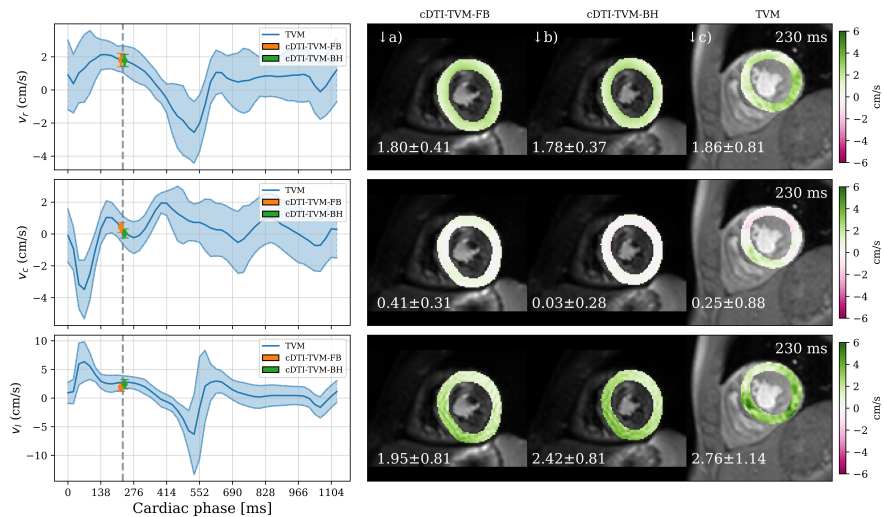


FIGURE 8.6: Comparison of radial (top), circumferential (mid) and longitudinal (bottom) velocities estimated by a) cDTI-TVM-FB (free-breathing), b) cDTI-TVM-BH (breath-hold) and c) conventional TVM for the mid-ventricular slice of a single volunteer. On the left, curves for mean and standard deviation of tissue velocities obtained by TVM are shown by the blue lines. The orange and green error bars illustrate the mean and standard deviation of the tissue velocities estimated from cDTI-TVM data corresponding to the chosen trigger delay (the markers for FB and BH are staggered for better visibility). The numbers in the velocity plots (a-c) state the mean and standard deviation of the corresponding velocity over all LV pixels in cm s^{-1} .

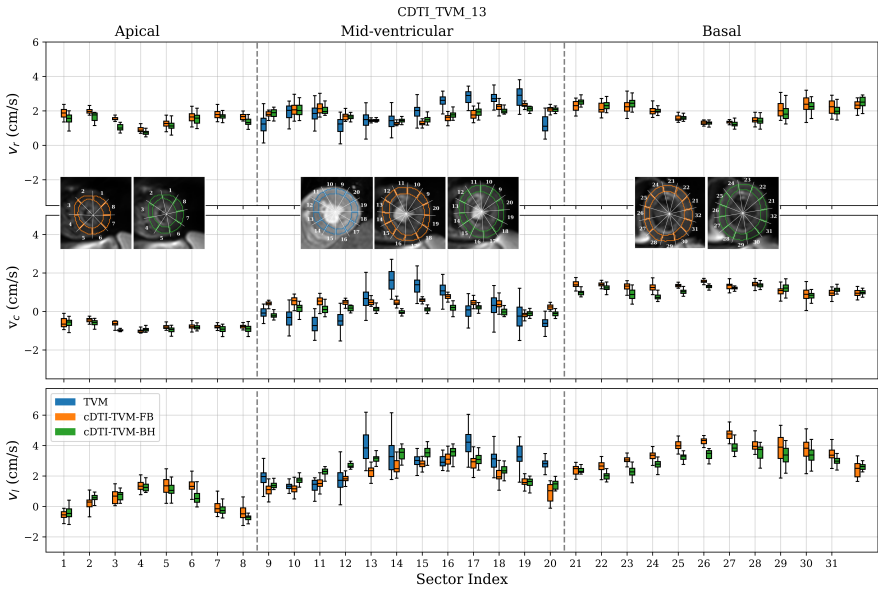


FIGURE 8.7: Boxplot of velocity projections per mask sector at a trigger delay of 230 ms for a single volunteer. The top row shows radial velocity while the mid/bottom rows show circumferential and longitudinal velocities. Sector indices plotted on the x-axis correspond to the mask areas shown in the inlets which are derived from the AHA definition for twice the number of sectors. The vertical gray dashed lines subdivide the apical/mid-ventricular/basal sector indices. Line colors in the inlets adhere to the box plot colors: conventional TVM (blue), cDTI-TVM-FB (free breathing) (orange) and cDTI-TVM-BH (breath-hold) (green). For the reference TVM (blue), only the mid-ventricular slice reference was acquired, therefore no blue boxes are shown for the apical and basal slices.

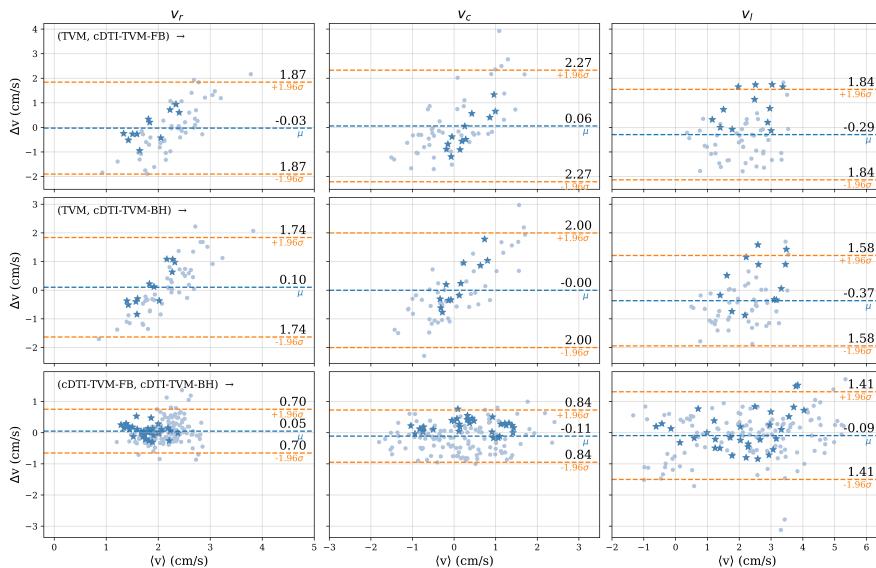


FIGURE 8.8: Bland-Altman plots comparing mean of local velocity vectors per sector for the three estimates: conventional TVM, cDTI-TVM-FB and cDTI-TVM-BH. The top row shows the comparison of TVM and cDTI-TVM-FB; the middle row shows the comparison of TVM and cDTI-TVM-BH and the bottom row provides the comparison of cDTI-TVM-FB and cDTI-TVM-BH. Columns from left to right correspond to radial (v_r), circumferential (v_c) and longitudinal (v_l) velocity. Mean and 1.96 times the standard deviation of the difference distributions are shown as blue and orange dashed lines as well as numbers in cm s^{-1} . Star-shaped markers of the scatter plot correspond to the values of the exemplary volunteer as presented in Figure 8.7.

velocity encoding gradients was shown to have negligible impact on the estimation of diffusion metrics (cf. Figure S1 left column) while measured tissue velocities agreed well with data obtained with the standard TVM method.

In-vivo, estimated diffusion metrics showed the expected characteristics such as approximately linear transmural progression of HA¹². Also, MD and FA values derived in our work were close to those reported elsewhere ($1.41 \pm 0.07 \cdot 10^{-3} \text{mm}^2 \text{s}^{-1}$ and 0.35 ± 0.03^{24} or $1.48 \pm 0.11 \cdot 10^{-3} \text{mm}^2 \text{s}^{-1}$ and 0.37 ± 0.04^{163}). Furthermore, areas of high off resonance were found close to the posterior vein of the left ventricle as well as the great cardiac vein (especially in the apical slice) in accordance to previous findings⁷⁶.

Tissue velocity estimates obtained from cDTI-TVM data showed no bias for in-plane velocities and a small global overestimation for through-plane velocities when compared to TVM as shown in Figure 8.8 and Table 2. The global averages obtained with TVM at 60% peak systole for the mid-ventricular slice agreed with published results^{122-124,270}. Furthermore, global averages of radial and circumferential velocities obtained with TVM were close to values obtained with the proposed cDTI-TVM sequence as shown in Table 2. The counter-rotation of base and apex was well captured by cDTI-TVM (Table 2, Figure S4-S8). The global average velocities in the apical and basal slices obtained with cDTI-TVM (Table 2) agreed well with literature values^{122,123,270}.

The presented simulations successfully captured effects of motion-induced phase accrual during diffusion encoding. Accordingly, the trigger delay dependency of MD due to the cardiac strain could be demonstrated (Figure 8.3). Moreover, the simulated data showed a residual motion-induced phase due to diffusion encoding²⁷⁹. As the main component of contractile motion was radial and the diffusion encoding linear, the residual motion phase resulted in a phase ramp. To this end, the proposed signal model used in the fitting algorithm contained a phase ramp term, resulting in low errors on the order of 1mm s^{-1} (cf. Figure S2).

In relation to the Bland-Altman analysis (Figure 8.8) it is noted that the variance of velocity differences between TVM and cDTI-TVM was relatively high with a single standard deviation being on the order of 1cm s^{-1} . Additionally, the cDTI-TVM velocity maps appeared smoother than the TVM maps resulting in the apparent linear dependency in Figure 8.8 for radial and circumferential velocities. In contrast to the smooth cDTI-TVM maps, the TVM maps exhibited areas of high spatial variation which could have been caused by insufficient motion correction (cf. sector 13 in Figure 8.6 and 8.7). The standard deviation for the global TVM velocity estimates were relatively high compared to the corresponding standard deviation reported in¹²³. This may in parts be attributed to imperfect masks used with TVM, causing high variance in diastole in Figure 8.7 as border regions may have created outliers (compare supplementary material S6). Excluding these outliers will likely reduce the apparent cDTI-TVM overestimation of through-plane velocity. On the other

hand, as computing strain-rate tensors involves evaluating the spatial derivative of velocities, the smoothing due to the averaging in the presented fitting approach for cDTI-TVM data might remove potential local spatial variations. Furthermore, the quality of spatial derivatives depends on image resolution, slice positions and gaps²⁷¹. Optimization of imaging parameters as well as increasing LV coverage^{98,280} (e.g. by using simultaneous multi-slice imaging^{11,281}) remains to be accomplished in future work.

A systematic underestimation of through-plane tissue velocities was noted for both simulated as well as in-vivo data. The calibration to address this effect was based on simulated data which was applied to the in-vivo data. However, the compensation remained incomplete. A potential cause for the residual underestimation is thought to be related to the non-integer and directionally-coupled phase aliasing of Hadamard encoding¹¹³. Therefore, if compatible with the sequence timing and system limits, it is advised to use a standard referenced 4-point velocity encoding scheme¹¹³ along the Cartesian axes. As the duration of the gap between excitation and start of the diffusion encoding waveforms is determined by the EPI train length^{12,29,30}, the addition of stronger bipolar velocity encoding gradients is feasible even when longer 2D spatial excitation pulses are used⁹⁸. However, when utilizing partial Fourier acquisition, the gap might become too short to insert standard referenced velocity encoding gradients.

Since Mo12-SE cDTI was performed in mid or late systole^{169,279}, tissue velocity estimates were obtained for a single time point of the cardiac cycle only, which is a limitation. Furthermore, a minimum of four averages is required for cDTI-TVM. When using the diffusion gradient inversion as proposed in this work, the minimum number of averages increases to eight, which can be limiting.

8.5 Conclusion

The proposed sequence allows for joint estimation of cardiac diffusion tensors and tissue velocities in the myocardium during free breathing. The approach offers a single-scan solution to study the interrelation of cardiac microstructure and macroscopic tissue motion.

Discussion and Outlook

This thesis has addressed potential improvements in three aspects of cDTI. As a variety of methodological advances in acquisition and processing of cDTI were proposed, finding a consensus on preferred methods is required^{5,6,26}. MR simulations are an established tool to facilitate a quantitative comparison of such methods⁵⁹ but are required to capture all relevant effects. A flexible representation of motion is key to simulating MR signals of the cardiovascular anatomy. To this end, an open-source framework for MR simulation, including complex motion and flow, has been developed. The framework's capability has been demonstrated using a simulation of turbulent flow encoding, as well as a SE-cDTI acquisition with contractile and respiratory motion. Both yielded relevant image characteristics such as motion-induced phase and magnitude modulation. Furthermore, the possibility of using a wide range of digital phantoms as inputs was demonstrated, allowing to use digital twins and tissue models from patient data^{189,243,282,283} as phantoms for MR simulations.

As reproducibility and extensibility of simulation experiments also depend on the usability and documentation of the software, emphasis was put on creating the framework according to software-quality standards. Alongside the simulation package, an MR-sequence definition package was implemented, to contribute vendor-agnostic MR-sequence development²¹¹⁻²¹⁴.

The application of deep learning algorithms helps reducing the computational burden for registration and LV segmentation⁵⁰. Furthermore, the increased robustness against noise and artifacts potentially facilitates the reduction of required data while preserving the accuracy of the estimated diffusion metrics^{51,53,54}. However, high-quality training data is scarce, and generalization between cDTI data from substantially different scan methods is not given. Utilizing in-silico data, this constraint on training networks for new methodology can be alleviated. To this end, the feasibility of training a deep learning model on synthetic data to infer diffusion tensors directly from unregistered cDTI data was shown. The pipeline of data generation includes a random sampling of diffusion tensors with spatial coherence, allowing to control the prior of MD and FA to be box-constrained within intervals corresponding to healthy and diseased tissue. The network showed superior tensor inference on simulated test data as well as in-vivo data. Moreover, the inclusion of spatial information allowed the network to maintain inference accuracy when the averages were reduced. However, the simulation did not include off-resonance or eddy current effects. Their inclusion will likely improve the estimation performance; therefore, obtaining a sufficiently accurate description of the off-resonance fields containing blood flow and motion could be beneficial^{177,178}.

As macroscopic cardiac function is directly determined by the state of cardiac microstructure, combined measurements of macroscopic strain and diffusion data are a promising approach to obtain clinically valuable biomarkers^{34,35,37,38}. However, acquiring multiple scans increases the total protocol duration, and co-registering can be challenging. Therefore, the last part of this thesis introduced a joint encoding strategy of tissue velocity and diffusion tensors based on SE-cDTI without prolonging TE. In the proposed sequence, a bipolar gradient was inserted to encode the tissue velocity in the image phase. The proposed fitting algorithm was calibrated and tested on simulated data and was capable of estimating the tissue velocities in-vivo in agreement with a reference acquisition. At the same time, the estimated diffusion metrics were unaffected by the sequence adaptation. By evaluating the spatial derivative of the velocity, tissue strain can be obtained, which directly offers clinically valuable information on local tissue contractility without extra scan duration. This requires sufficient slice coverage in the longitudinal direction as well as in-plane resolution. Therefore, the optimization of imaging parameters in cDTI acquisitions could also include the improvement of velocity encoding schemes and related parameters.

To solidify the foundation of cDTI, future work could be dedicated to implementing the variety of cDTI acquisition methods as vendor-agnostic sequences. Furthermore, standardized digital reference objects could be defined, to facilitate comparative simulation experiments. Furthermore, based on these simulations, the optimization of experimental design or even online scan-adaptation with i.e. reinforcement-learning could be pursued. Using the image phase to encode macroscopic motion could also be part of this approach. Another approach to potentially improve 3D coverage and leverage spatial information in cDTI is to use a digital twin model for tensor inference. On the personalized cardiac meshes, (graph-) neural networks could be used to fuse diffusion and strain-rate data with additional tissue characterization, thereby possibly becoming more data efficient in diagnosing cardiac diseases. Using synthetic data to train such models could facilitate this approach. Although higher field strength are able to improve spatial resolution and SNR, accessibility is favoured in lower field systems. In combination with an increased data efficiency using DL algorithms, this could offer a path to increase the usability of cDTI.

References

1. Ismail, T. F., Strugnell, W., Coletti, C., Božić-Iven, M., *et al.* Cardiac MR: From Theory to Practice. *Front. Cardiovasc. Med.* **9**, 826283. ISSN: 2297055X (2022).
2. Seetharam, K. & Lerakis, S. Cardiac magnetic resonance imaging: the future is bright. *F1000Research* **8**, 1636. ISSN: 2046-1402 (2019).
3. Kramer, C. M., Barkhausen, J., Bucciarelli-Ducci, C., Flamm, S. D., *et al.* Standardized cardiovascular magnetic resonance imaging (CMR) protocols: 2020 update. *J. Cardiovasc. Magn. Reson.* **22**, 1. ISSN: 1532429X (2020).
4. Russo, V., Lovato, L. & Ligabue, G. Cardiac MRI: technical basis. *Radiol. Medica* **125**, 1040. ISSN: 18266983 (2020).
5. Paddock, S., Tsampasian, V., Assadi, H., Mota, B. C., *et al.* Clinical Translation of Three-Dimensional Scar, Diffusion Tensor Imaging, Four-Dimensional Flow, and Quantitative Perfusion in Cardiac MRI: A Comprehensive Review. *Front. Cardiovasc. Med.* **8**, 682027. ISSN: 2297055X (2021).
6. Guo, R., Weingärtner, S., Šiurytė, P., T. Stoeck, C., *et al.* Emerging Techniques in Cardiac Magnetic Resonance Imaging. *J. Magn. Reson. Imaging* **55**, 1043. ISSN: 1522-2586 (2022).
7. Stikov, N. & Karakuzu, A. The relaxometry hype cycle. *Front. Physiol.* **14**, 1281147. ISSN: 1664042X (2023).
8. Scatteia, A., Baritussio, A. & Bucciarelli-Ducci, C. Strain imaging using cardiac magnetic resonance. *Heart Fail. Rev.* **22**, 465. ISSN: 1573-7322 (2017).
9. Rajiah, P. S., Kalisz, K., Broncano, J., Goerne, H., *et al.* Myocardial Strain Evaluation with Cardiovascular MRI: Physics, Principles, and Clinical Applications. *Radiographics* **42**, 968. ISSN: 15271323 (2022).
10. Chitiboi, T. & Axel, L. Magnetic resonance imaging of myocardial strain: A review of current approaches. *J. Magn. Reson. Imaging* **46**, 1263. ISSN: 1053-1807 (2017).
11. Mekkaoui, C., Reese, T. G., Jackowski, M. P., Bhat, H., *et al.* Diffusion MRI in the heart. *NMR Biomed* **30** (2017).
12. Nielles-Vallespin, S., Scott, A., Ferreira, P. F., Khalique, Z., *et al.* Cardiac Diffusion: Technique and Practical Applications. *J Magn Reson Imaging* (2019).
13. Teh, I., McClymont, D., Zdora, M. C., Whittington, H. J., *et al.* Validation of diffusion tensor MRI measurements of cardiac microstructure with structure tensor synchrotron radiation imaging. *J. Cardiovasc. Magn. Reson.* **19**, 1. ISSN: 1532429X (2017).
14. Chen, J., Song, S. K., Liu, W., McLean, M., *et al.* Remodeling of cardiac fiber structure after infarction in rats quantified with diffusion tensor MRI. *Am. J. Physiol. - Hear. Circ. Physiol.* **285**, 946. ISSN: 03636135 (2003).

15. Nielles-Vallespin, S., Khalique, Z., Ferreira, P. F., de Silva, R., *et al.* Assessment of Myocardial Microstructural Dynamics by In Vivo Diffusion Tensor Cardiac Magnetic Resonance. *J. Am. Coll. Cardiol.* **69**, 661. ISSN: 0735-1097 (2017).
16. McGill, L., Ferreira, P. F., Scott, A., Nielles-Vallespin, S., *et al.* Relationship between cardiac diffusion tensor imaging parameters and anthropometrics in healthy volunteers. *eng. J. Cardiovasc. Magn. Reson.* **18**, 2. ISSN: 1532-429X (2015).
17. Sosnovik, D. E., Mekkaoui, C., Huang, S., Chen, H. H., *et al.* Microstructural Impact of Ischemia and Bone Marrow-Derived Cell Therapy Revealed with Diffusion Tensor MRI Tractography of the Heart In Vivo. *Circulation* **129**, 1731. ISSN: 15244539 (2014).
18. Ferreira, P. F., Kilner, P. J., McGill, L. A., Nielles-Vallespin, S., *et al.* In vivo cardiovascular magnetic resonance diffusion tensor imaging shows evidence of abnormal myocardial laminar orientations and mobility in hypertrophic cardiomyopathy. *J Cardiovasc Magn Reson* **16**, 87 (2014).
19. Von Deuster, C., Sammut, E., Asner, L., Nordsletten, D., *et al.* Studying Dynamic Myofiber Aggregate Reorientation in Dilated Cardiomyopathy Using In Vivo Magnetic Resonance Diffusion Tensor Imaging. *Circ Cardiovasc Imaging* **9** (2016).
20. Das, A., Chowdhary, A., Kelly, C., Teh, I., *et al.* Insight Into Myocardial Microstructure of Athletes and Hypertrophic Cardiomyopathy Patients Using Diffusion Tensor Imaging. *J. Magn. Reson. Imaging* **53**, 73. ISSN: 1053-1807 (2021).
21. Das, A., Kelly, C., Teh, I., Stoeck, C. T., *et al.* Pathophysiology of LV Remodeling Following STEMI. *JACC Cardiovasc. Imaging* **16**, 159. ISSN: 1936878X (2023).
22. Sosnovik, D. E. Magnetic Resonance-Based Characterization of Myocardial Architecture. *Heart Fail. Clin.* **17**, 85. ISSN: 1551-7136 (2021).
23. Khalique, Z., Ferreira, P. F., Scott, A. D., Nielles-Vallespin, S., *et al.* Diffusion Tensor Cardiovascular Magnetic Resonance in Cardiac Amyloidosis. *Circ. Cardiovasc. Imaging* **13**, e009901. ISSN: 1941-9651 (2020).
24. Gotschy, A., Von Deuster, C., van Gorkum, R. J. H., Gastl, M., *et al.* Characterizing cardiac involvement in amyloidosis using cardiovascular magnetic resonance diffusion tensor imaging. *J. Cardiovasc. Magn. Reson.* **21**, 56. ISSN: 1532-429X (2019).
25. Gotschy, A., von Deuster, C., Weber, L., Gastl, M., *et al.* CMR Diffusion Tensor Imaging Provides Novel Imaging Markers of Adverse Myocardial Remodeling in Aortic Stenosis. *JACC Cardiovasc. Imaging* **14**, 1472. ISSN: 1936-878X (2021).
26. Khalique, Z., Ferreira, P. F., Scott, A. D., Nielles-Vallespin, S., *et al.* Diffusion Tensor Cardiovascular Magnetic Resonance Imaging. *JACC Cardiovasc. Imaging* **13**, 1235. ISSN: 1936878X (2020).
27. Novikov, D. S., Kiselev, V. G. & Jespersen, S. N. On modeling. *Magn Reson Med* **79**, 3172 (2018).

28. Soares, J. M., Marques, P., Alves, V. & Sousa, N. A hitchhiker's guide to diffusion tensor imaging. *Front. Neurosci.* **7**, 38546. ISSN: 16624548 (2013).
29. Stoeck, C. T., von Deuster, C., Genet, M., Atkinson, D., *et al.* Second-order motion-compensated spin echo diffusion tensor imaging of the human heart. *Magn Reson Med* **75**, 1669 (2016).
30. Welsh, C. L., Dibella, E. V. R. & Hsu, E. W. Higher-Order Motion-Compensation for In Vivo Cardiac Diffusion Tensor Imaging in Rats. *IEEE Trans. Med. Imaging* **34**, 1843. ISSN: 1558-254X VO - 34 (2015).
31. Van Gorkum, R. J., Guenther, C., Koethe, A., Stoeck, C. T., *et al.* Characterization and correction of diffusion gradient-induced eddy currents in second-order motion-compensated echo-planar and spiral cardiac DTI. *Magn. Reson. Med.* **88**, 2378. ISSN: 1522-2594 (2022).
32. Van Gorkum, R. J., von Deuster, C., Guenther, C., Stoeck, C. T., *et al.* Analysis and correction of off-resonance artifacts in echo-planar cardiac diffusion tensor imaging. *Magn. Reson. Med.* **84**, 2561. ISSN: 15222594 (2020).
33. Brady, B., King, G., Murphy, R. T. & Walsh, D. Myocardial strain: a clinical review. *Ir. J. Med. Sci.* **192**, 1. ISSN: 18634362 (2023).
34. Wilson, A. J., Sands, G. B., LeGrice, I. J., Young, A. A., *et al.* Myocardial Mesostructure and Mesofunction. *Am. J. Physiol. - Hear. Circ. Physiol.* **323**. ISSN: 15221539 (2022).
35. Moulin, K., Croisille, P., Viallon, M., Verzhbinsky, I. A., *et al.* Myofiber strain in healthy humans using DENSE and cDTI. *Magn. Reson. Med.* **86**, 277. ISSN: 1522-2594 (2021).
36. Perotti, L. E., Magrath, P., Verzhbinsky, I. A., Aliotta, E., *et al.* Microstructurally anchored cardiac kinematics by combining in vivo dense MRI and cDTI. *Lect. Notes Comput. Sci.* **10263**, 381. ISSN: 16113349 (2017).
37. Verzhbinsky, I. A., Perotti, L. E., Moulin, K., Cork, T. E., *et al.* Estimating Aggregate Cardiomyocyte Strain Using In-Vivo Diffusion and Displacement Encoded MRI. *IEEE Trans. Med. Imaging* **39**, 656. ISSN: 1558254X (2020).
38. Perotti, L. E., Verzhbinsky, I. A., Moulin, K., Cork, T. E., *et al.* Estimating cardiomyofiber strain in vivo by solving a computational model. *Med. Image Anal.* **68**, 101932. ISSN: 1361-8415 (2021).
39. Sermesant, M., Delingette, H., Cochet, H., Jais, P., *et al.* Applications of artificial intelligence in cardiovascular imaging. *Nat. Rev. Cardiol.* **2021 188 18**, 600. ISSN: 1759-5010 (2021).
40. Knoll, F., Hammernik, K., Zhang, C., Moeller, S., *et al.* Deep-Learning Methods for Parallel Magnetic Resonance Imaging Reconstruction: A Survey of the Current Approaches, Trends, and Issues. *IEEE Signal Process. Mag.* **37**, 128. ISSN: 15580792 (2020).
41. Haskins, G., Kruger, U. & Yan, P. Deep learning in medical image registration: a survey. *Mach. Vis. Appl.* **31**, 8 (2020).

42. Khan, M. Z., Gajendran, M. K., Lee, Y. & Khan, M. A. Deep Neural Architectures for Medical Image Semantic Segmentation: Review. *IEEE Access* **9**, 83002. ISSN: 21693536 (2021).
43. Du, G., Cao, X., Liang, J., Chen, X., *et al.* Medical image segmentation based on U-Net: A review. *J. Imaging Sci. Technol.* **64**. ISSN: 19433522 (2020).
44. Alabed, S., Maiter, A., Salehi, M., Mahmood, A., *et al.* Quality of reporting in AI cardiac MRI segmentation studies – A systematic review and recommendations for future studies. *Front. Cardiovasc. Med.* **9**, 956811. ISSN: 2297055X (2022).
45. Chen, C., Qin, C., Qiu, H., Tarroni, G., *et al.* Deep Learning for Cardiac Image Segmentation: A Review. *Front. Cardiovasc. Med.* **7**, 508599. ISSN: 2297055X. arXiv: 1911.03723 (2020).
46. Mishro, P. K., Agrawal, S., Panda, R. & Abraham, A. A Survey on State-of-the-Art Denoising Techniques for Brain Magnetic Resonance Images. *IEEE Rev. Biomed. Eng.* **15**, 184. ISSN: 19411189 (2022).
47. Faiyaz, A., Doyley, M. M., Schifitto, G. & Uddin, M. N. Artificial intelligence for diffusion MRI-based tissue microstructure estimation in the human brain: an overview. *Front. Neurol.* **14**, 1168833. ISSN: 16642295 (2023).
48. Loktyushin, A., Herz, K., Dang, N., Glang, F., *et al.* MRzero - Automated discovery of MRI sequences using supervised learning. *Magn. Reson. Med.* **86**, 709. ISSN: 1522-2594. arXiv: 2002.04265 (2021).
49. Perlman, O., Zhu, B., Zaiss, M., Rosen, M. S., *et al.* An end-to-end AI-based framework for automated discovery of rapid CEST/MT MRI acquisition protocols and molecular parameter quantification (AutoCEST). *Magn. Reson. Med.* **87**, 2792. ISSN: 1522-2594 (2022).
50. Ferreira, P. F., Martin, R. R., Scott, A. D., Khalique, Z., *et al.* Automating in vivo cardiac diffusion tensor postprocessing with deep learning-based segmentation. *Magn. Reson. Med.* **84**, 2801. ISSN: 15222594 (2020).
51. Ferreira, P. F., Banerjee, A., Scott, A. D., Khalique, Z., *et al.* Accelerating Cardiac Diffusion Tensor Imaging With a U-Net Based Model: Toward Single Breath-Hold (2022).
52. Aliotta, E., Nourzadeh, H., Sanders, J., Muller, D., *et al.* Highly accelerated, model-free diffusion tensor MRI reconstruction using neural networks. *Med. Phys.* **46**, 1581. ISSN: 0094-2405 (2019).
53. Phipps, K., Van De Boomen, M., Eder, R., Michelhaugh, S. A., *et al.* Accelerated in Vivo Cardiac Diffusion-Tensor MRI Using Residual Deep Learning-based Denoising in Participants with Obesity. *Radiol. Cardiothorac. Imaging* **3**, e200580. ISSN: 26386135 (2021).
54. Tänzer, M., Ferreira, P. F., Scott, A., Khalique, Z., *et al.* Faster Diffusion Cardiac MRI with Deep Learning-Based Breath Hold Reduction. *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)* **13413**, 101. ISSN: 16113349. arXiv: 2206.10543 (2022).

55. Shimron, E., Tamir, J. I., Wang, K. & Lustig, M. Implicit data crimes: Machine learning bias arising from misuse of public data. *Proc. Natl. Acad. Sci. U. S. A.* **119**, e2117203119. ISSN: 1091-6490 (2022).
56. Mårtensson, G., Ferreira, D., Granberg, T., Cavallin, L., *et al.* The reliability of a deep learning model in clinical out-of-distribution MRI data: A multicohort study. *Med. Image Anal.* **66**, 101714. ISSN: 1361-8415. arXiv: 1911.00515 (2020).
57. Kiselev, V. G. Fundamentals of diffusion MRI physics. *NMR Biomed.* **30**. ISSN: 10991492 (2017).
58. Chen, R. J., Lu, M. Y., Chen, T. Y., Williamson, D. F. K., *et al.* Synthetic data in machine learning for medicine and healthcare. *Nat. Biomed. Eng.* **5**, 493. ISSN: 2157-846X (2021).
59. Veldmann, M., Ehses, P., Chow, K., Nielsen, J.-F. F., *et al.* Open-Source MR Imaging and Reconstruction Workflow. *Magn. Reson. Med.* **88**, 2395. ISSN: 1522-2594 (2022).
60. Stöcker, T., Vahedipour, K., Pflugfelder, D. & Shah, N. J. High-performance computing MRI simulations. *Magn. Reson. Med.* **64**, 186. ISSN: 1522-2594 (2010).
61. Castillo-Passi, C., Coronado, R., Varela-Mattatall, G., Alberola-López, C., *et al.* KomaMRI.jl: An open-source framework for general MRI simulations with GPU acceleration. *Magn. Reson. Med.* **90**, 329. ISSN: 1522-2594. arXiv: 2301.02702 (2023).
62. Liu, F., Velikina, J. V., Block, W. F., Kijowski, R., *et al.* Fast Realistic MRI Simulations Based on Generalized Multi-Pool Exchange Tissue Model. *IEEE Trans. Med. Imaging* **36**, 527. ISSN: 1558254X (2017).
63. Xanthis, C. G., Venetis, I. E. & Aletras, A. H. High performance MRI simulations of motion on multi-GPU systems. *J. Cardiovasc. Magn. Reson.* **16**, 1. ISSN: 1532429X (2014).
64. Xanthis, C. G., Venetis, I. E., Chalkias, A. V. & Aletras, A. H. MRISIMUL: A GPU-based parallel approach to MRI simulations. *IEEE Trans. Med. Imaging* **33**, 607. ISSN: 02780062 (2014).
65. Benoit-Cattin, H., Collewet, G., Belaroussi, B., Saint-Jalmes, H., *et al.* The SIMRI project: a versatile and interactive MRI simulator. *J. Magn. Reson.* **173**, 97. ISSN: 1090-7807 (2005).
66. Jochimsen, T. H. & Von Mengershausen, M. ODIN-Object-oriented development interface for NMR.
67. Kose, R. & Kose, K. BlochSolver: A GPU-optimized fast 3D MRI simulator for experimentally compatible pulse sequences (2017).
68. Kose, R., Setoi, A. & Kose, K. A Fast GPU-optimized 3D MRI Simulator for Arbitrary k-space Sampling. *Magn. Reson. Med. Sci.* **18**, 208. ISSN: 1880-2206 (2019).
69. Drobnyak, I., Gavaghan, D., Süli, E., Pitt-Francis, J., *et al.* Development of a functional magnetic resonance imaging simulator for modeling realistic rigid-body motion artifacts. *Magn. Reson. Med.* **56**, 364. ISSN: 1522-2594 (2006).

70. Niederer, S. A., Aboelkassem, Y., Cantwell, C. D., Corrado, C., *et al.* Creation and application of virtual patient cohorts of heart models. *Philos. Trans. R. Soc. A* **378**. ISSN: 1364503X (2020).
71. Pesapane, F., Rotili, A., Penco, S., Nicosia, L., *et al.* Digital Twins in Radiology. *J. Clin. Med.* **11**, 6553. ISSN: 20770383 (2022).
72. Britannica, T. E. o. E. *Heart*
73. Cerqueira, M. D., Weissman, N. J., Dilsizian, V., Jacobs, A. K., *et al.* Standardized Myocardial Segmentation and Nomenclature for Tomographic Imaging of the Heart. *Circulation* **105**, 539. ISSN: 0009-7322 (2002).
74. Iuzzo, P. A. Handbook of cardiac anatomy, physiology, and devices: Second edition. *Handb. Card. Anatomy, Physiol. Devices Second Ed.*, 1 (2005).
75. Goodwill, A. G., Dick, G. M., Kiel, A. M. & Tune, J. D. in *Compr. Physiol.* **2**, 321 (Wiley, 2017).
76. Reeder, S. B., Faranesh, A. Z., Boxerman, J. L. & McVeigh, E. R. In vivo measurement of T2(*) and field inhomogeneity maps in the human heart at 1.5 T. *Magn. Reson. Med.* **39**, 988. ISSN: 07403194 (1998).
77. Pope, A. J., Sands, G. B., Smaill, B. H. & LeGrice, I. J. Three-dimensional transmural organization of perimysial collagen in the heart. *Am J Physiol Hear. Circ Physiol* **295**, H1243. ISSN: 03636135 (2008).
78. Satoh, H., Delbridge, L. M., Blatter, L. A. & Bers, D. M. Surface:volume relationship in cardiac myocytes studied with confocal microscopy and membrane capacitance measurements: species-dependence and developmental effects. *Biophys J* **70**, 1494 (1996).
79. Olivetti, G., Cigola, E., Maestri, R., Corradi, D., *et al.* Aging, cardiac hypertrophy and ischemic cardiomyopathy do not affect the proportion of mononucleated and multinucleated myocytes in the human heart. *J Mol Cell Cardiol* **28**, 1463 (1996).
80. Gilbert, S. H., Benson, A. P., Li, P. & Holden, A. V. Regional localisation of left ventricular sheet structure: integration with current models of cardiac fibre, sheet and band structure. *Eur. J. Cardio-Thoracic Surg.* **32**, 231. ISSN: 1010-7940 (2007).
81. Smerup, M., Nielsen, E., Agger, P., Frandsen, J., *et al.* The Three-Dimensional Arrangement of the Myocytes Aggregated Together Within the Mammalian Ventricular Myocardium. *Anat. Rec. Adv. Integr. Anat. Evol. Biol.* **292**, 1. ISSN: 1932-8494 (2009).
82. Sengupta, P. P., Jamil Tajik, A., Chandrasekaran, K. & Khandheria, B. K. *Twist Mechanics of the Left Ventricle Principles and Application* (2008).
83. Axel, L., Wedeen, V. J. & Ennis, D. B. Probing dynamic myocardial microstructure with cardiac magnetic resonance diffusion tensor imaging. *J Cardiovasc Magn Reson* **16**, 89 (2014).
84. Harrington, K. B., Rodriguez, F., Cheng, A., Langer, F., *et al.* Direct measurement of transmural laminar architecture in the anterolateral wall of the ovine

- left ventricle: new implications for wall thickening mechanics. *Am. J. Physiol. Heart Circ. Physiol.* **288**, H1324. ISSN: 03636135 (2005).
85. Kilner, P. J., McCarthy, K., Murillo, M., Ferreira, P. F., *et al.* Histology of human myocardial laminar microstructure and consideration of its cyclic deformations with respect to interpretation of in vivo cardiac diffusion tensor imaging. *J. Cardiovasc. Magn. Reson.* **17**, Q10. ISSN: 1532-429X (2015).
 86. Sánchez-Ferrero, G. V., Vega, A. T., Grande, L. C., de la Higuera, P. C., *et al.* in, 355 (Springer, London, 2009). ISBN: 978-1-84882-299-3.
 87. Nasiraei-Moghaddam, A. & Gharib, M. Evidence for the existence of a functional helical myocardial band. *Am. J. Physiol. - Hear. Circ. Physiol.* **296**, 127. ISSN: 03636135 (2009).
 88. Pedrizzetti, G., Kraigher-Krainer, E., De Luca, A., Caracciolo, G., *et al.* Functional Strain-Line Pattern in the Human Left Ventricle. *Phys. Rev. Lett.* **109**, 048103. ISSN: 0031-9007 (2012).
 89. Voigt, J. U. & Cvijic, M. 2- and 3-Dimensional Myocardial Strain in Cardiac Health and Disease. *JACC Cardiovasc. Imaging* **12**, 1849. ISSN: 1936-878X (2019).
 90. Van Heeswijk, R. B., Bonanno, G., Coppo, S., Coristine, A., *et al.* Motion Compensation Strategies in Magnetic Resonance Imaging. *Crit. Rev. TM Biomed. Eng.* **40**, 99 (2012).
 91. Firmin, D. & Keegan, J. Use of Navigator Echoes in Cardiovascular Magnetic Resonance and Factors Affecting Their Implementation. *Cardiovasc. Magn. Reson.*, 129 (2010).
 92. Wang, Y., Riederer, S. J. & Ehman, R. L. Respiratory Motion of the Heart: Kinematics and the Implications for the Spatial Resolution in Coronary Imaging. *Magn. Reson. Med.* **33**, 713. ISSN: 15222594 (1995).
 93. Nguyen, T. D., Nuval, A., Mulukutla, S. & Wang, Y. Direct monitoring of coronary artery motion with cardiac fat navigator echoes. *Magn. Reson. Med.* **50**, 235. ISSN: 1522-2594 (2003).
 94. Stehning, C., Börnert, P., Nehrke, K., Eggers, H., *et al.* Free-breathing whole-heart coronary MRA with 3D radial SSFP and self-navigated image reconstruction. *Magn. Reson. Med.* **54**, 476. ISSN: 1522-2594 (2005).
 95. Khalil, A., Ng, S. C., Liew, Y. M. & Lai, K. W. *An Overview on Image Registration Techniques for Cardiac Diagnosis and Treatment* 2018.
 96. Klein, S., Staring, M., Murphy, K., Viergever, M. A., *et al.* elastix: A Toolbox for Intensity-Based Medical Image Registration. *IEEE Trans. Med. Imaging* **29**, 196. ISSN: 1558-254X VO - 29 (2010).
 97. Vishnevskiy, V., Gass, T., Székely, G., Goksel, O., *et al.* *Total Variation Regularization of Displacements in Parametric Image Registration* in (Springer International Publishing, Cham, 2014), 211. ISBN: 978-3-319-13692-9.
 98. Nguyen, C. T., Christodoulou, A. G., Coll-Font, J., Ma, S., *et al.* Free-breathing diffusion tensor MRI of the whole left ventricle using second-order motion compensation and multitasking respiratory motion correction. *Magn. Reson. Med.* **85**, 2634. ISSN: 0740-3194 (2021).

99. Coll-Font, J., Chen, S., Eder, R., Fang, Y., *et al.* Manifold-based respiratory phase estimation enables motion and distortion correction of free-breathing cardiac diffusion tensor MRI. *Magn. Reson. Med.* **n/a**. ISSN: 0740-3194 (2021).
100. Gil, C. S., Meredith, S. D. & Curran, K. M. Full tensor registration of diffusion tensor magnetic resonance imaging for assessment of cardiac pathologies. *J. Cardiovasc. Magn. Reson.* **14**, W41. ISSN: 1532-429X (2012).
101. Aletras, A. H., Ding, S., Balaban, R. S. & Wen, H. DENSE: Displacement Encoding with Stimulated Echoes in Cardiac Functional MRI. *J. Magn. Reson.* **137**, 247. ISSN: 10907807 (1999).
102. Liu, Z., Liu, W. & Clarysse, P. *Cardiac strain measurement with DENSE MRI: a review in 2022 16th IEEE Int. Conf. Signal Process.* **2022-October** (IEEE, 2022), 474. ISBN: 978-1-6654-6056-9.
103. Aletras, A. H. & Wen, H. Mixed echo train acquisition displacement encoding with stimulated echoes: An optimized DENSE method for in vivo functional imaging of the human heart. *Magn. Reson. Med.* **46**, 523. ISSN: 0740-3194 (2001).
104. Ibrahim, E.-S. H. Myocardial tagging by Cardiovascular Magnetic Resonance: evolution of techniques-pulse sequences, analysis algorithms, and applications (2011).
105. Axel, L. & Dougherty, L. MR imaging of motion with spatial modulation of magnetization. *Radiology* **171**, 841. ISSN: 0033-8419 (1989).
106. Axel, L. & Dougherty, L. Heart wall motion: improved method of spatial modulation of magnetization for MR imaging. *Radiology* **172**, 349. ISSN: 0033-8419 (1989).
107. Mosher, T. J. & Smith, M. B. A DANTE tagging sequence for the evaluation of translational sample motion. *Magn. Reson. Med.* **15**, 334. ISSN: 1522-2594 (1990).
108. Fischer, S. E., McKinnon, G. C., Maier, S. E. & Boesiger, P. Improved myocardial tagging contrast. *Magn. Reson. Med.* **30**, 191. ISSN: 1522-2594 (1993).
109. Rutz, A. K., Ryf, S., Plein, S., Boesiger, P., *et al.* Accelerated whole-heart 3D CSPAMM for myocardial motion quantification. *Magn. Reson. Med.* **59**, 755. ISSN: 0740-3194 (2008).
110. Pedrizzetti, G., Claus, P., Kilner, P. J. & Nagel, E. Principles of cardiovascular magnetic resonance feature tracking and echocardiographic speckle tracking for informed clinical use. *J. Cardiovasc. Magn. Reson.* **18**, 51. ISSN: 1532-429X (2016).
111. Bernstein, M. A., King, K. F. & Zhou, X. J. *Handbook of MRI Pulse Sequences* ISBN: 9780080533124 (Elsevier Science & Technology, San Diego, UNITED STATES, 2004).
112. Bryant, D. J., Payne, J. A., Firmin, D. N. & Longmore, D. B. Measurement of Flow with NMR Imaging Using a Gradient Pulse and Phase Difference Technique. *J. Comput. Assist. Tomogr.* **8**, 588. ISSN: 0363-8715 (1984).

113. Pelc, N. J., Bernstein, M. A., Shimakawa, A. & Glover, G. H. Encoding strategies for three-direction phase-contrast MR imaging of flow. *J. Magn. Reson. Imaging* **1**, 405. ISSN: 1053-1807 (1991).
114. Dumoulin, C. L., Souza, S. P., Darrow, R. D., Pelc, N. J., *et al.* Simultaneous acquisition of phase-contrast angiograms and stationary-tissue images with Hadamard encoding of flow-induced phase shifts. *J. Magn. Reson. Imaging* **1**, 399. ISSN: 1053-1807 (1991).
115. Guenthner, C., Runge, J. H., Sinkus, R. & Kozerke, S. Analysis and improvement of motion encoding in magnetic resonance elastography. *NMR Biomed.* **31**. ISSN: 0952-3480 (2018).
116. Haase, A., Frahm, J., Matthaei, D. & Merboldt, K.-D. FLASH Imaging. Rapid NMR Imaging Using Low Flip-Angle Pulses. *J. Magn. Reson.* **67**, 258 (1986).
117. Hennig, J., Schneider, B., Peschl, S., Markl, M., *et al.* Analysis of myocardial motion based on velocity measurements with a black blood prepared segmented gradient-echo sequence: Methodology and applications to normal volunteers and patients. *J. Magn. Reson. Imaging* **8**, 868. ISSN: 10531807 (1998).
118. Markl, M., Schneider, B., Hennig, J., Peschl, S., *et al.* Cardiac phase contrast gradient echo MRI: Measurement of myocardial wall motion in healthy volunteers and patients. *Int. J. Card. Imaging* **15**, 441. ISSN: 01679899 (1999).
119. Markl, M. & Hennig, J. Phase contrast MRI with improved temporal resolution by view sharing: k-space related velocity mapping properties. *Magn. Reson. Imaging* **19**, 669. ISSN: 0730725X (2001).
120. Markl, M., Schneider, B. & Hennig, J. Fast phase contrast cardiac magnetic resonance imaging: Improved assessment and analysis of left ventricular wall motion. *J. Magn. Reson. Imaging* **15**, 642. ISSN: 1522-2586 (2002).
121. Jung, B. A., Kreher, B. W., Markl, M. & Hennig, J. Visualization of tissue velocity data from cardiac wall motion measurements with myocardial fiber tracking: principles and implications for cardiac fiber structures. *Eur. J. Cardio-Thoracic Surg.* **29**, S158. ISSN: 1010-7940 (2006).
122. Jung, B., Markl, M., Föll, D. & Hennig, J. Investigating myocardial motion by MRI using tissue phase mapping. *Eur. J. Cardio-Thoracic Surg.* **29**, S150. ISSN: 1010-7940 (2006).
123. Jung, B., Föll, D., Böttler, P., Petersen, S., *et al.* Detailed analysis of myocardial motion in volunteers and patients using high-temporal-resolution MR tissue phase mapping. *J. Magn. Reson. Imaging* **24**, 1033. ISSN: 1522-2586 (2006).
124. Codreanu, I., Robson, M. D., Golding, S. J., Jung, B. A., *et al.* Longitudinally and circumferentially directed movements of the left ventricle studied by cardiovascular magnetic resonance phase contrast velocity mapping. *J. Cardiovasc. Magn. Reson.* **12**, 48. ISSN: 1532429X (2010).
125. Jung, B., Zaitsev, M., Hennig, J. & Markl, M. Navigator gated high temporal resolution tissue phase mapping of myocardial motion. *Magn. Reson. Med.* **55**, 937. ISSN: 0740-3194 (2006).

126. Steeden, J. A., Knight, D. S., Bali, S., Atkinson, D., *et al.* Self-navigated tissue phase mapping using a golden-angle spiral acquisition—proof of concept in patients with pulmonary hypertension. *Magn. Reson. Med.* **71**, 145. ISSN: 0740-3194 (2014).
127. Tanner, J. E. Self-Diffusion of Water in Frog-Muscle. *Biophys. J.* **28**, 107 (1979).
128. Torrey, H. C. Bloch Equations with Diffusion Terms. *Phys. Rev.* **104**, 563 (1956).
129. Novikov, D. S., Fieremans, E., Jespersen, S. N. & Kiselev, V. G. Quantifying brain microstructure with diffusion MRI: Theory and parameter estimation. *NMR Biomed* **32**, e3998 (2019).
130. Kiselev, V. in, 152 (2010). ISBN: 9780195369779.
131. Kiselev, V. G. & Il'yasov, K. A. Is the "biexponential diffusion" biexponential? *Magn. Reson. Med.* **57**, 464. ISSN: 1522-2594 (2007).
132. Callaghan, P. T. *Principles of nuclear magnetic resonance microscopy* XVII, 492 S., 4 S. Taf. ISBN: 0-19-853944-4 (Clarendon Press, Oxford, 1991).
133. Kingsley, P. B. Introduction to diffusion tensor imaging mathematics: Part III. Tensor calculation, noise, simulations, and optimization. **28A**, 155 (2006).
134. Scott, A. D., Nielles-Vallespin, S., Ferreira, P. F., McGill, L.-A., *et al.* The effects of noise in cardiac diffusion tensor imaging and the benefits of averaging complex data. *NMR Biomed.* **29**, 588 (2016).
135. Veraart, J., Sijbers, J., Sunaert, S., Leemans, A., *et al.* Weighted linear least squares estimation of diffusion MRI parameters: Strengths, limitations, and pitfalls. *Neuroimage* **81**, 335. ISSN: 1053-8119 (2013).
136. Jones, D. K. & Cercignani, M. Twenty-five pitfalls in the analysis of diffusion MRI data. *NMR Biomed.* **23**, 803. ISSN: 1099-1492 (2010).
137. Collier, Q., Veraart, J., Jeurissen, B., Den Dekker, A. J., *et al.* Iterative reweighted linear least squares for accurate, fast, and robust estimation of diffusion magnetic resonance parameters. *Magn. Reson. Med.* **73**, 2174. ISSN: 1522-2594 (2015).
138. Chang, L. C., Walker, L. & Pierpaoli, C. Informed RESTORE: A method for robust estimation of diffusion tensor from low redundancy datasets in the presence of physiological noise artifacts. *Magn. Reson. Med.* **68**, 1654. ISSN: 15222594 (2012).
139. Holmes, A. A., Scollan, D. F. & Winslow, R. L. Direct Histological Validation of Diffusion Tensor MRI in Formaldehyde-Fixed Myocardium (2000).
140. Scollan, D. F., Holmes, A., Winslow, R. & Forder, J. Histological validation of myocardial microstructure obtained from diffusion tensor magnetic resonance imaging. *Am. J. Physiol. - Hear. Circ. Physiol.* **275**. ISSN: 03636135 (1998).
141. Kung, G. L., Nguyen, T. C., Itoh, A., Skare, S., *et al.* The presence of two local myocardial sheet populations confirmed by diffusion tensor MRI and histological validation. *J. Magn. Reson. Imaging* **34**, 1080. ISSN: 1522-2586 (2011).
142. Helm, P., Beg, M. F., Miller, M. I. & Winslow, R. L. Measuring and Mapping Cardiac Fiber and Lamellar Architecture Using Diffusion Tensor MR Imaging. *Ann. N. Y. Acad. Sci.* **1047**, 296. ISSN: 1749-6632 (2005).

143. Stoeck, C. T., Kalinowska, A., von Deuster, C., Harmer, J., *et al.* Dual-phase cardiac diffusion tensor imaging with strain correction. *PLoS One* **9**, e107159. ISSN: 1932-6203 (2014).
144. Sosnovik, D. E., Wang, R., Dai, G., Wang, T., *et al.* Diffusion Spectrum MRI Tractography Reveals the Presence of a Complex Network of Residual Myofibers in Infarcted Myocardium. *Circ. Cardiovasc. Imaging* **2**, 206. ISSN: 19420080 (2009).
145. Nguyen, C., Fan, Z., Xie, Y., Dawkins, J., *et al.* In vivo contrast free chronic myocardial infarction characterization using diffusion-weighted cardiovascular magnetic resonance. *J. Cardiovasc. Magn. Reson.* **16**, 1. ISSN: 1532429X (2014).
146. Amzulescu, M. S., De Craene, M., Langet, H., Pasquet, A., *et al.* Myocardial strain imaging: review of general principles, validation, and sources of discrepancies. *Eur. Hear. J. - Cardiovasc. Imaging* **20**, 605. ISSN: 2047-2404 (2019).
147. Wei, L., Dong, J. X., Jin, L. X., He, J., *et al.* Peak early diastolic strain rate improves prediction of adverse cardiovascular outcomes in patients with ST-elevation myocardial infarction. *Radiol. Med.* **128**, 1372. ISSN: 1826-6983 (2023).
148. Mekkaoui, C., Jackowski, M. P., Kostis, W. J., Stoeck, C. T., *et al.* Myocardial scar delineation using diffusion tensor magnetic resonance tractography. *J. Am. Heart Assoc.* **7**. ISSN: 20479980 (2018).
149. Kung, G. L., Vaseghi, M., Gahm, J. K., Shevtsov, J., *et al.* Microstructural Infarct Border Zone Remodeling in the Post-infarct Swine Heart Measured by Diffusion Tensor MRI. *Front Physiol* **9**, 826 (2018).
150. Das, A., Kelly, C., Teh, I., Stoeck, C. T., *et al.* Acute microstructural changes after ST-segment elevation myocardial infarction assessed with diffusion tensor imaging. *Radiology* **299**, 86. ISSN: 15271315 (2021).
151. Yip, G., Abraham, T., Belohlavek, M. & Khandheria, B. K. Clinical applications of strain rate imaging. *J. Am. Soc. Echocardiogr.* **16**, 1334. ISSN: 0894-7317 (2003).
152. Wu, M. T., Tseng, W. Y. I., Su, M. Y. M., Liu, C. P., *et al.* Diffusion Tensor Magnetic Resonance Imaging Mapping the Fiber Architecture Remodeling in Human Myocardium After Infarction. *Circulation* **114**, 1036. ISSN: 00097322 (2006).
153. Wu, M. T., Su, M. Y., Huang, Y. L., Chiou, K. R., *et al.* Sequential changes of myocardial microstructure in patients postmyocardial infarction by diffusion-tensor cardiac mr correlation with left ventricular structure and function. *Circ. Cardiovasc. Imaging* **2**, 32. ISSN: 19419651 (2009).
154. Konstam, M. A., Kramer, D. G., Patel, A. R., Maron, M. S., *et al.* Left ventricular remodeling in heart failure: current concepts in clinical significance and assessment. *JACC Cardiovasc Imaging* **4**, 98 (2011).
155. Strijkers, G. J., Bouts, A., Blanckesteijn, W. M., Peeters, T. H., *et al.* Diffusion tensor imaging of left ventricular remodeling in response to myocardial infarction in the mouse. *NMR Biomed* **22**, 182 (2009).

156. Ariga, R., Tunnicliffe, E. M., Manohar, S. G., Mahmood, M., *et al.* Identification of Myocardial Disarray in Patients With Hypertrophic Cardiomyopathy and Ventricular Arrhythmias. *J. Am. Coll. Cardiol.* **73**, 2493. ISSN: 0735-1097 (2019).
157. Li, A., Ruh, A., Berhane, H., Robinson, J. D., *et al.* Altered regional myocardial velocities by tissue phase mapping and feature tracking in pediatric patients with hypertrophic cardiomyopathy. *Pediatr. Radiol.* **50**, 168. ISSN: 14321998 (2020).
158. Carruth, E. D., Teh, I., Schneider, J. E., McCulloch, A. D., *et al.* Regional variations in ex-vivo diffusion tensor anisotropy are associated with cardiomyocyte remodeling in rats after left ventricular pressure overload. *J. Cardiovasc. Magn. Reson.* **22**, 1. ISSN: 1532429X (2020).
159. Rider, O. J., Ajufo, E., Ali, M. K., Petersen, S. E., *et al.* Myocardial tissue phase mapping reveals impaired myocardial tissue velocities in obesity. *Int. J. Cardiovasc. Imaging* **31**, 339. ISSN: 15730743 (2015).
160. Abdullah, O. M., Seidel, T., Dahl, M., Gomez, A. D., *et al.* Diffusion tensor imaging and histology of developing hearts. *NMR Biomed* **29**, 1338. ISSN: 1099-1492 (Electronic) 0952-3480 (Linking) (2016).
161. Tous, C., Gentles, T. L., Young, A. A. & Pontr , B. P. Ex vivo cardiovascular magnetic resonance diffusion weighted imaging in congenital heart disease, an insight into the microstructures of tetralogy of Fallot, biventricular and univentricular systemic right ventricle. *J. Cardiovasc. Magn. Reson.* **22**, 69. ISSN: 1532429X (2020).
162. Abbasi, M. A., Bruno, G., Di Stefano, C., Garcia Bello, L., *et al.* Detection of Early Myocardial Dysfunction by Imaging Biomarkers in Cancer Patients Undergoing Photon Beam vs. Proton Beam Radiotherapy: A Prospective Study. *J. Cardiovasc. Dev. Dis.* **10**, 418. ISSN: 2308-3425 (2023).
163. Von Deuster, C., Stoeck, C. T., Genet, M., Atkinson, D., *et al.* Spin echo versus stimulated echo diffusion tensor imaging of the in vivo human heart. *Magn. Reson. Med.* **76**, 862. ISSN: 0740-3194 (2016).
164. Stoeck, C. T., Scott, A. D., Ferreira, P. F., Tunnicliffe, E. M., *et al.* Motion-Induced Signal Loss in In Vivo Cardiac Diffusion-Weighted Imaging. *J. Magn. Reson. Imaging* (2019).
165. Reese, T. G., Wedeen, V. J. & Weisskoff, R. M. Measuring Diffusion in the Presence of Material Strain. *J. Magn. Reson. Ser. B* **112**, 253 (1996).
166. Gamper, U., Boesiger, P. & Kozerke, S. Diffusion imaging of the in vivo heart using spin echoes—considerations on bulk motion sensitivity. *Magn Reson Med* **57**, 331 (2007).
167. Tseng, W. Y., Reese, T. G., Weisskoff, R. M. & Wedeen, V. J. Cardiac diffusion tensor MRI in vivo without strain correction. *Magn Reson Med* **42**, 393 (1999).
168. Stoeck, C. T., von Deuster, C., Fleischmann, T., Lipiski, M., *et al.* Direct comparison of in vivo versus postmortem second-order motion-compensated cardiac diffusion tensor imaging. *Magn Reson Med* **79**, 2265 (2018).

169. Moulin, K., Verzhbinsky, I. A., Maforo, N. G., Perotti, L. E., *et al.* Probing cardiomyocyte mobility with multi-phase cardiac diffusion tensor MRI. *PLoS One* **15** (ed Lionetti, V.) e0241996. ISSN: 1932-6203 (2020).
170. Lasič, S., Szczepankiewicz, F., Dall'Armellina, E., Das, A., *et al.* Motion-compensated b-tensor encoding for in vivo cardiac diffusion-weighted imaging. *NMR Biomed* **33**, e4213. ISSN: 1099-1492 (2020).
171. Westin, C.-F., Knutsson, H., Pasternak, O., Szczepankiewicz, F., *et al.* Q-space trajectory imaging for multidimensional diffusion MRI of the human brain. *Neuroimage* **135**, 345. ISSN: 10538119 (2016).
172. Szczepankiewicz, F., Sjölund, J., Dall'Armellina, E., Plein, S., *et al.* Motion-compensated gradient waveforms for tensor-valued diffusion encoding by constrained numerical optimization. *Magn. Reson. Med.* **85**, 2117. ISSN: 0740-3194 (2021).
173. Szczepankiewicz, F., Westin, C. F. & Nilsson, M. Gradient waveform design for tensor-valued encoding in diffusion MRI. *J. Neurosci. Methods* **348**. ISSN: 1872-678X. arXiv: 2007.07631 (2021).
174. Teh, I., Shelley, D., Boyle, J. H., Zhou, F., *et al.* Cardiac q-space trajectory imaging by motion-compensated tensor-valued diffusion encoding in human heart in vivo. *Magn. Reson. Med.* **90**, 150. ISSN: 0740-3194 (2023).
175. Reischauer, C., Wilm, B. J., Froehlich, J. M., Gutzeit, A., *et al.* High-resolution diffusion tensor imaging of prostate cancer using a reduced FOV technique. *Eur. J. Radiol.* **80**, e34. ISSN: 0720048X (2011).
176. Zijlstra, F., Bouwman, J. G., Braškutė, I., Viergever, M. A., *et al.* Fast Fourier-based simulation of off-resonance artifacts in steady-state gradient echo MRI applied to metal object localization. *Magn. Reson. Med.* **78**, 2035. ISSN: 15222594 (2017).
177. Shang, Y., Theilenberg, S., Terekhov, M., Mattar, W., *et al.* High-resolution simulation of Bo field conditions in the human heart from segmented computed tomography images. *NMR Biomed.* **35**, e4739. ISSN: 1099-1492 (2022).
178. Shang, Y., Theilenberg, S., Gajdošik, M., Schreiber, L. M., *et al.* High resolution simulation and measurement demonstrate oscillatory spatiotemporal Bo fluctuations across the human cardiac cycle. *Magn. Reson. Med.* ISSN: 1522-2594 (2023).
179. Aliotta, E., Wu, H. H. & Ennis, D. B. Convex optimized diffusion encoding (`<sc>CODE</sc>`) gradient waveforms for minimum echo time and bulk motion-compensated diffusion-weighted `<sc>MRI</sc>`. *Magn. Reson. Med.* **77**, 717. ISSN: 0740-3194 (2017).
180. Aliotta, E., Moulin, K., Magrath, P. & Ennis, D. B. Quantifying precision in cardiac diffusion tensor imaging with second-order motion-compensated convex optimized diffusion encoding. *Magn. Reson. Med.* **80**, 1074. ISSN: 1522-2594 (2018).

181. Aliotta, E., Moulin, K. & Ennis, D. B. Eddy current-nulled convex optimized diffusion encoding (EN-CODE) for distortion-free diffusion tensor imaging with short echo times. *Magn. Reson. Med.* **79**, 663. ISSN: 15222594 (2018).
182. Loecher, M., Middione, M. J. & Ennis, D. B. A gradient optimization toolbox for general purpose time-optimal MRI gradient waveform design. *Magn. Reson. Med.* **84**, 3234. ISSN: 0740-3194 (2020).
183. Dirix, P., Buoso, S., Peper, E. S. & Kozerke, S. Synthesis of patient-specific multipoint 4D flow MRI data of turbulent aortic flow downstream of stenotic valves. *Sci. Reports* 2022 121 **12**, 1. ISSN: 2045-2322 (2022).
184. Li, J.-R., Nguyen, V.-D., Tran, T. N., Valdman, J., *et al.* SpinDoctor: A MATLAB toolbox for diffusion MRI simulation. *Neuroimage* **202**, 116120. ISSN: 10538119. arXiv: 1902.01025 (2019).
185. Buoso, S., Joyce, T. & Kozerke, S. Personalising left-ventricular biophysical models of the heart using parametric physics-informed neural networks. *Med. Image Anal.* **71**, 102066. ISSN: 13618415 (2021).
186. Lee, K. L., Doorly, D. J. & Firmin, D. N. Numerical simulations of phase contrast velocity mapping of complex flows in an anatomically realistic bypass graft geometry. *Med. Phys.* **33**, 2621. ISSN: 0094-2405 (2006).
187. Guerquin-Kern, M., Lejeune, L., Pruessmann, K. P. & Unser, M. Realistic Analytical Phantoms for Parallel Magnetic Resonance Imaging. *IEEE Trans. Med. Imaging* **31**, 626. ISSN: 0278-0062 (2012).
188. Graf, C., Rund, A., Aigner, C. S. & Stollberger, R. Accuracy and performance analysis for Bloch and Bloch-McConnell simulation methods. *J. Magn. Reson.* **329**, 107011. ISSN: 1090-7807 (2021).
189. Buoso, S., Joyce, T., Schulthess, N. & Kozerke, S. MRXCAT2.0: Synthesis of realistic numerical phantoms by combining left-ventricular shape learning, biophysical simulations and tissue texture generation. *J. Cardiovasc. Magn. Reson.* 2023 251 **25**, 1. ISSN: 1532-429X (2023).
190. Zubal, I. G., Harrell, C. R., Smith, E. O., Rattner, Z., *et al.* Computerized three-dimensional segmented human anatomy. *Med. Phys.* **21**, 299. ISSN: 0094-2405 (1994).
191. Kramer, R., Houry, H. J., Vieira, J. W., Loureiro, E. C. M., *et al.* All about FAX: a Female Adult voXel phantom for Monte Carlo calculation in radiation protection dosimetry. *Phys. Med. Biol.* **49**, 5203. ISSN: 0031-9155 (2004).
192. Segars, W. P. & Tsui, B. M. W. MCAT to XCAT: The Evolution of 4-D Computerized Phantoms for Imaging Research. *Proc. IEEE* **97**, 1954. ISSN: 0018-9219 (2009).
193. Segars, W. P., Sturgeon, G., Mendonca, S., Grimes, J., *et al.* 4D XCAT phantom for multimodality imaging research. *Med. Phys.* **37**, 4902. ISSN: 00942405 (2010).
194. Veress, A. I., Segars, W. P., Weiss, J. A., Tsui, B. M. W., *et al.* Normal and Pathological NCAT Image and Phantom Data Based on Physiologically Realistic Left Ventricle Finite-Element Models. *IEEE Trans. Med. Imaging* **25**, 1604. ISSN: 0278-0062 (2006).

195. Segars, W. P., Veress, A. I., Sturgeon, G. M. & Samei, E. Incorporation of the Living Heart Model Into the 4-D XCAT Phantom for Cardiac Imaging Research. *IEEE Trans. Radiat. Plasma Med. Sci.* **3**, 54. ISSN: 2469-7311 (2019).
196. Joyce, T., Buoso, S., Stoeck, C. T. & Kozerke, S. Rapid inference of personalised left-ventricular meshes by deformation-based differentiable mesh voxelization. *Med. Image Anal.* **79**, 102445. ISSN: 13618415 (2022).
197. Bai, W., Shi, W., de Marvao, A., Dawes, T. J., *et al.* A bi-ventricular cardiac atlas built from 1000+ high resolution MR images of healthy subjects and an analysis of shape and motion. *Med. Image Anal.* **26**, 133. ISSN: 13618415 (2015).
198. Suinesiaputra, A., Dhooge, J., Duchateau, N., Ehrhardt, J., *et al.* Statistical Shape Modeling of the Left Ventricle: Myocardial Infarct Classification Challenge. *IEEE J. Biomed. Heal. Informatics* **22**, 503. ISSN: 2168-2194 (2018).
199. Duan, J., Bello, G., Schlemper, J., Bai, W., *et al.* Automatic 3D Bi-Ventricular Segmentation of Cardiac Images by a Shape-Refined Multi- Task Deep Learning Approach. *IEEE Trans. Med. Imaging* **38**, 2151. ISSN: 0278-0062. arXiv: 1808.08578 (2019).
200. Lee, H. H., Fieremans, E. & Novikov, D. S. Realistic Microstructure Simulator (RMS): Monte Carlo simulations of diffusion in three-dimensional cell segmentations of microscopy images. *J. Neurosci. Methods* **350**, 109018. ISSN: 0165-0270. arXiv: 2012.06478 (2021).
201. Lihui Wang, Yuemin Zhu, Hongying Li, Wanyu Liu, *et al.* Multiscale Modeling and Simulation of the Cardiac Fiber Architecture for DMRI. *IEEE Trans. Biomed. Eng.* **59**, 16. ISSN: 0018-9294 (2012).
202. Hall, M. G. & Clark, C. A. Diffusion in hierarchical systems: A simulation study in models of healthy and diseased muscle tissue. *Magn Reson Med* **78**, 1187 (2017).
203. Jing, Y., Frindel, C. & Magnin, I. E. *Simulation of water diffusion through a simple virtual cardiac cell model: optimization of Monte Carlo parameters and observation by simulated MRI in 2020 15th IEEE Int. Conf. Signal Process. 2020-Decem* (IEEE, 2020), 527. ISBN: 978-1-7281-4479-5.
204. Wang, L., Hong, Y., Qin, Y. B., Cheng, X. Y., *et al.* Connecting macroscopic diffusion metrics of cardiac diffusion tensor imaging and microscopic myocardial structures based on simulation. *Med. Image Anal.* **77**, 102325. ISSN: 1361-8415 (2022).
205. Jing, Y., Magnin, I. E. & Frindel, C. Monte Carlo simulation of water diffusion through cardiac tissue models. *Med. Eng. Phys.* **120**, 104013. ISSN: 13504533 (2023).
206. Bates, J., Teh, I., McClymont, D., Kohl, P., *et al.* Monte Carlo Simulations of Diffusion Weighted MRI in Myocardium: Validation and Sensitivity Analysis. *IEEE Trans. Med. Imaging* **36**, 1316. ISSN: 0278-0062 (2017).
207. Lashgari, M., Ravikumar, N., Teh, I., Li, J.-R., *et al.* Three-dimensional micro-structurally informed in silico myocardium—Towards virtual imaging trials

- in cardiac diffusion weighted MRI. *Med. Image Anal.* **82**, 102592. ISSN: 13618415 (2022).
208. Rose, J. N., Nielles-Vallespin, S., Ferreira, P. F., Firmin, D. N., *et al.* Novel insights into in-vivo diffusion tensor cardiovascular magnetic resonance using computational modelling and a histology-based virtual microstructure. *Magn. Reson. Med.* **81**, 2759. ISSN: 0740-3194 (2019).
209. Stimm, J., Buoso, S., Berberoglu, E., Kozerke, S., *et al.* A 3D personalized cardiac myocyte aggregate orientation model using MRI data-driven low-rank basis functions. *Med. Image Anal.* **71**, 102064. ISSN: 13618415 (2021).
210. Weine, J., McGrath, C. & Kozerke, S. *CMRSeq - A Python package for intuitive sequence design* in *Proc. ISMRM 32nd Annu. Meet.* (Toronto, 2023).
211. Layton, K. J., Kroboth, S., Jia, F., Littin, S., *et al.* Pulseq: A rapid and hardware-independent pulse sequence prototyping framework. *Magn. Reson. Med.* **77**, 1544. ISSN: 15222594 (2017).
212. Sravan Ravi, K., Geethanath, S. & Thomas Vaughan Jr, J. PyPulseq: A Python Package for MRI Pulse Sequence Design. *J. Open Source Softw.* **4**, 1725. ISSN: 2475-9066 (2019).
213. Nielsen, J. F. & Noll, D. C. TOPPE: A framework for rapid prototyping of MR pulse sequences. *Magn. Reson. Med.* **79**, 3128. ISSN: 15222594 (2018).
214. Cordes, C., Konstandin, S., Porter, D. & Günther, M. Portable and platform-independent MR pulse sequence programs. *Magn. Reson. Med.* **83**, 1277. ISSN: 0740-3194 (2020).
215. Petersson, S., Dyverfeldt, P., Gårdhagen, R., Karlsson, M., *et al.* Simulation of phase contrast MRI of turbulent flow. *Magn. Reson. Med.* **64**, 1039. ISSN: 1522-2594 (2010).
216. Lorthois, S., Stroud-Rossman, J., Berger, S., Jou, L. D., *et al.* Numerical simulation of magnetic resonance angiographies of an anatomically realistic stenotic carotid bifurcation. *Ann. Biomed. Eng.* **33**, 270. ISSN: 00906964 (2005).
217. Marques, J. P., Meineke, J., Milovic, C., Bilgic, B., *et al.* QSM reconstruction challenge 2.0: A realistic in silico head phantom for MRI data simulation and evaluation of susceptibility mapping procedures. *Magn. Reson. Med.* **86**, 526. ISSN: 1522-2594 (2021).
218. Wissmann, L., Santelli, C., Segars, W. P. & Kozerke, S. MRXCAT: Realistic numerical phantoms for cardiovascular magnetic resonance. *J. Cardiovasc. Magn. Reson.* **16**, 63. ISSN: 1532429X (2014).
219. Cao, Z., Oh, S., Sica, C. T., McGarrity, J. M., *et al.* Bloch-based MRI system simulator considering realistic electromagnetic fields for calculation of signal, noise, and specific absorption rate. *Magn. Reson. Med.* **72**, 237. ISSN: 1522-2594 (2014).
220. Puiseux, T., Sewonu, A., Moreno, R., Mendez, S., *et al.* Numerical simulation of time-resolved 3D phase-contrast magnetic resonance imaging. *PLoS One* **16**, e0248816. ISSN: 1932-6203 (2021).

221. Jordan, S. P., Hu, S., Rozada, I., McGivney, D. F., *et al.* Automated design of pulse sequences for magnetic resonance fingerprinting using physics-inspired optimization. *Proc. Natl. Acad. Sci. U. S. A.* **118**, e2020516118. ISSN: 10916490. arXiv: 2106.04740 (2021).
222. Weine, J., van Gorkum, R. J., Stoeck, C. T., Vishnevskiy, V., *et al.* Synthetically trained convolutional neural networks for improved tensor estimation from free-breathing cardiac DTI. *Comput. Med. Imaging Graph.* **99**, 102075. ISSN: 0895-6111 (2022).
223. Wu, Z., Chen, W. & Nayak, K. S. Minimum Field Strength Simulator for Proton Density Weighted MRI. *PLoS One* **11**, e0154711. ISSN: 1932-6203 (2016).
224. Peretti, L., Cencini, M., Cecchi, P., Donatelli, G., *et al.* *PySynthMRI: An open-source Python tool for Synthetic MRI* in *Proc. ISMRM 31st Annu. Meet.* (London, 2022).
225. Paganelli, C., Summers, P., Gianoli, C., Bellomi, M., *et al.* A tool for validating MRI-guided strategies: a digital breathing CT/MRI phantom of the abdominal site. *Med. Biol. Eng. Comput.* **55**, 2001. ISSN: 17410444 (2017).
226. Van Gorkum, R. J. H., Weine, J. L., Segars, W. P., Stoeck, C. T., *et al.* MRXCAT-CDTI: A Numerical Cardiac Diffusion Tensor Imaging Phantom. arXiv: 2204.07966 (2022).
227. Bijmens, B., Cikes, M., Butakoff, C., Sitges, M., *et al.* Myocardial Motion and Deformation: What Does It Tell Us and How Does It Relate to Function? *Fetal Diagn. Ther.* **32**, 5. ISSN: 1015-3837 (2012).
228. Bijmens, B. H., Cikes, M., Claus, P. & Sutherland, G. R. Velocity and deformation imaging for the assessment of myocardial dysfunction. *Eur. J. Echocardiogr.* **10**, 216. ISSN: 1532-2114 (2009).
229. Van Tyen, R., Saloner, D., Jou, L.-D. & Berger, S. MR imaging of flow through tortuous vessels: A numerical simulation. *Magn. Reson. Med.* **31**, 184. ISSN: 0740-3194 (1994).
230. Steinman, D. A., Ethier, C. R. & Rutt, B. K. Combined analysis of spatial and velocity displacement artifacts in phase contrast measurements of complex flows. *J. Magn. Reson. Imaging* **7**, 339. ISSN: 1053-1807 (1997).
231. Siegel, J. M., Oshinski, J. N., Pettigrew, R. I. & Ku, D. N. Computational simulation of turbulent signal loss in 2D time-of-flight magnetic resonance angiograms. *Magn. Reson. Med.* **37**, 609. ISSN: 0740-3194 (1997).
232. Fortin, A., Salmon, S., Baruthio, J., Delbany, M., *et al.* Flow MRI simulation in complex 3D geometries: Application to the cerebral venous network. *Magn. Reson. Med.* **80**, 1655. ISSN: 1522-2594 (2018).
233. Klepaczko, A., Szczypiński, P., Strzelecki, M. & Stefańczyk, L. Simulation of phase contrast angiography for renal arterial models. *Biomed. Eng. Online* **17**, 41. ISSN: 1475-925X (2018).
234. Dillinger, H., McGrath, C., Guenther, C. & Kozerke, S. Fundamentals of turbulent flow spectrum imaging. *Magn. Reson. Med.* **87**, 1231. ISSN: 0740-3194 (2022).

235. Jou, L.-D., van Tyen, R., Berger, S. A. & Saloner, D. Calculation of the magnetization distribution for fluid flow in curved vessels. *Magn. Reson. Med.* **35**, 577. ISSN: 0740-3194 (1996).
236. Jou, L. & Saloner, D. A numerical study of magnetic resonance images of pulsatile flow in a two dimensional carotid bifurcation. *Med. Eng. Phys.* **20**, 643. ISSN: 13504533 (1998).
237. Marshall, I. Simulation of in-plane flow imaging. *Concepts Magn. Reson.* **11**, 379. ISSN: 1043-7347 (1999).
238. Jurczuk, K., Kretowski, M., Bellanger, J. J., Eliat, P. A., *et al.* Computational modeling of MR flow imaging by the lattice Boltzmann method and Bloch equation. *Magn. Reson. Imaging* **31**, 1163. ISSN: 0730725X (2013).
239. Jurczuk, K., Kretowski, M. & Bezy-Wendling, J. GPU-based computational modeling of magnetic resonance imaging of vascular structures. *Int. J. High Perform. Comput. Appl.* **32**, 496. ISSN: 17412846 (2018).
240. Hanson, H. M., Eiben, B., McClelland, J. R., van Herk, M., *et al.* Technical Note: Four-dimensional deformable digital phantom for MRI sequence development. *Med. Phys.* **48**, 5406. ISSN: 2473-4209 (2021).
241. Abadi, M., Barham, P., Chen, J., Chen, Z., *et al.* TensorFlow: A System for Large-Scale Machine Learning TensorFlow: A system for large-scale machine learning.
242. Sullivan, C. B. & Kaszynski, A. A. PyVista: 3D plotting and mesh analysis through a streamlined interface for the Visualization Toolkit (VTK). *J. Open Source Softw.* **4**, 1450. ISSN: 2475-9066 (2019).
243. Buoso, S., Manzoni, A., Alkadhi, H., Plass, A., *et al.* Reduced-order modeling of blood flow for noninvasive functional evaluation of coronary artery disease. *Biomech. Model. Mechanobiol.* **18**, 1867. ISSN: 16177940 (2019).
244. Buoso, S., Manzoni, A., Alkadhi, H. & Kurtcuoglu, V. Stabilized reduced-order models for unsteady incompressible flows in three-dimensional parametrized domains. *Comput. Fluids* **246**, 105604. ISSN: 0045-7930 (2022).
245. Lund, T. S., Wu, X. & Squires, K. D. On the Generation of Turbulent Inflow Conditions for Boundary Layer Simulations. *J. Comput. Phys.* **140**, 233. ISSN: 00219991 (1998).
246. Huilier, D. G. An Overview of the Lagrangian Dispersion Modeling of Heavy Particles in Homogeneous Isotropic Turbulence and Considerations on Related LES Simulations. *Fluids 2021, Vol. 6, Page 145* **6**, 145. ISSN: 2311-5521 (2021).
247. Jones, D., Horsfield, M. & Simmons, A. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magn. Reson. Med.* **42**, 515. ISSN: 0740-3194 (1999).
248. Uecker, M., Holme, C., Blumenthal, M., Wang, X., *et al.* BART Toolbox for Computational Magnetic Resonance Imaging (2021).
249. Alemany, I., Rose, J. N., Ferreira, P. F., Pennell, D. J., *et al.* Realistic numerical simulations of diffusion tensor cardiovascular magnetic resonance: The effects

- of perfusion and membrane permeability. *Magn. Reson. Med.* **90**, 1641. ISSN: 1522-2594 (2023).
250. Rose, J. N., Nielles-Vallespin, S., Ferreira, P. F., Firmin, D. N., *et al.* Novel insights into in-vivo diffusion tensor cardiovascular magnetic resonance using computational modelling and a histology-based virtual microstructure. *Magn. Reson. Med.* **81**, 2759. ISSN: 1522-2594 (2019).
 251. Wasserman, L. *All of Statistics: A Concise Course in Statistical Inference* ISBN: 1441923225 (Springer Publishing Company, Incorporated, 2004).
 252. Andersen, A. H. On the Rician distribution of noisy MRI data. *Magn Reson Med* **36**, 331 (1996).
 253. Orton, M. R., Collins, D. J., Koh, D.-M. & Leach, M. O. Improved intravoxel incoherent motion analysis of diffusion weighted imaging by data driven Bayesian modeling. *Magn. Reson. Med.* **71**, 411 (2014).
 254. Spinner, G. R., Federau, C. & Kozerke, S. Bayesian inference using hierarchical and spatial priors for intravoxel incoherent motion MR imaging in the brain: Analysis of cancer and acute stroke. *Med. Image Anal.* **73**, 102144. ISSN: 1361-8415 (2021).
 255. Spinner, G. R., von Deuster, C., Tezcan, K. C., Stoeck, C. T., *et al.* Bayesian intravoxel incoherent motion parameter mapping in the human heart. *J Cardiovasc Magn Reson* **19**, 85 (2017).
 256. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning* ISBN: 0262035618 (The MIT Press, 2016).
 257. Abdullah, A. A., Hassan, M. M. & Mustafa, Y. T. A Review on Bayesian Deep Learning in Healthcare: Applications and Challenges. *IEEE Access* **10**, 36538. ISSN: 2169-3536 (2022).
 258. Gao, C., Killeen, B. D., Hu, Y., Grupp, R. B., *et al.* Synthetic data accelerates the development of generalizable learning-based algorithms for X-ray image analysis. *Nat. Mach. Intell.* **5**, 294. ISSN: 2522-5839 (2023).
 259. Gyori, N. G., Palombo, M., Clark, C. A., Zhang, H., *et al.* Training Data Distribution Significantly Impacts the Estimation of Tissue Microstructure with Machine Learning. *bioRxiv*, 2021.04.13.439659 (2021).
 260. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., *et al.* Deep Learning Techniques for Automatic MRI Cardiac Multi-Structures Segmentation and Diagnosis: Is the Problem Solved? *IEEE Trans. Med. Imaging* **37**, 2514. ISSN: 0278-0062 (2018).
 261. Arsigny, V., Fillard, P., Pennec, X. & Ayache, N. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magn. Reson. Med.* **56**, 411. ISSN: 07403194 (2006).
 262. Szegedy, C., Ioffe, S., Vanhoucke, V. & Alemi, A. A. *Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning* in Proc. Thirty-First AAAI Conf. Artif. Intell. (AAAI Press, 2017), 4278.
 263. Ronneberger, O., Fischer, P. & Brox, T. *U-Net: Convolutional Networks for Biomedical Image Segmentation BT - Medical Image Computing and Computer-*

- Assisted Intervention – MICCAI 2015* in (eds Navab, N., Hornegger, J., Wells, W. M. & Frangi, A. F.) (Springer International Publishing, Cham, 2015), 234. ISBN: 978-3-319-24574-4.
264. Golkov, V., Dosovitskiy, A., Sperl, J. I., Menzel, M. I., *et al.* q-Space Deep Learning: Twelve-Fold Shorter and Model-Free Diffusion MRI Scans. *IEEE Trans Med Imaging* **35**, 1344 (2016).
265. Chartsias, A., Joyce, T., Papanastasiou, G., Semple, S., *et al.* Disentangled representation learning in cardiac image analysis. *Med. Image Anal.* **58**, 101535. ISSN: 1361-8415 (2019).
266. Flury, B. D. Acceptance-Rejection Sampling Made Easy. *SIAM Rev.* **32**, 474. ISSN: 00361445 (1990).
267. Scott, A. D., Nielles-Vallespin, S., Ferreira, P. F., Khaliq, Z., *et al.* An in-vivo comparison of stimulated-echo and motion compensated spin-echo sequences for 3 T diffusion tensor cardiovascular magnetic resonance at multiple cardiac phases. *J. Cardiovasc. Magn. Reson.* **20**, 1. ISSN: 1532-429X (2018).
268. Anand, I. S. & Florea, V. G. in, 247 (Springer International Publishing, 2017).
269. Sharrack, N., Das, A., Kelly, C., Teh, L., *et al.* The relationship between myocardial microstructure and strain in chronic infarction using cardiovascular magnetic resonance diffusion tensor imaging and feature tracking. *J. Cardiovasc. Magn. Reson.* **24**, 1. ISSN: 1532429X (2022).
270. Föll, D., Jung, B., Schilli, E., Staehle, F., *et al.* Magnetic Resonance Tissue Phase Mapping of Myocardial Motion. *Circ. Cardiovasc. Imaging* **3**, 54. ISSN: 1941-9651 (2010).
271. Wedeen, V. J. Magnetic resonance imaging of myocardial kinematics. technique to detect, localize, and quantify the strain rates of the active human myocardium. *Magn. Reson. Med.* **27**, 52. ISSN: 15222594 (1992).
272. Tseng, W. Y. I., Reese, T. G., Weisskoff, R. M., Brady, T. J., *et al.* Myocardial fiber shortening in humans: Initial results of MR imaging. *Radiology* **216**, 128. ISSN: 00338419 (2000).
273. Dou, J., Tseng, W.-Y. I., Reese, T. G. & Wedeen, V. J. Combined Diffusion and Strain MRI Reveals Structure and Function of Human Myocardial Lamellar Sheets In Vivo. *Magn. Reson. Med.* **50**, 107 (2003).
274. Bjerregaard, C. L., Olsen, F. J., Lassen, M. C. H., Svartstein, A. S. W., *et al.* Ratio of early transmitral inflow velocity to early diastolic strain rate predicts atrial fibrillation following acute myocardial infarction. *Int. J. Cardiovasc. Imaging.* ISSN: 1875-8312 (2023).
275. Zhang, H.-K., Shi, C.-Y., Liu, D.-T., Gao, H.-Q., *et al.* Dynamic changes in cardiac morphology, function, and diffuse myocardial fibrosis duration of diabetes in type 1 and type 2 diabetic mice models using 7.0 T CMR and echocardiography. *Front. Endocrinol. (Lausanne)*. **14**. ISSN: 1664-2392 (2023).
276. Inci, S., Gul, M., Elcik, D., Aktas, H., *et al.* Identification of subclinical myocardial dysfunction by Speckle Tracking Imaging in patients with myocardial

- infarction with non-occlusive coronary arteries (MINOCA). *Int. J. Cardiovasc. Imaging* **38**, 2099. ISSN: 1875-8312 (2022).
277. Nizhnikava, V., Reiter, U., Kovacs, G., Reiter, C., *et al.* Myocardial strain parameters in pulmonary hypertension are determined by changes in volumetric function rather than by hemodynamic alterations. *Eur. J. Radiol.* **170**, 111187. ISSN: 0720-048X (2024).
278. Weine, J., McGrath, C., Dirix, P., Buoso, S., *et al.* CMRsim –A python package for cardiovascular MR simulations incorporating complex motion and flow. *Magn. Reson. Med.* ISSN: 0740-3194 (2024).
279. Stoeck, C. T., von Deuster, C., van Gorkum, R. J., Kozerke, S., *et al.* Motion and eddy current-induced signal dephasing in in vivo cardiac DTI. *Magn. Reson. Med.* **84**, 277 (2020).
280. Moulin, K., Croisille, P., Feiweier, T., Delattre, B. M. A., *et al.* In vivo free-breathing DTI and IVIM of the whole human heart using a real-time slice-followed SE-EPI navigator-based sequence: A reproducibility study in healthy volunteers. *Magn. Reson. Med.* **76**, 70 (2016).
281. Lau, A. Z., Tunncliffe, E. M., Frost, R., Koopmans, P. J., *et al.* Accelerated human cardiac diffusion tensor imaging using simultaneous multislice imaging. *Magn. Reson. Med.* **73**, 995. ISSN: 1522-2594 (2015).
282. Coorey, G., Figtree, G. A., Fletcher, D. F., Snelson, V. J., *et al.* The health digital twin to tackle cardiovascular disease—a review of an emerging interdisciplinary field. *npj Digit. Med.* **5**, 126. ISSN: 2398-6352 (2022).
283. Baillargeon, B., Rebelo, N., Fox, D. D., Taylor, R. L., *et al.* The Living Heart Project: A robust and integrative simulator for human heart function. *Eur. J. Mech. A, Solids* **48**, 38. ISSN: 09977538 (2014).

List of Publications

Articles in peer-reviewed journals:

1. Weine, J., McGrath, C., Dirix, P., Buoso, S., *et al.* CMRsim—A python package for cardiovascular MR simulations incorporating complex motion and flow. *Magnetic Resonance in Medicine* **n/a**. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.30010>.
2. Weine, J., van Gorkum, R. J., Stoeck, C. T., Vishnevskiy, V., *et al.* Synthetically trained convolutional neural networks for improved tensor estimation from free-breathing cardiac DTI. *Comput. Med. Imaging Graph.* **99**, 102075. ISSN: 0895-6111 (2022).
3. Julius, K., Weine, J., Gao, M., Latarius, J., *et al.* Impact of Macromolecular Crowding and Compression on Protein–Protein Interactions and Liquid–Liquid Phase Separation Phenomena. *Macromolecules* **52**, 1772. eprint: <https://doi.org/10.1021/acs.macromol.8b02476> (2019).
4. Julius, K., Weine, J., Berghaus, M., König, N., *et al.* Water-Mediated Protein-Protein Interactions at High Pressures are Controlled by a Deep-Sea Osmolyte. *Phys. Rev. Lett.* **121**, 038101 (3 2018).
5. Schulze, J., Möller, J., Weine, J., Julius, K., *et al.* Phase behavior of lysozyme solutions in the liquid–liquid phase coexistence region at high hydrostatic pressures. *Phys. Chem. Chem. Phys.* **18**, 14252 (21 2016).

Conference contributions:

6. Weine, J., McGrath, C. & Kozerke, S. *CMRSeq - A Python package for intuitive sequence design in ISMRM2023* June 3–8, 2023 (Proceedings of ISMRM 32nd Annual Meeting, Toronto, Canada).
7. Weine, J., Buoso, S., McGrath, C. & Kozerke, S. *Joint cardiac tissue velocity and diffusion tensor mapping in ISMRM2023* June 3–8, 2023 (Proceedings of ISMRM 32nd Annual Meeting, Toronto, Canada).
8. Weine, J., McGrath, C. & Kozerke, S. *CMRsim - A Python package for MRI simulations incorporating complex organ motion and flow in ISMRM2023* June 3–8, 2023 (Proceedings of ISMRM 32nd Annual Meeting, Toronto, Canada).
9. Weine, J., Buoso, S. & Kozerke, S. *Theoretical Considerations on Joint Inference of Cardiac Diffusion and Strain Tensors from second-order motion compensated cDTI Data in ISMRM2022* May 7–12, 2022 (Proceedings of ISMRM 31st Annual Meeting, London, England).
10. Weine, J., van Gorkum, R. J. H., Stoeck, C. T., Vishnevskiy, V., *et al.* *Inferring Diffusion Tensors on Unregistered Cardiac DWI Using a Residual CNN and Implicitly Modelled Data Prior in ISMRM2021* May 15–20, 2021 (Proceedings of ISMRM 30th Annual Meeting, Online).

11. Vogel, R., Ritter, D., Weine, J., Faust, J., *et al.* *Deep Learning-Based Needle Tracking Trained on Bloch-Simulated Data and Evaluated on Clinical Real-Time bSSFP Images in ISMRM2021* May 15–20, 2021 (Proceedings of ISMRM 30th Annual Meeting, Online).
12. Weine, J., Schneider, R., Kägebein, U., Hensen, B., *et al.* *Interleaved White Marker Contrast with bSSFP Real-Time Imaging for Deep Learning based Needle Localization in MR-Guided Percutaneous Interventions in ISMRM2019* May 10–13, 2019 (Proceedings of ISMRM 27th Annual Meeting, Montreal, Canada).
13. Weine, J., Breton, E., Garnon, J., Gangi, A., *et al.* *Deep Learning Based Needle Localization on Real-Time MR Images of Patients Acquired During MR-guided Percutaneous Interventions in ISMRM2019* May 10–13, 2017 (Proceedings of ISMRM 27th Annual Meeting, Montreal, Canada).
14. Weine, J., Maier, F., Polak, D. & Umathum, R. *Continuous SWIFT: Analog Leakage Compensation Utilizing an Embedded System in ISMRM2019* Apr. 22–27, 2017 (Proceedings of ISMRM 25th Annual Meeting, Montreal, Canada).

Acknowledgements

First of all, I would like to thank Prof. Sebastian Kozerke for giving me the opportunity to undertake my PhD journey under his supervision. I deeply appreciate the open and honest communication we had and the trust you put in me. Your dedication to CMR and your willingness to always be available for discussing the challenges popping up along the way, profoundly impresses me.

I would like to thank Prof. Maxim Zaitsev for being referee of my thesis and a role model in open MR research.

Special thanks also go to Florian Maier, without whom my path maybe would have been very different.

Further, I especially would like to thank Stefano Buoso for offering me invaluable professional advice, proofreading, collaboration and foremost keeping me (almost) sane with his company. Moreover, I want to express my gratitude to Charles McGrath for jumping on the train of open-source development with me, as tackling our simulation project would have been to scary alone. Also, geeking out about computer-stuff and random internet-things was indispensable in battling the PhD-despair. For providing me help to get started with cDTI, I would like to thank Robbert van Gorkum and Christian Stoeck.

I am thankful for Hannes Dillinger, Christian Günthner, Tobias Hoh, Julia Bernreiter, Christian Stoeck, Valery Vishnevskiy, Jonas Wahlheim, Robbert van Gorkum, Sophie Peereboom, Maximilian Fütterer, Johanna Stimm kindly welcoming me in the CMR-group. Aside from sharing an office with me, I want to thank Hannes Dillinger for being an exceptional, inspiring person and Eva Peper for emitting seemingly inexhaustible amounts of energy (also for encouraging me to do things right for once!). For keeping up with all the disturbances coming from next door or next desk I would like to thank Pietro Dirix, Gloria Wolkerstorfer and Manuel Christanell. Furthermore, this extends to all my current and former colleagues in the CMR and MR-tech group. Your company was a great pleasure!

Thanks to Roger Lüchinger for his support in managing and maintaining all the computers and Martin Bühner for his help with navigating the raw data and license jungle. Also, thanks to Isabel Spiess for her invaluable administrative support. Furthermore, I would like to thank Maximilian Fütterer for his support in the final days of this journey.

Sophie, words cannot not describe how grateful I am for having you in my life. Thank you for your love, your support, and your trust, as all of this would not have happened without you taking the leap of faith to Zurich with me.

Dear Mom and Dad, thank you for always believing in me and your love along the rocky road leading to this very moment. Balthasar, I wish I could share this moment with you.