Imaging Vulnerable Coronary Plaques Using 3D Motion-Corrected Simultaneous PET/MR

vorgelegt von
M.Sc.
Johannes Mayer
ORCID: 0000-0002-2500-445X

von der Fakultät V - Verkehrs- und Maschinensysteme
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Naturwissenschaften
-Dr. rer. nat.-

genehmigte Dissertation

Promotionsausschuss:
Vorsitzender: Prof. Dr.-Ing. Marc Kraft
Gutachter: Prof. Dr. rer. nat. Tobias Schäffter
Gutachter: Prof. Dr. Sebastian Kozerke

Tag der wissenschaftlichen Aussprache: 30. September 2021

Berlin 2021

Zusammenfassung

Vor kurzem hat der radioaktive PET-Tracer [^{18}F]NaF aus der Knochenbildgebung eine Anwendung im Herzen gefunden: er verzeichnet Aufnahme in vulnerablen atherosklerotischen Plaques, die mit hoher Wahrscheinlichkeit abreißen und potenziell einen Herzinfarkt verursachen. Die PET-Aufnahme für [^{18}F]NaF -Scans der Koronararterien benötigen jedoch eine halbe Stunde oder länger, weshalb Herz- und Atembewegung die PET-Bildqualität verschlechtern können. Zu den auftretenden Bewegungsartefakten gehören eine Unschärfe in der Darstellung des Tracer sowie ein fehlerhafter Überlapp von PET-Daten und der Schwächungskarte. Herkömmliche Ansätze zur Bewegunsunterdrückung, wie beispielsweise das Gating, sind auf eine hohe Zählrate in den aufgenommenen PET-Daten oder eine starke lokale Ansammlung von PET-Tracer im Plaque angewiesen. In den Koronararterien sind diese jedoch klein. Daher werden neue Bewegungskorrekturverfahren benötigt, die alle aufgenommenen Daten nutzen können und diese in einen Zustand ohne Bewegung zurückversetzen.

Eine geeignete Modalität für die Umsetzung solcher Techniken ist PET/MR, ein Hybrid welcher die gleichzeitige Aufnahme schneller und hochaufgelöster anatomischer Bildgebung von MR und der komplementären Informationen über Stoffwechselprozesse durch PET ermöglicht. Das übergeordnete Ziel dieser Arbeit war die Verbesserung der Darstellung und Quantifizierung von Plaques in den Koronararterien durch eine MR-basierte Korrektur der Herz-Atembewegung in simultan aufgenommenen [^{18}F]NaF -PET/MR Daten. Dazu wurde eine MR-Sequenz zur Darstellung von Wasser- und Fettgewebe modifiziert, um den k-Raum entlang einer neuen Trajektorie abzutasten, welche besser für die bewegungsaufgelöste Bildgebung geeignet ist. Bewegungskorrektur wurde in eine hochauflösende 3D-Modell-basierte Fett-Wasser MR-Rekonstruktion integriert. Zudem wurde eine Simulationsumgebung für dynamische PET/MR-Daten mit “Ground-Truth” Informationen entwickelt. Diese wurde genutzt, um die Genauigkeit von synergistischen Bildregistrierung zu messen, welche zur Erstellung von Herz-Atmungsbewegungs-Modellen eingesetzt wurde.

Die entwickelten Methoden wurden zunächst in einer eigenständigen Herz-MR-Anwendung validiert und verbesserten die lokale Schärfe von 3D epikardialen Fettstrukturen durch Bewegungskorrektur um 40 %. Daraufhin wurden die entwickelten Techniken auf [^{18}F]NaF PET/MR Patienten mit koronarer Herzkrankheit angewandt. Bei den detektierten koronaren [^{18}F]NaF -positiven Plaques der Patienten führte die Anwendung der Herz-Atembeweguns-Korrektur zu einer Erhöhung des Signal-zu-Hintergrund Verhältnisses um bis zu 14 % und zu einer durchschnittlichen Reduktion der Plaqueweite um 23 %. Zudem konnten Artefakte aufgrund von fehlerhafter Schwächungskorrektur verringert werden.

Die im Rahmen dieser Arbeit entwickelten Techniken könnten ein Schritt zu einer zuverlässigeren und reproduzierbareren Bildgebung von Hochrisiko-Plaques in den Koronararterien mittels [^{18}F]NaF -PET/MR sein. Sie können einen Beitrag für eine verbesserte Diagnose und Therapiekontrolle liefern, und die individuelle Behandlung von Plaques ermöglichen.

Abstract

Recently, the radioactive PET tracer[^{18}F]NaF previously used in skeletal imaging found a cardiac application: it shows uptake in vulnerable high-risk atherosclerotic plaques likely to rupture and cause myocardial infarction. Subsequent studies revealed that coronary [^{18}F]NaF uptake provides a powerful predictor of myocardial infarction fatality in patients with coronary artery disease. PET acquisition times for [^{18}F]NaF scans of the coronary arteries, however, are of the order of half an hour and therefore cardiac and respiratory motion can strongly affect the PET image quality. These motion artefacts include blurring of the uptake, as well as misalignment of PET emission data and attenuation correction map. Motion compensation approaches such as motion-gating are limited to plaques with large uptakes as they rely on a high number of PET counts in the gated data - however, the uptakes in the coronary arteries are small. New approaches are needed to make use of all acquired data and motion-correct them to the right anatomical location while mitigating the artefacts.

A suitable candidate for the implementation of such techniques is PET/MR, a hybrid modality allowing the simultaneous acquisition of the fast, high-resolution anatomical images of MR and the complementary, metabolic information of PET. The overall aim of this thesis was to use MR-based, cardio-respiratory motion correction to improve the visualisation and quantification of coronary plaques imaged with simultaneous [^{18}F]NaF PET/MR. To enhance the quality of the motion-models, a 3D fat-water separation sequence was extended to sample k-space along a novel trajectory that is better suited for cardiac and respiratory resolved imaging. Motion-correction was incorporated into a high-resolution 3D model-based fat-water separated MR reconstruction. Furthermore, a simulation framework for dynamic PET/MR data with motion ground truth information was developed. It was used to assess the accuracy of synergistic image registration that were employed to generate cardiac and respiratory motion models.

The developed techniques were first validated in a stand-alone cardiac MR application where they could improve the local sharpness of 3D epicardial fat structures in patients by 40 % using motion correction. Afterwards, they were applied to [^{18}F]NaF PET/MR patients suffering from atherosclerosis. In the detected coronary [^{18}F]NaF -positive lesions in patients, the use of cardio-respiratory motion correction yielded an increase in target-to-background ratio of up to 14 % and an average reduction of plaque width of 23 %. Furthermore, artefacts due to attenuation misalignment could be mitigated. The techniques developed in the scope of this thesis could be a step towards more reliable and more reproducible imaging of high-risk coronary plaques using [^{18}F]NaF PET/MR. They can contribute to an improved diagnostic quality of coronary artery disease and allow for its patient-specific treatment.

Acknowledgements

This work could have not been completed if it was not for the continuous support I experienced from my colleagues, friends and family, to all of whom I feel greatly obliged.

I want to thank my supervisors, Christoph Kolbitsch and Tobias Schäffter for their invaluable input over the course of this project.

Furthermore, I am grateful to the esteemed people from CCP SyneRBI for teaching me a lot, and from Charité Berlin and BIOQIC for supporting me over the course of this project with their time and effort.

My heartfelt gratitude goes to my colleagues at the PTB, not only for their professional support but also for making me look forward to coming to work again every morning.

Finally, I want to thank my friends and family, who made me look forward to leaving work again every evening, and made me an incredibly rich man on a student's salary.

Table of Contents

List of Figures

List of Tables

Abbreviations

TV total variation 58

TVT total variation in time 58

2D two-dimensional 9

3D three -dimensional 12

4D four -dimensional xiii, 23

5D five -dimensional xiii, 16

AC attenuation correction xiv, 2

AIF arterial input function 31

AVG motion-averaged image reconstruction 60

BART Berkley Advanced Reconstruction Toolbox 38

CAD coronary artery disease 1

CBR contrast-to-background ratio 78

cMR cardiac magnetic resonance 3

cMRA coronary magnetic resonance angiography 7

CNN convolutional neural network 94

CNR contrast-to-noise ratio xiii, 32

cr-MCIR cardio-respiratory motion-corrected image reconstruction 59

CS compressed sensing 30

CT x-ray computed tomography 1

CVD cardiovascular disease 1

D plaque width 83

DCE dynamic contrast enhancement x, xiii, xv, 21

ECG electrocardiogram 12

FDG fluorodeoxyglucose 2

FFD free-form deformation 29

FFT Fast Fourier transform 93

FOV field of view 42

GT ground-truth 20

ISMRMRD International Society of Magnetic Resonance in Medicine Raw Data format 20

LGE late gadolinium enhancement 9

LM lipotamous metaplasia 8

LOR line of response 23

MCIR motion-corrected image reconstruction 3

MI myocardial infarction 1

MLEM maximum likelihood expectation maximisation 14

MoCo motion correction xiv, 3

MR magnetic resonance xiv, 2

MRI magnetic resonance imaging 1

MVF motion vector field xiii, 23

NMR nuclear magnetic resonance 9

NUFFT non-uniform Fast Fourier Transform 93

OCT optical coherence tomography 1

OSEM ordered subset expectation maximisation 15

PCA principal component analysis 16

PET positron emission tomography xiii–xv, 2

PET/CT sequential hybrid positron emission and x-ray computed tomography xiii, 2

PET/MR hybrid positron emission and magnetic resonance tomography i, ix, xiii–xv, 2

PVC partial volume corrections 18

r-MCIR respiratory motion-corrected image reconstruction 59

RF radio frequency 17

ROI region of interest 26

RPE radial phase-encoding 4

RTA reconstruct-transform-average 14

SIRF Synergistic Image Reconstruction Framework 20

SNR signal-to-noise ratio 22

STIR Software for Tomographic Image Reconstruction 20

SUV standardised uptake value 83

TBR target-to-background ratio 18

TOF time-of-flight 18

US ultrasound 7

XCAT extended cardiac-torso phantom 21

1

Introduction

In the modern world with its ageing population, cardiovascular disease (CVD) is the leading cause of death worldwide[1–3]. While there are preventive measures both at the population and the individual level to reduce the prevalence of the disease, (e.g. through sufficient physical activity or a healthy diet) [4] ischaemic heart disease and stroke combined are still responsible for over 15 million deaths in 2016 [5].

Ultimately, the disease leads to an occlusion of parts of the vascular system restricting blood supply to regions of the brain causing a stroke, or the heart leading to acute coronary syndrome and likely myocardial infarction (MI). Also in Germany, despite a long-term positive trend reducing the number of deaths associated with cardiac ischaemic disease by 40 % since 1990 [6] coronary artery disease (CAD) remains the most lethal disease. With a predicted increase of CVD prevalence, this indicates that ever-improving technologies are required for an earlier diagnosis and treatment of the underlying pathological processes.

The pathological condition underlying CVD is atherosclerosis [7, 8]. Fatty deposits, so-called atheroma or plaques, form on the interior walls of the arterial vascular tree, in particular in the coronary system. Over years and decades, in an inflammatory process, they develop fibrous structures and necrotic cores, which can calcify leading to ossification of part of the atheroma. The final stage of atherosclerosis is that a thin fibrous cap on the surface of these plaques ruptures and a thrombus occurs inside the vessel, often with fatal consequences.

Imaging plaques in the vascular system of the heart demands a high spatial resolution as the coronary arteries are at most a few millimetres in diameter. Intravascular optical coherence tomography (OCT) can detect individual coronary plaque substructures and identify vulnerable, thin fibrous caps and necrotic plaque cores but the procedure is highly invasive. There are also other, non-invasive tomographic modalities available to image atherosclerosis and its effects. Cardiac x-ray computed tomography (CT) can spatially resolve coronary arteries and show calcification of the vessel, assigning the so-called calcium score. Also with cardiac magnetic resonance imaging (MRI), angiographies can be performed that yield sub-millimetre resolution images of the coronary arteries [9]. However, in standard qualitative MRI, plaques are visualised as a narrowing of the vessel with no information

on its substructure, and novel contrast agents for plaque classification [10] are still in a research phase. Cardiac positron emission tomography (PET) on the other hand can provide information on plaque metabolism and hence potentially plaque composition. While [^{18}F]-fluorodeoxyglucose (FDG) is used in clinical routine to image MI using perfusion it could in principle be employed to image its uptake in plaques due to inflammation. However, FDG is not used in clinical routine, because the large tracer uptake in the adjacent myocardium usually completely obscures any uptake in the vessel wall[11].

Recently, the radiotracer [^{18}F]sodium fluoride ([^{18}F]NaF) usually used for skeletal imaging or to detect osseous prostate cancer metastasis imaging [12] has found novel applications in vascular PET imaging. It could be shown in a first study using [^{18}F]NaF sequential hybrid positron emission and x-ray computed tomography (PET/CT) that the tracer uptake was significantly higher in coronary plaques meeting the risk criteria for vulnerability compared to non-vulnerable plaques [13]. It was also found that ruptured, culprit plaques in patients with MI showed the highest uptake compared to the non-culprit plaques. Furthermore, a recent study has shown a high prognostic value [14] in predicting a fatal outcome of MI. Hence, despite that the underlying physiological processes are not entirely understood, [^{18}F]NaF -PET offers the first possibility to non-invasively identify high-risk coronary plaques. Its precise uptake quantification could lead to a more patient, and plaque-specific treatment of CAD.

Yet, a reliable and reproducible quantification is still challenging as coronary [^{18}F]NaF uptakes are small and at the limit of the spatial resolution of PET imaging systems. Additionally, as the PET data acquisition process is in the order of 30 minutes physiological motion of the coronaries due to respiration and the heartbeat impairs the quality of the resulting images. Motion does not only introduce blurring into the PET images, it can also compromise an accurate attenuation correction (AC) for coronary arteries and hence lead to erroneous quantification. As the vessels run close to the interfaces between cardiac, fat and lung tissue they are vulnerable to a motion-induced emission-attenuation-data mismatch [15].

The modality simultaneous hybrid positron emission and magnetic resonance tomography (PET/MR) [16, 17] has the potential to address these challenges. It allows the simultaneous data acquisition of both modalities which is not possible for other, sequential hybrids such as PET/CT. Hence, the complementary information provided by magnetic resonance (MR) is not limited to serve merely as an anatomical reference of the uptake but can be used to improve the PET image quality. MRI provides high-resolution 3D anatomical information with flexible soft-tissue contrast, while also the ability to measure cardiac and respiratory motion. These can be utilised to create a patient-specific model of cardio-respiratory motion, and a high-resolution, tissue-detailed, and motion-synchronised attenuation map.

Eventually, this synergy could enable the PET/MR modality to generate fully-motion corrected, 3D high-resolution maps of high-risk coronary plaques in PET images, supported by matching MRI anatomical information from one non-invasive data acquisition.

Scope of the Thesis

This thesis aims to develop cardiac and respiratory non-rigid motion correction for uptakes in the coronary system measured with simultaneous [^{18}F]NaF PET/MR. The individual components dedicated and required to yielding a high-quality PET/MR motion correction are:

Numerical PET/MR Simulation Image registration-based motion models generated from patient data suffer from the unavailability of ground truth motion information and their precision can not be evaluated. Hence, a numerical simulation of simultaneous PET/MR with available ground-truth motion information is required to serve as a testing ground for the generation of motion-models on which their accuracy could be assessed. Furthermore, it would allow optimising image registration hyperparameters on a simulated dataset and transfer them onto patient data to achieve an optimal motion model in-vivo.

Fat-Water Separated MRI A high-resolution MRI acquisition as a source for both motion and attenuation information is required. A fat-water separated imaging approach unifies multiple advantages: the additional positive fat contrast could be beneficial to the motion-estimation, while at the same time providing a detailed MR-derived AC.

[^{18}F]NaF PET motion correction (MoCo) Finally, simultaneous [^{18}F]NaF PET/MR patient data needs to be acquired. To assess the effect of cardio-respiratory motion correction on [^{18}F]NaF uptake strength and localisation the acquired data need to be reconstructed in a motion-corrected image reconstruction (MCIR). Motion and AC can be derived from the MRI, while the PET MCIR should include motion-modelling for both AC and emission data to address both sources of motion-induced artefacts.

Thesis Outline

The following chapters address the components and challenges listed above and discuss developments concerning the state of the art:

Chapter 2 The physiological cause of atherosclerosis and its inherent risks are briefly introduced. Afterwards, an overview of the research on imaging atherosclerosis using [^{18}F]NaF PET is given. The chapter is completed by discussing cardiac magnetic resonance (cMR) imaging techniques for depicting fat and water and advances in motion-corrected PET/MR imaging.

Chapter 3 The development, implementation and application of a numerical simulation framework for dynamic PET/MR data are presented. The framework is designed to simulate a simultaneous PET/MR examination with the ground-truth motion information readily available. The input and output are based on existing patient raw data, such that the simulations integrate seamlessly into existing reconstruction pipelines. The framework is tested in a series of PET and MR applications including the evaluation of the accuracy of synergistic registration algorithms. This chapter has been published in parts as [J1] and [J4] and were accepted for publication [J3].

Chapter 4 In this chapter a free-breathing, 3D fat-water separated acquisition-reconstruction framework for cardiac MR is developed. It is aimed at providing a comprehensive MR examination dedicated to image small structures in the heart, while simultaneously being able to supply anatomic and motion information as well as serve as a basis for PET attenuation correction. As a first step, a novel phase-encoding pattern for a 3D radial phase-encoding (RPE) trajectory is presented. Subsequently, a model-based reconstruction combining fat-water separation and motion correction as well as compressed sensing is developed. A synergistic PET/MR image registration algorithm is applied to fat-water separated MR reconstructions to generate motion models for the respiratory and cardiac motion of the heart. The numerical simulation framework introduced in the previous chapter is applied to assess the accuracy of the motion models.

Chapter 5 The methods outlined in chapter 4 are validated for cardiac adipose tissue imaging. The fat-water separated reconstruction is first validated against images acquired in a clinical routine protocol. Afterwards, data from nine patients are evaluated concerning the quality of motion-resolved and motion-corrected images. A local sharpness metric is used to quantify the improvement of the sharpness of fat structures in the heart. Finally, possibilities to accelerate the acquisition process are discussed.

Chapter 6 This chapter describes the application of the methods developed in the scope of this thesis to simultaneous [^{18}F]NaF PET/MR data acquired in ten patients. The fat-water separated image reconstruction is applied to simultaneous PET/MR data. The generated motion models are used to correct respiratory and cardiac motion for the fat-water separated MR data. This minimises motion artefacts and provides a high-resolution 3D image of the patients' anatomy. These fully motion-corrected images are segmented into an attenuation map for the PET reconstruction. At the same time, the motion models are utilised in the PET image reconstruction to also minimise cardiac and respiratory motion artefacts. The effect of motion correction on focal hotspots in the coronary arteries in the motion-corrected PET reconstructions is assessed quantitatively. Parts of this chapter were published as [J2].

Chapter 7 A summary of the challenges faced and the associated solutions developed in the scope of this thesis is given. The results are discussed concerning state-of-art limitations and potential further developments. A short outlook for future publications is given.

2

Background

This chapter gives an overview of atherosclerosis as well as of previous studies of atherosclerosis imaging using [^{18}F]NaF PET. Subsequently, pathologies associated with myocardial fat infiltration are discussed and fat imaging using MR is introduced to the reader. The chapter is completed by an overview of current motion compensation methods and image reconstruction techniques employed in clinical routine cMR and cardiac PET/MR as well as state of the art developments. The following sections are based on [8], [18] and [19].

2.1 Atherosclerosis

Plaque Progression

Atherosclerosis is a multifactorial disease which is responsible for most deaths worldwide [20]. While the exact causal mechanisms of plaque development are not understood, an increase of low-density lipoproteins in the blood can already lead to the development of clinically relevant atheroma, commonly referred to as plaques. However, multiple risk factors such as a poor diet, smoking and a genetic predisposition can increase the risk of disease development [21] as they affect the interaction between aterial wall and proteins. Plaques initially form in reproducible positions in the arterial vascular system that are related to shear stress generated by blood flow [22]. These include in particular the bifurcations of the coronary arteries and the carotids. With time, however, the disease spreads to adjacent sites as well. In elderly patients most of the coronary tree can eventually be affected. The different stages of plaque formation are depicted schematically in Fig. 2.1. This disease progression usually stretches over several decades. Plaque formation is initiated when lipoproteins bind to the intimal vessel cells and cause a reaction of the immune system leading to inflammation. The inflamed sites will subsequently form foam cells from present macrophages which serve as deposits for the lipids. Until this point, the disease is still reversible if the stimuli causing them disappear, e.g. through a diet change. However, a formation of a necrotic core leads to the formation of a fibroatheroma which is not reversible anymore. Apoptosis inside the plaque and additional necrosis of foam cells and smooth muscle cells in the vessel wall can lead to the development of

a necrotic core and potential non-functioning removal of apoptotic cells accelerates its growth. The plaque content will become more and more fibrotic, in particular, the boundary with the lumen forms a potentially unstable fibrotic cap over the plaque interior. In the last stages of plaque development, calcification can set in that can cover the whole necrotic core. A progressive occlusion of the vessel will eventually lead to stenosis.

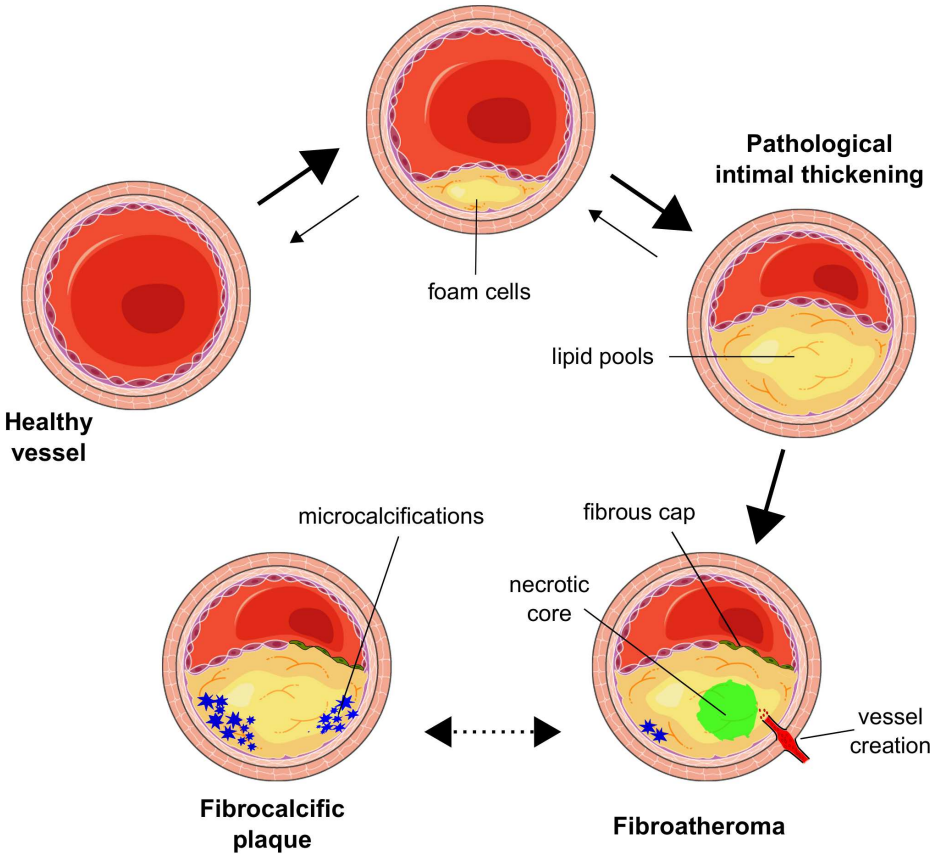

Figure 2.1: Atherosclerotic disease progression. Starting from a healthy vessel foam cells and lipid pools accumulate leading to a pathological intimal thickening. After a fibroatheroma is created the process is no longer reversible. Apoptosis creates a necrotic core and angiogenesis leads to intra-plaque haemorrhage. The thin fibrous cap makes the structure unstable. After a calcification of the necrotic core, a fibrocalcific plaque is created. Figure based on [8], Figure 1. This figure was created using graphics from [23].

A more dangerous event is so-called plaque rupture. It occurs if the fibrous cap is damaged. This is depicted schematically in Fig. 2.2. The contents of fibroatheromas are highly thrombogenic such that when they spill into the lumen a clot forms. This thrombus can be transported to a more narrow part of the vessel tree where it can cause a stroke or MI with often fatal consequences.

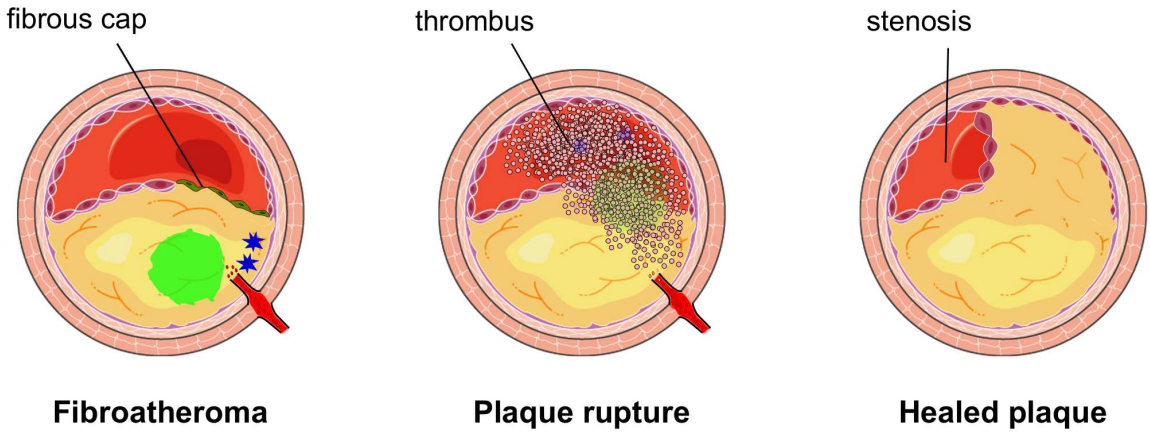

Figure 2.2: Plaque rupture. Left: plaque prone to rupture with thin fibrous cap, necrotic core intra-plaque haemorrhage and calcifications. Middle: the fibrous cap ruptures and the thrombogenic plaque core accesses the lumen. Right: healed plaque leading to stenosis. Figure based on [8], Figure 3. This figure was created using graphics from [23].

[^{18}F]NaF Atherosclerosis Imaging

Diagnostic methods that can give information on the plaque structure and composition and allow to quantify the risk of plaque rupture are of central interest when for atherosclerosis treatment planning [24]. Invasive methods such as intravascular OCT or ultrasound (US) can resolve plaque substructures and identify fibrotic caps and other risk factors. However, they are highly invasive procedures and can cause a high level of patient stress.

While CT and MR can non-invasively create high-resolution images of the heart, their access to plaque substructures is limited. The so-called calcium score can be computed based on CT images showing macrocalcifications in the vessels that correlate with cardiovascular event risk and mortality [25]. In the coronaries due to the challenges posed by motion and the small vessel size, coronary magnetic resonance angiography (cMRA) can only be applied to image stenosis qualitatively as long as dedicated contrast agents are still in early research stages [10].

PET is a non-invasive imaging modality that allows to image metabolism, usually acquired as a combined modality PET/CT or PET/MR. There are several tracers which can provide information about plaques. [^{18}F]FDG can be used to image the inflammatory processes occurring in the plaque but in studies this tracer did not show increased uptake in patients with atherosclerosis [13, 26]. An additional disadvantage is that the large myocardial activity makes quantification of uptake in the adjacent coronaries difficult [26].

[^{18}F]NaF is a PET tracer which can be used to image calcification processes. Initially used for skeletal imaging it has been used in oncology to detect osseos bone, prostate and breast cancer metastasis [12, 27]. Since risk factors strongly vary based on the composition of plaques, [^{18}F]NaF adds clinical value in the analysis of plaque risk over global detection methods such as blood markers or inferring disease progression based on non-invasive US of the carotids [8]. [^{18}F]NaF has been used in prospective clinical trials to evaluate its potential using PET/CT. The uptake in the coronaries was measured and correlated with both the calcium score that indicates the degree of macrocalcifications as well as the occurrence of atherosclerosis [26]. It

was found that 40% of the patients with increased calcium score did not show any significant [^{18}F]NaF uptake. This suggests that the tracer labels active micro-calcifications instead of calcified macro-structures. This was confirmed for patients with aortic stenosis [28].
An detailed analysis of [^{18}F]NaF uptake in coronary plaques using multiple imaging modalities has between presented in [13]. Patients with stable angina or MI were examined using [^{18}F]NaF PET/CT and intravascular US and histological data were used as a reference for the plaque substructure. All ruptured carotid plaques showed [^{18}F]NaF uptake. In addition, uptake in coronary plaques correlated well with high-risk characteristics obtained with US. Further studies confirmed these findings using additional data acquired by intravascular OCT [29].
A subsequent in-depth analysis of the tracer distribution and molecular mechanisms involved in [^{18}F]NaF binding in plaques was performed in multiple studies [30, 31]. In [30] histological analysis was combined with μPET and PET imaging showing how [^{18}F]NaF binds to the calcification surfaces and can distinguish between micro-and macrostructures. This is relevant because while the occurence of macrocalcifications is associated with higher patient risk they can act as a stabilising structure to individual plaques. In [31] the authors concluded that the [^{18}F]NaF uptake is a sign of early plaque calcification while it decreases with increased deposition of calcium in the arterial walls. Recently its use as a prognostic predictor for fatal or non-fatal MI outcomes [14] where it has been shown that [^{18}F]NaF outperformed other well-established clinical markers.
When using [^{18}F]NaF PET/MR, it was found that there is a strong correlation between the detected plaque locations with the ones obtained from PET/CT [32]. However, it was highlighted that MR-based AC map generation can be hindered by susceptibility artefacts caused by coronary stents [32, 33].
Physiological motion poses a major challenge when imaging the coronaries. While in the initial studies physiological motion was not corrected for, PET motion compensation techniques in the heart were addressed in PET/CT for cardiac motion of the coronaries [34, 35] and in PET/MR of the aortic valves [36]. They showed that correcting for motion leads to a better and more reliable quantification of the tracer. However, for PET/CT CT and PET have to be acquired one after the other instead of simultaneously (c.f. Fig. 2.7) the used motion models were PET based. The motion estimation from [^{18}F]NaF PET frames is challenging since they do not contain anatomical information except for bones whose motion does not correlate with the motion of the heart. The acquired CT data could only be used to support the registration by improving the alignment of anatomical features. For [^{18}F]FDG PET/MR multiple motion correction approaches obtain motion information from simultaneously acquired high-resolution MR data. An overview of the current state of PET/MR MoCo is given at the end of this chapter.

2.2 Adipose Tissue Imaging in Cardiac MR

Fat in the Heart

Fat is a major tissue component of the human body and naturally occurs in the healthy heart as pericardial and epicardial adipose tissue. However, the infiltration of the myocardium with adipose cells, known as lipotamous metaplasia (LM) is associated with risk factors for

cardiovascular diseases as has been shown in autopsy studies [37]. LM is associated with poor prognosis after MI [38], and sudden cardiac death [39, 40] and is also a significant predictor for all-cause mortality, sustained ventricular arrhythmia and heart failure [41]. Furthermore, an excess in epicardial adipose tissue is associated with atrial fibrillation and ventricular arrhythmias due to the local changes in electrical conductivity [42–47]. To image fat in the heart with cMR, mainly fat-water separation techniques have been used [48] which will be discussed in more details in the next section. High prevalence of fat deposition in healed MI could be detected using two-dimensional (2D) cardiac fat-water separation [49]. Furthermore, fat-water separation could show the association of LM with dilated cardiomyopathy [50]. Also, cMR could be employed to measure the volume of epicardial fat and showed it to be correlated with increased myocardial fat, interstitial myocardial fibrosis and damaging effects on myocardial contractile function [51]. In similar studies, cMR was used to measure functional parameters showing for patients with metabolic syndrome that epicardial and pericardial fat depositions are linked to left-ventricular diastolic functional deficits, such as a reduced peak-filling-rate [52]. A worsening of cardiac function could be due to the replacement of compact scar by compressible adipocytes [53]. While fat often impairs cardiac function, lately, it has also been found that patients suffering from heart failure with a preserved ejection fraction have significantly more adipose cardiac tissue when compared to patients with reduced ejection fraction [54].

Magnetic Resonance Fat-Water Imaging

CMR serves as a non-invasive tool to image cardiac fat [38, 55] and has been employed to assess myocardial fat deposits. The physical basis of imaging different chemical components is the so-called chemical shift. The molecular environment of hydrogen bound in lipids causes their spins to have a different Larmor frequency compared to hydrogen in water. The chemical frequency shift depends on the external magnetic field B_0:

$$\Delta_\omega(Hz) = \Delta_\omega(ppm) \cdot \frac{\gamma}{2\pi} B_0 \tag{2.1}$$

where $\Delta_\omega(ppm)$ is the chemical shift relative to a reference and γ is the gyromagnetic ratio. The different chemical components can observed in a nuclear magnetic resonance (NMR) spectrum, schematically displayed for fat and water in Fig. 2.3.

While the frequency offset gives MR flexibility in imaging the different components [18] it is also associated with multiple artefacts: The fat voxels will accumulate additional phase depending on the bandwidth encoded in the readout. This leads to a global displacement of fat in readout direction. Furthermore, partial volume effects lead to the black boundary artefact due to signal cancellations. Finally, fat typically has a T_1 value similar to those achieved in blood after injection of a T_1 contrast agent. Hence, for example in T_1 weighted post-contrast cMRA the high fat signal surrounding coronary arteries will obscure the blood signal or can be mistaken for fibrosis in late gadolinium enhancement (LGE) imaging.
There are multiple techniques dedicated to nulling the fat signal to avoid any artefacts. These include fat saturation, water excitation and Short T_1 Inversion Recovery (STIR).

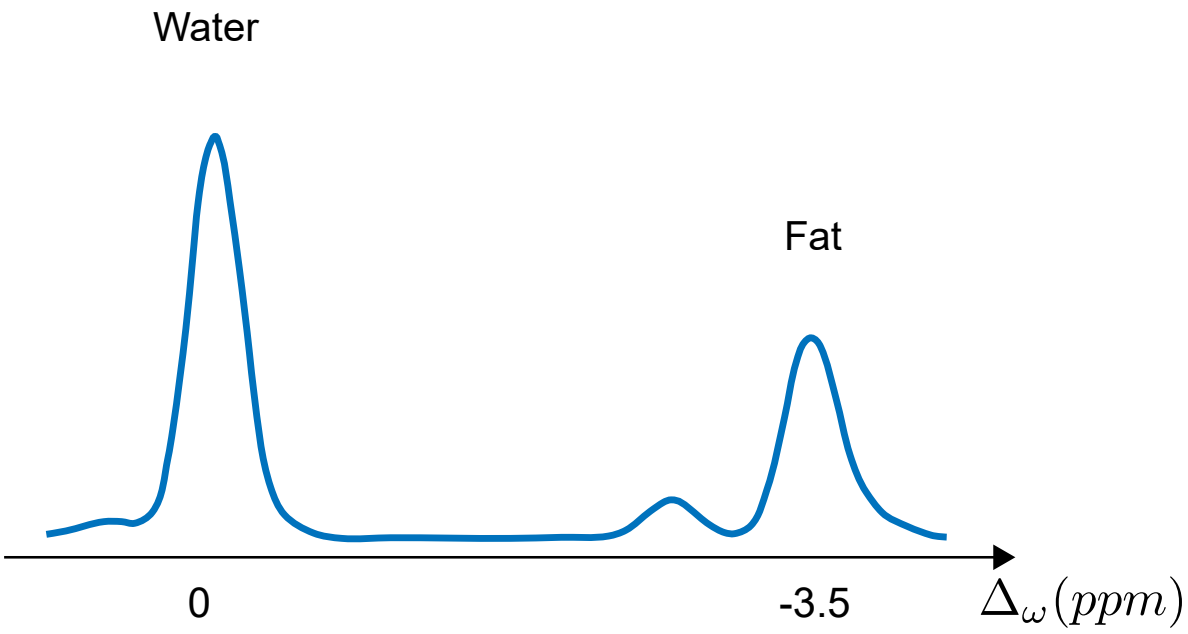

Figure 2.3: Fat and water NMR spectrum. Schematic depiction of an NMR spectrum containing both fat and water. The spectrum is normed to the water peak at 0 ppm while the fat-peaks are shifted in frequency. The main fat peak lies at approximately −3.5 ppm relative to water. The offset in frequency can be computed using Eq. (2.1). Figure based on [18, 56].

While these methods are effective to suppress the fat and are used in clinical routine, they can be impaired by B_0 and B_1 inhomogeneties which lead to incomplete fat suprression or even supression of water signals [57]. As discussed above, imaging fat is of central interest in cMR as it is associated with multiple cardiomyopathies. Since the above signal suppression techniques are only providing negative fat contrast [58–60] they only provide little diagnostic information about cardiac fat. Chemical-shift based separation approaches on the other hand, yield both fat and water with positive contrast. They use multiple echo acquisitions and are referred to as Dixon methods (Fig. 2.4). During data acquisition voxels containing fat F acquire a phase relative to the water signal W such that the complex signal $s(t)$ at time t consists of:

$$s(t) = W + e^{2\pi i \Delta_\omega \cdot t} F. \tag{2.2}$$

Choosing the correct timing for the MR echo time during data acquisition and a subsequent combination of the images at the different echo times can be used to separate the fat from the water component. To ensure data consistency the different images are not acquired one after another, but instead, multiple echoes are acquired during each T_R interval.

As the chemical shift depends on the external magnetic field, local inhomogeneities, i.e. $B_0 = B_0(\boldsymbol{r}) = B_0 + \Delta B_0(\boldsymbol{r})$ will lead to a spatially dependent chemical shift. This can be accounted for by a three-point Dixon methods [61] in which a third echo is acquired and the off-resonance map $\Phi = \frac{\gamma}{2\pi}\Delta B_0(\boldsymbol{r})$ is computed. In order to accurately calculate fat, water and an off-resonance map from (almost) arbitrary echo-times, iterative approaches have been proposed [62–65]. A comprehensive overview of techniques employed in clinical routine is given in [66].

These methods can also be incorporated into a model-based separation of fat and water [67,

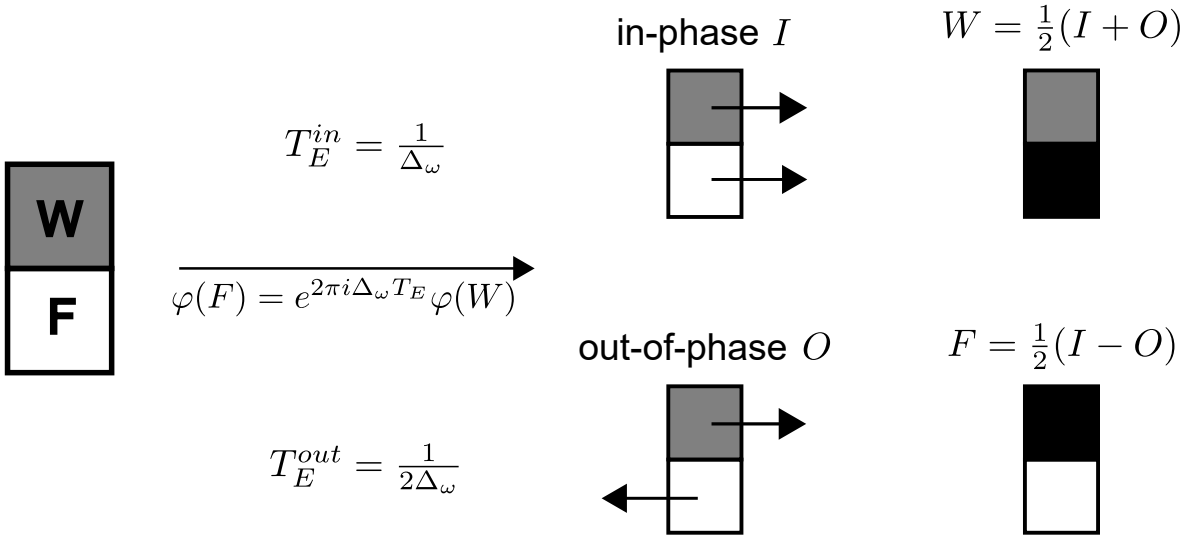

Figure 2.4: Dixon imaging. Principle of chemical shift-based fat-water separation (Dixon imaging). Left: a two-pixel image containing water (W) and fat (F) is acquired at two different echo times. Center: reconstructed images for individual echoes. Arrows depict the direction of the complex signal phase. Right: computation of the images separated into water and fat by combining the in- and out-of-phase reconstructions.

68] which skips the reconstruction of images at individual echo times. The forward operator projects water and fat (W, F) to k-space for different echo times T_E:

$$\begin{aligned} k_c(T_E) &= \hat{E}_{(c,T_E)}(W, F) \\ &= \mathcal{F}(C_c e^{2\pi i \Phi T_E} \cdot W) + D(T_E)\mathcal{F}(C_c e^{2\pi i \Phi T_E} \cdot F) \end{aligned}$$

where C_c are the receiver channel sensitivities, Φ is the off-resonance map and $D(T_E)$ a chemical shift model. To invert such a forward model an iterative optimisation scheme is required. This allows for regularisation on the final fat and water images instead of the intermediate indivdiual echo images. A detailed discussion of this approach and an extension to include physiological motion is given in chapter 5.

2.3 Motion Compensation in Image Reconstruction

For both cMR and cardiac PET/MR physiological motion is one of the major challenges. Respiration and heartbeat impair the image quality of both modalities. Motion introduces a blurring into the reconstructed images or can cause more complex artefacts in MR such as ghosting. In PET, a mismatch between emission data and AC map can severely impair the images if uptake is located near a tissue boundary with large differences in attenuation values, such as between soft tissue and lung. Hence, there is the need to address motion during the acquisition and reconstruction. General aspects of image reconstruction will be discussed in the following including retrospective MoCo and the currently available techniques in both cMR and cardiac PET/MR.

PET and MR image reconstruction is an inverse problem: the acquired raw data are the result of a known physical measurement applied to an unknown tissue and tracer distribution in the

human body. The task of the reconstruction is to infer the underlying distributions from the acquired data, effectively inverting the measurement process. The measurement process of both PET and MR can be described by:

$$\boldsymbol{y} = \hat{E}\boldsymbol{x} + \boldsymbol{b} + \boldsymbol{n}, \tag{2.3}$$

where $\boldsymbol{y}$ is the measured data as a column vector, $\hat{E}$ is the acquisition model, i.e. usually containing a transformation such as the Fourier or Radon transform, $\boldsymbol{x}$ is the image as a column vector, $\boldsymbol{b}$ is some additive background term and $\boldsymbol{n}$ is the measurement noise. The individual components for both modalities are depicted in Fig. 2.5. Solving the reconstruction

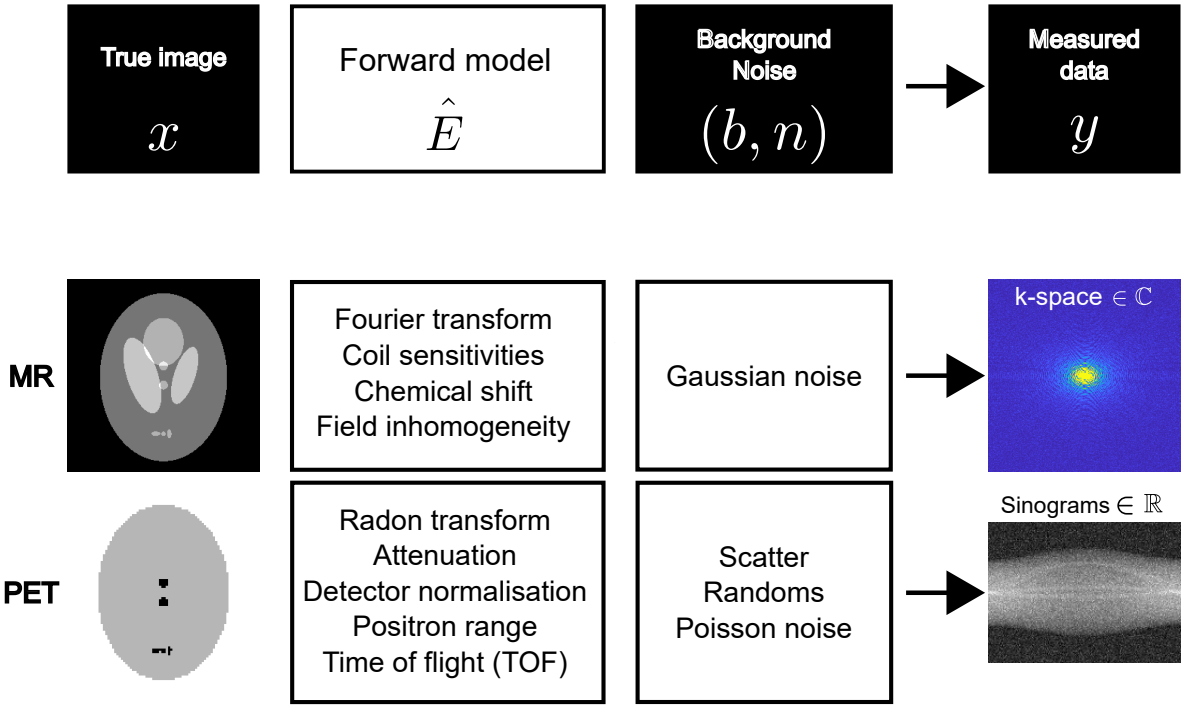

Figure 2.5: Components of PET/MR image reconstruction. Top: overview of involved components during PET/MR data acquisition as in Eq. (2.3). Bottom: non-exhaustive overview of the individual components for PET and MR reconstruction. Potential image contrasts are exemplified by Shepp-Logan phantoms.

problem involves determining $\boldsymbol{x}$ given $\hat{E}$, $\boldsymbol{y}$, $\boldsymbol{b}$, and a model for $\boldsymbol{n}$. The underlying assumption is that the image is a function of space only and not of time: $\boldsymbol{x} = \boldsymbol{x}(\boldsymbol{r})$. Motion, however, breaks this assumption. The acquired raw data are inconsistent for different time points, and solving the reconstruction problem with a time-independent acquisition model $\hat{E} \neq \hat{E}(t)$ will yield artefacts in the resulting images. For thoraic PET and three -dimensional (3D) MR this can be the case due to physiological motion. A common strategy is to exploit the periodic nature of the respiratory and cardiac motion. To correlate the acquired data with the motion there must be some surrogate signal $sig(t)$ available. For respiration, this can either be acquired using a respiratory belt [69], extracted from the data using MR navigators [70], or PET listmode data [71], while for cardiac motion usually an electrocardiogram (ECG) signal is used. Then the time-dependent image can be expressed as $\boldsymbol{x}(\boldsymbol{r}, t) = \boldsymbol{x}(\boldsymbol{r}, sig(t))$. As the same motion states will approximately reoccur in each motion cycle it is possible to acquire the necessary data over many cycles. This is depicted in Fig. 2.6.

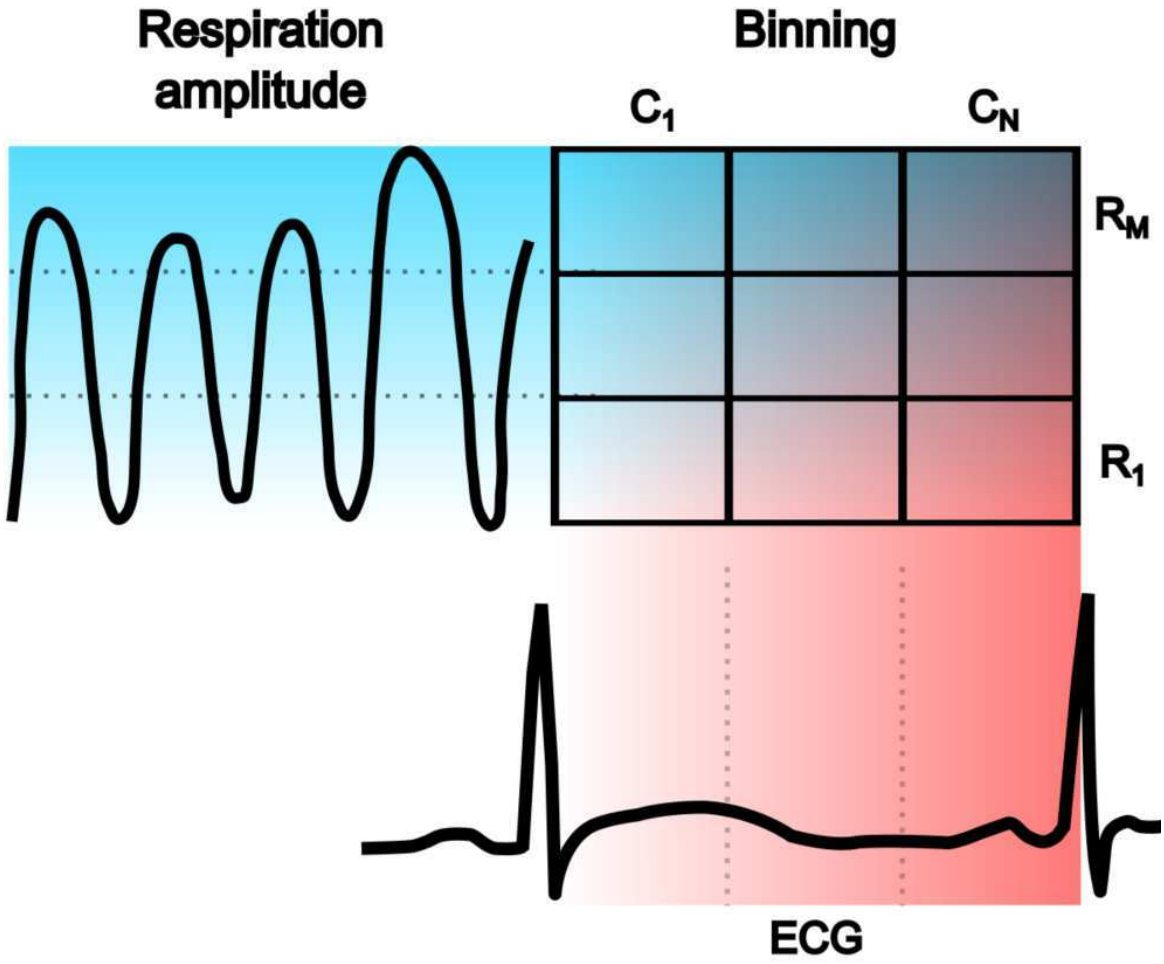

Figure 2.6: Schematic overview of the influence of motion gating. Left: respiratory amplitude, the increased amplitude is shaded from white to cyan and partitioned into M respiratory bins from $R_1 \dots R_M$. Bottom: ECG with increase in cardiac cycle shaded from white to magenta and partitioned into N cardiac bins from $C_1 \dots C_N$. The overlayed area under the grid represents the time available for data acquisition per cardio-respiratory motion cycle. Taking only a subset of the data will mitigate motion artefacts, however, at the cost of increasing the required scan time for the same number of MR data points and PET counts.

The straight-forward motion compensation approach is to only acquire the raw data when the surrogate signal sig is in a narrow window $sig \in (s-\delta, s+\delta)$. If data acquistion is only started for a fixed amount of time if $sig \in (s-\delta, s+\delta)$, then this is often called triggering. If data are acquired irrespectively of the value of sig, but then reacquired until $sig \in (s-\delta, s+\delta)$, this is commonly referred to as gating [72].

Motion Corrected Image Reconstruction (MCIR)

While gating is very reliable, effectively freezing the image, it comes at the cost of increased data acquisition time. In order to overcome this problem of long scan times, MoCo has been proposed. Here, data is acquired in different motion states and then the acquisition model is adapted to include a motion model.[1]

In the following, only periodic, non-rigid motion due to physiological motion in thoracic

[1] In MR there are also prospective approaches using so called slice-tracking [73–75]. Rigid or affine retrospective MoCo (e.g. head MoCo) can be carried out directly in the raw data space[76, 77]

applications is discussed. To resolve motion, based on the surrogate signal a retrospective binning of the data into different motion state $s \in 1, \ldots, S$ can be performed: $\boldsymbol{y} \rightarrow \boldsymbol{y}_s$. A motion model to be included in the acquisition model is defined as the transformation:

$$\boldsymbol{x}_s(\boldsymbol{r}) = (\boldsymbol{x}_0 \circ \boldsymbol{m}_s)(\boldsymbol{r}), \tag{2.4}$$

where x_s is the image in motion state s, x_0 the image in the reference state $s = 0$, and $\boldsymbol{m}$ is a coordinate transform with $\boldsymbol{m} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$. Usually, this transformation is written as:

$$\boldsymbol{x}_s = \mathbb{M}_s \boldsymbol{x}_0, \tag{2.5}$$

where the image $\boldsymbol{x}_0$ is a column vector and $\mathbb{M}_s$ is a matrix containing the respective interpolation weights. This, however, is more a question of notational convenience. An actual storing of this matrix would not be computationally feasible, and its application as in Eq. (2.5) is commonly carried out using Eq. (2.4).

To obtain the transformation in thoraic imaging it is usually generated using image registration. There are different available methods that are B-spline based. These are used to maximise image similarity (e.g. using mutual information or normalised cross-correlation) and possess intrinsic regularisations [78, 79]. Others optmise a pixel-wise sum-of-squared differences to register images, add regularisation by enforcing diffeomorphic motion fields and periodic boundary conditions on the motion or include the estimation into the reconstruction iterations to improve data consistency [80, 81]. Also, multi-modality registrations are available that estimate the motion simultaneously from PET and MR images [82]. This transformation is assumed to be known in the following.

Once a motion transform is available there are two ways to include it in a MoCo, either the so-called reconstruct-transform-average (RTA) approach or a full MCIR[83]. RTA reconstructs the data of the individual bins separately and then averages them:

$$\boldsymbol{x} = \langle \mathbb{M}_s^{-1} \boldsymbol{x}_s \rangle_s. \tag{2.6}$$

While for this approach the motion does not have to be included into the forward model it has some disadvantages. In particular in PET the individual bins might not contain enough counts to obtain a sufficient image quality. However, this approach is still useful if PET quantification is not the main objective [83]. A full MCIR on the other hand makes use of all acquired data simultaneously but requires to extend the reconstruction's encoding model. As an analytic inversion is typically used to invert a forward model that contains more than either the Fourier or Radon transform an iterative inversion is employed to perform MCIR.

PET MCIR

The following maximum likelihood expectation maximisation (MLEM) reconstruction formula serves as an exemplary reference for iterative PET reconstruction. It originates in the assumption that PET data contain Poisson noise and is given without proof. The image

MLEM algorithm uses the iterative update:

$$\boldsymbol{x}^{(k+1)} = \frac{\boldsymbol{x}^{(k)}}{\hat{E}^\intercal \mathbb{1}} \hat{E}^\intercal \frac{\boldsymbol{y}}{\hat{E}\boldsymbol{x}^{(k)} + \boldsymbol{b}} \tag{2.7}$$

where $x^{(k)}$ is the image estimate after k iterations of the algorithm, The initial image is initialised with as positive $x^{(0)} > 0$, and division is understood as element-wise. The denominator of the first term $\hat{E}^\intercal \mathbb{1}$ is the so-called sensitivity term which is the adjoint forward model applied to a unit projection set and ensures attenuation correction and detector normalisation. The second part backprojects the ratio between the acquired data and the prediction of the forward model based on the current image estimate. While this algorithm is not computationally feasible and multiple improved and accelerated version are available, including the ordered subset expectation maximisation (OSEM) algorithm [84], it can serve as a simple example in how to include motion into a PET reconstruction. To extend the MLEM algorithm by a motion model $\mathbb{M}$, the update step would be modified as [85]:

$$\boldsymbol{x}^{(k+1)} = \frac{\boldsymbol{x}^{(k)}}{\sum_s \mathbb{M}_s^\intercal \hat{E}^\intercal \mathbb{1}} \sum_s \mathbb{M}_s^\intercal \hat{E}^\intercal \frac{\boldsymbol{y}_s}{\hat{E}\mathbb{M}_s\boldsymbol{x}^{(k)} + \boldsymbol{b}}. \tag{2.8}$$

where s sums over all motion states.
This reconstruction approach yields the reconstructed image in the reference state $s = 0$ and the forward model deforms it into the respective motion states during reconstruction. This allows using all available data without the need for gating. In the same fashion, motion models can be included in an OSEM algorithm as well [83, 86].

MR MCIR

In contrast to PET, MR data contains Gaussian noise leading to an L_2 minimisation problem:

$$\hat{\boldsymbol{x}} = \min_{\boldsymbol{x}} \|\hat{E}\boldsymbol{x} - \boldsymbol{y}\|_2. \tag{2.9}$$

This problem can be iteratively solved with a range of optimisation algorithms, depending on the acquisition model $\hat{E}$ and potential applications such as parallel imaging and compressed sensing or fat-water imaging [67, 87, 88].
The inclusion of motion into a model-based MR reconstruction is performed analogously as to the PET case by concatenating the model with a known motion transformation $\mathbb{M}$:

$$\hat{E} \rightarrow \hat{E} \circ \mathbb{M} \tag{2.10}$$

A successful iterative inversion will automatically lead to an MCIR [89]. This is further discussed in chapter 5.

High Resolution Motion-Compensated Cardiac MR and PET/MR

Commonly in clinical routine cardiac imaging is carried out as multiple slices with high in-plane resolution but a large slice thickness of 4 mm to 8 mm. While the cardiac motion is usually compensated using cardiac triggering, the respiratory motion is addressed by acquiring the data during one or over multiple breath-holds.

To achieve an improved spatial coverage high-resolution 3D cMR imaging has been developed. Yet, its implementation in clinical routine suffers from long acquisition times. As breath-holds become difficult [90] motion compensation techniques for respiratory motion are required. The usual approach to respiratory motion compensation is the application of respiratory gating which leads to the rejection of a large fraction of the data (c.f. Fig. 2.6) and unpredictable scan times that depend on the cardiac frequency and a patient's breathing pattern [91–95]. Also, the correlation between diaphragm navigators and actual cardiac motion may change over time [96].

Instead of compensating for respiratory motion, compressed sensing reconstructions [88] can be used to resolve the respiratory motion in free-breathing 3D imaging applications from undersampled respiratory motion states. Exploiting sparsity along the respiratory direction has been suggested for radial (XD-GRASP) [97] and also extended towards cartesian trajectories [98]. XD-GRASP has also been used to ommit cardiac triggerng and resolve cardio-respiratory motion [99] with radial 3D encoding [100], providing five -dimensional (5D) high-resolution whole-heart coverage in acquisition times of the order of 15 min. A similar sequence has also been employed to achieve high-resolution whole-heart coverage [101] with a complete extraction of cardio-respiratory surrogate information from the MR data. As radial trajectories sample the k-space centre during every readout the cardiac signal can be extracted using principal component analysis (PCA) with a similar performance as an external ECG.

As resolving the motion relies on reconstructing images from heavily undersampled k-space data there are MoCo approaches that make use of all acquired data by including a motion model for respiration in the reconstruction. In cMRAs this has been addressed by using image-based navigators [98, 102–107] to estimate the motion. Low-resolution images can be acquired during the start-up pulses of steady-state sequences. From these images, a rigid motion model for the respiration can be extracted and applied to correct for motion between heartbeats. This acquisition-reconstruction approach mitigates most respiratory motion artefacts and provides high-resolution 3D images. But cardiac triggering restricts the acquisition period to a short window during each heartbeat such that still long scan times are required.

Instead of obtaining a rigid or affine respiratory motion model from image navigators, non-rigid motion models can also be estimated retrospectively by binning the acquired data, carrying out an image registration and applying the obtained motion model in MCIR [89, 108–111] as described above. Lately, regularisation methods in the form of so-called low-rank regularisation allowed for the reconstruction of high-resolution images from strongly undersampled data acquired in shorter scan times. This has been used in cMR and been combined with MoCo techniques [9, 112, 113] to produce high-resolution 3D images of the heart. This approach was also applied to 3D cardiac fat-water separated LGE imaging [114]. So far, however, there has not been a combination of physiological motion compensation with fat-water model-based reconstruction. Furthermore, in the above approaches cardiac motion

is not corrected for but compensated using cardiac triggering to acquire data in mid-diastole. While the above-mentioned approaches are applied to qualitative cMR there are of course substantial efforts towards quantitative parameter mapping. To name only a few, some of the above methods were applied for 3D T_1 and T_2 mapping [115, 116], and 2D low-rank methods were used in multi-tasking approaches [117, 118], or cMR fingerprinting [119, 120].
The above techniques were applied to cMR. Cardiac motion was mainly compensated by cardiac triggering or gating. However, for PET cardiac gating might not be feasible as insufficient counts strongly increase the image noise. Hence, in PET/MR there are other approaches in particular aimed at correcting cardiac motion. In Fig. 2.7 the layout of the two possible clinical systems for PET data acquisition, i.e. PET/CT and PET/MR is displayed. While in the PET/CT the data of the two modalities are acquired sequentially the hardware of a PET/MR system enables actual simultaneous data acquisition. This makes PET/MR a true hybrid modality opening up possibilities to exploit structural similarities and ultimately joint, synergistic image reconstructions [82, 121–124]. In research as well as clinically, it has

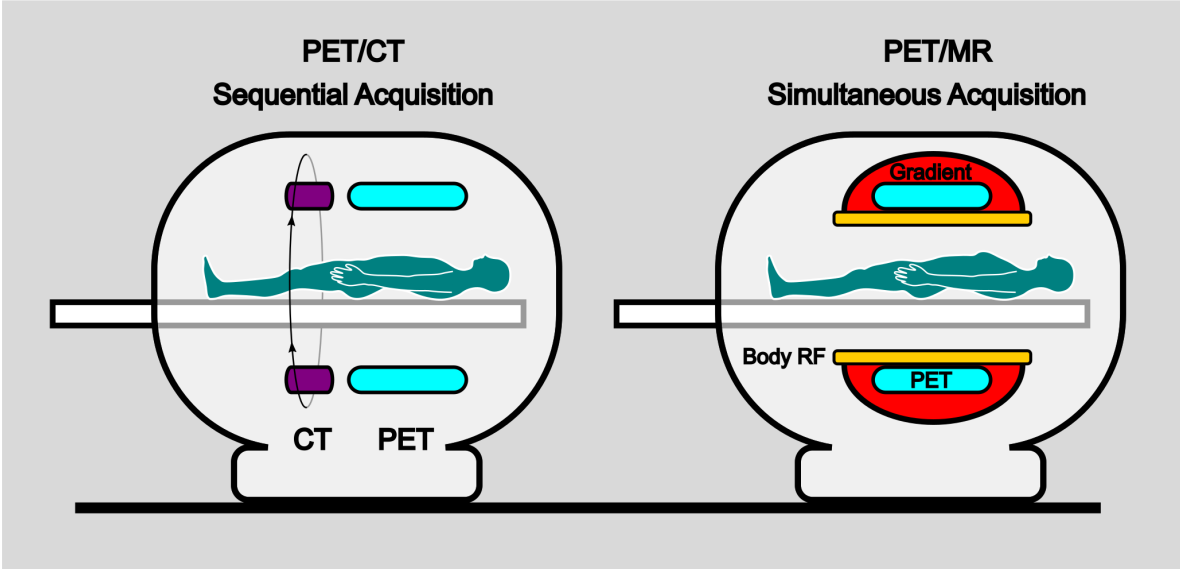

Figure 2.7: Comparison between PET/CT (left) and PET/MR (right). PET/CT: CT (purple) and PET (cyan) hardware are incorporated into one case, but the data are acquired in a sequential fashion. PET/MR: gradient coils (red), body radio frequency (RF) coils (yellow) and PET hardware are all combined. Figure adapted from [125] and [126].

been shown for multiple tracers used in oncology that respiratory MoCo strongly improves the detectability of lesions in simultaneous PET/MR [127–131].
Usually, as MR provides high resolution anatomical information, the motion models are generated from MR and applied subsequently to PET using an approach as described in Eq. 2.8. An overview of employed techniques is given in [132], usually using the tracer [^{18}F]FDG. Also, 2D multislice MR has been used purely dedicated to extracting a motion model and shown to be effective while keeping the additional MR scan time to generate the motion model in the order of one minute [133].
Respiratory MoCo while simultaneously providing high-quality and valuable anatomical information has been employed simultaneously with a 3D cMRA [134] and extended by estimating cardiac motion from a dual-phase cMRA [135]. Also separate retrospectively motion-binned MR acquisitions for cardiac and respiratory motion have been successfully

combined to improve PET image quality [136]. Also, holistic approaches have been used, estimating multi-phase cardiac and respiratory motion from a single MR acquisition [137, 138] while simultaneously extracting surrogate signals and providing high-quality 3D anatomical information, and have been extended to include respiratory-resolved attenuation correction [139].

Recently, also real-time MR respiratory motion models able to cope with irregular breathing patterns have been applied in abdominal imaging [140]. A positive effect of respiratory MoCo has been shown for myocardial perfusion PET/MR using [^{18}F]Flurpiridaz in porcine animal models [141]. This effect will potentially be confirmed in patients once the tracer is clinically available.

An overview of the use of PET/MR dedicated to imaging atherosclerotic plaques (not necessarily using [^{18}F]NaF) is given in [142]. As discussed above, MoCo techniques have also been applied to [^{18}F]NaF PET/CT imaging [34–36]. However, these addressed cardiac motion only. One study included respiration, and examined the effect of MoCo and PET partial volume correcition based on a numerical phantom study and four patients [143]. It was found that both cardiac and respiratory MoCo, and even more, partial volume corrections (PVC) lead to a strong increase in the target-to-background ratio (TBR) of the detected lesions. Additionally, there was no MCIR employed but instead relied on an RTA approach.

With respect to motion estimation PET/MR can improve with respect to PET/CT as the both modalities are acquired simultaneously. The methods developed for cardiac [^{18}F]FDG PET/MR are discussed in the final section of this chapter, and show how the high resolution anatomical MR data can be used to generate high-precision cardio-respiratory motion models and motion-resolved attenuation maps for both cardiac and respiratory motion. So far, respiratory motion has not been corrected for in [^{18}F]NaF PET/MR and also detrimental effects of AC mismatches due to motion have not been addressed so far but will be in the scope of this thesis.

Aside from an improved motion estimation, PET/MR has more advantages over PET/CT in [^{18}F]NaF imaging: the limiting factor in terms of acquisition time is the PET data. This neutralises advantage posed by the overall shorter time for data acquisition of CT compared to MR. MR can be used during the whole PET acquisition to acquire for example quantitative cardiac data. So far, PET/MR studies have not employed time-of-flight (TOF) information in their reconstructions which was available for the PET/CT systems. This could lead to a further improvement in resolution if data were acquired with a TOF PET/MR system.

Please note, that this discussion was focussed on "classical" imaging approaches. It omitted the advances in image reconstruction involving machine learning, which will most certainly allow for stronger data undersampling, leading to further decreases in scan and reconstruction times.

3

Flexible Numerical Simulation Framework for Dynamic PET/MR Data

Parts of this chapter have been published as J1, J3, J4, and J7.

3.1 Background

Positron-emission tomography (PET) and magnetic resonance imaging MRI are two commonly used medical imaging techniques. MRI yields high-resolution anatomical images and functional information such as tissue perfusion. PET, on the other hand, provides metabolic information with excellent sensitivity. Recently, they have been combined into the hybrid modality PET/MR allowing simultaneous data acquisition of complementary diagnostic information [16].

A major challenge for 3D high-resolution MR and especially PET in applications in the abdomen or thorax is physiological motion [34, 144–146], which impacts data quality during the long acquisition time of both modalities in the order of several minutes. Simultaneous PET/MR offers the possibility to obtain physiological motion information and utilise it in order to minimise motion artefacts. Several approaches use high-resolution anatomical MR images or dedicated MR sequences for accurate motion correction of PET data [129, 135, 138, 143, 147–149]. Other techniques combine both MR and PET information for image registration and use this to improve both MR and PET image quality [82, 137]. One of the main challenges of these motion correction schemes is the evaluation of their accuracy. Commonly there is no ground truth physiological motion information available. In addition, acquisition of ground truth image information without motion artefacts (e.g. using gated acquisitions) is often not feasible due to long acquisition times. Evaluation of motion correction schemes is therefore often limited to assessing the improvement in image quality between uncorrected and motion-corrected images. Although this gives an indication of the

motion-correction scheme's functionality, it does not necessarily provide information on how accurate it is.

To overcome this challenge numerical simulations are employed. While many existing simulation frameworks are dedicated to realistic simulations of a single modality [150–156] they are not primarily aimed at simulating simultaneous PET/MR data acquisition. There are PET/MR simulations which also incorporate physiological motion [83, 157, 158]. Nevertheless, these methods are limited to respiratory motion and cannot easily be extended to other types of motion or other dynamic processes such as contrast agent uptake. The design is tailored to certain applications making it challenging to adapt for other purposes. In particular, the frameworks' emphases are put onto the data generation to create more realistic anatomical input rather than utilising available solutions [151, 159, 160]. It is not possible to supply custom segmentations, motion models and surrogate signals and their output cannot be easily integrated into different reconstruction pipelines. Furthermore, simultaneous multidimensional motion simulations (i.e. combining both cardiac and respiratory motion) in a PET/MR framework remain unavailable so far.

In this chapter, a flexible framework to simulate dynamic simultaneous PET/MR raw data is presented. It allows for the combined simulation of a range of different motion types such as breathing or heartbeat. The framework provides ground-truth (GT) motion information for the simulated dynamic processes. Thereby, an evaluation of spatial accuracy of image registration and motion correction schemes can be performed. A detailed signal model is used allowing the simulation of image contrast for different MR sequence parameters. Additional dynamic processes such as uptake of MR contrast agents and PET tracers can be flexibly included, enabling a wide range of different PET/MR applications. The input of MR and PET acquisition parameters and output of MR and PET raw data files uses standardised formats (International Society of Magnetic Resonance in Medicine Raw Data format (ISMRMRD) for MR [161] and Interfile for PET [162]). This makes it compatible with standard reconstruction frameworks [163–165]. The framework is implemented in C++ and incorporated as a submodule of the Synergistic Image Reconstruction Framework (SIRF) [166], an open-source PET/MR image reconstruction framework aimed specifically at synergistic PET/MR applications. Its PET and MR capabilities are based on the powerful tomographic image reconstruction toolboxes Gadgetron and Software for Tomographic Image Reconstruction (STIR) [164, 165].

3.2 Methods

In the following an overview of the design of the framework is given, detailing its functionality and use when simulating dynamic processes, such as cardiac and respiratory motion, contrast agent uptake and PET tracer kinetics. Furthermore, the performance of the framework is demonstrated for three different MR and PET/MR applications: (a) optimisation of image registration parameters for cardiac and respiratory motion estimation in PET and PET/MR, (b) comparison of 5D (cardiac and respiratory) motion-corrected image reconstruction of PET/MR data and (c) assessment of the effect of inaccuracies in the motion estimation on quantitative biophysical parameters obtained from motion-corrected free-breathing abdominal

dynamic contrast enhancement (DCE) MR. All input parameters for the simulation were taken from patient scans carried out on a simultaneous PET/MR scanner.

Simulation Input and Output

An overview of the framework design is given in Figure 3.1. Input and output for the simulation is provided in standardised community format (ISMRMRD [161] format (https://github.com/ismrmrd/ismrmrd) for MRI[1] , and Interfile for PET [162]). The acquisition parameters (e.g. T_E , T_R , flip angle, sequence type, image resolution, number of receiver coils and k-space trajectory for MR or detector geometry for PET) are automatically set based on the header information of MR and PET raw data files. This enables and reproduction of a phantom or in-vivo scan carried out on a PET/MR scanner using the exact same hardware- and sequence-related parameters in the simulation framework or the user can choose to overwrite them.
During the simulation, the data part of the input file is replaced by the generated simulation data yielding a raw data output file, which is fully compatible with any image reconstruction framework capable of dealing with ISMRMRD and/or Interfile raw data. In this manner, the simulation can emulate the acquisition of already available in-vivo data while simultaneously providing GT information. Furthermore, this allows using the simulation output in already existing reconstruction workflows without modifications otherwise necessary to read and process the simulation output.
In addition to the raw data file, an anatomical segmentation combined with an XML descriptor detailing the tissue parameters for each voxel in the segmentation must be supplied. Each tissue type is characterised by T_1 , T_2 , spin density, chemical shift, PET tracer radioactive activity concentration and PET attenuation value which describes its MR and PET behaviour. In this study, the extended cardiac-torso phantom (XCAT) [159] phantom was employed to generate a tissue segmentation and respiratory and cardiac motion models. However, the input format for the segmentation is vendor-independent and arbitrary voxelised anatomical segmentations and matching motion models can be used. The MR and PET scan parameters are then utilised to simulate realistic MR and PET signal intensities for the different tissue types present in the anatomical segmentation. The MR signal for each voxel is computed based on tissue parameters and sequence type using an effective signal model [167]. Currently, only a low flip angle FLASH contrast is available, but extensions to other sequences with a signal model available can be easily made. The PET signal assigns the activity concentration (defined in the XML descriptor) to each voxel and transforms it into specific activity (Bq) by multiplication with the voxel volume.

Signal Encoding

The signal encoding process is performed after the generation of PET and MR signal intensities. The MR raw data is generated by sampling the k-space of the 3D signal. The default k-space shape and acquisition parameters are defined by the information contained in the header of the input file, but can also be changed by the user to simulate modifications to the sampling

[1] Conversion apps from vendor raw data formats of Philips, GE and Siemens to ISMRMRD are provided and maintained by ISMRMRD.

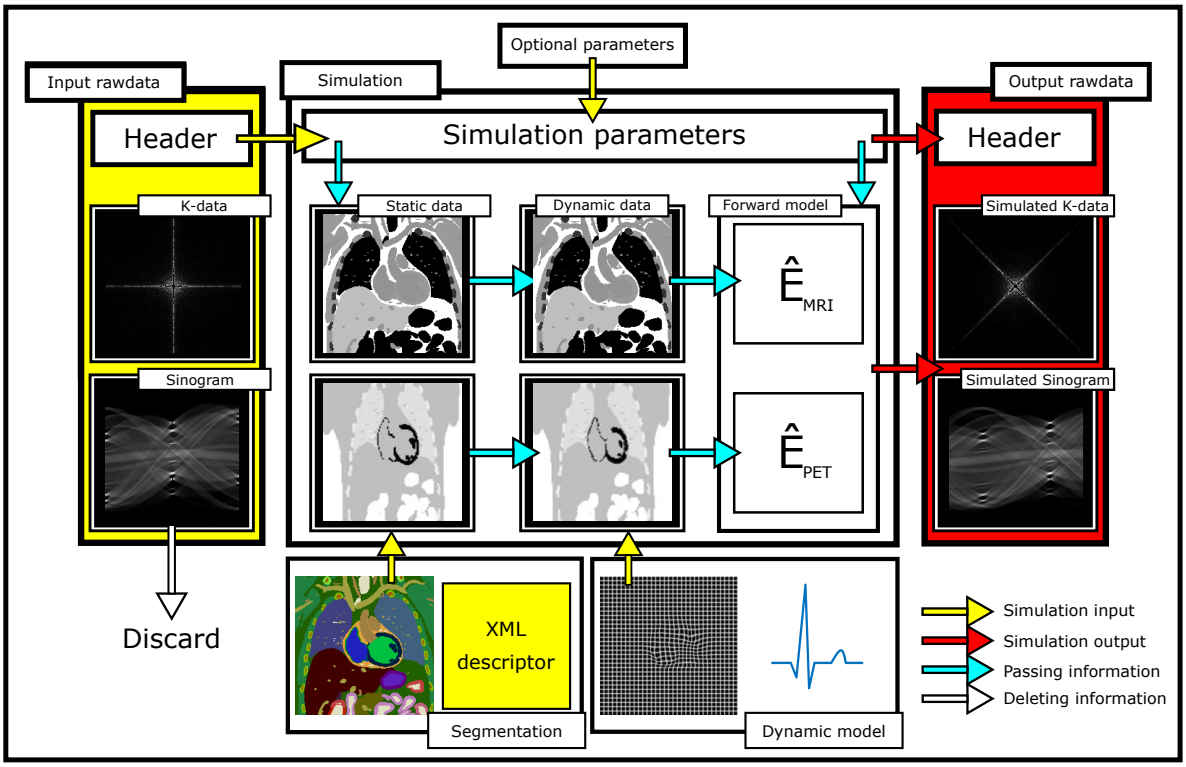

Figure 3.1: Schematic overview of the simulation framework. The default scan parameters of the simulation, including MR and PET acquisition details, hardware layout, are determined by the header information of the template data. These parameters can be overwritten by the user to further customise the simulation. A voxelised tissue segmentation in combination with an XML descriptor must be supplied as the second input. The XML file details the MR and PET specific tissue parameters. Thirdly, a motion model and/or temporal variation of contrast and tracer concentration can be specified in order to simulate dynamic processes. The simulation applies the signal and dynamic models and simulates data sampling based on the parameters defined in the input file. Finally, the simulated data are written together with the simulation parameters into a raw data file. This figure was published in J1.

procedure. Furthermore, MR signal reception is simulated using multiple receiver coils to allow parallel imaging [87]. Each sampled readout line is subsequently stored in the ISMRMRD data structure and written to file.

In contrast to MR, the parameters of PET raw data are defined by the geometry of the detector. This is provided by a template raw data file in Interfile format describing the desired scanner, which is currently limited to ring-shaped designs with possible future extensions to prototype architectures [168]. The forward operator $\hat{E}_{\mathrm{PET}}$ maps the activity distribution into sinogram space and includes tissue attenuation. During a simulation including motion, the motion model is applied simultaneously to the activity distribution and the attenuation map to ensure they are aligned for all simulated time points.

Noise Simulation

For MR, Gaussian noise is added to the real and imaginary part of the simulated k-space data based on an signal-to-noise ratio (SNR) parameter defined in the image as SNR = $\mathcal{S}/\sigma_n$, where $\mathcal{S}$ is the simulated signal and σ_n is the standard deviation of a complex Gaussian distribution in image space. The signal $\mathcal{S}$ in the MR image depends on sequence but also tissue-specific parameters (T_1 , T_2 , spin density). Therefore the user can set a reference tissue type in which

the SNR is defined. This enables to simulate data with a predefined SNR in one region of the image, e.g. SNR = 15 in the liver. The noise is then added to the k-space data in Fourier space as a Gaussian distribution with a scaled width. The SNR is computed assuming data acquisition is carried out with the body coil. Therefore, using the simulation with multiple receiver coils the data will yield data with increased SNR compared to the simulation SNR parameter due to noise averaging. This has to be accounted for by the user.
The sinograms for the PET raw data given by the forward operation $\hat{E}_{\mathrm{PET}}$ applied to the activity distribution $A(\boldsymbol{r})$ describe the decay events along each line of response (LOR) per unit time. Poisson noise is added to them based on the acquisition time $\mathrm{T}_{\mathrm{acq}}$ defined by the user.

Motion models

The framework is set up to flexibly accept arbitrarily many independent dynamic processes as input, in the following referred to as modes, which are combined into one simultaneous process during the simulation. Hence, the framework allows the inclusion of multiple four -dimensional (4D) motion models, e.g. 4D respiratory and cardiac motion as independent motion modes, that result in a data set containing cardiorespiratory motion.
To fully describe one such motion mode during the simulation, a set of motion vector fields (MVFs) describing the anatomical deformations and a motion control signal describing its temporal progression are required. The motion control signal can be passed in the most generic format as a set of pairs (t, s) = (time-point, signal amplitude) with $s \in [0, 1]$, which are interpolated onto continuous time during the simulation. Deformations must define the displacement of each voxel during the motion relative to the reference motion state given by the segmented image. The motion model must encompass all the affected organs, e.g. the entire abdomen for respiration, and must be self-consistent, such that different organs move correctly relative to each other. Anatomy and motion models available for individual organs must be combined and interpolated onto a grid matching the segmentation prior to using them as an input and must be supplied as a set of dense MVFs: $\boldsymbol{m} = \{\boldsymbol{m}_n\}, n \in \{1, \ldots, N\}$ where $\boldsymbol{m}_1 = \boldsymbol{m}(s = 0)$ and $\boldsymbol{m}_N = \boldsymbol{m}(s = 1)$ respectively correspond to the borders of the interval on which the motion control signal is defined. However, the number of supplied MVFs does not limit the number of simulated motion states. The simulation performs a continuous simulation of the entire motion range interpolating between the discrete motion states given by $\{\boldsymbol{m}_n\}$: for example, assume $N = 3$ with $\{\boldsymbol{m}_1 = \boldsymbol{m}(s = 0), \boldsymbol{m}_2 = \boldsymbol{m}(s = 0.5), \boldsymbol{m}_3 = \boldsymbol{m}(s = 1)\}$. Then the state $s = 0.75$ is generated using $\boldsymbol{m}(s = 0.75) = 0.5 \cdot (\boldsymbol{m}_2 + \boldsymbol{m}_3)$.

Dynamic Contrast Models

While motion deforms the anatomy, dynamic contrast modes can be included to modify the tissue parameters as a function of time, e.g. T_1 or activity concentration, translating into signal variation over time. Independent contrast modes can be included flexibly and are combined during the simulation to simultaneous contrast variations. As for motion, a contrast mode requires a contrast control signal describing the temporal variation of the tissue parameters. For each tissue type θ, (e.g. θ = liver) the user must define two border tissue parameter states, $\theta_0 = \theta(s = 0)$ and $\theta_1 = \theta(s = 1)$ and the normalised contrast control signal

as pairs $(t, s), s \in [0, 1]$. Every tissue parameter θ_s can be described by a linear combination between the two border condition tissue states θ_0 and θ_1:

$$\theta_s = (1 - s) \cdot \theta_0 + s \cdot \theta_1. \tag{3.1}$$

For example to simulate the contrast changes in the liver upon injection of contrast agent, $\theta = \text{liver}$. $s = 0$ means no contrast agent is in the tissue and $s = 1$ corresponds to the maximum concentration. For a contrast agent effecting the T_1 relaxation time the liver parameter in the simulation could be set up as $\text{liver}_0.T_1 = 900\,\text{ms}$ and $\text{liver}_1.T_1 = 300\,\text{ms}$. E.g. at $s = 0.3$ this yields $\text{liver}_{0.3}.T_1 = 0.7 \cdot \text{liver}_0.T_1 + 0.3 \cdot \text{liver}_1.T_1 = 720\,\text{ms}$. A dynamic behaviour is generated since $s(t)$ is a function of time. Aside from the T_1 value, the contrast modes can for example also model the decay of activity over time or changes in the attenuation map linked to respiration in PET. The independent contrast dynamic modes are not limited to one tissue type, but each tissue type (i.e. liver parenchyma, blood, lesion etc.) can be assigned its own contrast mode independently without computational overhead. Naturally, the simulation combines the added contrast modes with added motion modes, yielding their combined dynamic behaviour.

Data Binning

For MR, the time axis underlying the simulation process is defined by the timestamps of the MR template raw data. During a dynamic simulation, it must be guaranteed that each readout line is acquired in the correct dynamic state defined by the simulation control signal at the readout's timestamp. In principle, the above approach of interpolating the dynamic models onto a continuous time axis allows simulations with an arbitrary temporal resolution. In practice, however, only a finite number of dynamic states are simulated for computational reasons. This means that multiple time points are grouped together into a bin with a certain temporal width for contrast modes and with a certain motion amplitude for motion states. Readouts at these time points are acquired in the same dynamic state. Eventually, only one k-space is stored, whose individual points are sampled at different contrast and motion states. In most cases, motion leads to periodic changes in the anatomy which is different for contrast modes. Therefore, two alternative binning schemes are used (Fig. 3.2):

- For contrast modes the time axis is binned directly by grouping together a set of subsequent time points. The dynamic state is then defined by sampling the contrast control signal in the bin centre. This yields a unique partition of the time axis independent of the contrast control signal. Therefore contrast variations in many different tissue types can be simulated simultaneously without the computational overhead.

- For motion modes, instead of the time axis, the motion control signal axis is split into bins and the dynamic state is defined directly at their centre. The time axis binning results from grouping together time points from the same dynamic state interval. Since the same motion state occurs repeatedly throughout the simulation it is much more efficient to only apply the motion model a minimum number of times and sample the necessary k-space points depending on their timestamps. However, for

each motion dynamic, an individual time axis partition must be computed. All time points corresponding to one signal bin can be simulated with a single application of the motion model [169], avoiding redundant motion simulations.

For each independent mode, their individual time axis partition is computed independently and intersected subsequently with the others to generate the simultaneous dynamic behaviour of the simulation.

For PET, however, each dynamic state is stored as an independent sinogram. The time spent in each bin defines how many counts contribute to each sinogram and hence its noise level which therefore differs between dynamic states.

Generally, for each mode, the number of simulated motion states should be higher than the number of reconstructed motion states in order to accurately simulate effects such as intra–bin blurring. The suitable number of simulated states depends on the motion amplitude and should be such that the maximum amplitude in voxels divided by the number of simulated phases is below half the voxel size.

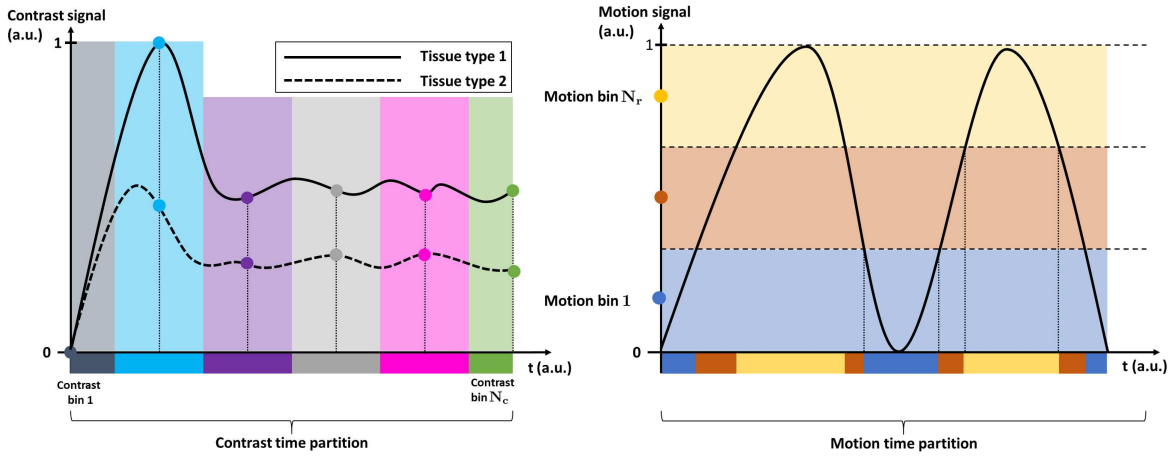

Figure 3.2: Time axis binning based on simulation control signals. Left: example simulation contrast control signal. It is sampled at equidistant time points and the time axis is partitioned into N_c bins. Right: motion dynamic performs an equidistant binning of the motion control signal axis and converts it into a time axis partition with N_r bins. This figure was published in J1.

Motion Ground Truth

The MVFs $\boldsymbol{m}_a^b$ are defined to transform images I between motion states a and b:

$$\boldsymbol{m}_a^b : \mathbb{R}^3 \rightarrow \mathbb{R}^3; I_a \mapsto I_b \tag{3.2}$$

$$I_b(\boldsymbol{r}) = I_a(\boldsymbol{m}_a^b(\boldsymbol{r})) \tag{3.3}$$

The framework yields one MVF for each of the N_{sim} simulated motion states. The set of GT MVFs are defined relative to a reference phase I_0, given by the segmentation input, and point to the simulated N_{sim} motion states:

$$\mathcal{GT}'_{sim} = \{\boldsymbol{m}_0^s\} \tag{3.4}$$

$$s \in [0,1],\ |\mathcal{GT}'_{sim}| = N_{sim}.$$

The image reconstruction can yield a different number of dynamic states N_{reco} compared to the simulated ones N_{sim} and can also yield a different reference phase which is used in the image registration. In order to be able to compare the MVFs obtained from the reconstructed images, the GT motion information output $\mathcal{GT}'_{sim}$ must be post-processed.

Post-processing of Motion Ground Truth

In order to be able to compare MVFs obtained with an image registration algorithm from the reconstructed images of the simulated data to the GT, two important things need to be considered:

- the reconstructed motion states are not the simulated motion states (black and red dots in Fig. 3.3).
- the reference coordinate system for the ground truth differs from the one used in image registration (blue offset vector in Fig. 3.3).

The GT information supplied by the framework consists of N_{sim} MVFs defined in Eq. (3.4): $\{\boldsymbol{m}_0^{\varsigma}\}$, where $\{\varsigma\}$, $\varsigma \in [0.1]$ are the simulated motion states. Reconstruction, however, generates an image series of size N_{reco}: $\{I_s\}, s \in [0,1], |\{I_s\}| = N_{reco}$. Generally, the reconstructed dynamic states are not identical to the ones simulated: $\{s\} \neq \{\varsigma\}$ and $N_{reco} \neq N_{sim}$. Hence, based on the surrogate-based data binning during reconstruction the GT must be interpolated onto the reconstructed states. In the scope of this work a weight matrix $W \in \mathbb{M}(N_{reco}, N_{sim})$ was computed and used for linear interpolation:

$$\mathcal{GT}'_{reco} = \{\boldsymbol{m}_0^s\} = \left\{ \sum_{\varsigma} W_{s,\varsigma} \cdot \boldsymbol{m}_0^{\varsigma} \right\}. \tag{3.5}$$

An algorithm registering a reference image $I_{s_0} \in \{I_s\}$ to the whole series $\{I_s\}$ yields the set of MVFs:

$$\mathcal{REG} = \{\tilde{\boldsymbol{m}}_{s_0}^s\} \tag{3.6}$$

Hence, the MVFs generated by the registration algorithm are relative to one of the reconstructed images I_{s_0}, but the MVFs from $\mathcal{GT}'_{reco}$ are defined relative to the reference phase I_0. There is a systematic mismatch between $\mathcal{GT}'_{reco}$ and $\mathcal{REG}$ as their coordinate systems are different ($\boldsymbol{m}_0^s \neq \boldsymbol{m}_{s_0}^s$). This offset was corrected for by:

$$\mathcal{GT} = \boldsymbol{m}_{s_0}^0 \circ \mathcal{GT}'_{reco} = \{\boldsymbol{m}_{s_0}^0 \circ \boldsymbol{m}_0^s\} = \{\boldsymbol{m}_{s_0}^s\}, \tag{3.7}$$

where $\boldsymbol{m}_{s_0}^0$ was computed by numerical inversion of $\boldsymbol{m}_0^{s_0}$.

Computation of Registration Error

For most applications the precision of the registration output $\mathcal{REG}$ is only of interest in a certain region (region of interest (ROI)). However, in a MVF the motion of a voxel is not encoded in the voxel itself but rather at the location where it is going to be deformed to

(c.f. Eq. (3.3)). Hence, the inverse MVF encodes the motion of the voxel. To compute the registration error both $\mathcal{GT}$ and $\mathcal{REG}$ must be inverted and evaluated in the ROI:

$$\Delta_s^{reg}(\mathrm{ROI}) = \|(\boldsymbol{m}_s^{s_0} - \tilde{\boldsymbol{m}}_s^{s_0})(\mathrm{ROI})\|_2. \tag{3.8}$$

The precision of the inversion algorithm can be set arbitrarily small, such that its effect is negligible on the accuracy of the registration output.[2] Since the registration output must be inverted eventually, in practice the inverse is registered directly: $\mathcal{REG} = \{\tilde{\boldsymbol{m}}_s^{s_0}\}$. Note that the ROI must be deformed from the reference phase of the segmentation into the motion state to which the dynamic images were registered: $\mathrm{ROI}_{s_0}(\boldsymbol{r}) = \mathrm{ROI}_0(\boldsymbol{m}_0^{s_0}(\boldsymbol{r}))$.

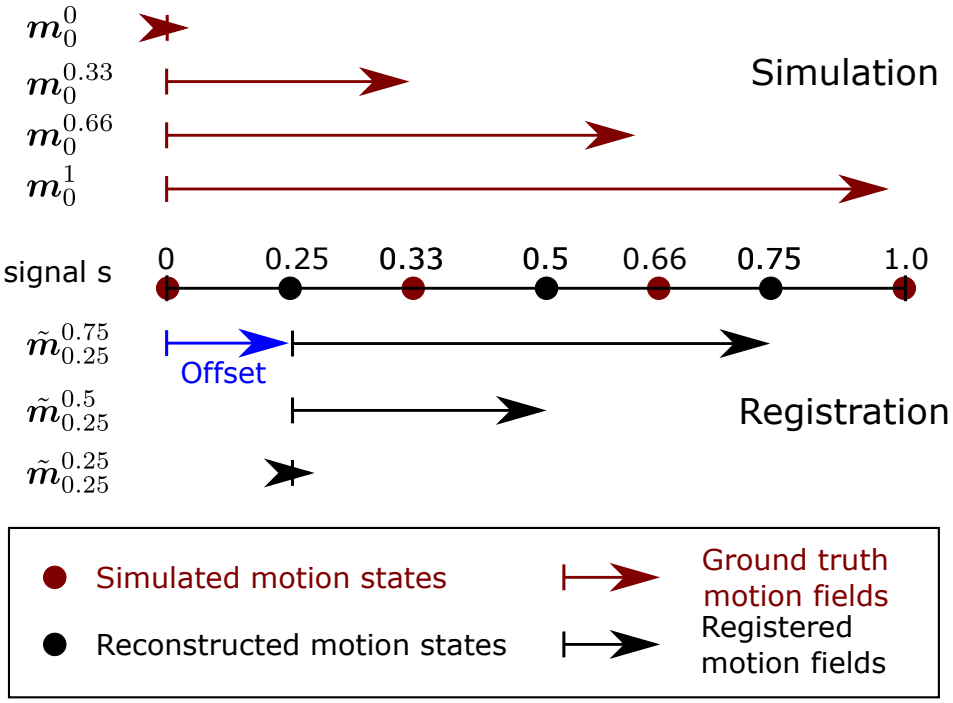

Figure 3.3: Schematic overview of relationship between registration and ground truth outputs. $N_{sim} = 4$ red dots show the simulated motion states, $N_{reco} = 3$ black dots show the reconstructed states. Red arrows symbolise the MVFs provided by the simulation, black arrows show the registration output against which the GT must be compared. The arrows' starting and endpoint symbolise in which coordinate system they are defined and to which they point. The post-processing steps discussed above would correspond to the following operations: first the red arrows are weighted to create three arrows which point from $s = 0$ to each of the black dots. To adjust their starting coordinate system their starting points is shifted by the blue offset vector is added s.t. they each depart from the black dot at $s = 0.25$ instead of $s = 0$. Then GT and registration coordinate systems match and their content can be meaningfully compared. This figure was published in J1.

To evaluate the registration output the simulation framework is able to generate ROIs based on the underlying tissue segmentation for the different motion states.

[2]This assumes that the output of the registration algorithm is invertible. However, the two features of physiological motion, tissues not passing through each other and motion cyclicality strongly suggest that registration output should be invertible to describe motion.

3.3 Experiments

As mentioned above, we evaluated the framework on three different applications. In the first, we optimised the parameters of a non-rigid image registration algorithm for cardiac-resolved PET data and of a synergistic motion estimation approach to estimate cardiac and respiratory motion using both MR and PET data. In the second application, 5D non-rigid motion fields were utilised in a motion-corrected image reconstruction to compensate for cardiac and respiratory motion. Finally, the effect of inaccuracies in the motion fields on the final quantitative MR parameters obtained from DCE MR was assessed.

PET/MR Simulation Parameters

All tissue segmentation and motion models were generated by the XCAT software [159]. A simulation of a PET/MR exam on a 3T Siemens Biograph mMR was performed using raw data files from a patient data examination with the patient's self-navigator and ECG signal as motion control signal input (Fig. 3.4). Continuous MR data acquisition during free-breathing was simulated for a triple-echo prototype Dixon-based GRE Golden angle Radial Phase Encoding sequence (TR = 5.9ms, TE = 1.2/2.7/4.2ms, FA = 10°, #PE points = $125\frac{\text{pts}}{\text{angle}} \times$ 256 angles)[139, 170]. The spatial resolution of MR was $1.9\,\text{mm}^3 \times 3.2\,\text{mm}^3 \times 3.2\,\text{mm}^3$. The simulated coil sensitivities consisted of a simple setup of four 3D varying Gaussian profiles without phase variations between the different RF receive channels. However, more sophisticated coil map simulations or measurements from an existing scanner could be easily included and passed as a simulation parameter. The PET sinograms were simulated based on the same segmentation input as the MRI data. The forward model encoded the spatial resolution of a Biograph mMR detector using ray tracing. The individual PET images were reconstructed at different resolutions depending on the application. This is described below. Attenuation was included in the simulation but scatter and randoms effects were so-far omitted. The total acquisition time was set to match the MR scan at 3.2 min. These parameters were used for all subsequent PET simulations.

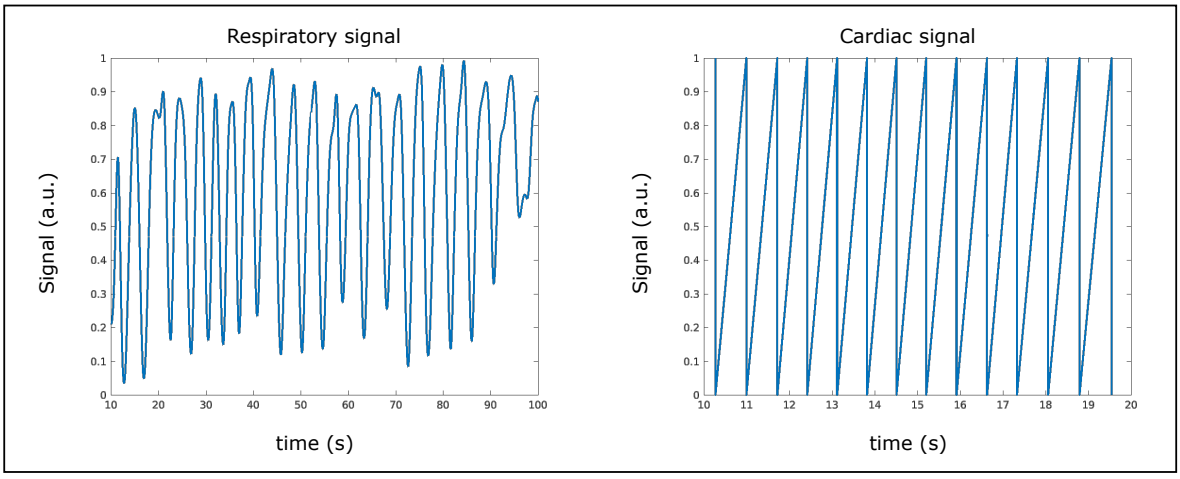

Figure 3.4: Respiratory and cardiac motion control signal used as 5D PET/MR simulation input. The respiratory signal changes smoothly over time. The cardiac signal has the shape of a sawtooth function increasing linearly from R-peak to R-peak. This figure was published in J1.

Quantitative Evaluation of Image Registration

The output of the simulation framework was reconstructed into different cardiac and respiratory motion states resulting in a series of 3D images $\{I_s\}, s \in [0,1]$. To generate a motion model, this image series was registered to a reference phase I_{s_0} evaluated against the GT in an ROI containing the left ventricle myocardium as described in Eq. (3.5).

Using Eq. (3.8), the maximum error over all motion phases, i.e.

$$\Delta^{reg} = \max_s \left(\left\langle \Delta_s^{reg}(\boldsymbol{r}) \right\rangle_{\mathrm{ROI}} \right), \tag{3.9}$$

is computed to quantify the spatial accuracy of the registration output. Δ^{reg} will be referred to as the registration error in the following. When evaluating the registration error it was compared to the maximum amplitude of the GT MVFs which is the error made in neglecting motion and reconstructing the average motion. The application of a motion-compensated reconstruction can only expect to improve image quality if the motion model is more accurate than this threshold. This comparison is depicted in Fig. 3.6 and Fig. 3.7.

Image Registration Using Cardiac-resolved 4D PET

To quantitatively evaluate a registration algorithm a 4D PET dataset was simulated in 16 cardiac motion states. Eight motion states $\{I_t\}$ were reconstructed at the signal locations $t = \{0.125, 0.25, \ldots 1.0\}$ at a spatial resolution matching the MR reconstructions at $1.92\,\mathrm{mm}^3 \times 3.2\,\mathrm{mm}^3 \times 3.2\,\mathrm{mm}^3$. Attenuation correction was omitted to avoid the imprint of a static attenuation map veiling the actual motion content of the images at the cost of an apparently increased uptake in the lungs. The registration algorithm optimised an energy functional:

$$\boldsymbol{m}_s^{s_0} = \min_{\hat{\boldsymbol{m}}_s^{s_0}} \mathcal{E}(\hat{\boldsymbol{m}}_s^{s_0}) = \min_{\hat{\boldsymbol{m}}_s^{s_0}} \mathcal{S}(I_s \circ \hat{\boldsymbol{m}}_s^{s_0}, I_{s_0}) \tag{3.10}$$

where $\mathcal{S}$ is a similarity measure. In the case of this study, normalised mutual information was used as a similarity metric. The transformation was modelled as a free-form deformation (FFD) using spatial B-spline interpolation [78, 171] without explicit regularisation because the intrinsic regularisation ensures a smooth motion field on the interpolation scale. The reconstructed images were registered for different spline control point distances ΔB to assess the influence of the intrinsic regularisation on the registration quality. Motion correction was performed using the registration-transform-add (RTA) approach, i.e.

$$I_{corr}(\boldsymbol{r}) = \sum_{\sigma \in \{s\}} (I_{s_0}^{\sigma} \circ \boldsymbol{m}_{\sigma}^{s_0})(\boldsymbol{r}) \tag{3.11}$$

This motion-corrected image was used to assess the influence of the registration output on the image quality.

Synergistic registration Using 4D PET/MR

Further application of the framework was the evaluation of an advanced synergistic motion estimation algorithm [82]. To this end two PET/MR datasets were simulated containing 4D cardiac and respiratory motion, respectively. The MR images were reconstructed using an iterative SENSE approach [87] for respiration and compressed sensing (CS) [88] for cardiac motion. Reconstructions yielded image series at the same motion states and with the same spatial resolution as the 4D PET data in section 3.3. An extended version of the image registration from Eq. (3.10) with energy functional:

$$\mathcal{E}(\boldsymbol{m}_s^{s_0}) = \lambda \cdot \mathcal{S}(I^{MR} \circ \boldsymbol{m}_s^{s_0}, I_{s_0}^{MR}) + (1-\lambda) \cdot \mathcal{S}(I^{PET} \circ \boldsymbol{m}_s^{s_0}, I_{s_0}^{PET}) \tag{3.12}$$

$$\tag{3.13}$$

with $I^{MR/PET}$ as the PET and MR image series relative to a phase $I_{s_0}^{MR/PET}$, and the synergy weight $\lambda \in [0,1]$ balancing the PET and MR contribution to the cost function. The extremes $\lambda = 0$ corresponds to PET, and $\lambda = 1$ corresponds to MR input only. The synergy weighting parameter was varied in the range $\lambda \in \{0, 0.1, \ldots, 0.9, 1.0\}$ at $\Delta B = 12$ for cardiac motion and $\Delta B = 24$ for respiration.

Motion Correction of 5D Cardiac PET/MR

All simulations were carried out using in-vivo PET/MR data as template input. Cardiac and respiratory motion were included as independent motion modes with $N_r = N_c = 8$ states respectively, yielding simultaneous cardio-respiratory motion in 64 combined motion states. This corresponded to a cardiac time interval of 125 ms at 60 bpm. Based on the cardiac and respiratory input signal the simulation output was double-binned into $N_r \times N_c = 4 \times 4$ motion states for both modalities, corresponding to a temporal resolution 250 ms at a motion frequency of 60 bpm. The reconstructions were performed with and without motion correction [89] using the ground truth motion fields. Subsequently, the resulting images were compared to corresponding reconstructions of the in-vivo patient data to verify that the proposed framework yields realistic output data. The PET images were reconstructed at a resolution of $2\,\mathrm{mm}^3 \times 2.1\,\mathrm{mm}^3 \times 2.1\,\mathrm{mm}^3$ matching the patient data used in the comparison. The MR simulation data were reconstructed using an iterative SENSE approach [87]. For PET AC during reconstruction a static attenuation map used as the simulation input was employed. For a more detailed description of the employed motion-correction scheme, please refer to [172].

Dynamic Contrast Enhanced (DCE) Abdominal MRI

Data acquisition

An abdominal DCE MR exam during the injection of an MR contrast agent (gadoxeate disodium) was simulated with the input data also acquired on a 3T Siemens Biograph mMR and using a GRE Golden angle Radial Phase Encoding [173] sequence (T_R / T_E = 3.29/1.36 ms, FA = 11°, #PE points = $120\frac{\text{pts}}{\text{angle}} \times 640$ angles) at an isotropic spatial resolution

of $1.5\,\mathrm{mm}^3 \times 1.5\,\mathrm{mm}^3 \times 1.5\,\mathrm{mm}^3$. Fat suppression was simulated by setting the spin density of fat to 10% of the water value.

Quantitative Parameter Estimation

An extended Toft's model for the temporal evolution of the gadoxeate disodium concentration in the tissue of interest (toi) $C_{toi}(t)$ was fitted based on an arterial input function (AIF) extracted from the hepatic artery in the template data [173]:

$$C_{toi}(t) = v_p \cdot C_p(t) + k_{trans} \cdot C_p(t - \Delta T) * e^{-\frac{k_{trans}}{v_e} t}, \tag{3.14}$$

with C_p the blood plasma contrast uptake, v_p the volume fraction of blood plasma, v_e the extravascular extracellular fractional volume, k_{trans} the volume transfer constant, ΔT the time delay of tissue enhancement with respect to the arterial concentration in plasma and $*$ representing the convolution operator. Every parameter of the right-hand side of equation (3.14) depends on the tissue of interest except C_p. The simulated tissues of interest were healthy liver tissue and a hepatic lesion with a necrotic core. The parameters used as simulation input can be found in table 3.1. The necrotic core was simulated as lesion tissue without any contrast agent kinetics. The resulting Toft's model and simulated T_1 variations are depicted in Fig. 3.5. The concentration of contrast agent was converted into a temporal

Table 3.1: Input parameters employed in the Toft's model described in Eq. (3.14). This table was published in J1.

Tissue type	Healthy Liver	Lesion	Lesion Core
k_{trans}	0.041	0.135	0
v_p	0.025	0.141	0
v_e	0.087	0.022	1
ΔT[min]	0.19	0.20	0

evolution of T_1 using $T_1(t) = \frac{1}{R_1(t)} = \frac{1}{R_1 + r \cdot C_{toi}(t)}$ where $r = 6.2\,\mathrm{mmol\,s^{-1}}$ is the relaxivity of gadoxeate disodium at 3T [174]. The simulated Toft's model and the calculated change in the T_1 parameter extracted from patient data used as simulation input is depicted in Fig. 3.5. The DCE simulation was performed for $N_r = 16$ respiratory states and $N_c = 49$ contrast states.

Image Reconstruction and Impact of Motion Field Inaccuracies

From the simulation output a k-t-SENSE reconstruction was performed with and without respiratory motion correction [89]. In order to assess, how inaccuracies in the motion estimation would impact the assessment of the quantitative parameters such as k_{trans}, a bias of different strength σ was added to the GT motion fields using:

$$\mathcal{GT}_\sigma = \{\boldsymbol{m}^t_\sigma\} = \boldsymbol{m}^t_0 + \sigma \cdot \frac{t}{N_{reco} - 1} \cdot \boldsymbol{n} \tag{3.15}$$

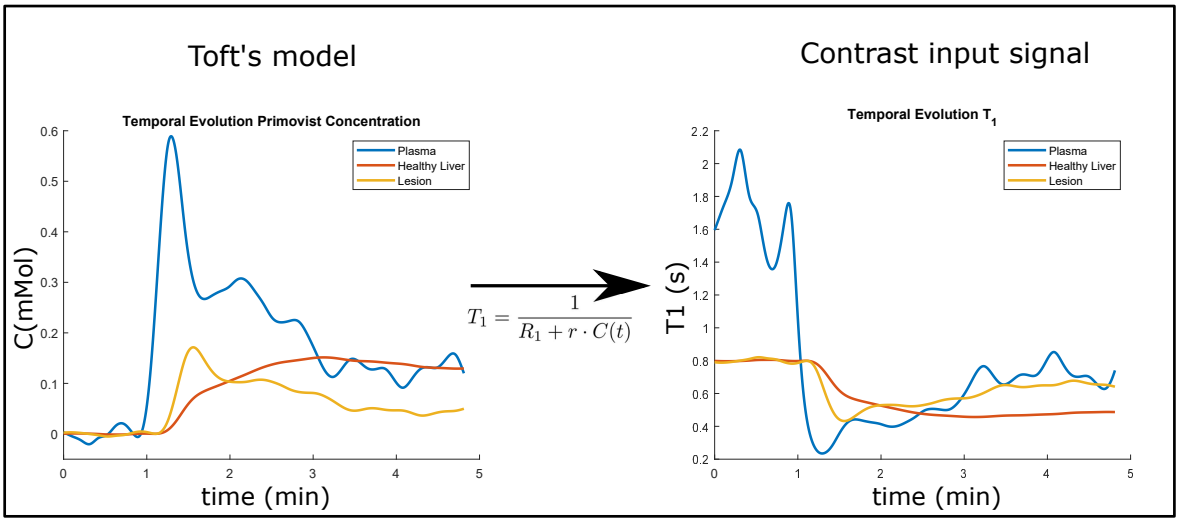

Figure 3.5: Depiction of motion control signal used as DCE simulation input. Left: Toft's model simulated using the parameters from Table 3.1 for three different tissue types. Right: Toft's model converted into T_1 variations and respiration signal passed to the simulation as contrast control signal based on which tissue contrast and respiratory state were changed. One can see how a contrast agent uptake leads to a drop in T_1 which generates contrast between the lesion and healthy tissue in late contrast phases.

where $\boldsymbol{m}_0^t$ are the GT MVFs for reconstructed motion phase $t \in \{0 \dots N_{reco} - 1\}$, σ is the bias strength, $\sigma \in \{0, \dots, 6\}$ and $\boldsymbol{n}$ a spatially varying (but temporally constant) random unit vector field which fixes the direction of the bias for each voxel. This random unit vector was not drawn for each voxel independently, but on a larger grid and subsequently spline-interpolated to the resolution of the motion fields. This procedure took into consideration that the image registration algorithm used here is spline-based, and hence errors in the motion estimation are expected to vary smoothly between neighbouring voxels. The increase of bias with each a larger motion phase number included potential accumulation of error for larger amplitudes in the registration input. Each of the motion field sets $\mathcal{GT}_\sigma$ were then used in the motion-corrected reconstruction and k_{trans} maps were computed from the dynamic contrast-resolved image time series. The contrast-to-noise ratio (CNR) between lesion and necrotic core for uncorrected and motion-corrected (with different values for σ) images was compared. CNR was computed as CNR $= \frac{\mathcal{S}_{\text{lesion}} - \mathcal{S}_{\text{core}}}{\sigma_{\text{lesion}}}$ where $\mathcal{S}_i$ and σ_i are the mean and standard deviation of the k_{trans} maps in the respective tissue regions. These were evaluated in a 2D axial slice through the lesion centre and segmented by hand.

3.4 Results

4D PET/MR Motion Estimation Evaluation

Image Registration Using Cardiac-resolved 4D PET

The results of the 4D PET image registration for different registration parameters are shown in Fig. 3.6. Coronal views of the motion-corrected PET image data are displayed, using a registration output for two different parameter settings, as well as motion correction using the GT motion model. The displayed images are corrected into end-diastole as in Eq. (3.11).

The figure also shows the respective transformations which transform end-diastole images into systole. The visualisations of the motion fields show that the motion fields with $\Delta B = 8$ contain many areas with non-physiological motion vectors. Motion fields obtained with $\Delta B = 18$ have a much better correspondence to the GT motion fields. Nevertheless, the motion-corrected PET images have similar image quality and appear to be accurately corrected for cardiac motion for both $\Delta B = 8$ and $\Delta B = 18$. The registration error evaluated in the left ventricle is depicted as a function of the spline point distance ΔB. Note that only amplitudes are plotted but the deviations between registration output and GT are computed vectorially. The effect of motion correction using the differently regularised registration outputs can be seen in the animation in supplementary figure S1. The proposed framework provides the GT motion fields and allows for the calculation of a motion field error in order to quantitatively assess the performance of different registration parameters.

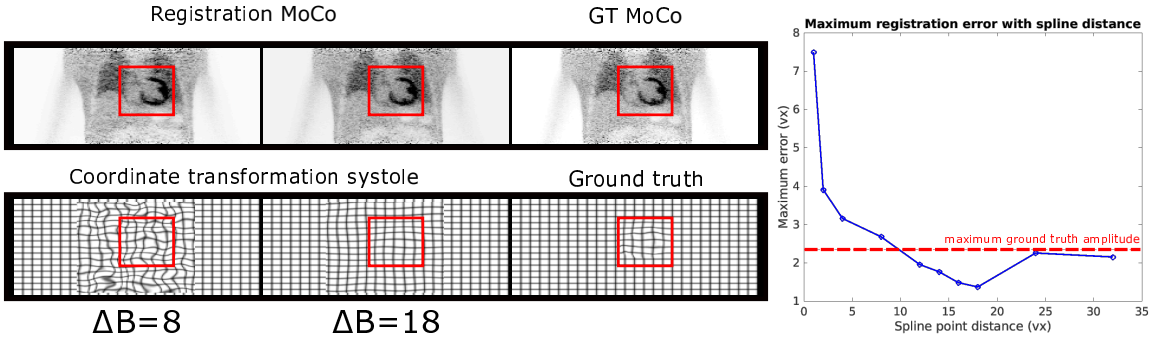

Figure 3.6: Quantitative evaluation of 4D PET image registration. Left: Cardiac motion-corrected PET images (top row) and corresponding motion information used for the correction (bottom row). All images were motion-corrected to end-diastole. The red box indicates the heart region for better comparison of images and motion information. A similar effect on the image quality can be seen for both poorly ($\Delta B = 8$) and well ($\Delta B = 18$) regularised image registration. Only the motion field transforming the data from systole to diastole is shown here. Right: Maximum error over all motion phases between estimated motion and ground truth in the left myocardium as a function of different spline point distances ΔB of the registration algorithm. The red line indicates the maximum ground truth amplitude. A minimum can be found at $\Delta B = 18$. This figure was published in J1.

Synergistic Registration Using 4D PET/MR

The results of the quantitative assessment of synergistic motion estimation is depicted in Fig. 3.7. Systolic and diastolic and exhale and inhale images are shown for cardiac and respiratory resolved PET and MR images, respectively. Markers indicate cardiac (top) and respiratory (bottom) motion contained in the image data. PET reconstructions show large uptake in the lungs in the absence of attenuation correction. The right column shows the registration error for different synergistic weighting λ. The red line is the amplitude of the GT motion and hence corresponds to the error when neglecting motion. For both respiration and cardiac motion using both MR and PET image data for the registration (i.e. $0 < \lambda < 1$) improves the registration accuracy compared to using either only MR or only PET image data separately. For respiration one can see that the maximum error is always much lower than the maximum amplitude of the ground truth motion whereas for cardiac motion the λ parameter needs to

be fine-tuned to achieve a registration error below the GT amplitude and hence a registration results which leads to an improvement in image quality when using motion correction.

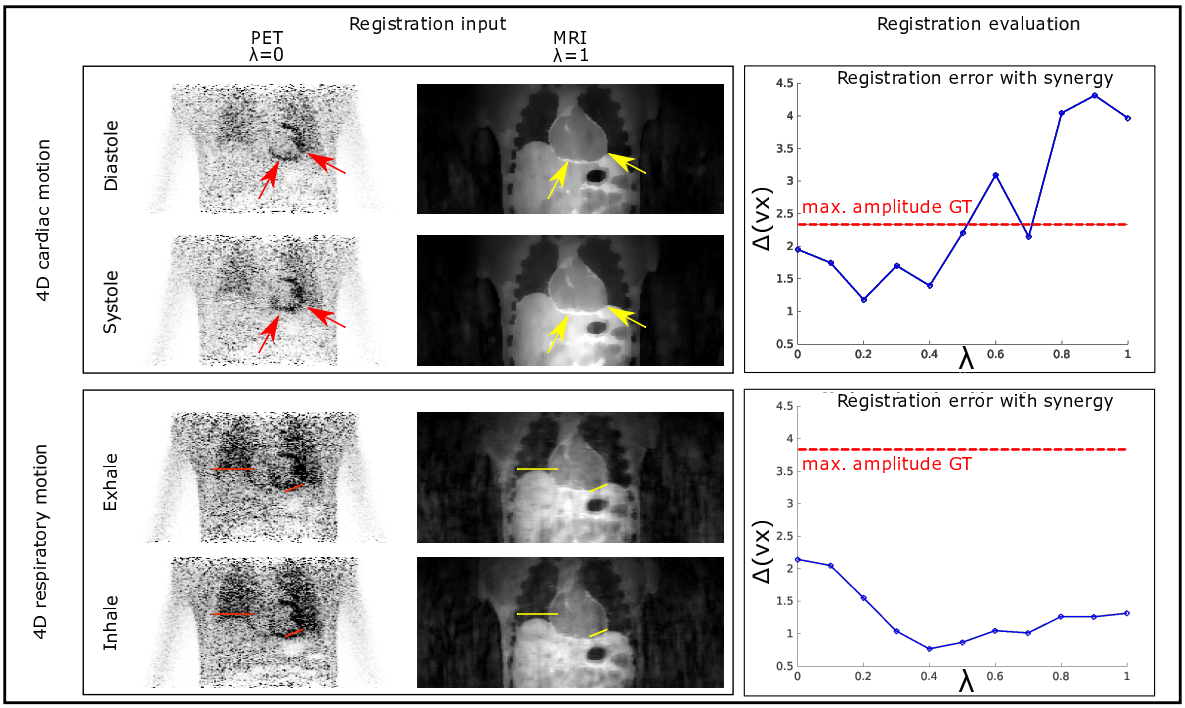

Figure 3.7: Quantitative assessment of joint PET/MR 4D motion estimation. Left: MR and PET images used as the input for 4D respiratory and cardiac motion estimation. Myocardial contraction between diastole and systole (red arrows) and translation between in- and exhale (red lines) can be clearly seen. The depicted motion states correspond to the respective maximum amplitude of the GT motion fields for both motion types. Right: quantitative evaluation of the registration. The registration error is plotted in blue, with the maximum amplitude of the GTMVFs in red. Minima can be found in both curves yielding the best choice of registration parameter λ. This figure was published in J1.

Motion Correction of 5D cardiac PET/MR

Images reconstructed from the simulated PET/MR raw data are compared to the corresponding in-vivo patient data in Fig. 3.8. Respiratory and cardiac motion leads to similar motion artefacts in both datasets for PET and MR. In the MR images blurring of anatomical structures such as the heart or the liver can be clearly seen for both in-vivo data and numerical simulations. The PET images show blurring of the myocardium and papillary muscles. Motion correction of respiratory and cardiac motion in the in-vivo data improves the image quality, leading to a better depiction of the anatomy comparable to the motion-free reference image obtained with the simulation framework.

DCE Abdominal MRI

The reconstructed DCE images are pictured in Fig. 3.9 for both simulation and patient data. For the two datasets both respiration averaged and motion compensated k-t-reconstructions are compared in selected phases. An ROI containing a lesion with a necrotic core is zoomed in on. The slice showing the respiration average has been selected to show the same anatomical region as the respiratory-compensated images. For the selected ROI the k_{trans} parameter of

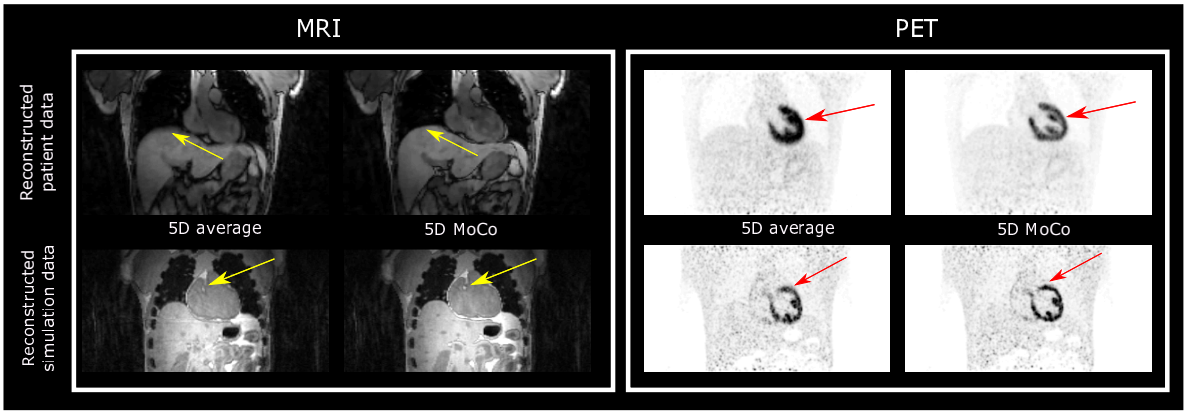

Figure 3.8: Reconstructions of in-vivo and simulated PET/MR data. The simulation and acquisition occured during free-breathing and over multiple cardiac cycles. Left: MR reconstructions. The top row shows the patient images for motion-average (5D average) and cardiac and respiratory motion-correction (5D MoCo), and the bottom row displays the corresponding reconstructions of the simulation output, where 5D MoCo was performed using the ground truth motion. Patient data and simulation have a comparable image contrast and quality. Motion correction reduces the blurring of anatomical structures indicated by the yellow arrows. Right: corresponding PET reconstructions. Again, the image quality and contrast of simulation and patient data reconstructions are comparable. Motion correction in PET yields a reduced blurring of the uptake in the heart indicated by red arrows. This figure was published in J1.

a Toft's model fit to the DCE images is depicted. In simulation and patient data motion correction improves the visualisation of the necrotic core of the lesion with lower k_{trans} value compared to the lesion itself. In the k_{trans} maps computed from motion-averaged cases the necrotic core is strongly blurred. This finding is confirmed quantitatively by the results shown in Fig. 3.10. The left plot shows the deviation from the true motion for the different σ values of the maximum respiratory amplitude. Only the configuration of $\sigma = 6$ has a deviation larger than the GT amplitude marked by the dotted line. Upper right: the CNR in the k_{trans} fits between lesion and core as a function of MVF bias. A non-linear relationship is revealed resulting in a decrease in CNR between lesion and necrotic core in the k_{trans}. A ROI-averaged bias larger than 4 mm yields a reduced CNR, in which the core is not visible any more. Deviations larger than 5 mm yield a CNR of the same magnitude as if no motion correction had been employed (indicated by the dotted line). The k_{trans} maps for the different bias strengths in which the CNR was evaluated are shown below the quantitative plots. While one can see a degradation in contrast for growing bias, between $\sigma = 0$ and $\sigma = 3$ the core stays clearly outlined.

3.5 Discussion

This paper presents the design and application of a simulation framework used to generate realistic dynamic PET/MR raw data. The framework is specifically targeted at providing dynamic ground truth information for different types of physiological motion. In addition, other effects such as contrast or tracer dynamics and chemical shift between fat and water in MRI can be included. It aims at supporting the development of new - and validation and improvement of currently employed reconstruction methods and, more specifically, PET/MR motion correction approaches.

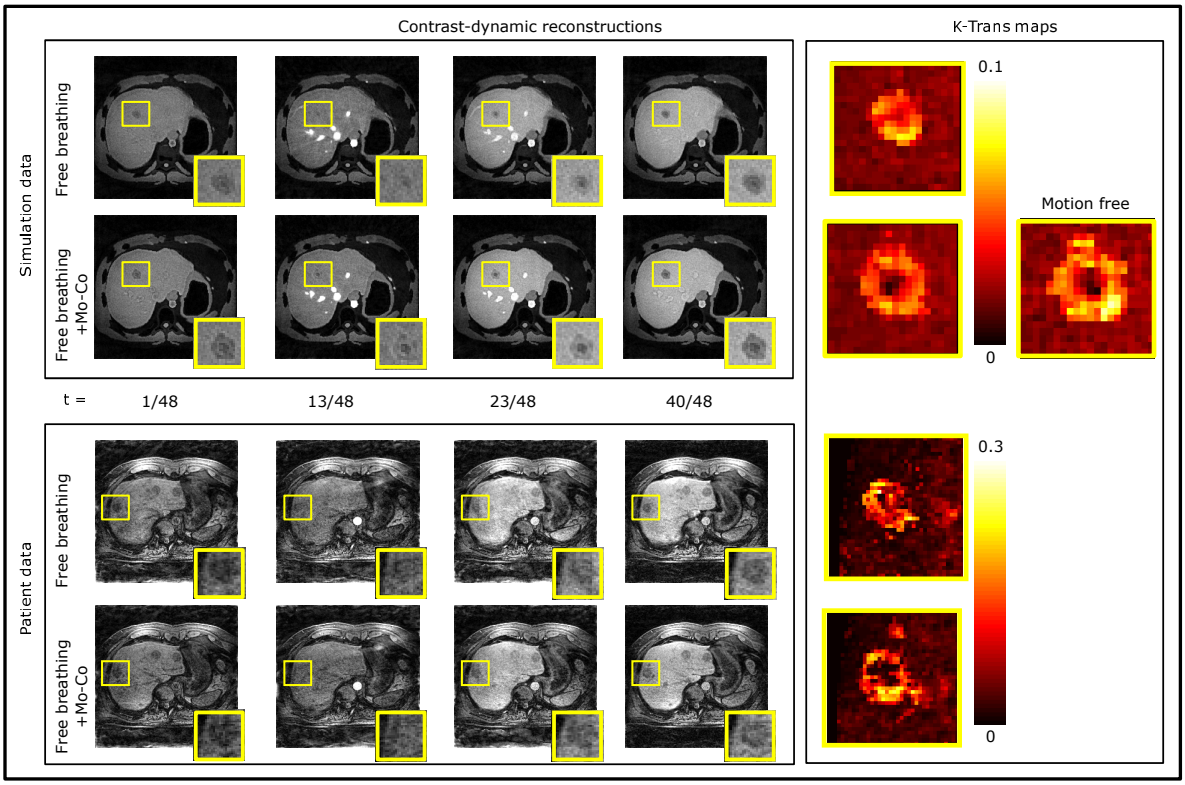

Figure 3.9: Comparison of simulated DCE and patient data. Top images show simulated data, bottom images show patient data. Motion-corrupted (Free breathing) and respiratory motion-corrected (Mo-Co) image reconstructions are shown in selected DCE images for both simulation and patient data. The right row shows the k_{trans} map resulting from a fit over all 48 contrast states. For the simulation a k_{trans} map of a motion free simulation is available, displayed next to the k_{trans} fit of the motion-corrected data. This figure was published in J1.

The success of a motion-corrected reconstruction depends on the accuracy of the estimated motion model from image registration. However, an improvement in image quality might be an insufficient indicator for successfully modelling motion. Figure 3.6 shows that also motion fields which are not physiologically plausible can lead to high image quality. Non-rigid image registration offers many degrees of freedom [70, 79, 175], but it can also yield strong improvements in image quality by computing an unrealistic and inaccurate motion model [176]. The proposed approach offers here the possibility to simulate not just dynamic data but also GT motion information which can serve as a reference standard, similar to previously proposed MR-only approaches [171, 177, 178]. This allows for a comprehensive evaluation of image registration techniques without the need for any manual annotations [179]. The realism of numerical simulations of human anatomy is often limited by the accuracy of the organ and tissue segmentation, assigned tissue parameters and signal and acquisition models. The proposed method is not developed for a specific anatomical model, but can be combined with any pixel-based model of the human anatomy [83, 155, 157, 159]. The presented framework builds itself on flexible input possibilities allowing the inclusion of currently existing and future human phantoms in a generic way. The output in standardised format allows direct integration of the simulation results into existing workflows. The computational demand is kept low, by directly computing effective MR and PET signals, and omitting Bloch and Monte-Carlo simulations. However, this comes at the cost of reduced realism, neglecting intravoxel

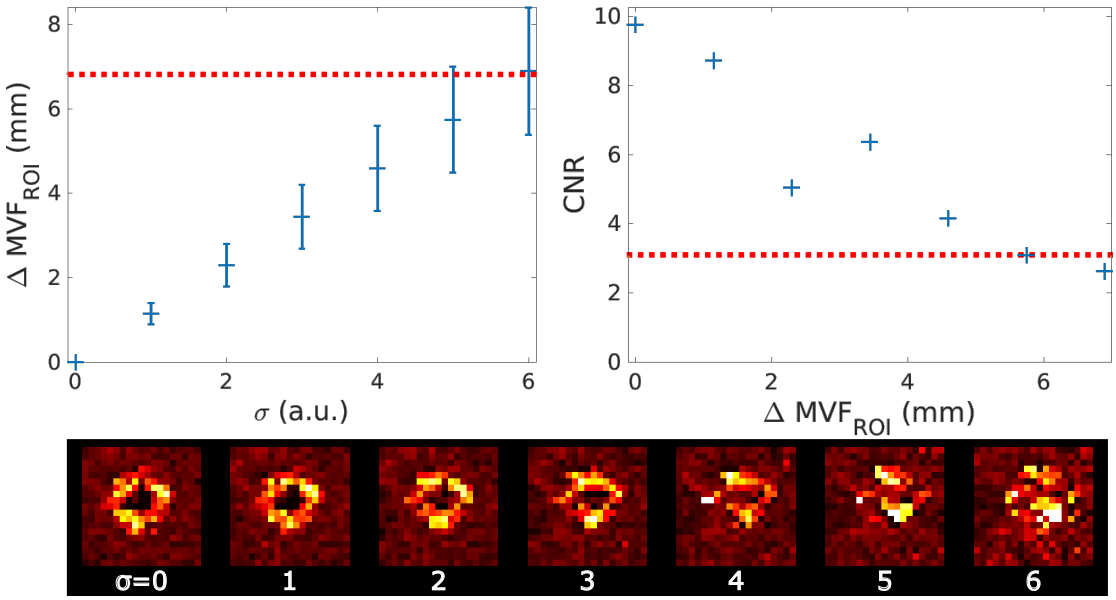

Figure 3.10: Artificial bias in vector field. Upper left: absolute value of the bias for different σ as in Eq. (3.15) computed as the mean bias averaged over all motion phases in an ROI including only the lesion and the necrotic core. Error bars correspond to the standard deviation in the ROI. $\sigma = 0$ yields the GT MVF with no bias; the dotted line indicates the mean motion amplitude used in the simulation. Only the largest bias $\sigma = 6$ exceeds the average motion amplitude. Right: variation of the CNR in lesion and core with different MVF bias. Upper dotted line: CNR without including motion in the simulation. Dotted line: resulting CNR without including motion correction in the image reconstruction. It can be observed that the CNR drops to the level of neglecting motion correction for deviations larger than 4 mm. Bottom: series of k_{trans} maps resulting from different bias strengths. A degradation in contrast for increasing bias starting from $\sigma = 3$ is visible. The CNR evaluation was performed on manual segmentations of these images This figure was published in J1.

effects and yielding identical contrast over regions of the same tissue. Despite this limitation, the presented framework is able to generate data which is very similar to real patient data as shown in Fig. 3.8. Nevertheless, more complex signal models could be included in the framework. Currently randoms and scatter effects, positron range and partial volume effects are not part of the PET simulation but this could also be added in the future. Following this could also yield applications in training data generation for machine learning.

All simulations were performed on a Virtual Machine running on a workstation with access to 58 GB of RAM and 8 CPUs. Using this setup and a 208x208x208 voxel segmentation the time to simulate 4 motion states was 25 minutes for PET using a Siemens Biograph mMR detector consisting of 11 segments and 4 minutes for MRI simulating a 96000 phase-encoding points on a 3D GRPE trajectory acquisition. The simulation time scales linearly with the number of dynamic states included. The current implementation only allows equal-sized CC bins and CC is sampled at the centre of each bin. Hence to avoid non-linear intra-bin CC behaviour the number of simulated states must be high. Further improvement could be achieved by computing the average CC in each bin instead and perform automatic bin sizing based on the slope of the CC signal to ensure proper signal sampling while reducing computational costs.

In addition to comparing the image quality of the simulation to the template data, the framework was used to evaluate the accuracy of registration algorithms. With the GT information the optimal registration parameters for 4D motion in PET, shown in Fig. 3.6,

MR images or weighted combinations thereof, as in Fig. 3.7 could be determined. For a strong weighting of the MR image information, image registration is affected by the undersampling artefacts in the MR images and hence yields inaccurate motion fields. The results for the optimal weighting between MR and PET image information obtained from the simulated data was $\lambda = 0.4$ and 0.2 ($\lambda = (0, 1)$=(PET,MR)) for respiratory and cardiac motion. This is in good agreement with the values reported from the in-vivo data of [82]. The framework was also shown to be feasible in standalone MR applications, such as the DCE simulation presented in 3.4. In this case, the number of simulated and reconstructed dynamic phases coincided for contrast as well as motion. The framework was able to produce data with realistic content and the same dynamic behaviour as observed in the patient data acquisitions. Furthermore, it enabled a Toft's model fit and reproduced the k_{trans}-maps found in the template input data. An improvement in CNR between lesion and necrotic core when applying respiratory motion correction could be visually detected for the simulation results where the motion-compensated fit showed a clear delineation of the core. To quantify this effect different degrees of bias were added to the MVFs in a fashion very similar to errors produced by artefacts and noise. Motion-compensated k_{trans} maps yielded a well-delineated core for ROI-averaged deviations of the MVFs from the correct motion in the maximum amplitude motion phase up to the limit of approximately twice the voxel size. Deviations larger than 4 mm, however, showed a CNR comparable to completely neglecting motion. The direction of bias was fixed to one random configuration and scaled to increase the deviation from the true motion. While the importance of precise image registration to obtain high-quality k_{trans} could be shown, to determine the exact behaviour of the CNR evaluations with different directions of bias (i.e. different noise realisations) would have to be carried out. In this study, the number of used template datasets was limited to two, one from a PET/MR and from a free-breathing DCE abdominal exam. The focus was set to prove the feasibility of the framework in multiple different tasks. It was not examined how the degree of realism changes with different template data or different segmentation input. Furthermore, for the cardiac PET/MR simulation, no pathology (e.g. fibrotic tissue in the myocardium) was added. Future extensions in order to improve the realism of the simulation framework will include scatter and randoms simulation and the application of a positron range kernel for PET, and more realistic coil map simulations for MR.

The output in standardised raw data format enables reconstruction with open-source reconstruction packages such as Berkley Advanced Reconstruction Toolbox (BART), Gadgetron, STIR and other frameworks able to read ISMRMRD or Interfile format. The framework allowed an assessment of image registration accuracy to be efficiently performed on data simulated with parameters as would be used in-vivo.

In this study, we used a spline-based model for image registration. Also, the GT motion fields used as input to the simulation were based on the same model. This could lead to a bias in the evaluation of image registration algorithms and requires further investigation.

The framework is not limited to periodic motion. The motion model can include global rotations and translations, which in combination with a non-continuous CM (e.g. a step function) enable the simulation of sudden patient motion.

Simulations often serve as a mean to validate newly developed methods, however, they should

also become a useful tool to allow others to reproduce previous results. Supplying the software as part of SIRF (https://www.ccppetmr.ac.uk/) ensures it will be available as part of a well-maintained framework with an active PET/MR community. Implemented in C++, Matlab and Python interfaces will enable easy use and both simulation and output can be shared among the community. Additionally, the simulation and all its dependencies as part of the SIRF software will be supplied in a virtual machine or Docker container.

3.6 Conclusion

A novel framework to simulate dynamic PET/MR data with motion and contrast changes were presented. It allows for flexible user input and generates output in standardised raw data format, integrating into existing workflows. Its performance was evaluated in different applications and its output compared to patient data, showing that realistic data was generated. The value of the framework for the evaluation and optimisation of registration algorithms was assessed and demonstrated for PET , MR and synergistic PET/MR image registration approaches of cardiac and respiratory motion estimation. Especially for complex reconstruction pipelines (e.g. motion-corrected image reconstruction with subsequent model-fitting to obtain quantitative biophysical parameters), this framework offers the possibility to evaluate how inaccuracies in motion estimation are propagated through the pipeline to the final diagnostic parameters.

This work is the basis for simulations performed in subsequent chapters. Especially, the chemical shift of the MR signal model which has not yet been used allows evaluating fat-water separation techniques. These are discussed in detail in the following and will be the foundation of the acquisition of simultaneous [^{18}F]NaF PET/MR patient data. When eventually applying the subsequently developed MoCo methods to them the simulations will be able to provide confidence in the generation of motion models when GT motion information is no longer available.

4

MR Acquisition-Reconstruction Framework for Improved Motion Estimation

The following chapter describes the development of MR acquisition-reconstruction techniques that allow for the estimation of high-resolution physiological motion models, dynamic attenuation correction maps and 3D anatomical reference images for PET uptake from one comprehensive MR examination. While the presented methods are MR imaging specific they were conceived to be als employed in simultaneous [^{18}F]NaF PET/MR patient data MCIR. Parts of this chapter have been submitted as J8.

4.1 Background

The acquisition of MR data in simultaneous [^{18}F]NaF PET/MR can serve multiple purposes: basis for AC map calculation, anatomical reference for uptake locations and source of motion information for MoCo. While all these data could be acquired in separate acquisitions [71] this reduces scan efficiency. Hence, it is more practical if data are acquired in one comprehensive MR sequence that allows for the extraction of all parameters simultaneously. To meet these requirements there are several demands to the employed MR sequence. The MR image must have a high spatial 3D resolution and sufficient contrast to locate plaques in the coronary system. Also, the sequence must be able to provide both cardiac and respiratory motion information to correct for both motion types. Finally, it requires to classify different tissues for AC generation. However, the MR data need to be acquired in a feasible scan time. Retrospective binning approaches hence are bound to yield undersampled data in each bin which requires the employed k-space trajectory to be motion-robust.
The approach taken in this work builds on a 3D T_1 weighted Dixon acquisition that has previously been used for cardio-respiratory PET/MR MoCo. The advantages of using a Dixon sequence and employing fat-water separated reconstruction are three-fold. First, the

fat around the coronaries can be effectively suppressed in anatomical reference images [57]. Secondly, if both a positive water and fat contrast are available it could be benefitial to use both in estimating motion. Lastly, a fat-water separated reconstruction automatically provides fat as a tissue for an AC map. A model-based fat-water separation framework was implemented and combined with MCIR that allows compressed sensing regularisation on the fat-water separated images instead of individual echo reconstructions. The previously assessed PET/MR synergistic registration is applied to simulated fat-water separated MR images and its performance is evaluated using the framework presented in chapter 3. Furthermore, a novel phase encoding pattern for a 3D RPE sequence [170, 180] was developed providing an ideal distribution of phase encoding points and ensuring the same image quality independent of the periodicity of cardiac and respiratory motion.

4.2 Methods

Sunflower Trajectory

K-space data acquired in this work were sampled with an adaptation of the radial phase encoding (RPE) trajectory [170, 180]. The readout occurs along a fixed direction k_x on a Cartesian grid. All readouts are performed parallel to each other. The phase encoding steps in the ky-kz-plane are arranged radially. In the following, radial lines refer to lines of phase encoding points sampled along a constant angle of the $ky - kz$ plane and do not refer to a radial readout. This is depicted in Fig. 4.1.

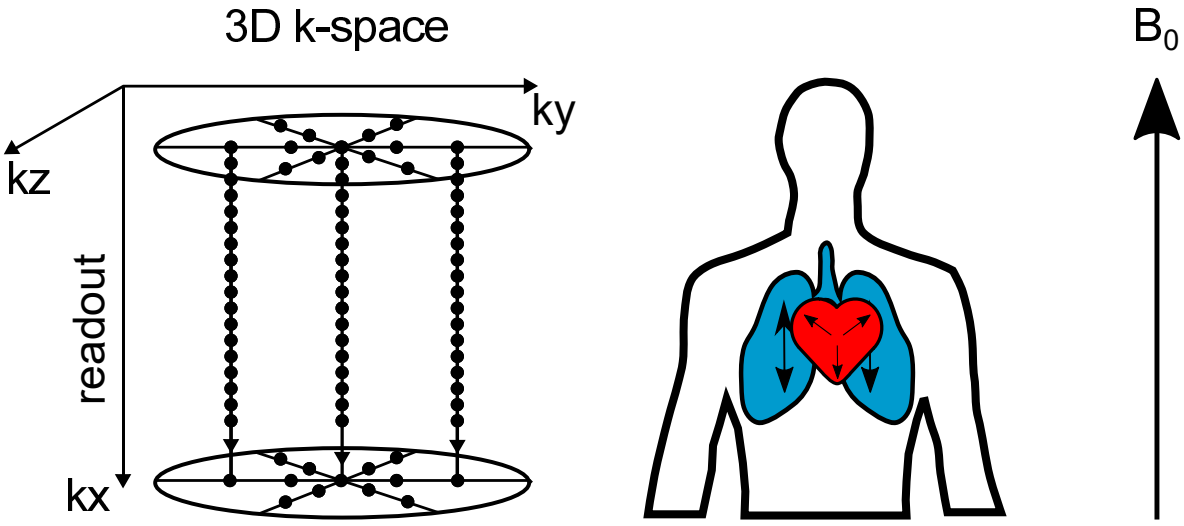

Figure 4.1: Radial phase encoding. Left: schematic k-space sampling pattern of 3D radial. Cartesian readouts along k_x are aligned in parallel and ordered in a non-cartesian fashion in the $k_y - k_z$ phase encoding plane. Readouts are usually performed along the head-foot direction to encode the largest field of view (FOV) along the main patient axis. Parts of this figure have been published in J7.

The trajectory in the phase encoding plane can be computed using Eq. (4.1):

$$\begin{pmatrix} k_y \\ k_z \end{pmatrix} = \begin{pmatrix} \mathfrak{Re} \\ \mathfrak{Im} \end{pmatrix} \left((n_r dr + \Delta(n_\phi)) \cdot e^{i \cdot n_\phi d\phi} \right), \tag{4.1}$$

with $n_r = -\frac{N_r}{2}, \ldots, \frac{N_r}{2} - 1$, $n_\phi = 0, \ldots, N_\phi - 1$ with N_r the total number of radial points per phase encoding line and N_ϕ the total number of angles, $\Delta(n_\phi)$ is an angle-dependent shift, and $\mathfrak{Re}$ and $\mathfrak{Im}$ are the real and imaginary part of a complex number. The step-wise increments in k-space in the radial and angular direction are given by dr and $d\phi$, respectively. Additional variability in the sampling patterns is introduced by the angle-dependent shift, displacing the complete set of points along this angle by $\Delta(n_\phi)$. Omitting an additional shift by setting $\Delta(n_\phi) = 0 \; \forall \; n_\phi$ yields concentric circles with a radial distance of dr between adjacent rings. As this distance defines the radial FOV, a reduction of radial folding artefacts can be achieved by a reduction of the distance between rings without increasing the number of sampled phase encoding points. Previously, the RPE trajectory has been used with 4-fold interleaved circles [170] with an angle-dependent shift in Eq. (4.2)

$$\Delta(n_\phi) = \frac{dr}{4} \cdot \pi_{BR}\big(\mathrm{mod}(n_\phi, 4)\big) = \Delta_{BR}(n_\phi) \tag{4.2}$$

$$\pi_{BR} : \begin{pmatrix} 0 & 1 & 2 & 3 \end{pmatrix} \mapsto \begin{pmatrix} 0 & 2 & 1 & 3 \end{pmatrix} \tag{4.3}$$

where mod represents the modulo operator and π_{BR} the bit-reversal permutation defined in Eq. (4.3). Each concentric ring is split into 4 sub-rings decreasing the effective radial gaps between them.
The trajectory proposed in this work is generated by the angle-dependent shift described in Eq. (4.4):

$$\Delta_{SFL}(n_\phi) = \frac{dr}{2}\big(2 \cdot \mathrm{mod}(n_\phi, \Phi) - 1\big), \tag{4.4}$$

where $\Phi = \frac{1+\sqrt{5}}{2}$ is the golden ratio. This applies the principle of golden-ratio based angular increments [181] to the radial sampling direction of polar sampling. The pattern generated through more and more phase encoding points is displayed in Fig. 4.2 and has the appearance of seeds assembled in a sunflower, similar to a spiral readout trajectory based on Fibonacci spirals [182].
The angle-dependent shift Δ_{SFL} from Eq. 4.4 assumes a linear angular increment $d\phi = \frac{\pi}{N_\phi}$. The trajectory in this work was used with a golden angle increment $d\phi = \pi \cdot (\Phi - 1)$, where when the final number of angles was known, assuming a constant angular spacing between them and the shift for each radial line was applied in the correct order and with the appropriate sign.

One challenge encountered in standard RPE encoding is that the acquisition time of the phase encoding points along one radial line is determined by how many points lie along it. Ultimately this number is constrained by the FOV that needs to be encoded in the radial direction. Subsets generated by retrospective binning of the data into motion states hence often yields connected phase encoding points along individual radial lines. For the shift Δ_{SFL}, however, such infinitely interleaved circles generated by (4.4) for $N_{r/\phi} \cdot d(r/\phi) = const$ allow to undersample the radial direction $N_r \mapsto \frac{N_r}{F_\phi}$ while simultaneously oversampling the angular direction $N_\phi \mapsto F_\phi \cdot N_\phi$. Hence, the encoded radial FOV is not depending on the number of points along one radial line. Upon retrospective binning this simultaneous radial under- and angular oversampling generates undersampling patterns with more homogeneous distribution

and hence more incoherent artefacts [183]. The shift described by (4.4) depends on the radial distance and hence the resolution increases at the same time by a factor of $\frac{\hat{N}_r - F_\phi}{\hat{N}_r - 1}$ compared to the fully sampled case where $\hat{N}_r$ is the number of radial points per radial line for the full FOV and F_ϕ the radial undersampling factor. This has to be taken into account during image reconstruction. The scaling yields 0.94 for $F_\phi = 16$ which would lead to approximately a 6% shrinkage of the object in the reconstructed images if not corrected for.

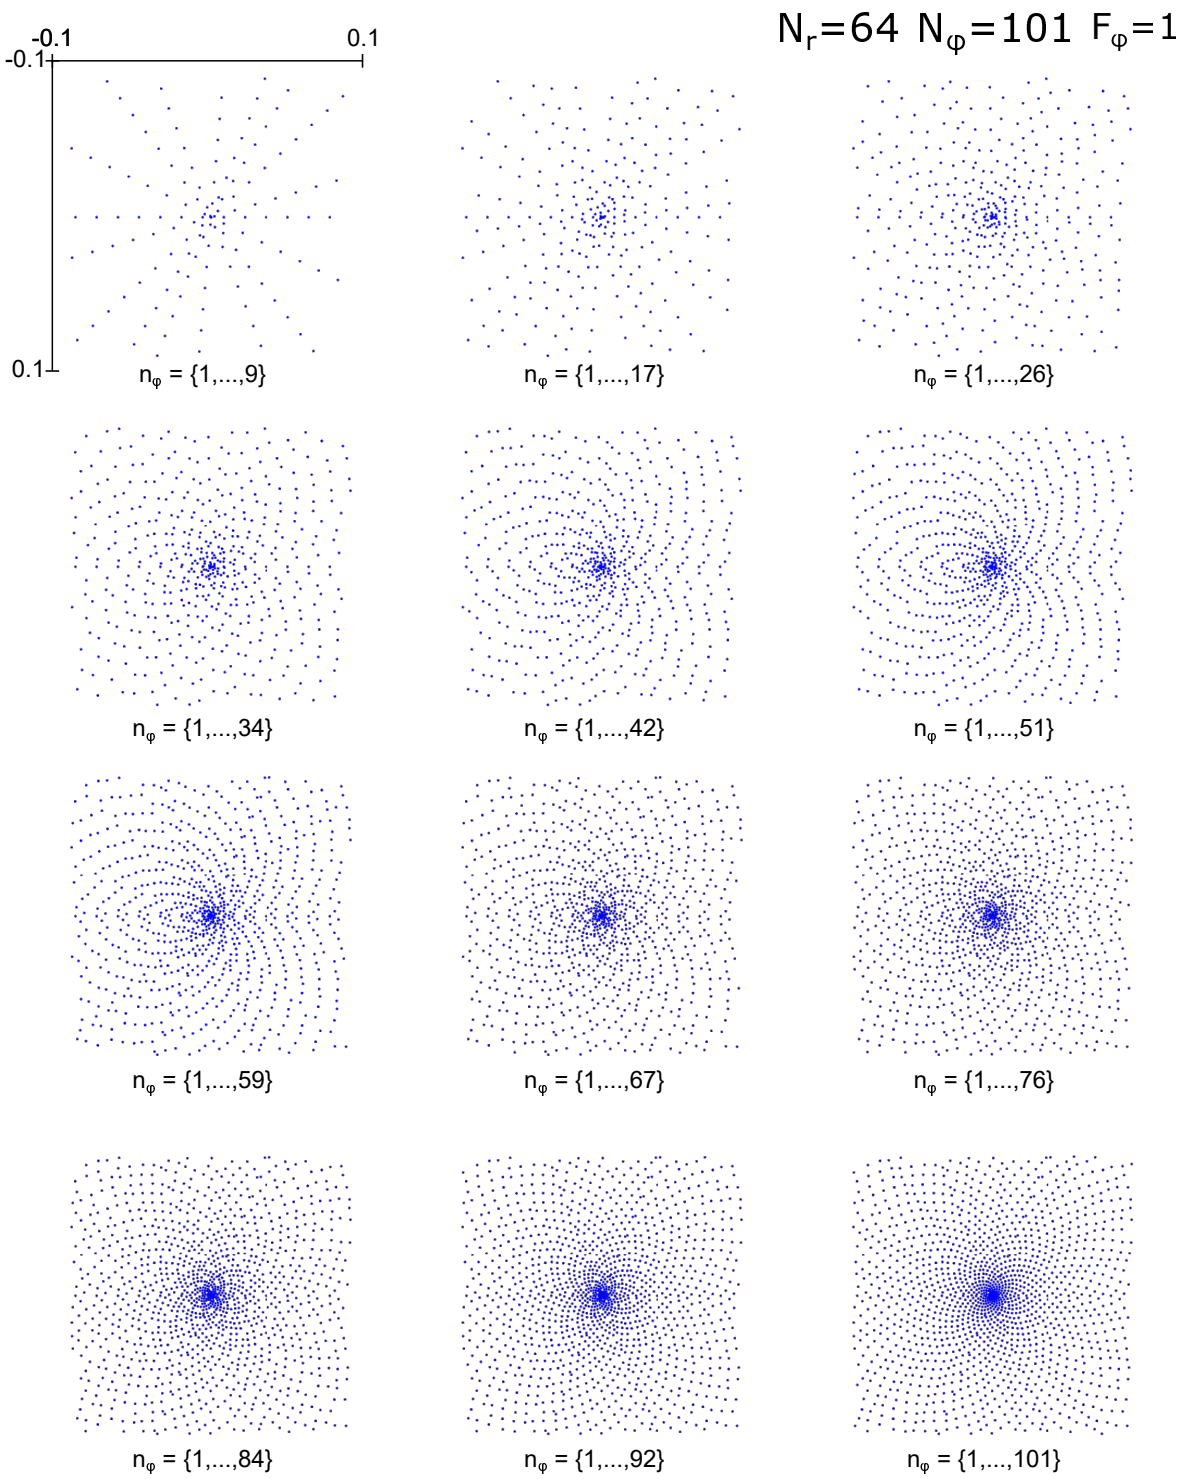

Figure 4.2: Sunflower pattern creation. With each frame the number of angles is increased for $N_r = 64$ radial points and $N_\phi = \frac{\pi}{2} \cdot Nr = 101$ angles. It can be seen, how the gaps in the pattern are filled in such that finally, no two points have the same distance from the centre. The depicted range is [-0.1, 0.1] while the trajectory is normed to [-0.5,0.5].

Model-based Fat-Water Separation

A model-based iterative fat-water separation was implemented [67, 68], and extended by a model describing physiological motion [89] (e.g., respiratory or cardiac motion). The image is a 4D water-fat doublet $(W, F)(x, y, z, m) = (W, F)(\boldsymbol{r}, m) = (W_m, F_m)$ where x, y, z or $\boldsymbol{r}$ are the spatial directions and m labels the motion state. Physiological motion is defined as a map $\hat{T}_m : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ where such that $(W, F)(\boldsymbol{r}, 0) = (W, F)(\hat{T}_m \circ \boldsymbol{r}, m)$ where $m = 0$ is an arbitrary reference motion state. The acquisition model is described by the action of encoding operator $\hat{E}$ for echo time T_E, receiver channel c and motion state m by:

$$k_c(T_E, m) = \hat{E}_{(c,T_E,m)}(W_0, F_0) \tag{4.5}$$

$$= \mathbb{P}_m\big(\mathcal{F}(C_c e^{2\pi i \Phi T_E} \cdot W_0 \circ \hat{T}_m^{-1}) + D(T_E)\mathcal{F}(C_c e^{2\pi i \Phi T_E} \cdot F_0 \circ \hat{T}_m^{-1})\big) \tag{4.6}$$

where C_c is the normalised coil sensitivity profile, $\mathcal{F}$ the Fourier transform, Φ an off-resonance frequency map due to local field inhomogeneities $\Delta B_0(\boldsymbol{r})$, $\mathbb{P}_m$ projects the data to the motion bin m, and $D(T_E)$ is the chemical shift between water and fat. The fat-water shift is approximated by a six-peak model [67, 68, 184] of the form:

$$D(T_E) = \sum_{p=1}^{6} \alpha_p \cdot e^{i\gamma\Delta_p T_E} \tag{4.7}$$

in which γ is the gyromagnetic ratio α_p is a relative amplitude and Δ_p the chemical shift of peak p relative to the water signal. The adjoint action is defined by:

$$(\tilde{W}_0, \tilde{F}_0) = \sum_{c,T_E,m} \hat{E}^H_{(c,T_E,m)} k_c(T_E, m) \tag{4.8}$$

$$= \sum_{c,T_E,m} e^{-2\pi i \Phi T_E} \cdot \hat{T}_m \circ C_c^\star \mathcal{F}^{-1}\big(k_c(T_E, m), D^\star(T_E) \cdot k_c(T_E, m)\big) \tag{4.9}$$

where the $\tilde{}$ in $(\tilde{W}_0, \tilde{F}_0)$ signifies that these are only the result of the application of the adjoint and not the actual inverse. A direct inverse of the model described in Eq. (4.6) does not exist such that the inversion must be found iteratively.

An inversion of the described model in Eq. 4.6 yields a fat-water separated MCIR in the state $m = 0$. However, this assumes that the motion $\hat{T}_m$ is known which is usually not the case but a motion-model must be generated from the data. To this end, a motion-resolved reconstruction is required, created based on motion-binning using a surrogate signal as described in chapter 2. The solution of the reconstruction problem will yield a 4D motion-resolved reconstruction if the application of a motion model and a summation over all motion bins in Eq.(4.6) and Eq.(4.9) are omitted.

The iterative inversion is formulated as the inversion of the cost function $C(W, F)$ given by:

$$C(W, F) = \|\hat{E}(W, F) - k\|_2 + \lambda_{TV}\|TV(W, F)\|_1 + \lambda_{TVT}\|TVT(W, F)\|_1 \tag{4.10}$$

where λ_{TV} and λ_{TVT} are the respective weights associated to the spatial total variation (TV) and total variation along the motion direction (TVT) regularisations. This was optimised using non-linear conjugate gradient descent with a line search optimisation [88, 185]. This

way the model-based inversion incorporates parallel imaging [87] and MoCo [89] as well as compressed sensing [88]. While in [67, 68] the off-resonance map Φ was a reconstructed parameter and part of the optimisation, in this study its estimate was computed before the optimisation [65] based on the three echoes and was subsequently kept constant and used as a model parameter. A schematic overview of this process is depicted in Fig. 4.3.

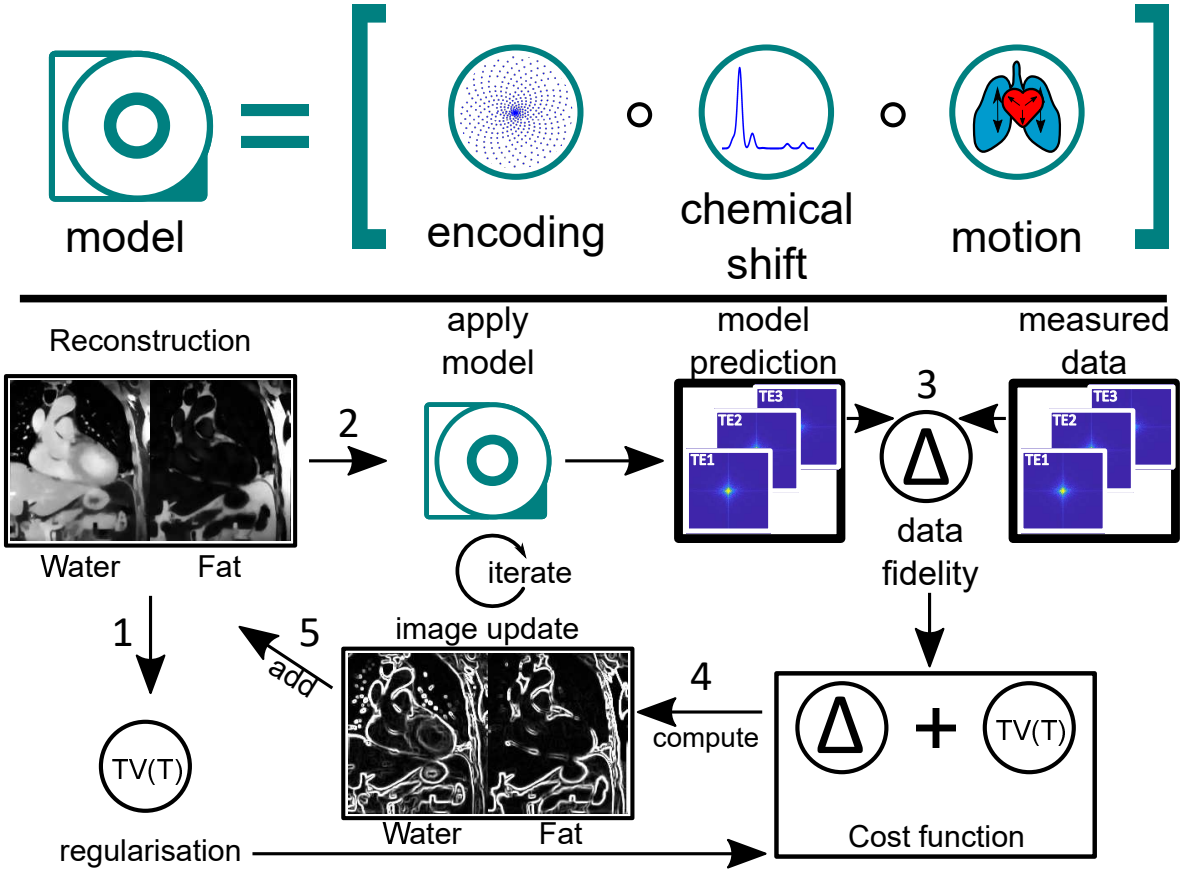

Figure 4.3: Schematic overview of model-based reconstruction framework. Top: components of employed acquisition model, symbolically referred to as an MR scanner. Fat and water images are first deformed by the motion model, then a chemical shift is applied by adding an echo time-dependent complex phase, and eventually, Fourier encoding including coil-sensitivities is performed. Bottom: iterative inversion of the acquisition model. Based on the current image estimate the regularisation is computed (1). Then the acquisition model is applied to the current image estimate (2) and its output (model prediction) is compared to the measured data by computing the data fidelity (3). Based on the cost function, comprised of the regularisation and data fidelity term, an image update for both water and fat is computed (4) and added to the current image estimate (5). This process is iterated until convergence.

Fat-Water Dual Image Registration

In the forward model underlying the reconstruction, the motion transformation $\hat{T}_m$ in Eq. (4.6) is assumed to be known. One way to obtain this information is to use an image registration algorithm on a set of motion-resolved images (W_m, F_m). The registration algorithm employed in this work is based on an algorithm previously used in dual-modality image registration in simultaneous cardiac PET/MR [82]. It was shown that combined use of both PET and MR images in the registration leads to a more robust motion estimation t han if both data streams

are used separately. This principle was transferred to the fat-water separated images using fat as complementary information to the water images. The registration algorithm is formulated as an optimisation problem with the associated cost function:

$$C(\hat{T}_m) = \lambda \cdot \mathcal{S}(W_m \circ \hat{T}_m, W_0) + (1-\lambda) \cdot \mathcal{S}(F_m \circ \hat{T}_m, F_0) + \rho \cdot \mathcal{R}(\hat{T}_m) \tag{4.11}$$

where $\mathcal{S}$ is a similarity metric, λ weighting the contribution of water relative to fat (i.e. $\lambda = 1$ corresponds to water only, $\lambda = 0$ to fat only), and ρ the weight associated to a regularising term $\mathcal{R}$. The parametrisation of the transformation $\hat{T}_m$ with B-splines [78] added further implicit regularisation. In this work, $\mathcal{S}$ was normalised mutual information, and the explicit term $\mathcal{R}$ was set to the so-called bending energy (i.e. the absolute value of the motion vector field's second spatial derivative) of the transformation [78, 186]. Larger bending energy forces the computed $\hat{T}_m$ towards a more linear transformation.

Quantitative Assessment of Image Sharpness

To assess the effect of MoCo on the image quality, the sharpness Σ of structures visible in the fat images was evaluated [187–190]. To this end, two coronal and two sagittal slices were extracted from the fat image of each reconstruction. In the sagittal slices the apex and the base, and the coronal slices, the pericardium at the left myocardium and right atrium were manually delineated with a smooth curve. These evaluated edge positions are schematically displayed in Fig. 4.4.

Figure 4.4: Sharpness computation locations in the numerical simulation. The four locations are shown here for the ground truth motion-corrected fat image as the red dotted lines.

A canny edge detection algorithm was applied to the fat images yielding edge information. The values of both fat and edge images were extracted along the delineated line ℓ with a width of 16 pixels and subsequently averaged over the line width yielding i_ℓ and e_ℓ for the image and edge information respectively. Finally, the edge sharpness was defined as

$$\Sigma_\ell = \frac{\max e_\ell}{\max i_\ell} \tag{4.12}$$

ΔB	$10^3 \cdot \rho$	λ
(4,6,8,10,12)	(0,1,2,3,4,5)	(0, 0.5, 1)

Table 4.1: Parameters used in the numerical registration evaluation.

4.3 Experiments

Sunflower Trajectory

The different patterns generated by $\Delta(n_\phi) = 0$, Δ_{BR} and Δ_{SFL} are generated for under-oversampling factors $F_\phi = 1, 2, 4$. Furthermore, phase encoding point distributions generated by Δ_{SFL} and $F_\phi = 4$ and Δ_{SFL} and $F_\phi = 1$ are binned retrospectively using a respiratory belt and ECG signal from patient data for $N_{rad} = 192$ and $N_\phi = 512$. The number of bins is $N_{resp} = 6$ for respiratory amplitude and $N_{card} = 12$ for cardiac phase binning.
In all cases, the patterns are compared and assessed visually.

Simulation for Fat-Water Registration Evaluation

The numerical simulation framework described in chapter 3 was used to simulate dynamic MR data containing either 4D respiratory or 4D cardiac motion [166, 191]. The underlying anatomy segmentation and matching ground-truth motion fields were generated by the XCAT [159]. Coil sensitivities were simulated using Gaussian distributions and no field inhomogeneities were added. Data were reconstructed using the same fat-water reconstruction as were applied to the patient data. A schematic overview of the employed simulation framework is given in Fig. 4.5. The registration can be performed for multiple sets of parameters and for each registration output, the error to the ground truth MVF can be computed in an ROI. For this evaluation, the ROI consisted of the pericardium and the myocardium as highlighted in Fig. 4.5.

Evaluation Registration Error

The free parameters of the registration are the triplet spline distance, bending energy weight, and water-weight $= (\Delta B, \rho, \lambda)$. The registration was run for different parameter spacing combinations that are displayed in Tab. 4.1.

Evaluation Image Sharpness Metric

Each registration was used to motion-correct the input images using the RTA method: $(W, F)_{MoCo} = \langle (W, F)(\hat{T}_m \circ \boldsymbol{r}, m) \rangle_m$. Σ defined in Eq. (4.12) was computed and for each parameter set it was related to the registration error of the MVF used for MoCo.

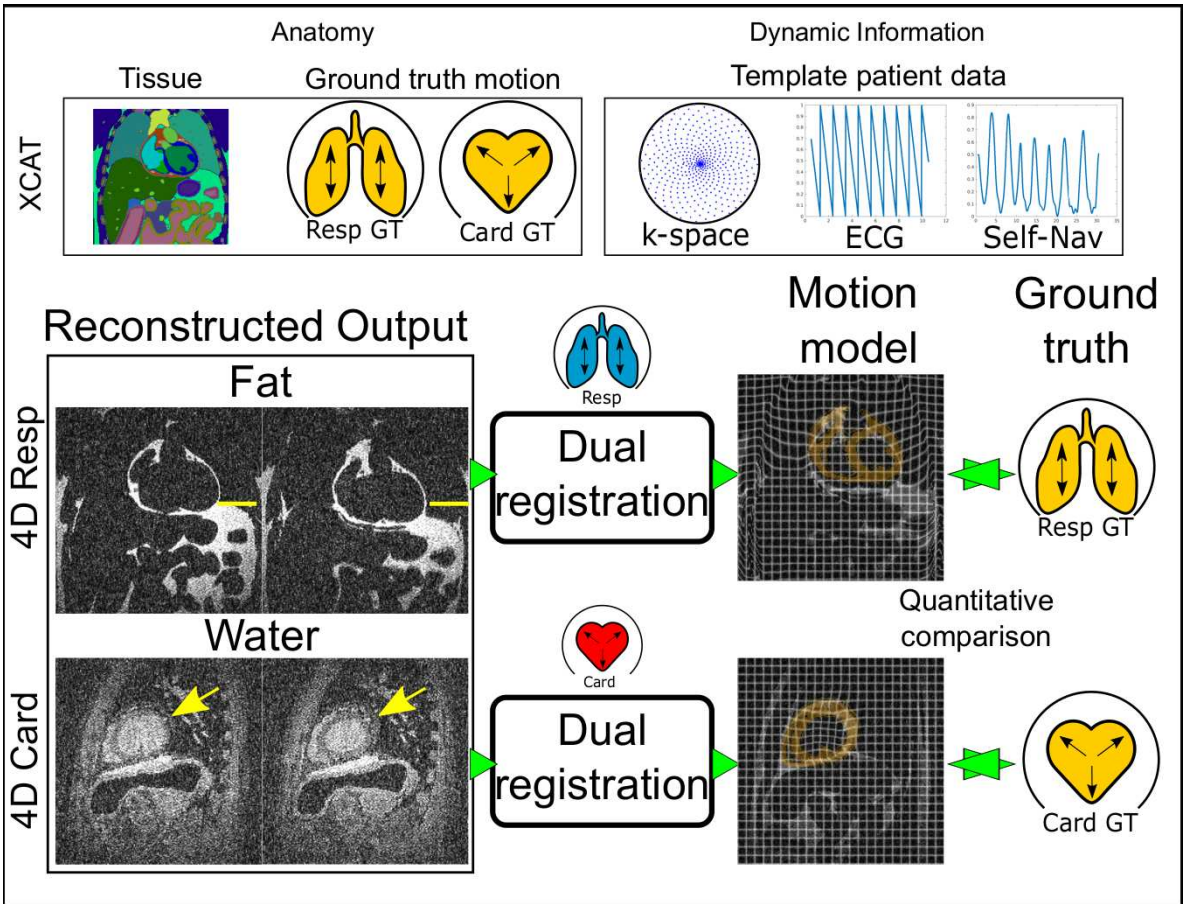

Figure 4.5: Schematic overview of simulations underlying the registration evaluation. Top: simulation input as an XCAT anatomy segmentation and corresponding ground truth motion fields, and dynamic information in the form of template MR k-space data and surrogate signals extracted from one patient. Bottom: reconstruction output displayed for the cardiac and respiratory motion for both fat and water images. Exemplary fat is displayed for respiratory motion, and water for cardiac motion. Yellow lines and arrows show the motion amplitude contained. Registrations were performed on data where water and fat were weighted by $1 - \lambda$ and λ respectively. Eventually, the registration output is compared to the ground truth motion in the ROI indicated by the orange overlay.

4.4 Results

Sunflower Trajectory

The resulting trajectories for simultaneous radial under- and angular oversampling are depicted in Fig. 4.6. The green H marker shows the radial k-space distance necessary to encode the FOV. It should be noted, however, if this FOV were used for encoding it was already insufficient to suppress all radial folding artefacts that appear as a ring around an acquired subject. For acceptable image quality, this radial distance between adjacent rings is not suited. As can be seen, only for Δ_{SFL} the newly sampled points are located in between the previously sampled ones, ensure no large gaps occur between the acquired k-space points for any F_ϕ.

This allows adapting the temporal resolution of one radial line (i.e., the time to acquire all phase encoding points along one line) almost arbitrarily. The effect of retrospective binning is depicted in Fig. 4.7. The motion gating does not generate subsets of single radial lines in the

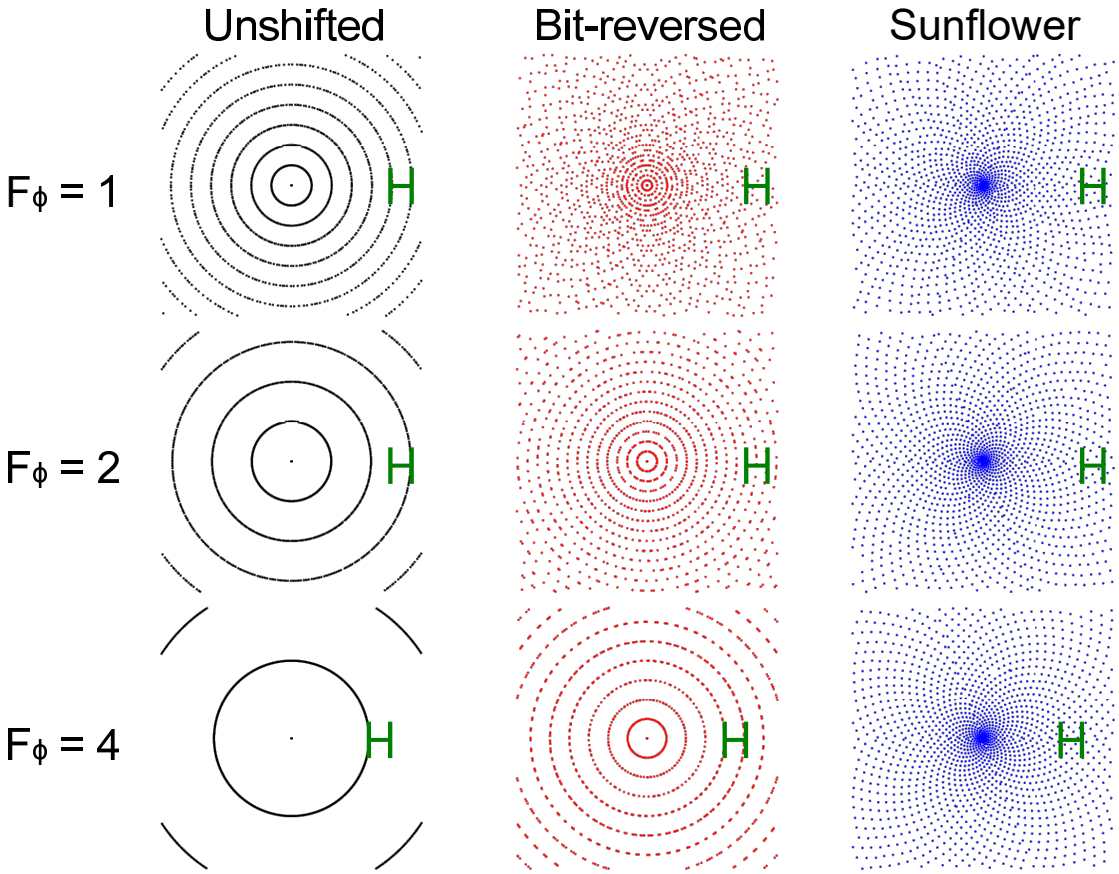

Figure 4.6: RPE undersampling-oversampling patterns. Three phase encoding trajectory patterns generated by $\Delta(n_\phi) = 0$ (left), $\Delta_{BR}(n_\phi)$ (centre), and $\Delta_{SFL}(n_\phi)$ (right). The green spacing defines the distance in k-space that defines the encoded FOV by dr. A simultaneous radial under- and angular over-sampling reduces the FOV for both unshifted and bit-reversed patterns but leaves the sunflower pattern invariant. Also, one can visually detect that the sunflower pattern has the smallest mean distance between adjacent k-space points and no concentric circles appear as there are no two points with the same distance from the centre.

phase encoding plane depending on motion periodicity and the regions without any sampled phase encoding points (black arrows) are smaller.

Finally, since the k-space centre is sampled with every radial line, a temporal resolution of a self-navigator is $\Delta t = N_r \cdot T_R$. Hence, the temporal resolution is improved to $\Delta t = \frac{N_r}{F_\phi} \cdot T_R$. This would also allow the sampling of a respiratory self-navigator if T_R becomes large.

Fat-water Synergistic Registration

The motion field errors for different registration parameters are plotted in Fig. 4.8. One can see that over a large range of parameters a registration error below 1 voxel can be achieved. For cardiac motion, when no additional regularisation ($\rho = 0$) is applied the only water-based ($\lambda = 1$) registration has an error larger than the motion amplitude.

The quantitative analysis of the registration error is summarised in Tab. 4.2. For both respiratory and cardiac motion, the minimum error is the smallest for a $\lambda = 0.5$ at 0.52 and 0.74 voxel, respectively. The average error for both motion types is below one voxel for

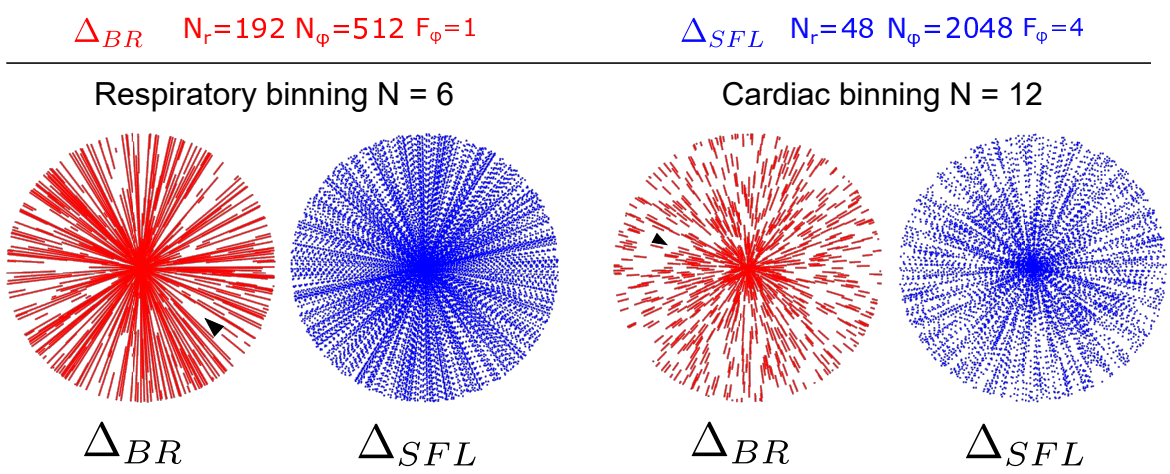

Figure 4.7: RPE patterns after retrospective binning. The pattern generated by Δ_{BR} (red) shows larger regions in which no points are sampled (black arrows) compared to Δ_{SFL} with $F_\phi = 4$ (blue). The overall distribution is more homogeneous for the patterns generated by Δ_{SFL}

respiration and slightly above one voxel for cardiac motion. The standard deviation of the error is similar for all weights.

Motion Type	λ	Minimum error	Average error	Std. error
Cardiac	0.0	0.57	0.69	0.08
	0.5	0.52	0.70	0.13
	1.0	0.62	0.81	0.21
Respiratory	0.0	0.82	1.16	0.29
	0.5	0.74	1.04	0.26
	1.0	0.78	1.09	0.24

Table 4.2: Quantitative analysis of the registration errors. The minimum, average and standard deviation are computed over the set of all performed registration parameter combinations. All errors are displayed in voxel sizes.

In Fig. 4.9 the results of the quantitative analysis of the registration are presented. One can see that the best registration can follow the motion amplitude accurately across all motion states where the average error is reduced to below one voxel. For the best registration parameters the dependency of the registration error λ is small, however, the optimal registration consists of both water and fat images weighted equally. The scatter plots of Σ versus the registration accuracy show that there is a clear difference between neglecting motion and using the ground truth to correct it. However, for the registrations, there is no visible correlation between a precise image registration and the resulting Σ for any λ. The largest error for cardiac motion is of the order of not correcting motion and scores a higher Σ than the ground truth MoCo.

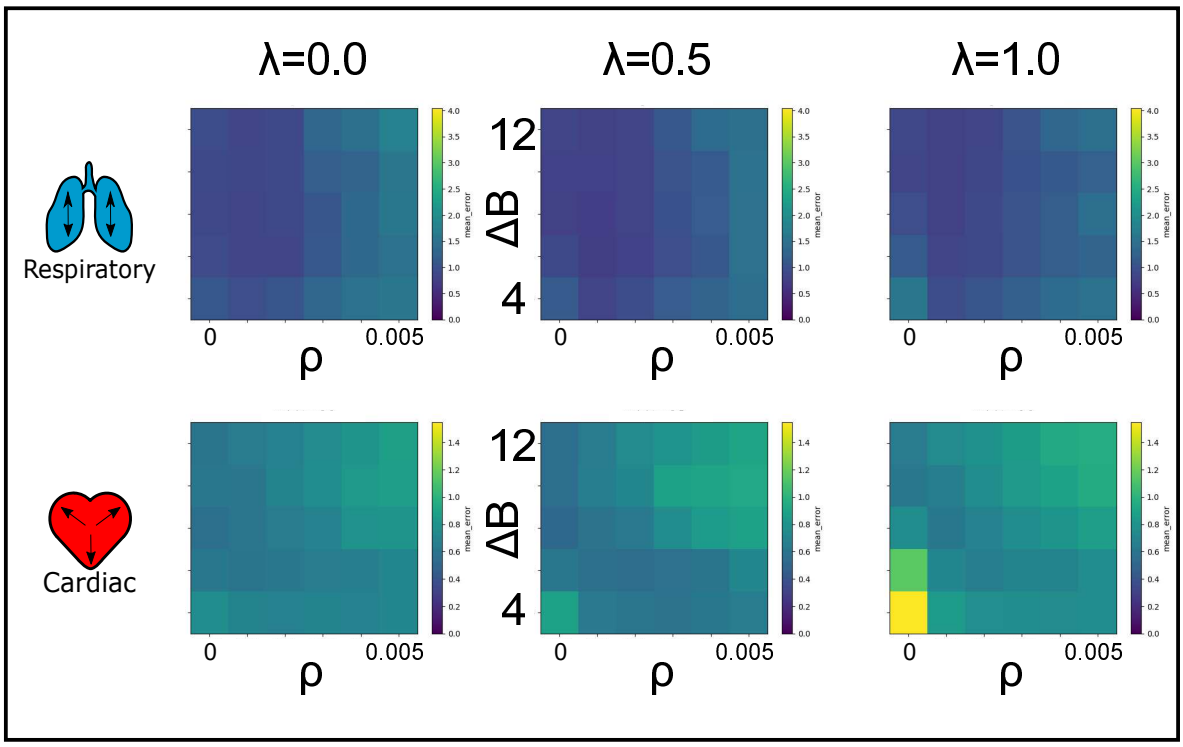

Figure 4.8: Results of the numerical evaluation of the registration accuracy. Top: respiratory motion. Bottom: cardiac motion. Left to right: $\lambda = 0$ (fat only), $\lambda = 0.5$ and $\lambda = 1$ (water only). The minimum of the colour bar corresponds to zero error, the maximum to the mean amplitude of the ground truth motion. Both motion types can be registered with an accuracy of up to 1 voxel size for a variety of parameters. Only cardiac motion reaches errors of the same size as the ground truth amplitude if no regularisation is used.

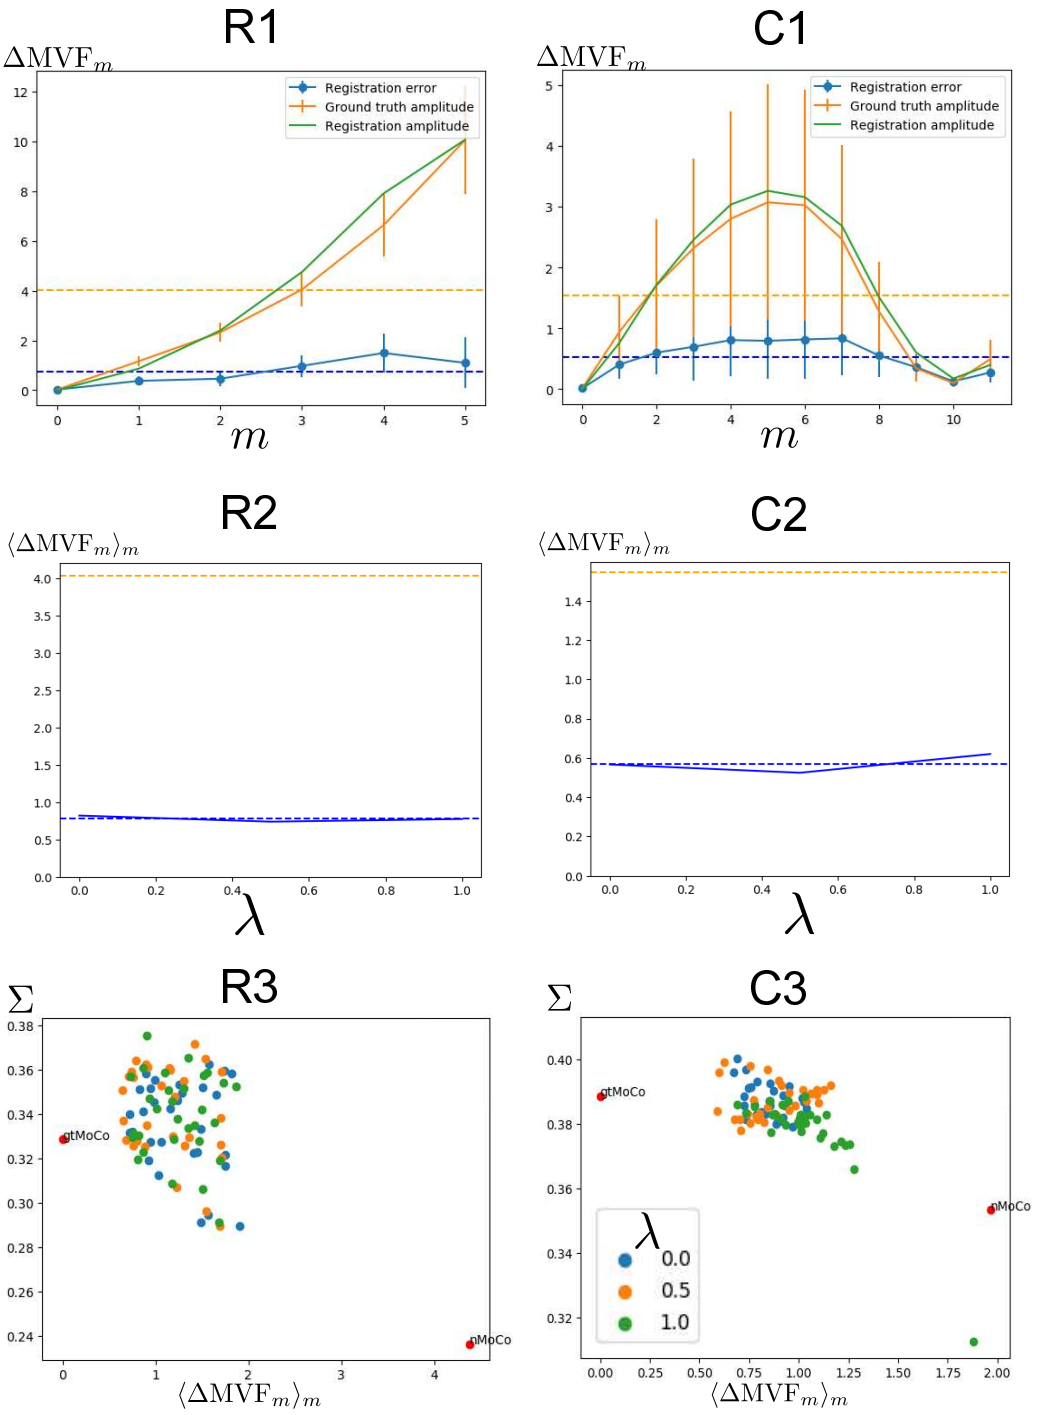

Figure 4.9: Evaluation of registration error. Left column: respiration (R1-R3), right column: cardiac motion (C1-C3). R1/ C1 show the motion-state resolved ground truth amplitude in green, the registration amplitude in orange, and the registration error in blue. Error bars correspond to one standard deviation. In both cases, the motion amplitudes are well captured and the residual errors have only a small dependency on the motion amplitude. R2/C2 show the relationship of the registration error for the best (ΔB, ρ) pair determined by the analysis of the registration errors. A minimum occurs in both curves for the $\lambda = 0.5$. However, the actual improvement compared to $\lambda = 0$ and $\lambda = 1$ is small. R3/C3: scatter plot showing the distributions registration error and image sharpness Σ. The two red points show an improvement from neglecting motion (nMoCo) to correcting with the GT motion (gtMoCo). The cardiac registrations indicate a correlation between $\langle\Delta MVF_M\rangle_m$ and Σ (C3). Respiratory registration shows little correlation (R3).

4.5 Discussion

In this chapter, a novel k-space trajectory, as well as a motion-corrected model-based fat-water, separated reconstruction framework was presented. Also, a simulation-based assessment of the registration accuracy of synergistic registration applied to water-fat separated images was performed for both cardiac and respiratory motion.

The novel phase encoding pattern in RPE encoding allowed for simultaneous radial under- and angular oversampling of the radial lines in the phase encoding plane. This resulted in a more homogeneous distribution of retrospectively motion-binned data. Large angular gaps in the phase encoding space lead to strong, coherent streaking artefacts as encountered for example in stack-of-stars sequences [81]. The size of these gaps space could be reduced compared to a bit-reversed shift of radial lines. The central region of k-space was still densely sampled post motion-binning, yielding patterns more suitable to compressed sensing applications. Furthermore, it increases the temporal resolution of the k-space center update and hence of the self-navigator. However, the frequency of k-space center samples per time did not allow for an extraction of cardiac surrogate signals as possible for example with 3D radial readouts [101].

The evaluation of the synergistic registration showed that the employed registration framework can successfully register respiratory and cardiac motion and reduce the motion errors to below the image resolution as shown in Fig. 4.8. It could also be seen, that the chosen image sharpness metric Σ improved when correcting with the ground truth motion compared to not correcting motion for both motion types. In the case of respiration, however, there was little correlation between respiratory registration precision and examined image quality metric Σ. For cardiac motion, where the error relative to the ground truth amplitude was larger there was a stronger correlation than for respiration. The outlier with the large registration error presented in Fig. 4.9 C3 showed that a large error in registration produces a low Σ. But for both motion types registrations were performed where a Σ higher than for the ground truth MoCo was scored despite having a non-zero registration error. This shows that while the correlation between registration error and Σ indicates that correcting for motion leads to an increase in Σ, this statement can not be inverted. An increase in Σ does not induce a good registration. An additional visual inspection is required to assess the quality of the estimated motion model to ensure that it follows the underlying image data, and provides a correct model of the physiological motion.

The dependence of the registration error for different weights of the water image λ was weak in the case of the reconstructed simulation data. The numerical phantom analysis was, however, limited to the single anatomy and motion provided by the underlying XCAT segmentation. Also, the similarity between reconstructed simulation data and patient data was limited. This is likely due to the use of unrealistic coil sensitivity profiles. The actual patient data shows a stronger drop-off of sensitivity, weighting the chest and back area stronger and leading to more coherent artefacts which could be achieved using a more sophisticated coil sensitivity computation [151]. Stronger dependencies between fat and water weighting λ and registration precision could well be possible when applied to patient data. This, however, could not be tested due to a lack of ground-truth motion information for patient data.

The performance of the motion-corrected model-based fat-water separation was not assessed here and merely applied to reconstruct the simulation output. However, a thorough analysis of its application to patient data is performed in chapter 5.

4.6 Conclusion

In this chapter, a comprehensive 3D high-resolution cardiac MR scan was developed aimed at cardiac PET-MR applications. To provide high image quality, accurate motion information and allow for AC calculation, a novel phase-encoding pattern for 3D RPE k-space sampling, a synergistic registration framework for fat-water separated image data, and a motion-corrected model-based fat-water reconstruction framework were developed. Although these methods were aimed to improve coronary plaque imaging for NaF-PET/MR, these methods are in general well-suited to image small cardiac structures. In the next chapter the framework developed here is validated for patient data by visualising fat infiltrations in the myocardium at 1.5 T . Chapter 6 then presents the application of the framework to NaF PET/MR in 10 patients suffering from atherosclerosis.

5

Model-based Fat-Water Separation in Cardiac Fat Quantification

In this chapter the methods developed in the previous chapter are applied in patients to image small structures in the heart. It describes an independent cardiac MR (cMR) application at 1.5 T which yields high-resolution fat and water images with diagnostic quality. Parts of this chapter are under revision as J8.

5.1 Background

Adipose tissue can infiltrate the myocardium, either isolated or paired with fibrosis, both of which are associated with a range of ischaemic and non-ischaemic cardiomyopathies [38]. These include chronic myocardial infarction (MI), muscular dystrophy and arrhythmogenic right ventricular cardiomyopathies (ARVC) [49, 192, 193]. CMR imaging can characterise chemical tissue components non-invasively and therefore can provide important diagnostic information [18].

Fat can be visualised by suppressing the water signal (spectral suppression) or by utilising the chemical shift to separate the fat and the water signal from multiple acquisitions at different echo times (Dixon imaging) [18]. In cardiac applications, multi-echo fat-water separation techniques feature multiple advantages over spectral suppression methods. Direct identification of a positive fat contrast automatically mitigates any incomplete suppression experienced with spectral methods [63, 194]. Also, multiple echoes enable taking B_0 field inhomogeneities into account [65, 195]. Separation techniques are also able to resolve ambiguities in late gadolinium enhancement (LGE) imaging when fatty infiltrations can be mistaken as fibrotic tissue in T_1 -weighted sequences because both have a low post-contrast T_1 value [194]. Furthermore, chemical shift artefact induced misregistrations between water and fat are suppressed in multi-point fat-water separated imaging [66, 196].

Commonly in clinical routine fat-water imaging is carried out as multiple slices with high in-plane resolution but a large slice thickness of 4 mm to 8 mm during one or over multiple

breath-holds. Images are either acquired as cine [64], resolving the cardiac motion or are ECG triggered such that data are only acquired in mid-diastole to mitigate artefacts due to motion caused by the heartbeat.
For fat-water imaging, images are usually first reconstructed at different echo times and then separated into fat and water content. Model-based fat water reconstruction on the other hand directly obtains a fat and water image from the acquired multi-echo k-space data. This allows a direct regularisation of the final fat and water images instead of regularising only of the multi-echo images and ensures that the obtained fat and water images are consistent with the obtained k-space data. This approach has been combined with compressed sensing [67] or with XD-GRASP for applications in the abdomen [68, 97].
Fat infiltration in the heart can be very small and irregularly distributed in the myocardium. Therefore, they can be poorly resolved or missed with multi-slice 2D imaging. High-resolution 3D fat-water imaging has been proposed [197], but the main challenge is long scan time because artefacts due to respiratory and cardiac motion need to be minimised using gating and triggering, respectively [49, 90, 103, 106, 198]. Respiratory motion correction has been proposed for fat-water imaging [107, 114], but cardiac triggering still reduces scan efficiency and can lead to scan times which are longer than clinically feasible. Similar technquies using fat have also been used for vessel imaging using a fat-selective excitation to correct for 3D translational motion using epicardial fat tracking [199, 200].
In this chapter, the development of a combined cardiac and respiratory motion-corrected 3D model-based fat water reconstruction is presented. The data acquisition is performed with a non-triggered and free-breathing sequence which covers the whole heart with a radial phase encoding (RPE) trajectory [180]. Images are reconstructed with a regularised model-based fat water framework. Respiratory and cardiac motion is estimated from the data and then utilised in a subsequent cardio-respiratory MCIR to minimise motion artefacts and improve the accuracy of fat visualisation. The motion modelling is based on simultaneous non-rigid registration of both water and fat images as analysed in the previous chapter. The model-based reconstruction is validated on data acquired from 9 patients using clincial 2D images as a reference. The effect of cardio-respiratory MoCo on adipose cardiac tissue is assessed with a local edge sharpness metric applied to the reconstructed fat images.

5.2 Methods

Two-Step Motion-Corrected Image Reconstruction

A schematic overview of the employed acquisition and reconstruction workflow is given in Fig, 5.1. In a first step, the respiratory self-navigator is extracted and used to bin the data into $N_{\text{resp}} = 6$ respiratory motion states based on respiratory amplitude with a 10 % overlap between adjacent bins. This yielded an undersampling factor of 3 for the respiratory reconstruction. The binned data are reconstructed into water and fat content using a total variation (TV) and total variation in time (TVT) regularisation strength of $\lambda_{TV} = 0.05$ and $\lambda_{TVT} = 0$. The resulting image series are subsequently registered using the dual fat-water registration with $(\lambda, \Delta B, \rho) = (0.7, 8, 0)$ yielding MVFs that describe the respiratory motion model. In a second step, the data are binned into $N_{\text{card}} = 12$ cardiac motion states based

on cardiac phase using the ECG signal. For cardiac binning, a 0 % overlap between adjacent bins was used. This yielded an undersampling factor of 6 for the cardiac reconstruction. The resulting data are reconstructed into water and fat modes using a regularisation strength of $\lambda_{TV} = 0.025$ and $\lambda_{TVT} = 0.05$ while simultaneously applying the previously generated respiration model for respiratory MoCo. The resulting image series resolves the heartbeat motion, and each image itself is already free of artefacts caused by respiration. The image registration with $(\lambda, \Delta B, \rho = (0.5, 2, 2 \cdot 10^{-3})$ was applied to determine the cardiac MVFs. Eventually, both motion models were combined into a cardio-respiratory motion model and applied to reconstruct a single, motion-free fat and water separated image. Finally, three reconstructions were compared: the motion-averaged reconstruction (AVG), respiratory motion-corrected image reconstruction (r-MCIR), and cardio-respiratory motion-corrected image reconstruction (cr-MCIR). Each was *TV* regularised with $\lambda_{TV} = 0.015$ (as there was only one reconstructed image *TVT* regularisation could not be performed, $\lambda_{TVT} = 0.0$).
Further details about the employed reconstruction framework can also be found in chapter 4.

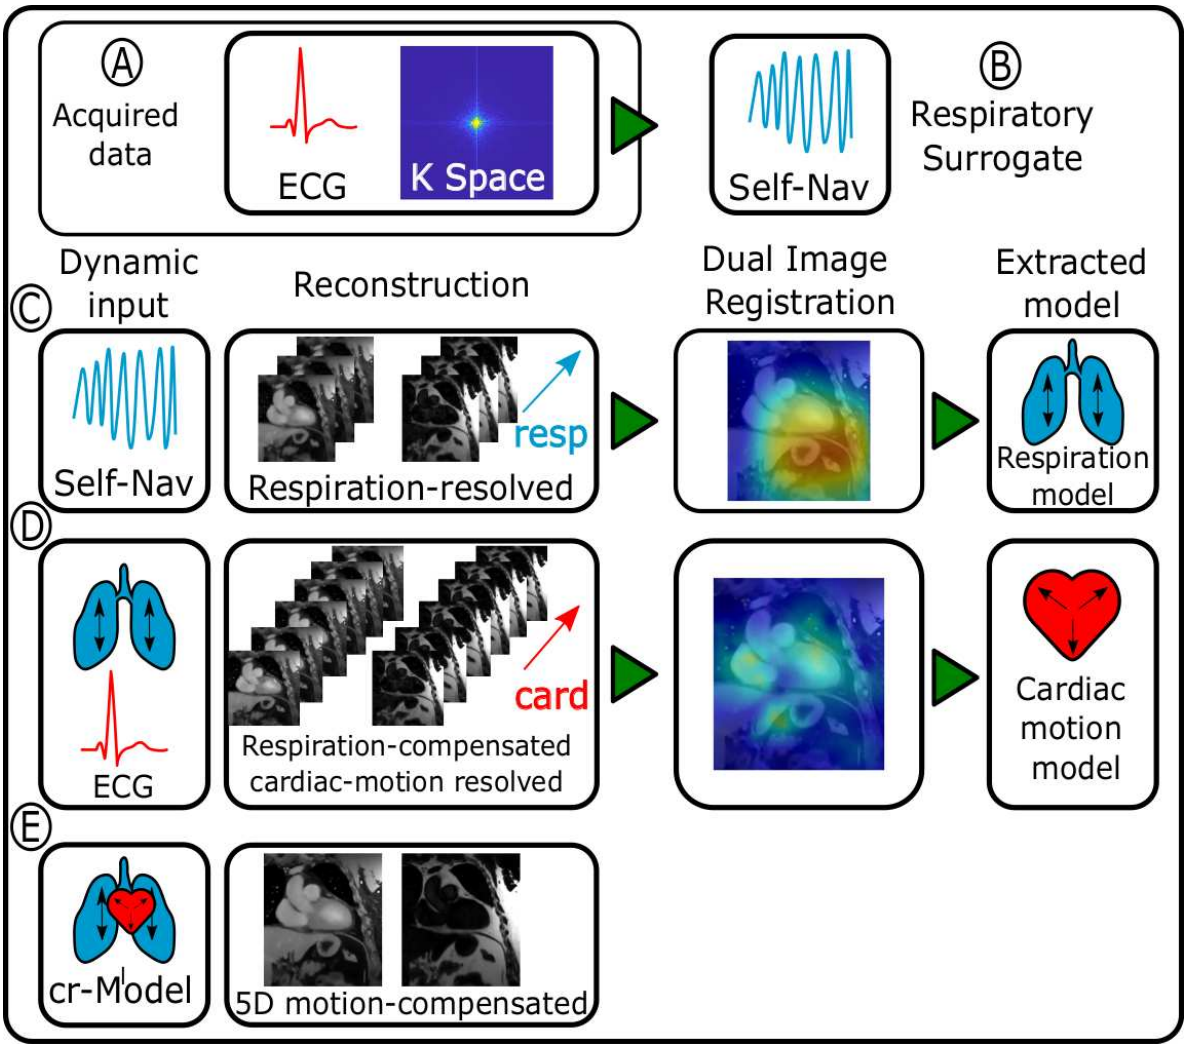

Figure 5.1: Overview of reconstruction workflow. A self-navigator (B) is extracted from the acquired data (A). It is used to bin the data into respiratory states (C) and extract a respiration model from the motion-resolved image series. Subsequently, it is combined with the ECG (D) to reconstruct a respiration-corrected and cardiac motion resolved image series. From this, a cardiac motion model is extracted, combined with the respiration model (E) and used to reconstruct a fully 5D motion-corrected image. This figure was submitted as J8.

5.3 Experiments

Patient Data Acquisition

Nine patients were scanned on a 1.5 T Siemens Avanto using a 32-channel cardiac phased coil array. Patients included in the study were suffering from ischaemic heart disease. Each patient provided informed written consent. Data was acquired with a 3D 3-echo Dixon (T_R = 8.2 ms, T_E = (2.90, 4.48, 6.06 ms), α=15°, T_A =13 min43 s) FLASH sequence. The reconstructed FOV in each direction was 288 mm with a 1.5 mm isotropic resolution, and covered the entire thorax. Phase and slice encoding were performed using the sunflower trajectory presented described in detail in section 4.4. The radial direction was undersampled by F_r = 4 and the angular direction was oversampled by F_ϕ = 10.67 yielding a net 1.7-fold oversampled phase encoding space. Data acquisition was carried out during free-breathing without cardiac triggering. Patients' ECG signals were recorded. Each patient was given a T_1 contrast agent prior to data acquisition.
Fat-water separated images were acquired in 2, 3, and 4 chamber views using a standard ECG-triggered 2D Cartesian acquisition scheme in 4 patients. The proposed 3D acquisition was then reformatted to these orientations and compared.

Quantitative Assessment of Image Sharpness

The local image sharpness Σ of the epicardial fat structures was computed as described in chapter 4 in four locations in the heart. They are displayed in Fig. 5.7 for one of the patients. The obtained four sharpness values for each patient were averaged, yielding one Σ value for each patient. Wilcoxon signed-rank tests were performed between the three different reconstruction groups motion-averaged image reconstruction (AVG), r-MCIR and cr-MCIR. A p-value $p \leq 0.05$ was viewed as statistically significant.

Systolic Gating to Correct Residual Cardiac Motion

Cardiac motion is particularly challenging to correct for. The temporal resolution of each cardiac motion bin is potentially not sufficient to resolve the fast contraction of the heart during systole. Additionally, the estimated motion model for the cardiac motion can be inaccurate for the highly non-linear cardiac motion with large amplitudes and the optimal regularisation parameters are potentially patient-specific. Hence, residual motion artefacts might still be present post MCIR. To investigate if gating of the systole in addtion to MoCo can further improve image quality, we compared r-MCIR and cr-MCIR with and without systolic gating in a patient with large cardiac motion amplitude. r-MCIR and cr-MCIR images were reconstructed using cardiac data of only 6 of the 12 cardiac bins that contained the least motion. These images were compared to the reconstructions using the entire data.

Retrospective Data Undersampling

To assess the potential impact on image quality by reducing the acquisition time the same analysis was performed after retrospectively undersampling, i.e. taking only the first N_ϕ^{US} = 1024 radial spokes. This corresponded to a reduction of the scan time by half from T_A =

13 min43 s to 6 min51 s. The cardiac registration parameters were altered to $(\lambda, \Delta B, \rho) = (0.5, 4, 2 \cdot 10^{-3})$, i.e. a larger ΔB was chosen to address increased undersampling artefacts. All other parameters were kept constant.

5.4 Results

Qualitative Evaluation of Fat-Water Reconstructions

Motion-resolved reconstructions for two patients are displayed in Fig. 5.2. For both patients water and fat images in end-exhale, end-inhale, diastole and systole are displayed. One patient is shown in sagittal and one in coronal view to highlight the 3D isotropic resolution of the acquired data.
For both patients, changes in the anatomy due to respiratory motion are clearly visible in water and fat reconstructions. In the cardiac-resolved images, the respiration has been corrected as indicated by the position of the liver and diaphragm. For patient 21 the contraction of the heart from diastolic to systolic phase is visible both in water and fat images.

Motion-corrected images for two patients are shown in Fig. 5.3. Both fat and water images are shown for AVG, r-MCIR, and cr-MCIR from left to right. Regions where a coronary vessel highlighted and depicted in a seperate zoomed-in square. Cyan arrows indicate locations where motion-blurring in the epicardial fat structures is visibly reduced by the application of the proposed motion-corrected image reconstruction. Patient 18 shows a reduction of motion blurring after the application of the respiratory model, especially in the abdominal region. Motion blurring due to the heartbeat can be further reduced upon an additional application of the cardiac motion model in both epicardial fat structures as well as for the coronary vessel. Patient 24 shows very strong blurring in the motion-averaged case, of which the abdominal fat and the structure at the apex could be improved with respiratory MoCo. Cardiac MoCo could further improve the basal fat structure. A similar improvement can be also seen in the complementary water image where the coronary arteries become clearly visible using cr-MCIR. In addition to a reduction in blurring also respiratory motion artefacts in the ventricle are reduced with MoCo.

Fat-water separated images reformatted to showing the left coronary artery are displayed in Fig. 5.4. Correction of both cardiac and respiratory motion using cr-MCIR leads to an improved depiction of the coronary arteries and the surrounding fat structures.

A comparison between 2D clinical acquisitions in 2, 3 and 4 chamber view and 3D fat cr-MCIR reconstructions reformatted to the same views are displayed for two patients in Fig. 5.5. The depiction of fat structures acquired is comparable between both scans. Nevertheless, it is important to note that the 2D acquisition was acquired during a breathhold using cardiac triggering and hence might show a different motion state than cr-MCIR.

A patient with fat-infiltration in the septum is displayed in Fig. 5.6. The reconstructed fat images are displayed in axial and sagittal view for both AVG and cr-MCIR. It can be observed that MoCo led to an improved depiction of the fat in the axial view. The sagittal slice shows the fine structure of the fat infiltration strongly blurred for the AVG reconstruction. In addition, there are many motion-artefacts which make it difficult to clearly identify these small structures. The proposed cr-MCIR reduced motion artefacts and blurring and strongly

improved the visibility of the fat infiltration. This can also be seen in a fat-water overlay of the sagittal slice.

Quantitative Evaluation of Image Sharpness

The measured Σ is displayed in Fig. 5.7. On the left, The locations ℓ where Σ is evaluated are shown for patient 22 on cr-MCIR fat images. On the right, the values of $\langle \Sigma_\ell \rangle_\ell$ are displayed for AVG, r-MCIR and cr-MCIR. The result of the statistical tests between AVG and r-MCIR of as well as cr-MCIR led to p-values $p_{\text{rMCIR}} = 0.008$ and $p_{\text{crMCIR}} = 0.008$, between r-MCIR and cr-MCIR the p-value was $p_{\text{card}} = 0.01$. This suggests that correcting for both the respiratory and the cardiac motion yields a significant improvement in the fat structure Σ of the overall heart.

The edge sharpness of the fat structures evaluated in the patient data is also summarised in Tab. 5.1. The effect of r-MCIR leads to an increase in fat structure Σ for every patient, with improvements ranging between 12 and 66 % with 33 ± 17 % on average. The effect of additional cardiac MoCo is highly patient-specific. For some patients it can further improve Σ by 20 % for other patients it does not lead to any further improvement. On average cardiac MoCo increase Σ by 8 ± 8 %.

MocoType Patientname	Σ AVG	Σ rMCIR	Σ crMCIR	ΔResp (%)	ΔCard (%)	ΔCaRe (%)
patient18	0.26	0.34	0.37	30	10	43
patient19	0.27	0.33	0.35	23	6	30
patient21	0.35	0.48	0.52	37	7	46
patient22	0.27	0.32	0.33	20	3	23
patient23	0.34	0.38	0.39	12	2	15
patient24	0.28	0.35	0.42	27	19	52
patient25	0.21	0.28	0.28	38	−2	35
patient26	0.26	0.39	0.47	50	20	80
patient29	0.23	0.38	0.39	66	4	73
Mean	0.27(5)	0.36(6)	0.39(7)	34(17) **	8(8) *	44(22) **

* $p < 0.05$, ** $p < 0.01$

Table 5.1: Quantitative analysis of Σ evaluated for each patient. ΔResp describes the increase in Σ between r-MCIR and AVG. ΔCard describes the increase in Σ between cr-MCIR and r-MCIR. ΔCaRe describes the increase in Σ between cr-MCIR and AVG. Statistical significance was evaluated between for all three differences. Standard deviation is given in parenthesis in significant digits. This table was submitted as J8

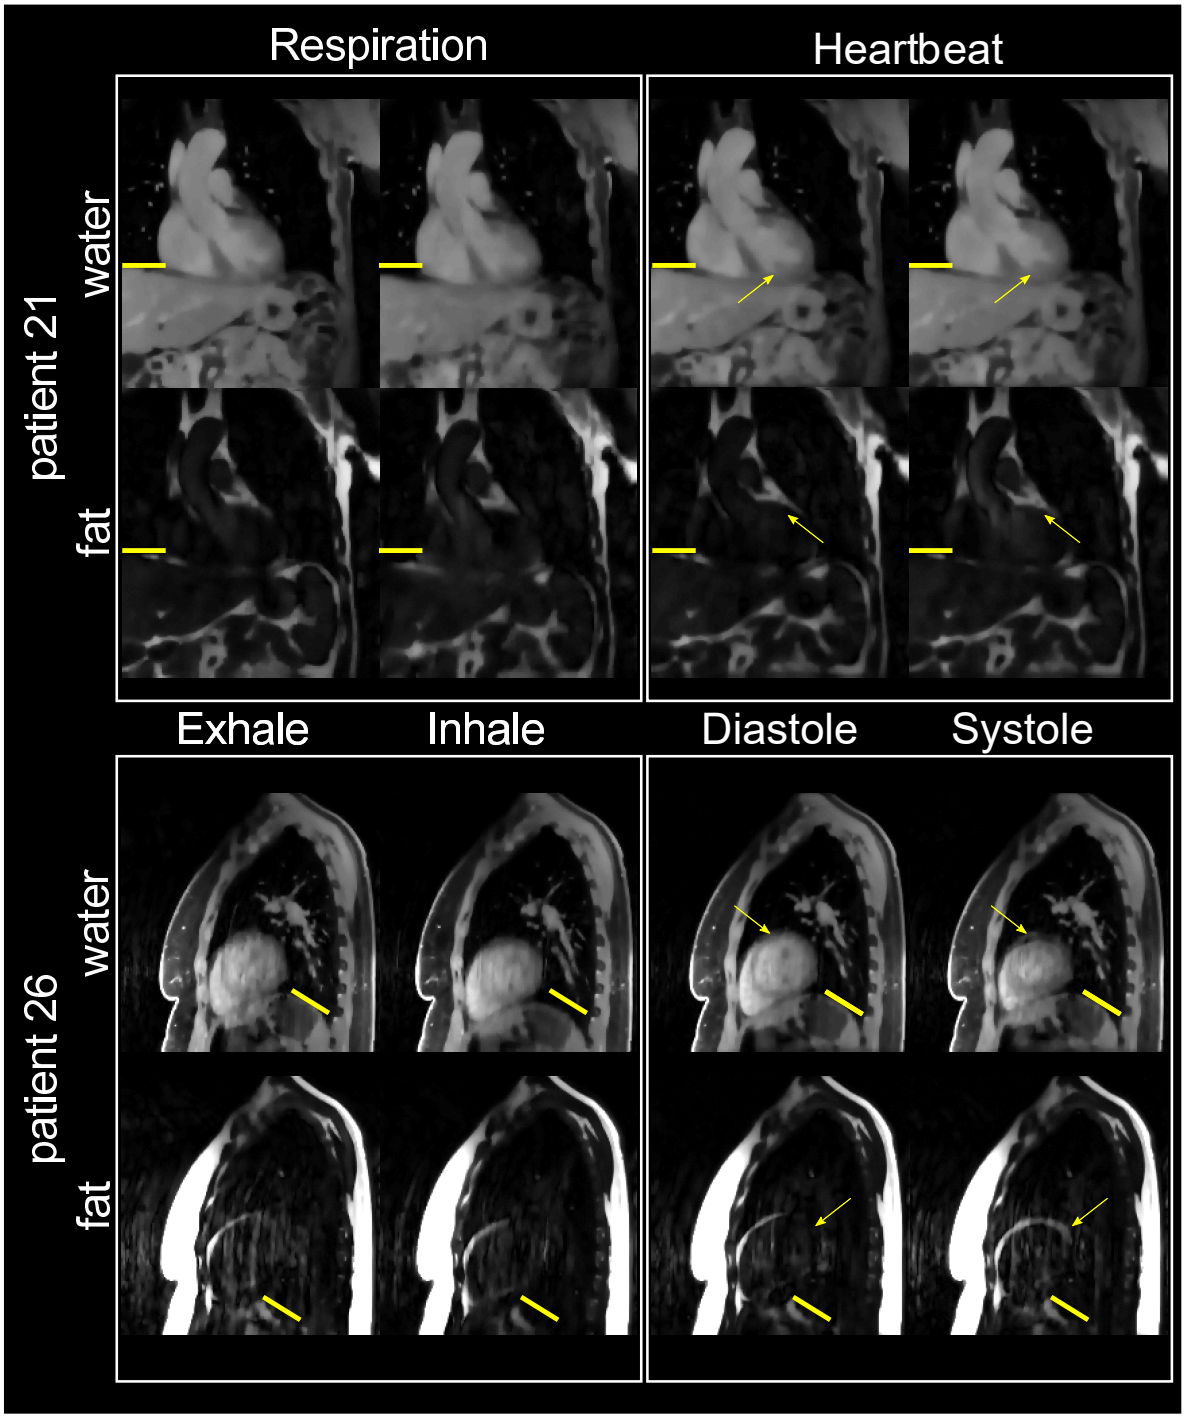

Figure 5.2: Motion-resolved patient data reconstructions as fat and water images. Two example patients 21 (top) and 26 (bottom) are depicted. The intensity of undersampling artefacts is larger in the reconstructions for patient 26, while cardiac and respiratory resolved images show similar intensities of undersampling artefacts. Fat and water are displayed for both respiratory (right) and cardiac (right) motion in the respective states of ex- and inhale, and diastole and systole. Yellow bars indicate the exhale position to which the cardiac-resolved images are already corrected to. Yellow arrows indicate cardiac motion which mainly leads to a contraction of the heart and a thickening of the myocardium. This figure was submitted as J8

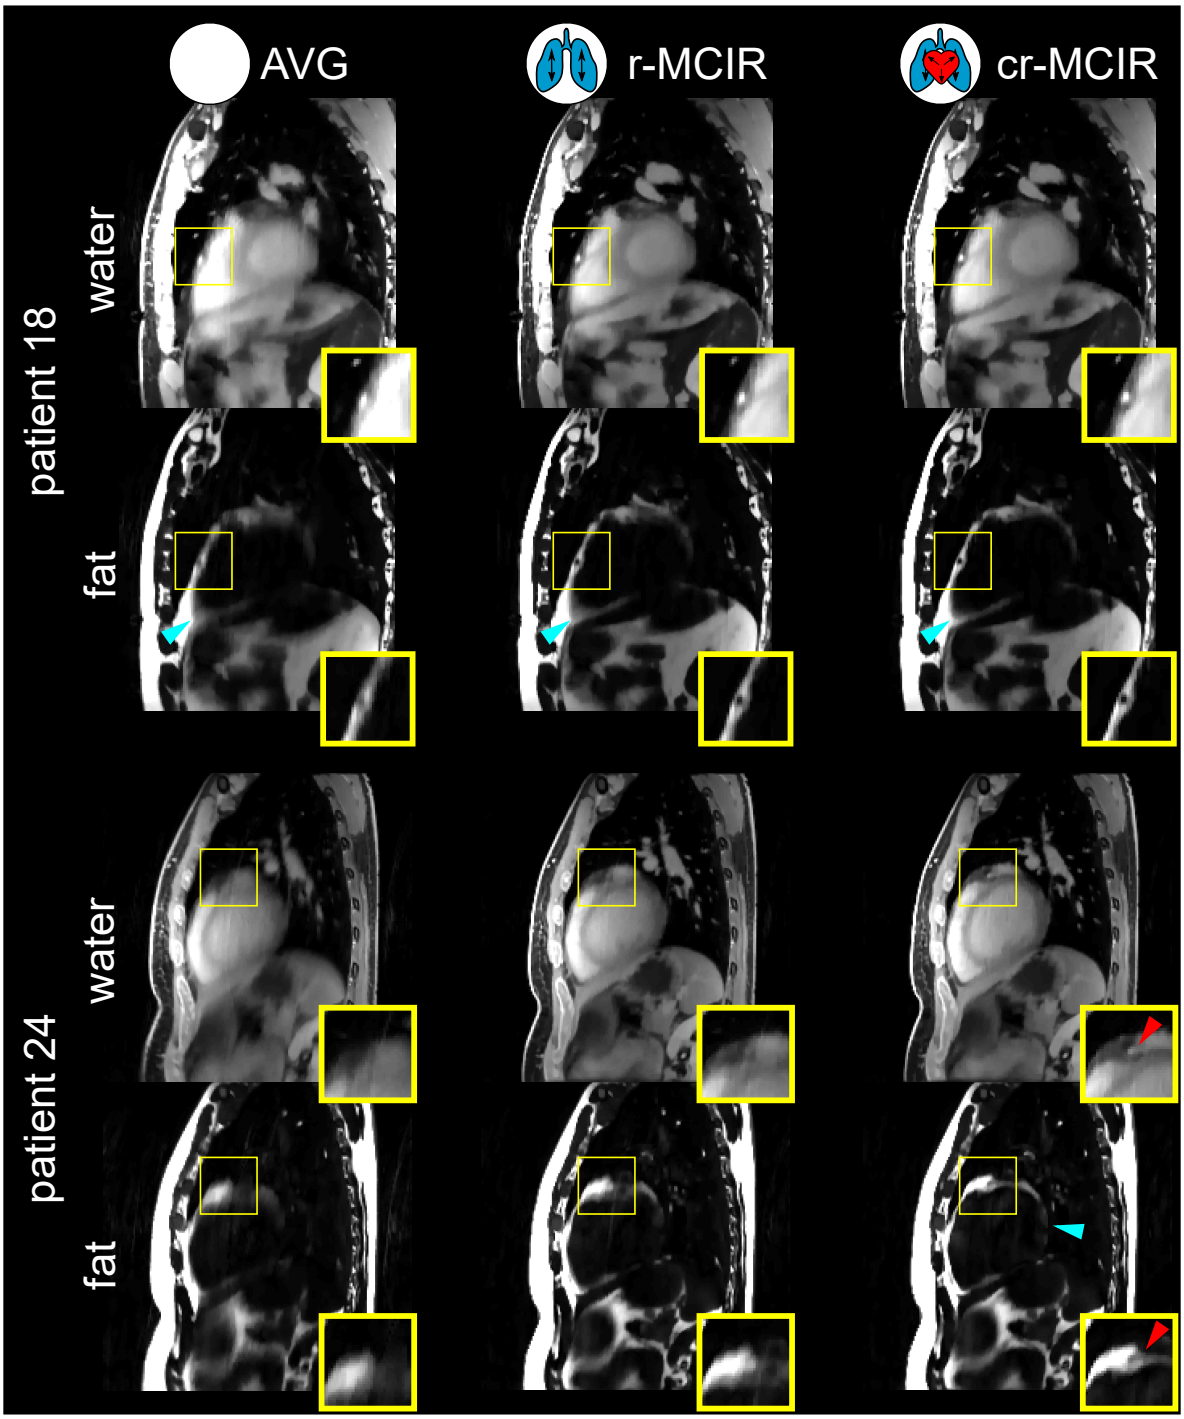

Figure 5.3: Motion-corrected patient data. Motion-corrected fat-water-separated image reconstruction of two patients 18 (top) and 24 (bottom). Cyan arrows highlight areas where respiratory (rMCIR) and subsequently cardio-respiratory (cr-MCIR) MoCo improve visualisation of epicardial fat. Yellow squares are zoomed in on to show delineation of coronary vessels with more motion correction in both water and compelmentary fat image. This figure was submitted as J8

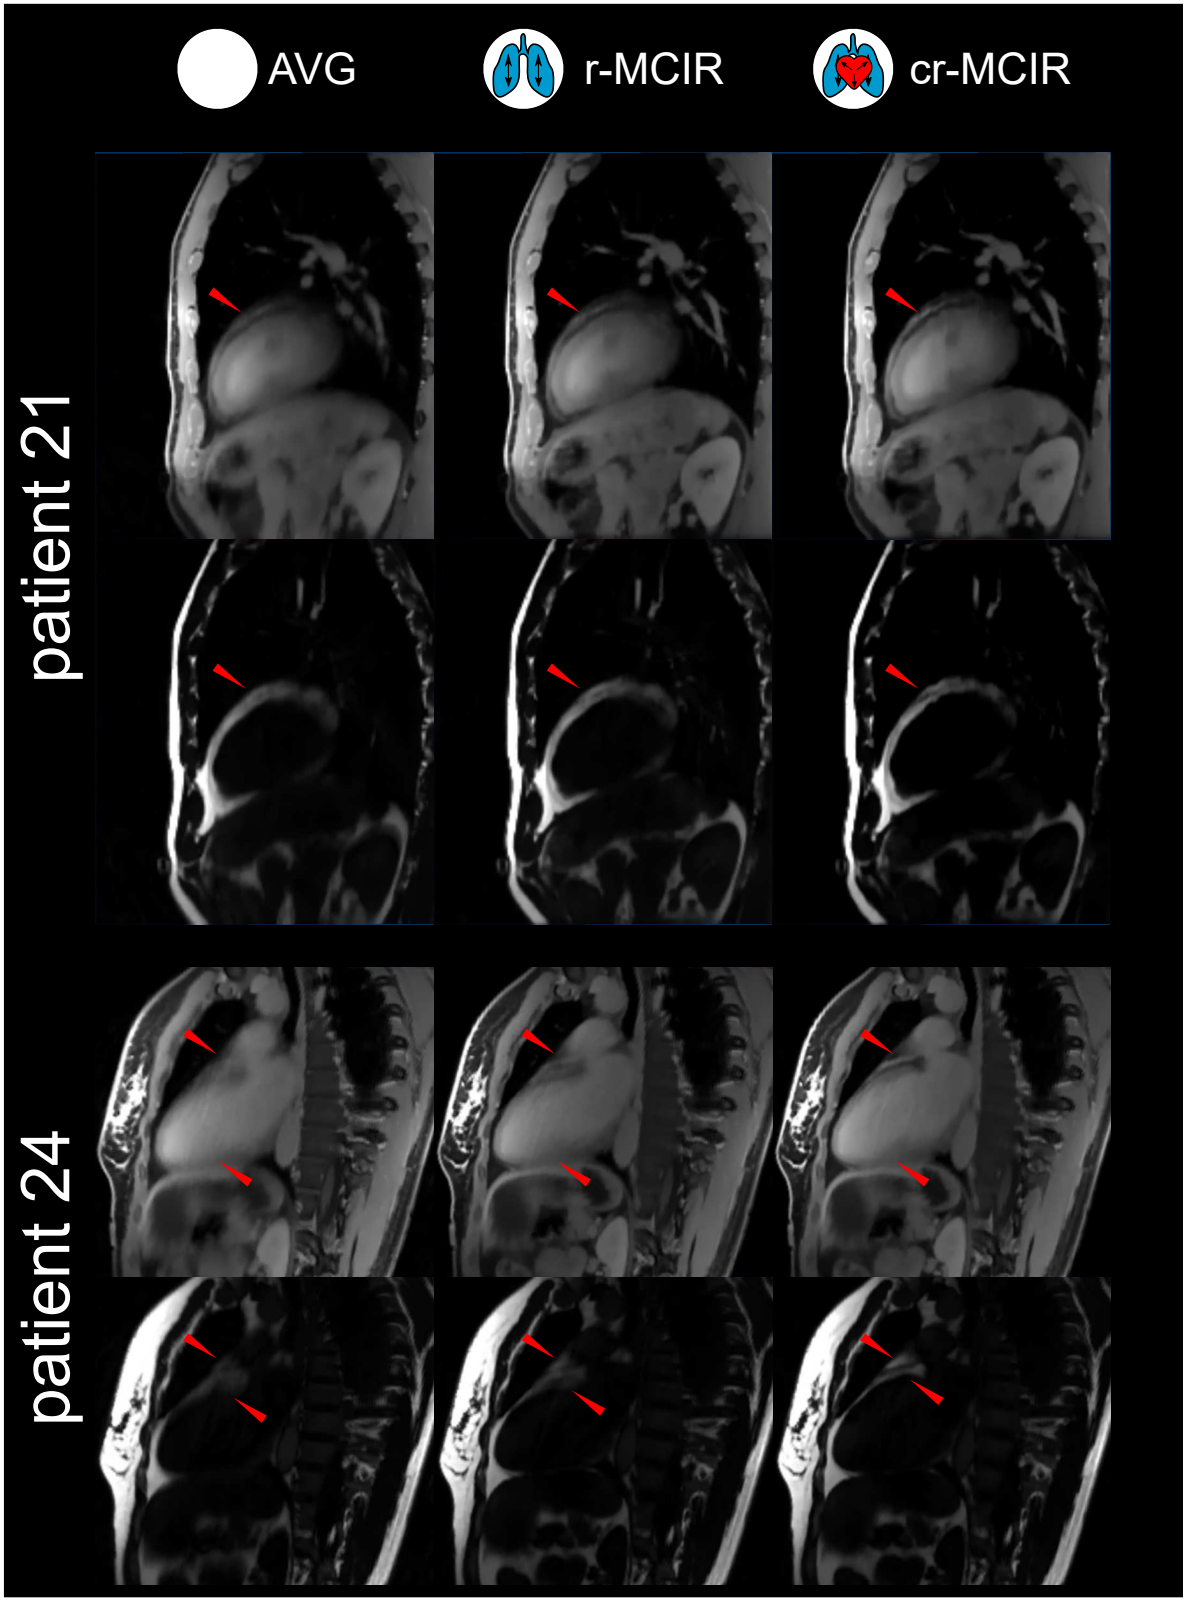

Figure 5.4: Effect of MCIR on coronaries. Reformatted motion-corrected fat-water-separated image reconstruction of two patients 21 (top) and 24 (bottom). Left to right: AVG, r-CMIR and cr-MCIR. Red arrows highlight improvements in the depiction of the vessels and surrounding fat.

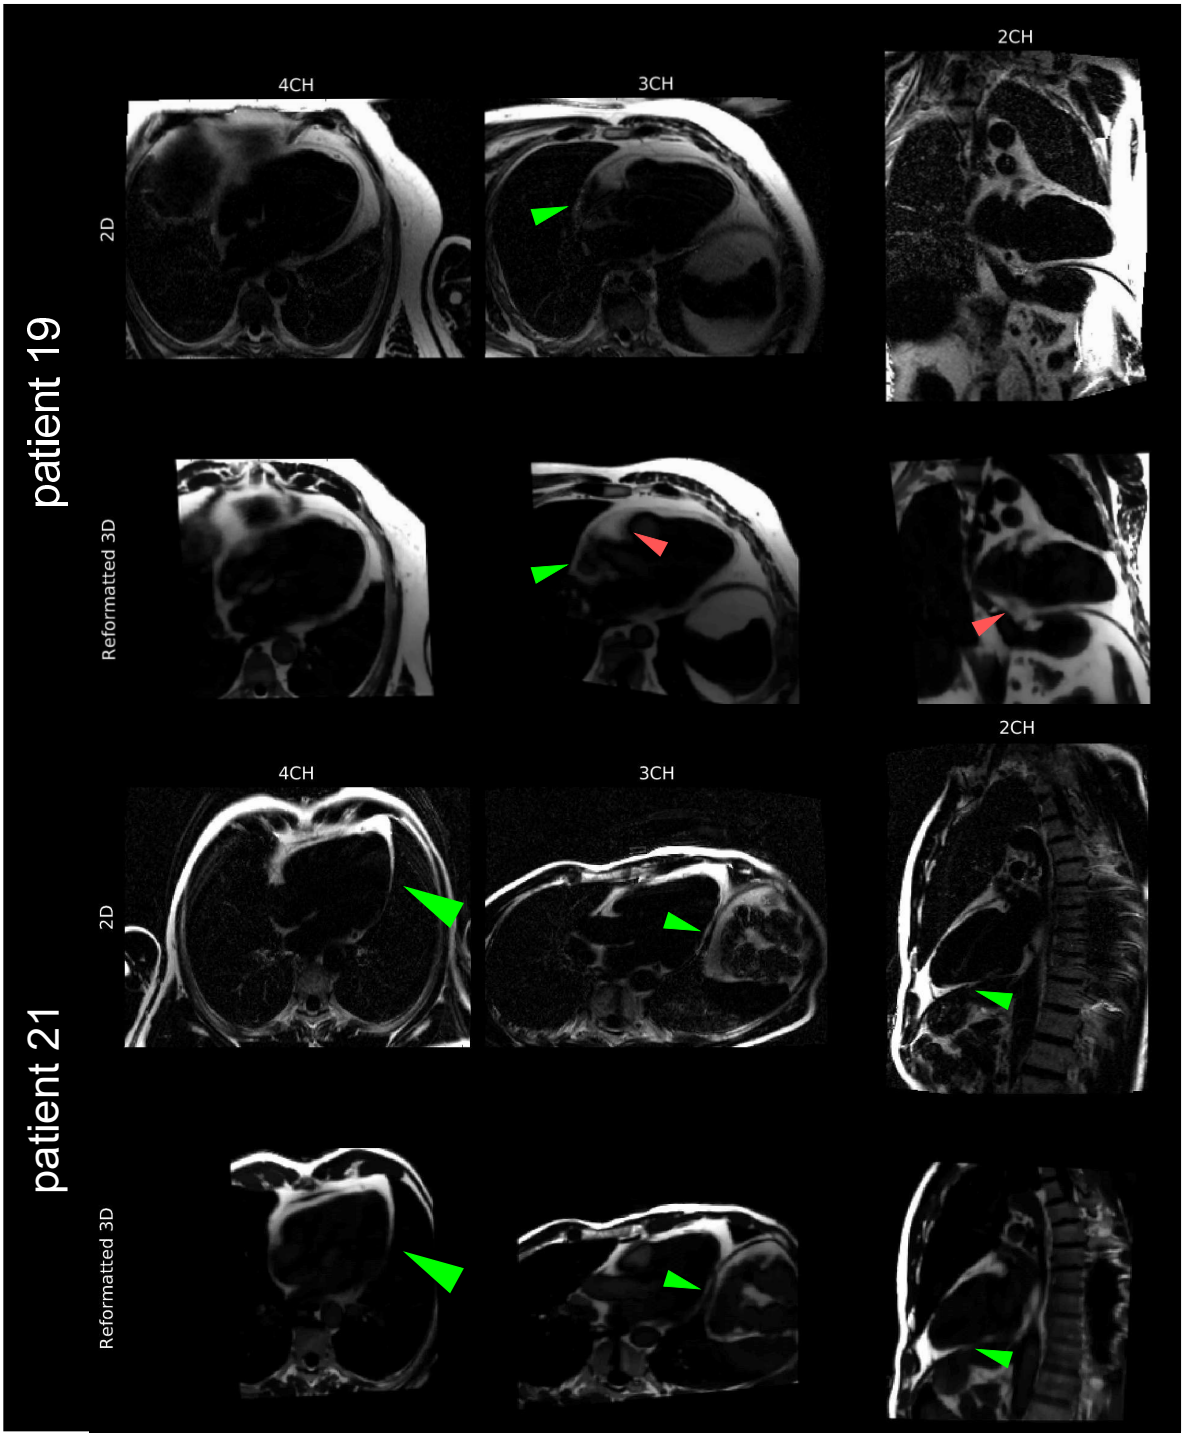

Figure 5.5: Comparison with clinical standard reference. Images of 3D fat cr-MCIR reformatted to clinical routine 4, 3 and 2 chamber view (left to right) for two patients. The top row for each patient shows the clinical standard cardiac triggered reference exam obtained during breath-hold. The bottom row shows the reformatted images generated from the free-breathing motion-corrected cr-MCIR 3D data. Green arrows point out fine fat structures that the 3D data can resolve. Red arrows show residual motion blurring due to imperfect cardiac MoCo. This figure was submitted as J8

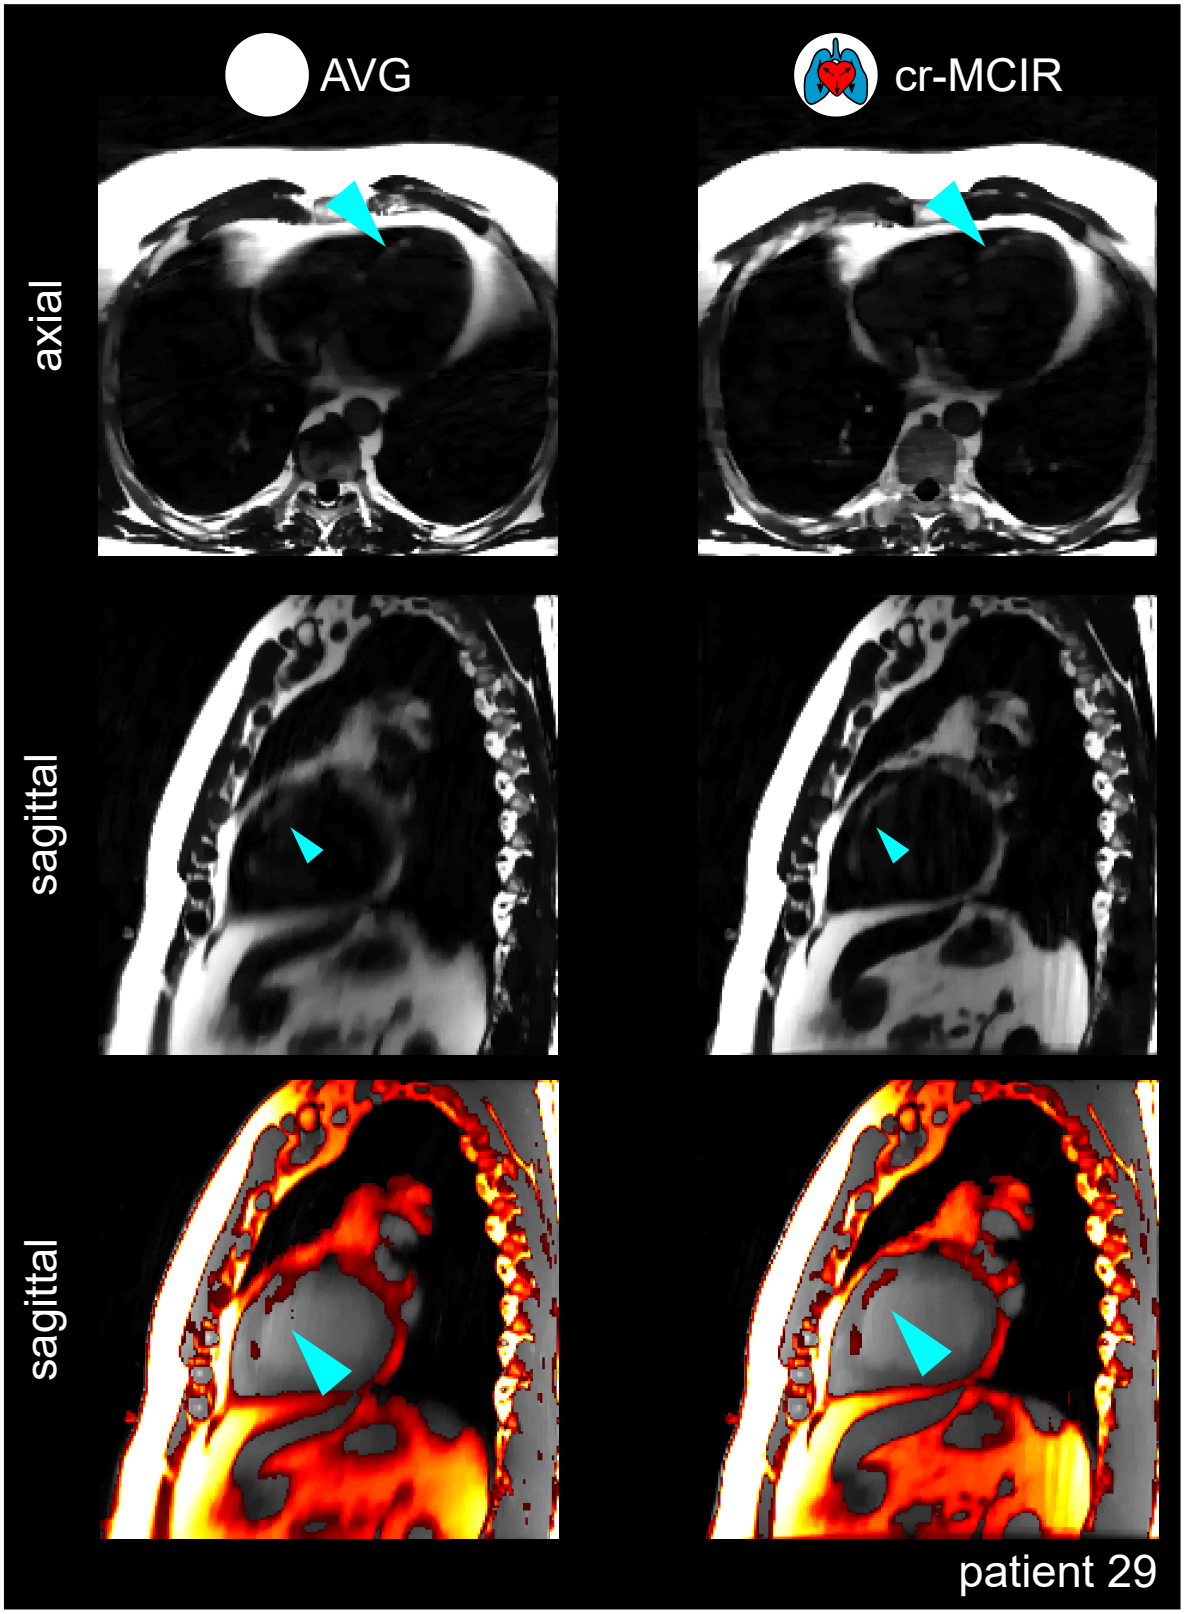

Figure 5.6: Patient with myocardial fat infiltration. Comparison of AVG and cr-MCIR for a patient with myocardial fat-infiltrations. Axial (top) and sagittal (centre) slices are depicted. The fat-infiltration is indicated by a blue arrow. The proposed cr-MCIR fat-water separated image reconstruction reduces motion artefacts and blurring, improving the visibility of the fat infiltration. This is also visible in a fat-water overlay (bottom). This figure was submitted as J8

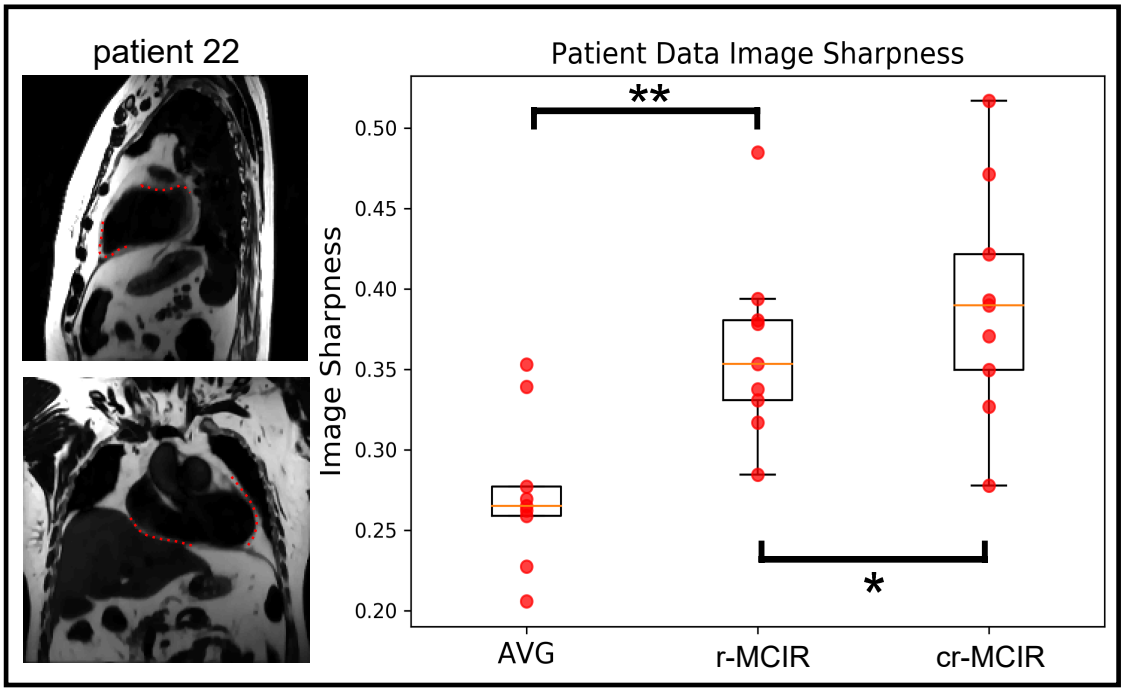

Figure 5.7: Σ computed for patient data reconstructions. Left: cr-MCIR of a patient with the spline curves ℓ along which the metric was evaluated marked in dashed red. Right: distribution and boxplot of the fat structure Σ metric for AVG (left), r-MCIR (centre), and cr-MCIR (right). Orange bars indicate the median, the box includes second and third quartile, and the whiskers span 1.5 times the box height from the median. The r-MCIR and cr-MCIR have a higher median than the AVG data set. This figure was submitted as J8

Systolic Gating to Correct Residual Cardiac Motion

Fig. 5.8 compares AVG, r-MCIR and cr-MCIR with and without cardiac gating. Cardiac gating strongly improves the image quality for r-MCIR images leading to a better depiction of the coronary artery in both the water and fat image. However, for cr-MCIR the improvements in image quality due to gating are very small. Cardiac gating does lead to additional undersampling artefacts (marked by the cyan arrow).

Retrospective Data Undersampling

Images reconstructed using all acquired data and from data undersampling by 1/2 are displayed in Fig. 5.9. As expected, SNR is reduced and undersampling artefacts become visible in the fat image due to the retrospective undersampling, but the overall image quality is comparable. Also, the motion-corrected framework can achieve the same improvements in image quality when applied to retrospectively undersampled data.

In Fig. 5.10 the effect of retrospective undersampling for the patient with septal fat infiltration is depicted. Image quality is also comparable and the fat infiltration can be visualised equally well using only half the k-space data.

The quantitative evaluation of Σ in the same regions as for the full dataset is presented in Tab. 5.2. The improvements of Σ are similar to using the full k-space data (Tab. 5.1). Fig. 5.11 shows Σ for the different MoCo schemes for a data fraction of 1/2. Although, cr-MCIR shows an improvement over r-MCIR in some patients, the median over all patients is only slightly improved.

MocoType Patientname	Σ AVG	Σ rMCIR	Σ crMCIR	ΔResp (%)	ΔCard (%)	ΔCaRe (%)
patient18	0.25	0.31	0.32	23	4	28
patient19	0.27	0.34	0.34	24	1	26
patient21	0.38	0.45	0.50	21	10	34
patient22	0.26	0.33	0.32	25	−2	22
patient23	0.34	0.38	0.40	13	4	18
patient24	0.27	0.33	0.34	20	4	25
patient25	0.17	0.25	0.25	45	2	48
patient26	0.24	0.37	0.45	56	22	90
patient29	0.23	0.38	0.40	66	5	74
Mean	0.27(6)	0.35(6)	0.37(8)	33(18) **	6(7) *	41(25) **

* $p < 0.05$, ** $p < 0.01$

Table 5.2: Quantitative analysis of Σ evaluated for each patient applied to k-space data retrospective undersampled by a factor of 1/2. ΔResp describes the increase in Σ between r-MCIR and AVG. ΔCard describes the increase in Σ between cr-MCIR and r-MCIR. ΔCaRe describes the increase in Σ between cr-MCIR and AVG. Statistical significance was evaluated between for all three differences. Standard deviation is given in parenthesis in significant digits.

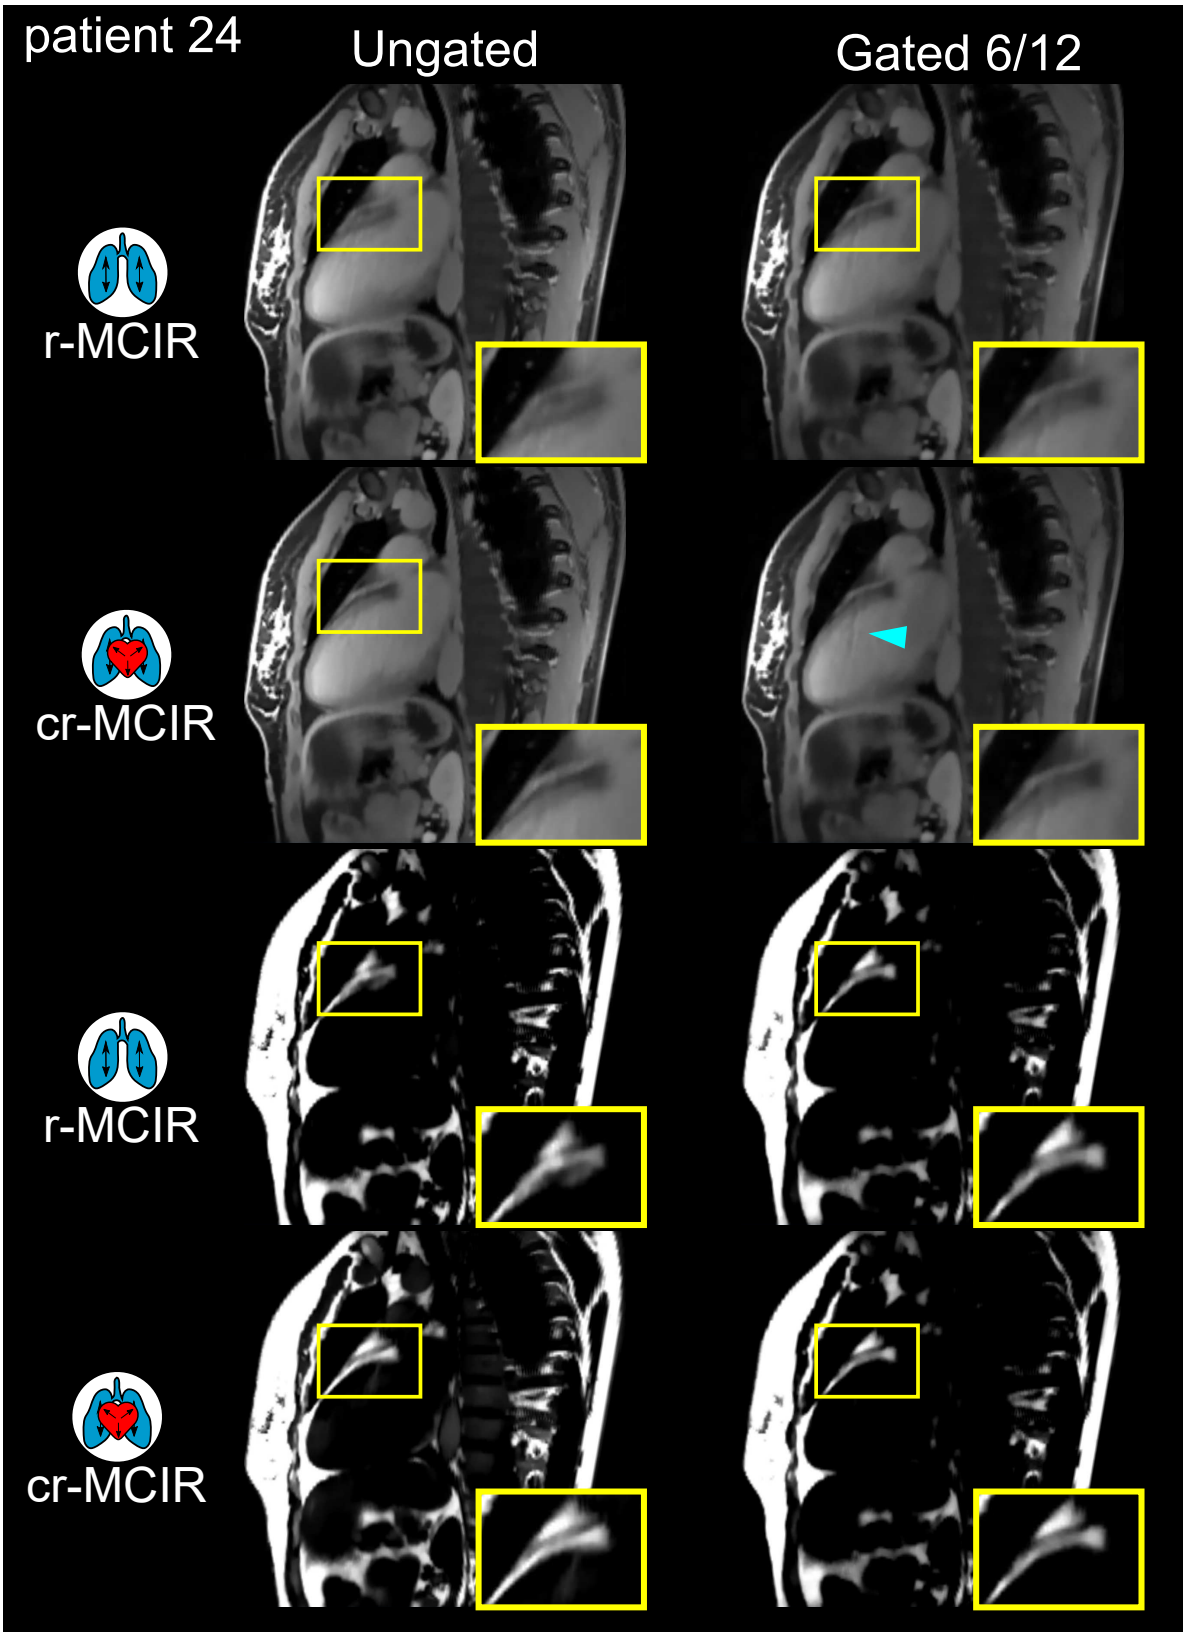

Figure 5.8: Effect of diastolic gating for MCIR. Top: water images, bottom fat images. The 2x2 matrices contain both types of motion compensation for the ungated and gated reconstruction. I.e. from left to right the effect of gating is visible, from top to bottom the effect of including cardiac motion correction can be seen. The coronary artery is zoomed in on in the yellow box. The cyan arrow points to a streak artefact.

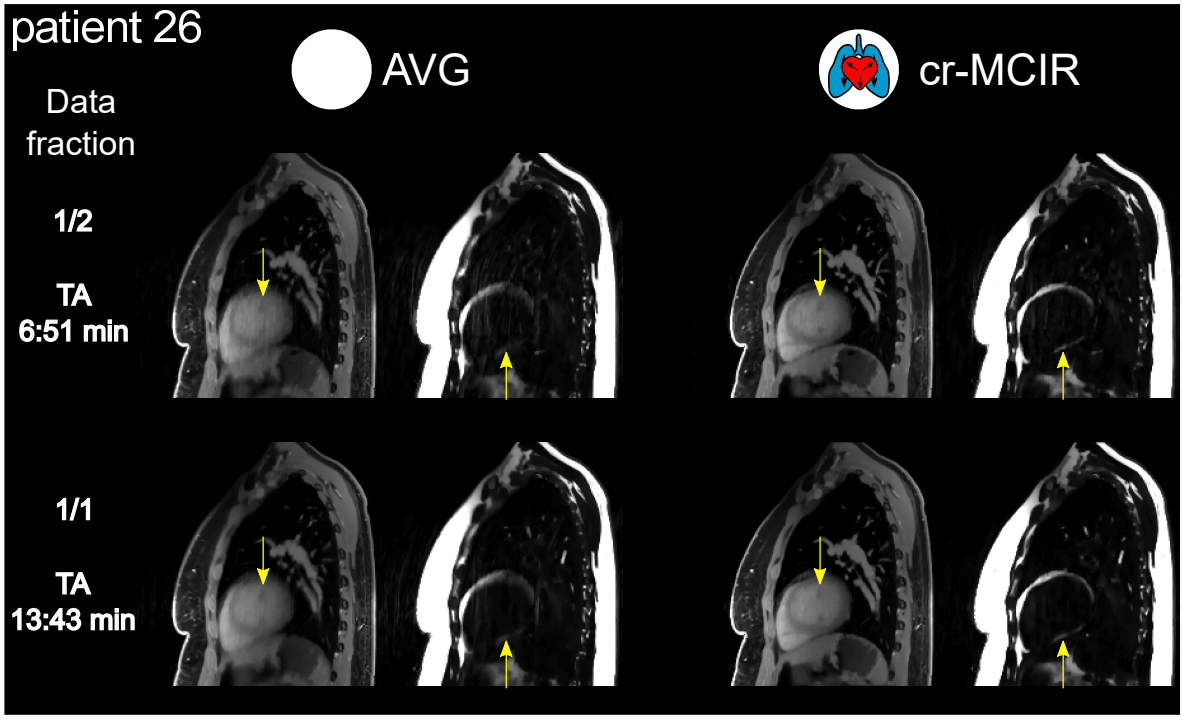

Figure 5.9: Retrospective undersampling applied to patient 26. Left: the reconstructed fraction of the total data and potential acquisition time T_A . Top: comparison AVG and cr-MCIR reconstructions for retrospectively undersampled k-space, using only half the data. Bottom: full dataset as a reference. Despite using only 1/2 of the k-space data, image quality is comparable between the two reconstructions and cr-MCIR leads to strong improvements in image quality in both cases (yellow arrows).

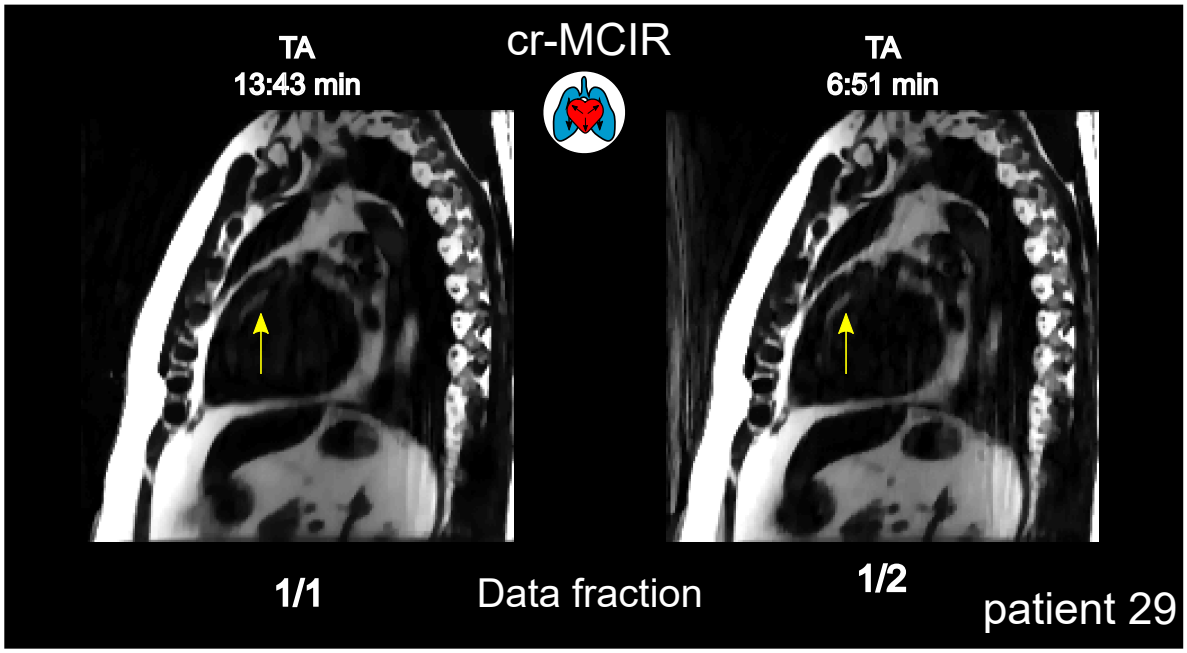

Figure 5.10: Depiction of myocardial fat infiltration. The cr-MCIR reconstructions are shown for the full (left) and retrospectively undersampled (right) data. The arrow points to the fat infiltration in the septum, which is clearly visible in both reconstructions.

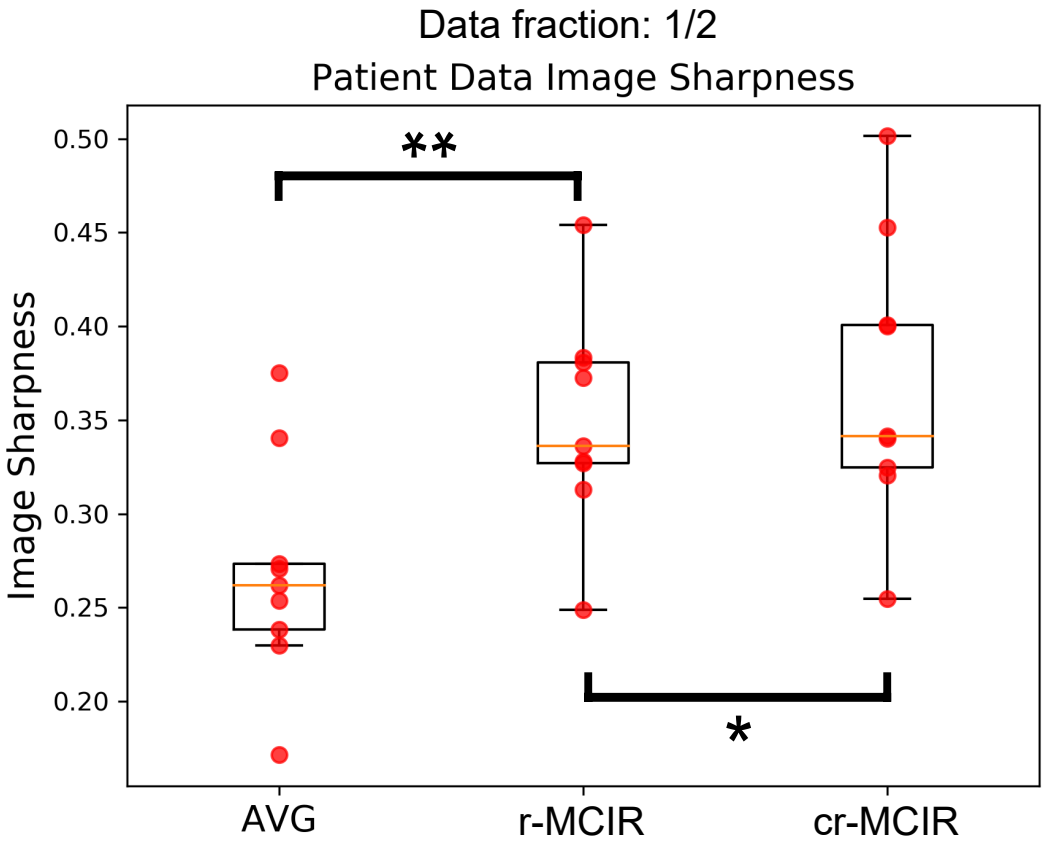

Figure 5.11: Σ computed for retrospectively undersampled patient data. The distribution and boxplot of the fat structure Σ metric for AVG (left), r-MCIR (centre), and cr-MCIR (right). Orange bars indicate the median, the box includes second and third quartile, and the whiskers span 1.5 times the box height from the median. The r-MCIR and cr-MCIR have a higher median than the AVG data set.

5.5 Discussion

In this chapter, we presented a model-based fat water reconstruction with cardiac and respiratory MoCo.

With the presented reconstruction framework compressed sensing regularisation could be successfully applied for motion-corrected image reconstruction of fat and water images. In Fig. 5.2 two example reconstructions were presented that showed the motion-resolved reconstruction of $N_{\text{resp}} = 6$ and $N_{\text{card}} = 12$ motion states. The *TV* and *TVT* regularisation were able to suppress artefacts while at the same time conserving motion amplitudes.

The motion resolved images could then be used to estimate and subsequently correct for respiratory and cardiac motion (Fig. 5.3). Image quality of both water and fat images was improved by reducing motion artefacts and blurring. This could be seen in both an improved delineation of of epicardial fat structures, as well an improved depiction of coronary vessels and their surrounding fat structures.

Quantitatively this was analysed by computing a local edge sharpness metric Σ for epicardial fat structures. The data showed a large improvement of $33 \pm 17\,\%$ of selected epicardial fat structures when including a respiratory motion model into the reconstruction process. Cardiac motion modelling further increased Σ by an additional $8 \pm 8\,\%$. Statistical tests showed that respiratory ($p < 0.01$) and cardiac ($p < 0.05$) MoCo led to a statistically significant increase in sharpness Σ. The increase in Σ is highly patient dependent, especially for cardiac motion (Fig. 5.7 and Tab. 5.1). This effect is expected as cardiac motion is much more patient-specific than respiratory motion. The impact of cardiac motion modelling depends on how long the systolic phase is relative to the whole cardiac cycle and how strongly the heart contracts during systole.

Cardiac motion is also more challenging to estimate than respiratory motion. While respiration is similar for most patients and occurs mainly in the head-foot direction the heartbeat leads to much more complex motion patterns which are more difficult to distinguish from any residual image artefacts. Finding a single set of patient-independent parameters for the registrations yielding high-quality registrations is very challenging. Adapting the registration to be more patient-specific could be supported by machine learning [201, 202]. Also, further temporal regularisation could be employed during the registration to improve the accuracy of motion estimation[80, 203, 204].

Fig. 5.6 shows the improvement in image quality and visualisation of fat structures infiltrating the myocardial tissue. The displayed patient's fat image was strongly deteriorated by motion blurring and other motion artefacts and showed the largest improvement in Σ for all patients using respiratory MoCo. In axial view both in AVG and cr-MCIR the structure was visible, however, only after MoCo, the structure can be identified as a coherent infiltration permeating a large part of the cardiac tissue.

Furthermore, we have demonstrated that using only 1/2 of the acquired k-space still leads to a high image quality and accurate visualisation of fat structures with a small decrease in SNR. We have also compared if cardiac gating can improve image quality by gating away 1/2 of the data containing the largest cardiac motion amplitudes. Our experiments showed that cardiac gating only led to small improvements in image quality but causes small undersampling

artefacts. A reduction of scan time by a factor of 2 can potentially justify a decrease in SNR to make an application more feasible in a clinical setting. However, the benefit of reduced motion artefacts as seen in the experiment due to gating does not outweigh the appearing undersampling artefacts.

Towards shorter scan time

In order to investigate the limits of the proposed motion-corrected model-based fat-water separation we retrospecitvely undersampled the k-space by a factor of 4 in one patient. We compared the glscrMCIR fat-water images using the fieldmap and motion fields from the reconstruction of the full k-space. In a second step, we investigated, how the image quality of the motion-resolved reconstructions could be improved to allow for accurate motion estimation despite being obtained from only 1/4 of the data.
In Fig. 5.12 cr-MCIR reconstructions from retrospective undersampling by a factor of 1 (≈14 min scan time), 2 (≈7 min scan time), and 4 (≈3.5 min scan time) are compared. Although the image quality for the 3.5 min scan time is still good, some fine structures in the fat image are not visible any more and fat and water cannot be separated succesfully anymore in challening structures, such as the moving blood in the aorta.

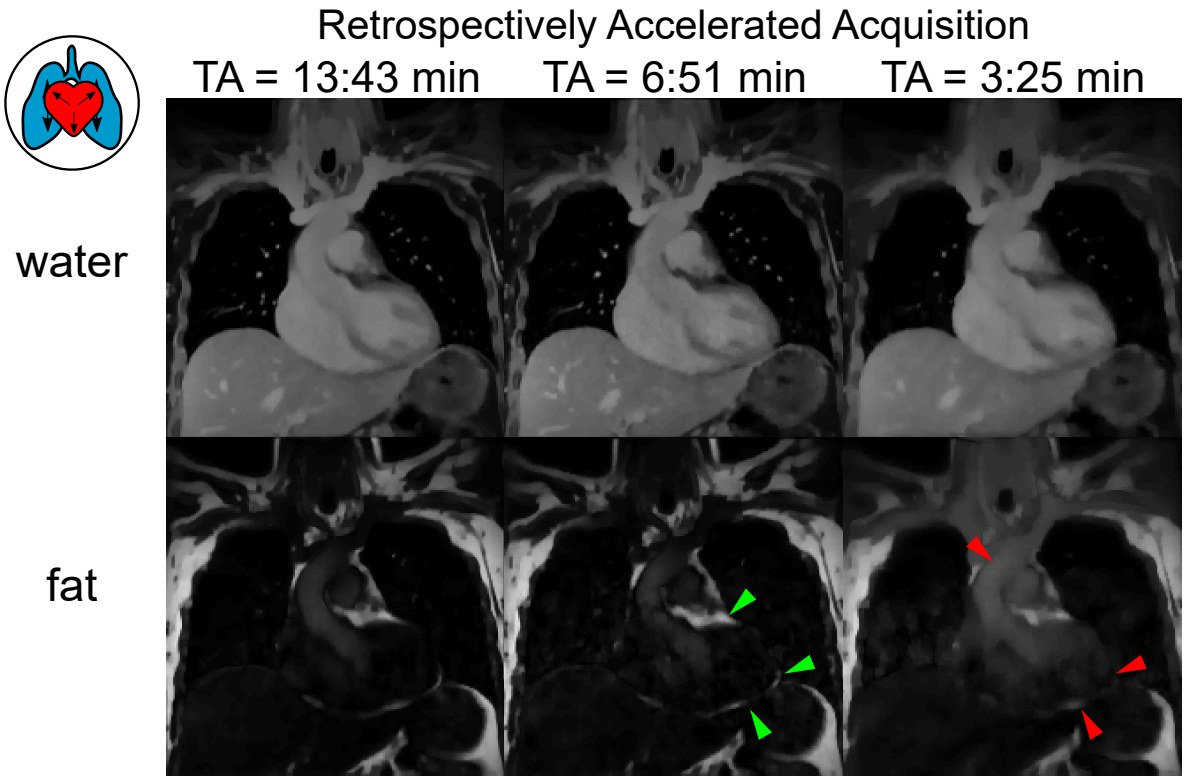

Figure 5.12: Effect of progressive retrospective undersampling on fat and water cr-MCIR. Left to right: full dataset, half the dataset to one-quarter of the dataset. By using only half the data the image quality is still good and fat structures can be seen clearly (green arrows). For using 1/4, smaller fat structures are not well depicted anymore and there is signal from the water mode in the fat image (red arrows).

The estimation of high-quality motion models for acquisition times faster than 5 minutes could potentially benefit from developments in artificial intelligence. Stronger undersampling

artefacts in cardiac-resolved reconstructions could be minimised by the application of convolutional neural networks. This is exemplified in Fig. 5.13. A network making use of spatio-temporal correlations that had previously been trained on 2D radial cardiac cine data [205] was applied to axial slices for the fat-water separated images individually. With this approach, the artefact level could be strongly reduced, especially for the fat images, while conserving the motion amplitude. These are promising results for fat-water imaging in less than 5 minutes, but further studies are required.

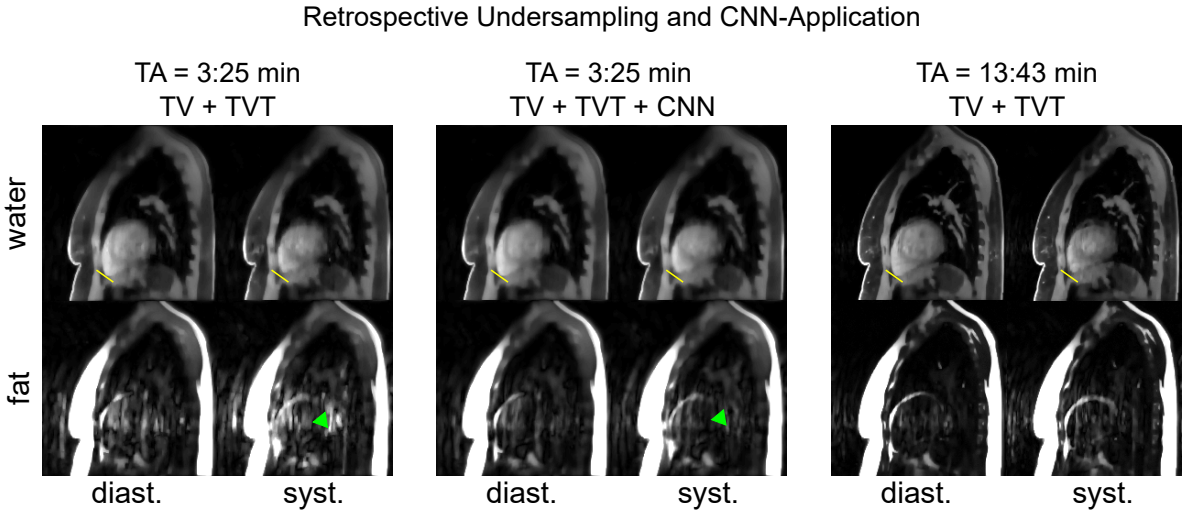

Figure 5.13: Potential use of a convolutional neural network (CNN) to reduce artefacts in cardiac-resolved image reconstructions. The motion-resolved images are displayed in diastole and systole (yellow lines indicate motion amplitude) and for water and fat. Left: cardiac resolved recon from $\frac{1}{4}$ of the data (retrospectively undersampled). Centre: CNN applied post reconstruction to the images from the left. Right: cardiac-resolved full dataset as a reference.

Although this analysis gives an idea of the image quality for shorter scan times, it is limited because the sunflower trajectory is not well suited for retrospective undersampling. As the phase encoding order is aimed to generate the sunflower pattern for a predefined undersampling factor, larger gaps in the acquired k-space data appear if the undersampling is increased retrospectively rather than prospectively (Fig. 5.14). Therefore, the image quality for images reconstructed from data acquired for 7 min and 3.3 min is expected to be higher for prospectively undersampled data acquisitions.

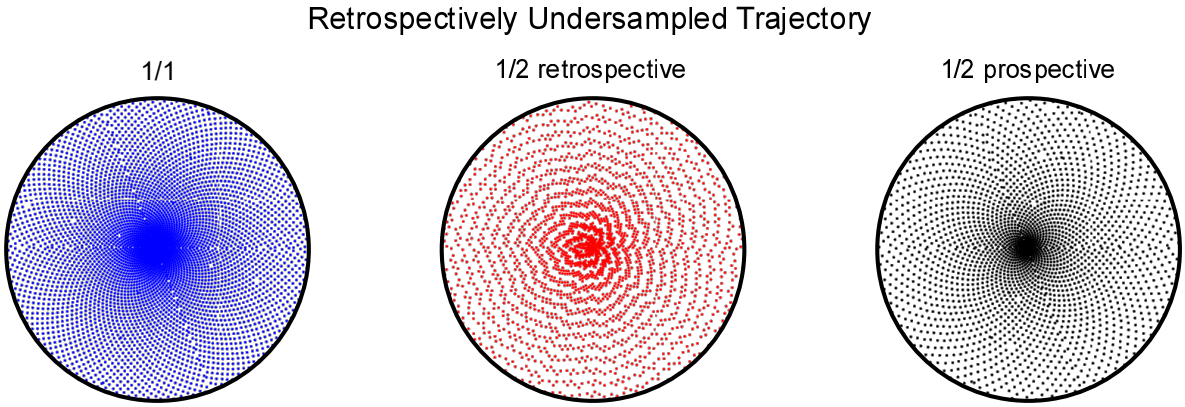

Figure 5.14: Effect of retrospective undersampling on the sunflower trajectory. Left: full k-space, centre: retrospectively undersampled k-space by factor $\frac{1}{2}$, right: prospectively acquiring only $\frac{1}{2}$ points in a sunflower trajectory. Prospective and retrospective undersampling lead to different undersampling patterns, with retrospective undersampling leading to larger gaps in k-space.

5.6 Conclusion

In this chapter, a cardio-respiratory motion-compensated and model-based fat-water separated image reconstruction framework was presented. 3D high-resolution fat-water separated images could be acquired during free-breathing and in a clinically feasible scan time. The effectiveness of the motion-correction approach could be shown for 9 patients.
Finally, the results presented in this chapter demonstrated that this framework is also well suited to address the challenges posed by [^{18}F]NaF PET/MR:

- Motion-corrected fat- and water images serve as a high-quality anatomical reference to locate [^{18}F]NaF uptake.
- Fat- and water images can be the basis for high-resolution attenuation maps resolved for fat, water and lung tissue.
- The framework can produce high-quality motion-models for both cardiac and respiratory motion. These can be used to minimise motion artefacts in simultaneously acquired PET data.

6

3D Motion Corrected Cardiac [^{18}F]NaF PET/MR

This chapter describes the extension of the cMR application developed at 1.5 T, to simultaneous [^{18}F]NaF PET/MR MCIR. It aims to develop a self-contained acquisition-reconstruction framework using MR-based high-quality physiological motion models to improve PET image quality, as well as high-resolution anatomical uptake localisation. Parts of this chapter have been published in the European Journal of Nuclear Medicine and Molecular Imaging (EJNMMI) as J2.

6.1 Background

In medical imaging, PET is used in a wide range of different cardiac applications, from the assessment of cardiac viability and perfusion to the recently developed atherosclerotic plaque imaging using [^{18}F]NaF [13, 24]. The latter can be used to highlight vulnerable plaques, that are likely to rupture and cause myocardial infarction by identifying pathological characteristics such as micro-calcifications and inflammation [26]. Reliable and reproducible quantification of [^{18}F]NaF uptake would be a step towards patient and plaque specific treatment planning. Both hybrid modalities PET/CT [13, 206, 207] and PET/MR [32, 208] have been shown to provide a good assessment of coronary [^{18}F]NaF uptake. A comprehensive comparison between both hybrid modalities[32] found equally successful plaque identification in aortic valves and coronary arteries. Yet, both hybrid modalities suffer from the main challenge of cardiac PET in clinical applications: the impairing effect of physiological heart motion due to both respiration and cardiac movement. The physiological motion of the heart leads to a PET uptake blurring, which is especially a challenge for coronary plaque imaging due to the small size of the plaques. Besides, motion can lead to a mismatch between attenuation correction (AC) maps and the PET emission data. The proximity of the coronary arteries to the lungs makes this especially a problem for coronary PET imaging. The major difference in attenuation values between cardiac tissue and lung can lead to severe artefacts. While

the straightforward motion-compensation approach of gating is very robust, only a small part of the acquired data is used for image reconstruction yielding low signal-to-noise ratio [13, 209]. Also, it does not address AC misalignment artefacts due to motion [208]. Motion-correction has been proposed to overcome these challenges for a range of other simultaneous PET/MR techniques [135, 139, 148, 190]. Improvements have been shown for MR-based ACs using different respiratory positions or allowing for free-breathing [210, 211]. So far motion-correction techniques in [^{18}F]NaF imaging have been PET/CT-based only. The application of motion correction showed an increase in uptake [35, 212], was extended to successfully improve the test-retest reproducibility of [^{18}F]NaF -PET plaque imaging [212] and combined with partial-volume corrections [143, 213]. However, the sequential PET/CT data acquisition restricts the motion estimation to be performed on motion-resolved low-resolution PET data only, limiting this approach to plaques with high uptake. Simultaneous PET/MR overcomes this challenge by allowing to estimate the motion from high-quality and anatomically detailed images: as of recently, a range of different motion correction schemes is available where motion information is extracted from simultaneously acquired MR data and utilised during PET image reconstruction [129, 138, 139, 172, 214]. Yet, so far there has been no application of cardio-respiratory PET/MR motion correction to [^{18}F]NaF imaging. In this work, we present for the first time MR-based cardio-respiratory motion-correction of simultaneous [^{18}F]NaF PET/MR imaging of the coronary arteries. A framework was developed combining advanced 3D MR acquisition with model-based reconstruction techniques and dedicated image registration for motion estimation. From the acquired MR data patient-specific motion models and dynamic AC maps were generated and applied during PET image reconstruction. This approach minimises motion artefacts in the emission data and ensures accurate alignment between PET and AC data. The framework yields 3D high-resolution and motion-compensated images for both PET and MR. The proposed methods were demonstrated in 10 patients. The effect of MCIR on both MR and PET image quality was assessed in uptake-positive plaques by comparing their target-to-background ratio (TBR), their contrast-to-background ratio (contrast-to-background ratio (CBR)), and the tracer visualisation.

6.2 Methods

PET/MR Data Acquisition

PET/MR data were acquired as part of the study "Molecular PET/MR–Imaging for detection and characterisation of vulnerable atherosclerotic plaques in coronary arteries" (EA4/052017) approved by the Charité ethics committee. It was performed in accordance with the Declaration of Helsinki. Before taking part in the study all patients provided written informed consent. The patient cohort consisted of 10 subjects (8 male), mean age 70 ± 7 years suffering from coronary artery disease. Data were acquired on a Siemens Biograph mMR hybrid PET/MR scanner. A dose of 169 ± 14 MBq [^{18}F]NaF tracer was administered intravenously 104 ± 26 min before starting the PET acquisition over 45 ± 20 min. Before the PET scan, an MR-AC scan provided by the vendor was carried out during breath-hold

(FOV=598 mm × 330 mm × 271 mm, $dx \times dy \times dz$=2.086 mm × 2.6 mm × 2.086 mm, TA=10.6 s). An overview of the timeline for patient 1 is given in Fig. 6.1.

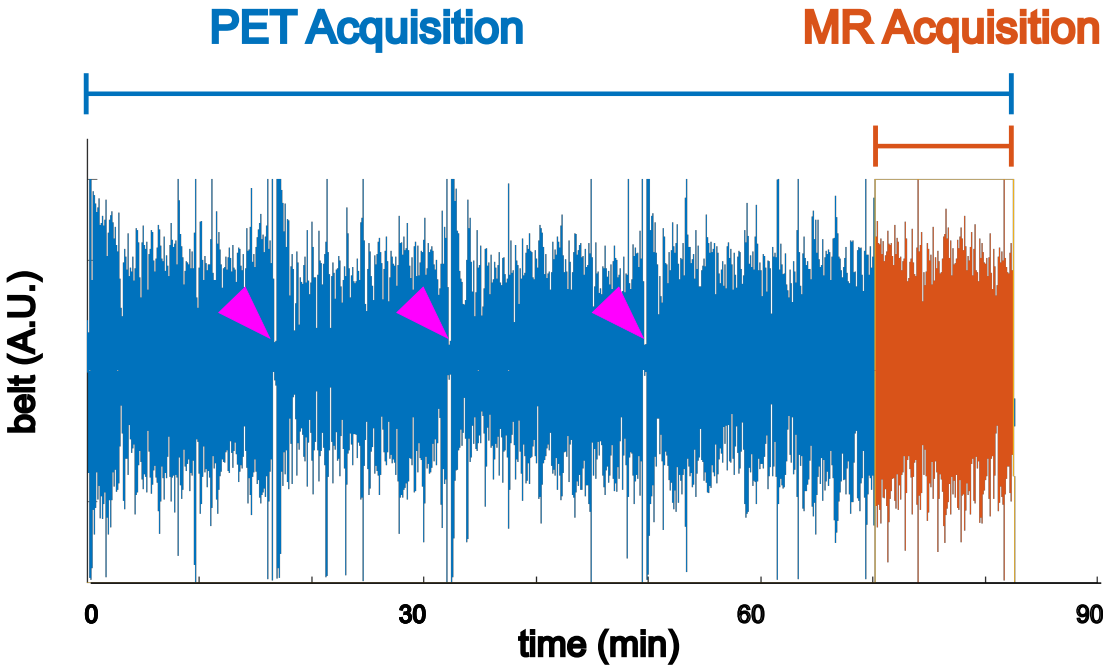

Figure 6.1: Timeline of performed PET/MR examination for an example patient. The displayed curve is the respiratory belt signal. PET data are acquired during the whole duration of the exam. The orange window shows the time window in which both PET and MR employed in this work were acquired simultaneously. Pink arrows indicate breath-holds when other, clinical MR examinations were performed.

The MR data were acquired for 12:25 minutes simultaneously with the PET data and used a T_1 -weighted, Dixon sequence (TR=7.57 ms, TE=2.62/4.13/5.64 ms, FA=15°) with a double-oversampled 3D RPE K-space sampling trajectory [180] with sunflower pattern and a field of view (FOV) covering the entire thorax (FOV=288 mm × 288 mm × 288 mm) at 1.5 mm isotropic resolution. ECG and respiratory signals were acquired simultaneously with the physiological monitoring unit. A T1-contrast agent (Gadovist) bolus of 0.2 mL/kg was administered before the MR exam. The contrast agent was used to increase the contrast between the blood pool and myocardium and improve the visualisation of the coronary arteries [214]. All PET/MR data were acquired as part of the study "Molecular PET/MR–Imaging for detection and characterisation of vulnerable atherosclerotic plaques in coronary arteries" (EA4/052017) approved by the Charité ethics committee. All procedures performed in studies involving human participants were in accordance with the ethical standards of the institutional and/or national research committee and with the 1964 Helsinki declaration and its later amendments or comparable ethical standards. Before taking part in the study all subjects provided written informed consent.

PET/MR Reconstruction Workflow

An overview of the reconstruction workflow is depicted in Fig. 6.2. The acquired data (A) consists of MR k-space, PET listmode data, and respiratory belt, and ECG signals as motion surrogates. Three consecutive fat-water separated MR reconstructions are performed to obtain a cardio-respiratory motion model, a 3D motion-compensated MR anatomical image, and an

AC map. First, using the belt signal the k-space data are binned into 6 different respiratory states (B), motion-resolved reconstructions of fat and water content are performed. While this eliminates the respiratory motion in each respiratory reconstructed state they still contain cardiac motion. Nevertheless, the motion artefacts due to cardiac motion are similar in all respiratory motion states and do not interfere with respiratory motion estimation. From these images, a respiratory motion model is generated using image registration. Subsequently, data are binned into 12 different cardiac states, (C) using the ECG signal, and motion-resolved fat-water reconstructions, including the application of the respiratory motion model, are performed. Hence each reconstructed cardiac state is free from respiratory motion artefacts. A model for cardiac motion is extracted analogously to respiration. Both motion models are combined in a third reconstruction (D) compensating for both motion types. Based on this motion-free 3D fat-water MR image an AC map is computed. For the PET reconstruction, the cardio-respiratory motion model is applied to the MR-based 4-tissue AC map [139] (E) as well as to the listmode data, yielding a cardio-respiratory motion-compensated (cr-MCIR) PET image (F).

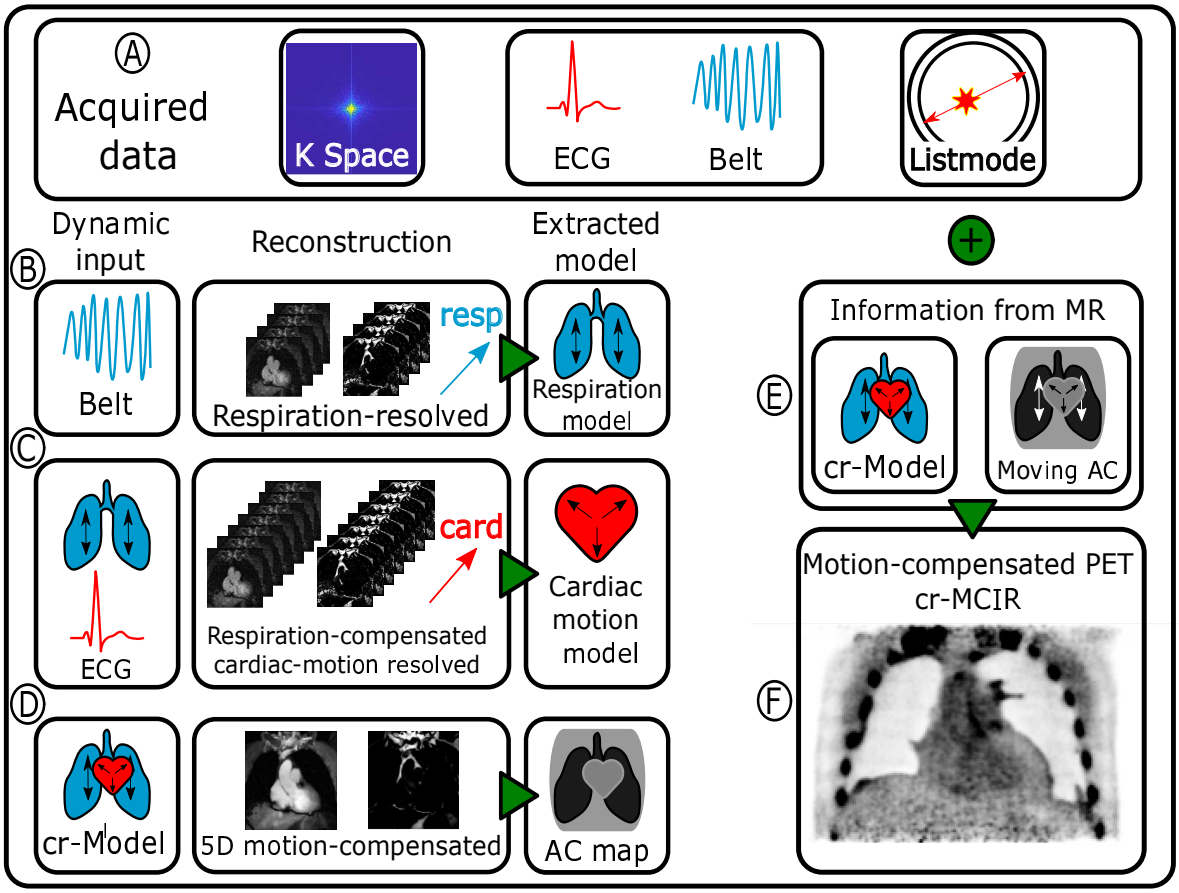

Figure 6.2: Overview of the reconstruction workflow. The acquired PET/MR data consist of listmode and k-space data, as well as the respiratory belt and ECG as surrogate signals for physiological motion. Three MR reconstructions (B-D) are performed from which motion information and an attenuation map are extracted. This information is incorporated into the PET reconstruction (E) compensating both emission data and attenuation map for motion yielding a cardio-respiratory motion-compensated (cr-MCIR) PET reconstruction (F). The cr-MCIR MR (D) and PET (F) reconstructions are hence in the same motion state and the MR anatomical image can be used to identify the anatomical location of the uptake. This figure was published in J2.

PET/MR Image Reconstruction

The MR data were reconstructed into fat and water content with an iterative model-based reconstruction [68] framework incorporating the effect of chemical shift, and using parallel imaging [87], compressed sensing [215], as well as motion information [89]. Motion-resolved images were reconstructed using the respiratory belt or ECG signal as a surrogate to bin data before reconstruction based on respiratory amplitude, and cardiac phase. Software for MR reconstruction was implemented in MATLAB (The MathWorks, Natick, MA) and Python. This has been described in detail in section 5. PET image reconstruction was performed with STIR (Software for Tomographic Image Reconstruction). A FOV of 718 mm × 718 mm × 258 mm was reconstructed at a resolution of 2.09 mm × 2.09 mm × 2.03 mm using an iterative 3D ordered subset expectation-maximisation algorithm with 21 subsets, 3 full iterations, and a 4 mm isotropic 3D Gaussian post-filtering[164]. PET scatter radiation estimates were performed using SIRF [166] (Synergistic Image Reconstruction Framework). An AC image was generated from the cr-MCIR MR reconstruction and transformed into the motion states matching the emission data during PET reconstruction.

MR-based Motion Models and Attenuation Correction

Respiratory and cardiac motion models were generated using motion-resolved reconstructed MR images. A dedicated image registration algorithm [78, 139] was used to generate a deformable motion model based on both the MR fat and water modes. The fat-water dual image registration is discussed in more detail in section 4.2. Specifically, the underlying cost function is described in Eq. (4.11). The rationale behind using both image modes for the registration is that while the water image describes the motion of coronary arteries and heart muscle, the surrounding fat tissue provides complementary information with very positive contrast [102, 200]. Additionally, the undersampling artefacts due to motion-binning are different in both images such that using both image content acts as additional regularisation as assessed in chapter 4. For this study $\lambda = 0.5$ was used, weighting water and fat images the same. The motion model was generated by registering the individual motion-states to the reference phase. Both motion types were combined into a cardio-respiratory motion model concatenating the two individual transformations. The cardiac motion was corrected to yield images in end-diastole while respiration was corrected to either end-exhale or end-inhale depending on which state was more prevalent in the surrogate signal. PET data were acquired for longer than MR data. To extend the obtained motion models onto the entire duration of the PET scan, the respiratory belt was used as a respiratory motion surrogate. Nevertheless, an MR self-navigator was available and used to improve the correlation [138] of the belt signal with the respiratory motion of the heart: in the time-window during the MR-acquisition a time shift for the belt-signal maximising the cross-correlation between belt and self-navigator is determined. This shift is then applied globally to the belt signal. In each case, the computed shift was much smaller than any period of respiration encountered during the exams. The MR-based AC map was extracted from the MR cr-MCIR image and used for both reconstructing motion-averaged, i.e. without the application of any motion model during reconstruction (AVG) and cardiac and respiratory motion-compensated PET images (cr-MCIR). Using the MR cr-MCIR images to calculate the AC map for both AVG PET and

cr-MCIR PET ensures that the difference in image quality between these two reconstructions is predominantly due to physiological motion artefacts in the emission data. A k-means clustering algorithm [216] was used to segment fat, soft tissue, lung tissue and air. MR signal voids caused by stents were automatically inpainted using morphological operations. As the arms were not covered in the MR FOV these were inserted from the vendor AC map. This process is displayed in Fig. 6.3.

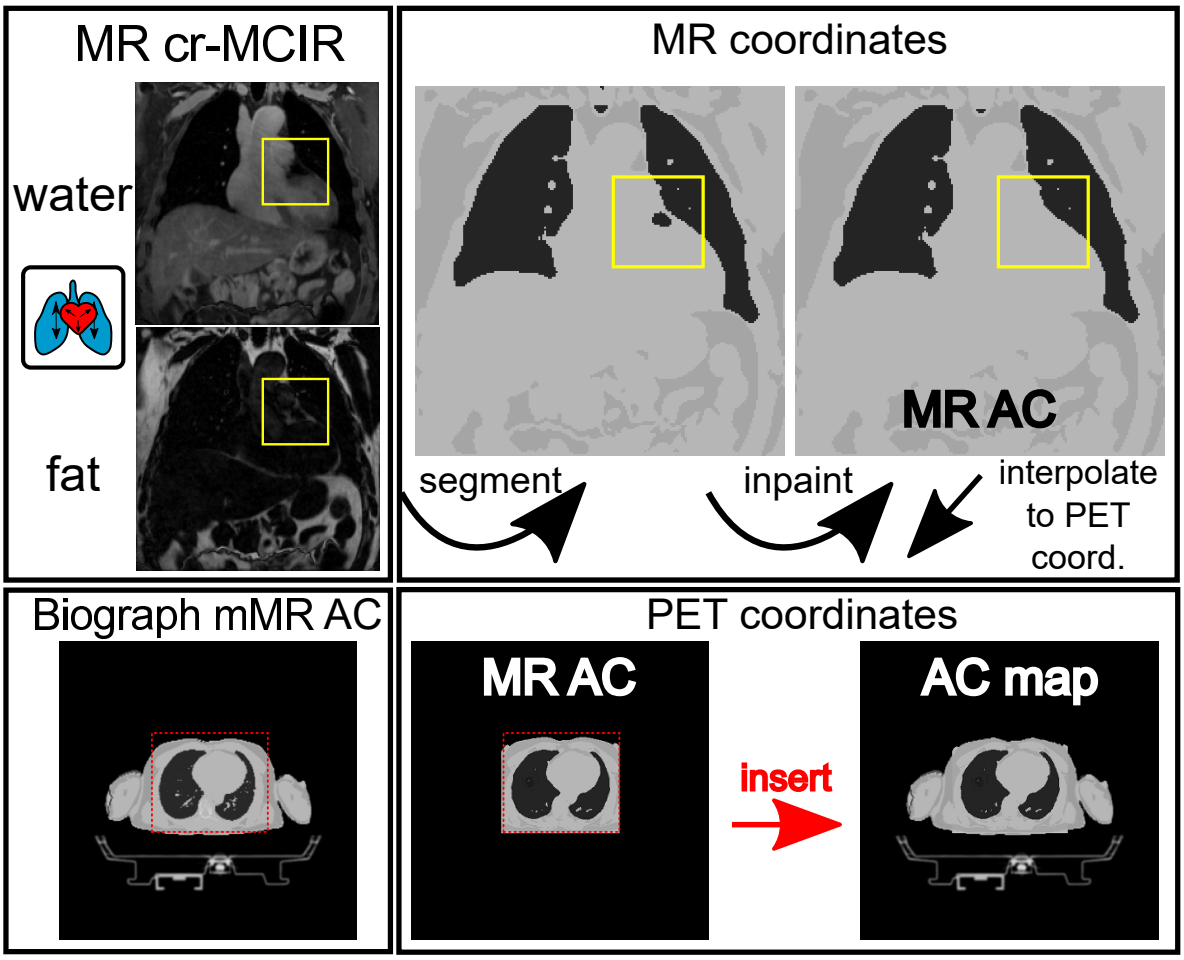

Figure 6.3: Example of AC-map generation. Top: water and fat cr-MCIR are segmented into fat, soft-tissue, lung tissue and outside air. Subsequently, the stent-caused signal void is inpainted (yellow box). Bottom: the generated AC map is interpolated into the PET coordinate system. Then the relevant window containing the thorax is inserted in the Biograph mMR AC (red box) to complete the AC map with the attenuation of the arms and hardware which are static during data acquisition.

PET Image Quality Assessment

The quality of MR images for both water and fat fraction was assessed visually between AVG, respiratory motion-compensated, i.e. only the respiratory motion model was applied during reconstruction (r-MCIR), and cr-MCIR reconstructions. An increase in coronary vessel sharpness and the position of the diaphragm was used to determine a successful generation of the cardio-respiratory motion model. All PET reconstructions were converted into SUV. The STIR open-source reconstructions were compared qualitatively to reconstructions provided by the Biograph mMR in terms of uptake position, noise level, and artefacts generated by AC

misalignment. The effect of motion-correction was evaluated quantitatively for all detected plaques based on their TBR, CBR, and diameter (D). Uptake-positive plaques were detected visually by an experienced observer. Plaques were identified as clearly visible, coherent tracer uptake structures following the coronary vessel tree. An independent method to confirm that the uptake was a plaque such as intra-vascular US or OCT was not available. Both 2D axial and coronal slice through the plaque centre were selected from both AVG and cr-MCIR images. Subsequently, the same plaque in all images was marked with a rectangular ROI.
The plaque signal s was in a 5x5 pixel patch around the maximum standardised uptake value (SUV) value in the ROI. To exclude background pixels from the patch an automated threshold value inside the patch was computed using Otsu's method [217]. This is depicted in Fig. 6.4. The background signal b was extracted from an independent coronary slice and measured as the blood signal average in an ROI in the left ventricle: $b = \langle SUV \rangle_{ROI\,LV}$. The TBR was calculated as $TBR = \frac{s}{b}$, and CBR as $CBR = \frac{\|s-b\|}{\sigma(b)}$, where $\sigma(b)$is the standard deviation of the background. TBR and CBR values generated from both the axial and coronal view of the plaque were averaged to include the effect of correcting motion both in and through transversal planes.

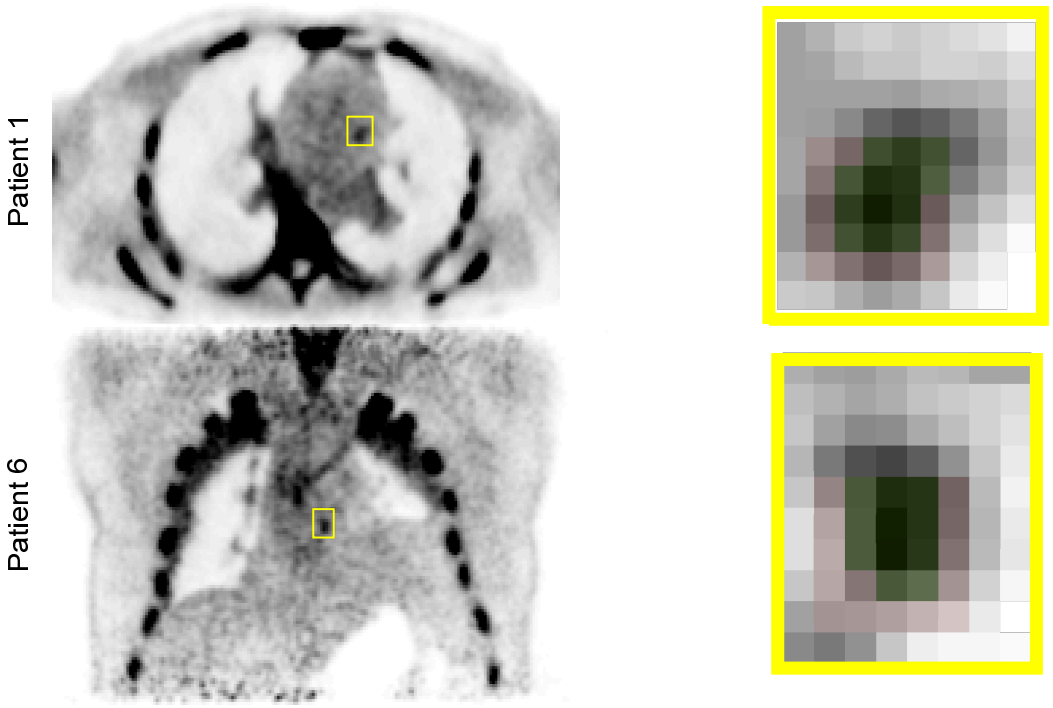

Figure 6.4: Patch-based computation of plaque signal. Yellow boxes show the ROI in which the plaque was localised. A zoomed window shows a 5x5 patch around the SUVmax value in the ROI in red shading with green pixels classified as signal by an automated segmentation in the patch using Otus's method. Top: axial slice of patient 1. Bottom: Coronal slice of patient 6.

To determine the diameter D, a line profile $l(x)$ perpendicular to the coronary artery was extracted from the coronary slices and fitted with a model containing a Gaussian for the plaque and a linear term and a constant locally describing the scatter background: $l(x) = a \cdot e^{(\frac{-(x-b)^2}{c^2})} + d \cdot x + e$. plaque width (D) was determined as the full-width at half maximum of the fitted Gaussian curve: $D = 2\sqrt{2\log 2} \cdot c$. The parameters were constrained to be positive and $d \in [-1, 1]$. Statistical analyses between AVG and cr-MCIR data were performed using

Python. To test the null hypothesis, that the values of both AVG and cr-MCIR both originate from the same underlying distribution for the case of normal data, a Wilcoxon signed-rank test was performed. P-values smaller than 0.05 were considered statistically significant.

6.3 Results

MR reconstruction and motion model generation

The results of the MR motion correction are displayed in Fig. 6.5. R-MCIR improved the visualisation of the coronary vessel, which can be seen especially for the fat image (small insert in Fig. 6.5). Nevertheless, there is still residual motion blurring which was further reduced using cr-MCIR.

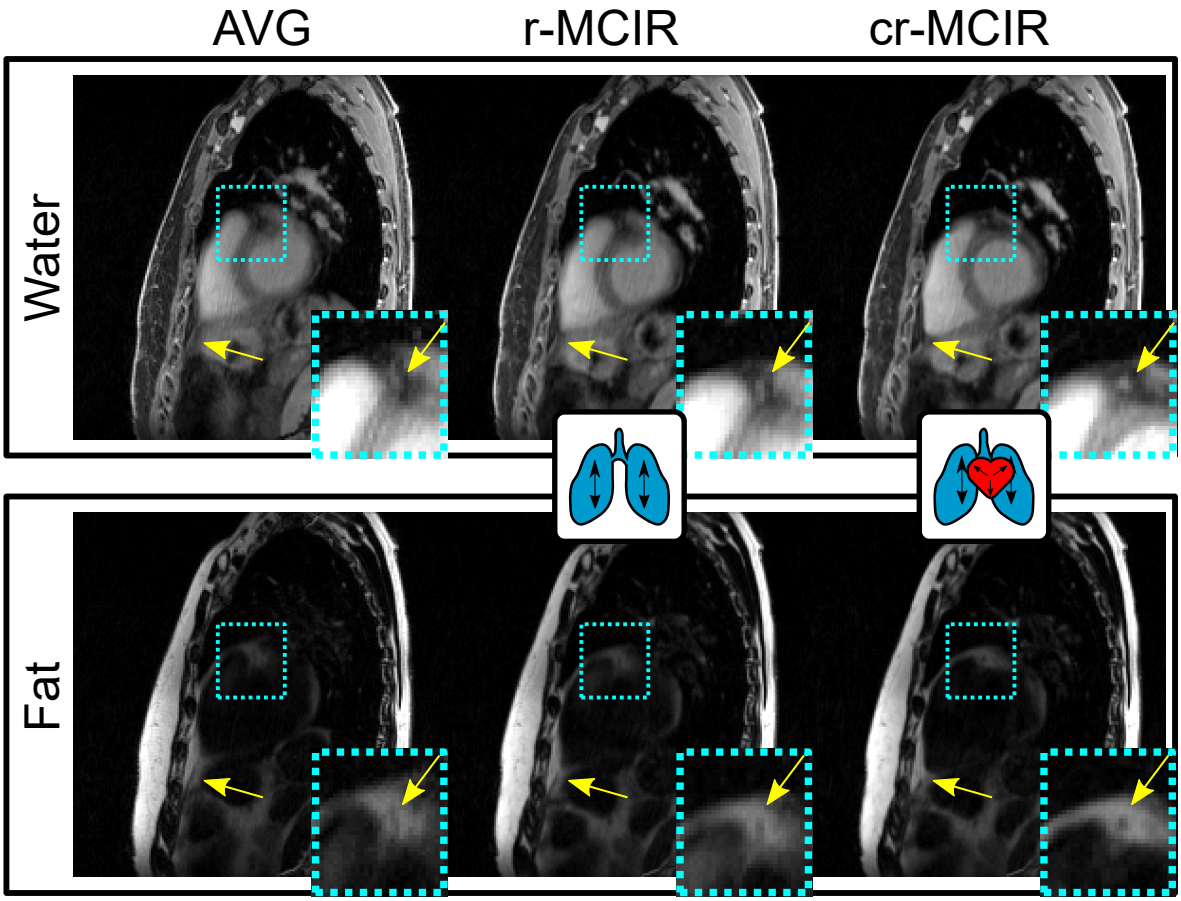

Figure 6.5: MR images in sagittal view of one exemplary dataset with reconstructed water (top) and fat (bottom) content. The motion-averaged reconstruction (AVG, left) is compared to respiratory- (r-MCIR, centre) and cardio-respiratory (cr-MCIR, right) reconstruction. Cyan boxes indicate enlarged areas depicted in the lower right corner reconstruction. Image quality increases with the inclusion of more motion information from AVG over r-MCIR to cr-MCIR for both the water and fat images especially in the coronary arteries and apex (yellow arrows). This figure was published in J2.

PET Motion-Correction

The effect of cr-MCIR in PET for three patients is depicted in Fig. 6.6. The AVG and cr-MCIR reconstructions are compared for three slices showing coronary uptake. Line profiles

in cyan for AVG and magenta for cr-MCIR are plotted and an increase in SUV_{max} is visible in the line profiles. The line profile fit based on which D is calculated is indicated as a black curve and for the three presented cases, D was reduced by 6 %, 44 %, and 34 % respectively. Furthermore, one can see that for patients 10 and 5 there was already a very good match between the emission data and the AC map as very little imprint artefacts are visible. The AVG reconstruction for patient 7, however, has an artefact indicated by a red arrow due to the misalignment between PET emission data and AC map that was strongly reduced by cr-MCIR.

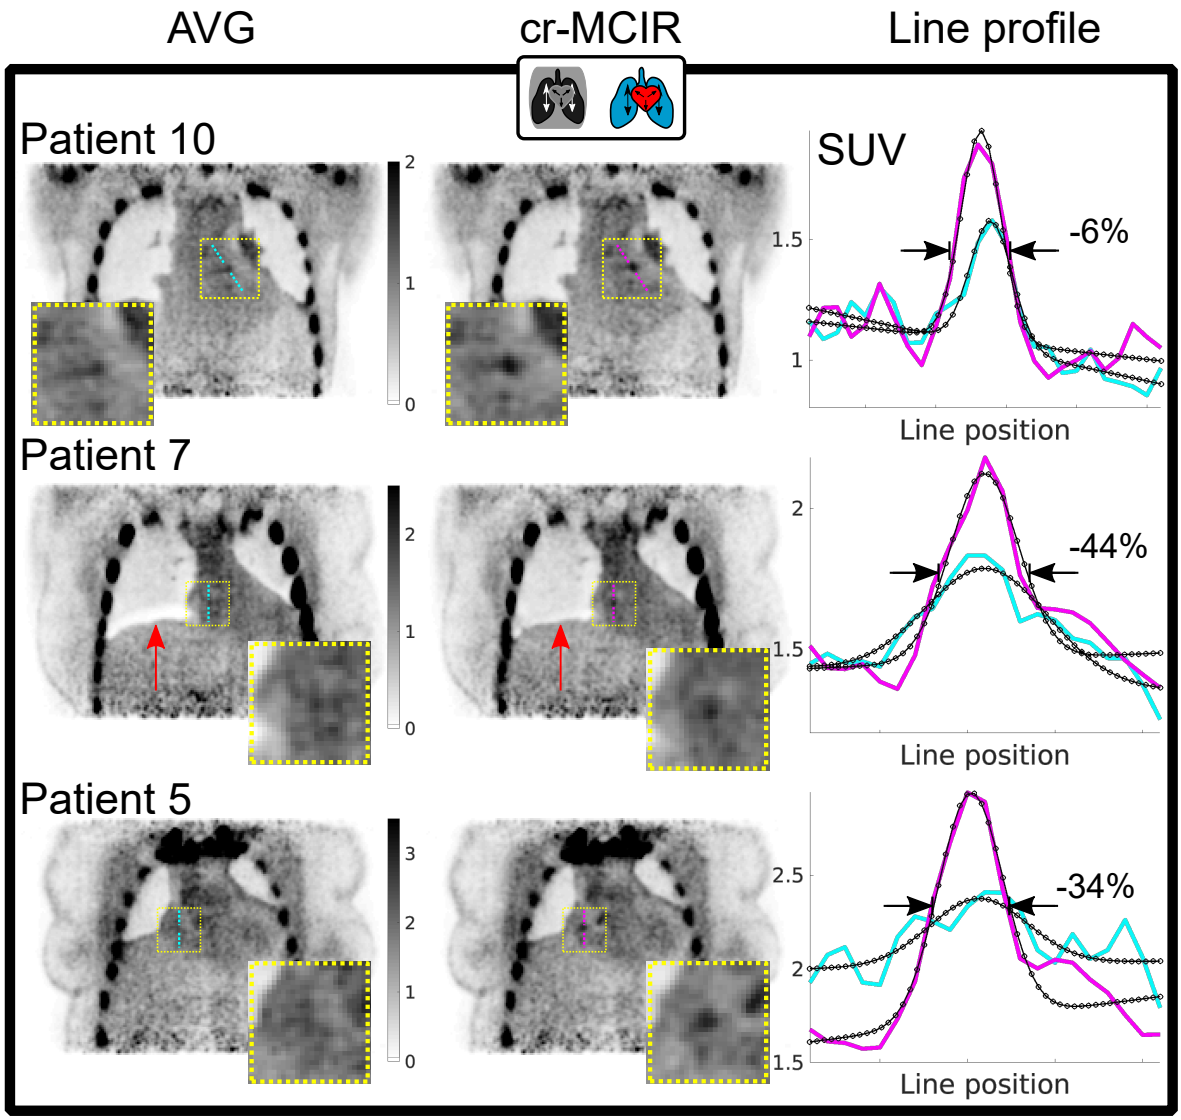

Figure 6.6: Comparison between PET AVG (left column) and cr-MCIR (center column) for patients 10, 7, and 5. Cyan and magenta lines indicate the position where line profiles (right column) were extracted. The fit is overlayed in black. Compared to AVG, cr-MCIR leads to an increase in the maximum uptake value and a decrease of plaque width D. The decrease ranged between −6 % and −44 % for the displayed cases. Patient 7 shows an artefact due to misalignment between AC map and PET emission data for AVG which is corrected by using cr-MCIR (red arrow). This figure was published in J2.

A comparison between the AVG, cr-MCIR, and Biograph mMR images is depicted in Fig. 6.7. Two different coronal positions are displayed, both of which contain uptake in

the left coronary artery. AVG and cr-MCIR reconstructions achieved a comparable quality of the PET uptake of the peripheric thorax region, visible when regarding uptake in the ribcage and shoulders. The overall background noise due to scatter radiation was lower in the Biograph reconstructions such that slightly different window settings were required to create the same visibility of the plaque uptake. The misalignment between the AC map and emission data is especially visible in the scanner reconstruction at the right hemidiaphragm (red arrow, "banana artefact"). Those impaired the visualisation of the coronary uptake which was located right at the edge of the heart and thereby obscured the actual extent of the uptake along the vessel indicated by the yellow arrows.

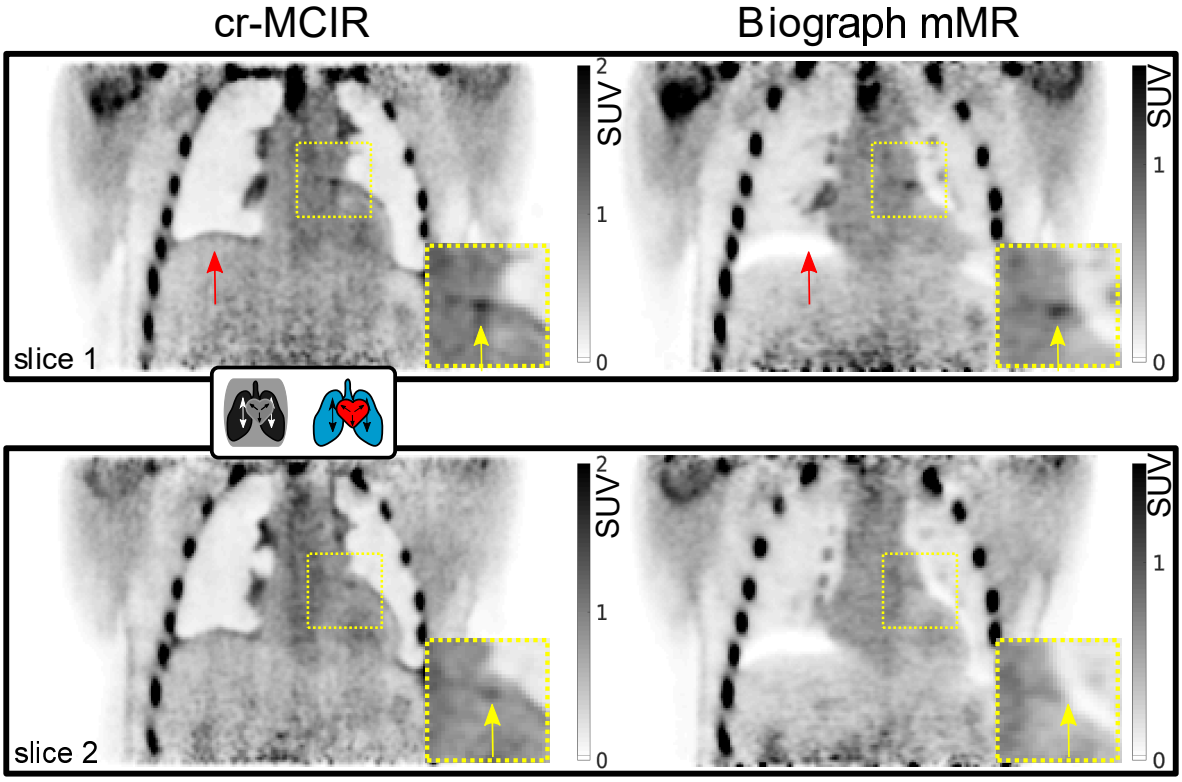

Figure 6.7: Result of the PET cr-MCIR application to patient 1. An enlarged ROI is marked by the dashed yellow square. Top: coronal slice 1, bottom: coronal slice 2. Left: cr-MCIR. Right: Biograph mMR vendor reconstruction. Yellow arrows highlight uptake in the left coronary artery; red arrows indicate artefacts due to an AC mismatch. The cr-MCIR shows a reduced AC mismatch compared to the Biograph mMR images. For the vendor reconstruction, the strong AC-data mismatch due to physiological motion means the uptake in the coronary plaque is not visible anymore in the second slice (bottom row). Scanner and STIR reconstructions required a different window setting for comparable contrast in the coronary uptake. This figure was published in J2.

An overlay of MR and PET is depicted in Fig. 6.8. For patients 5 and 8 a coronal slice is depicted where the PET image was reformatted to the MR coordinate system. Both patients showed uptake in a stent that caused an artefact in the MR image. Despite this void, a reliable MR-based AC map could be estimated allowing for accurate PET quantification and visualisation of plaque uptake. However, the exact anatomic position could not be determined in the MR image. Especially for patient 8 the signal void was much larger than the uptake.

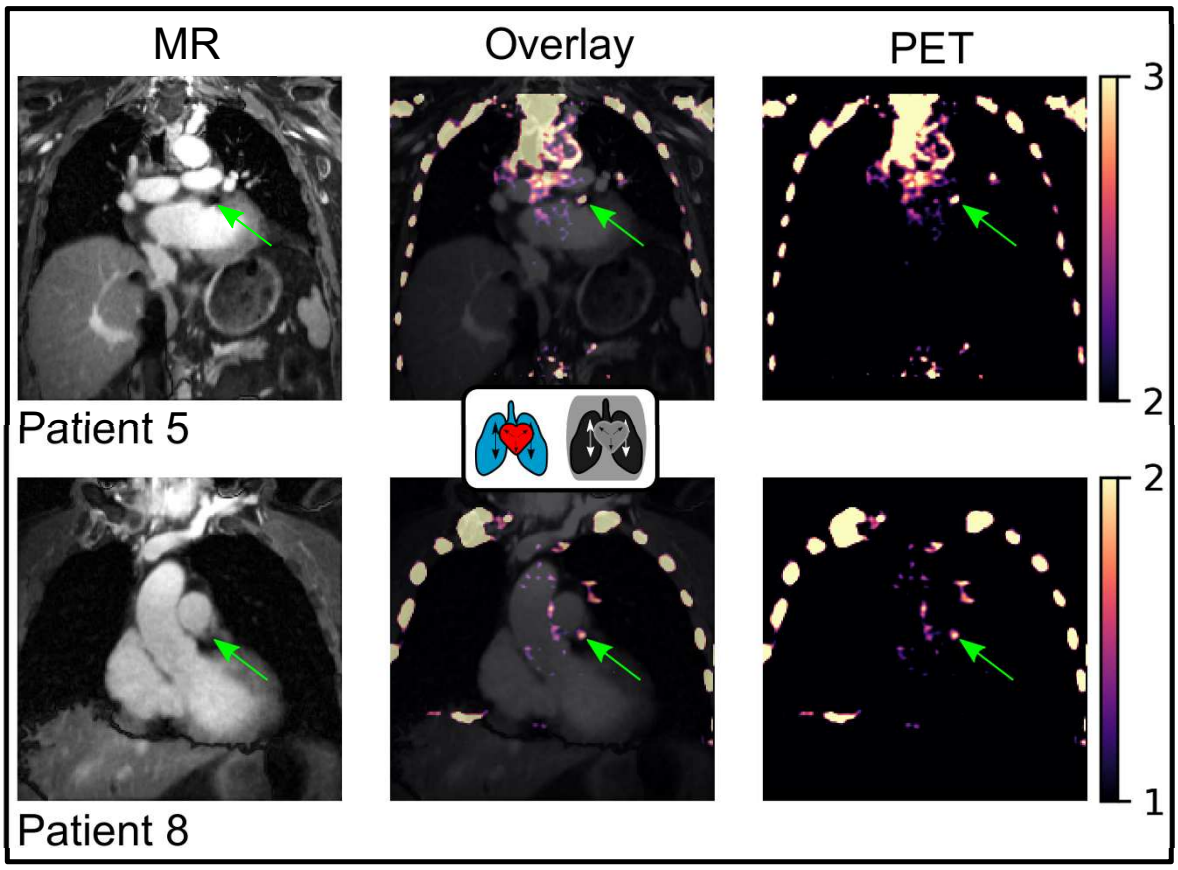

Figure 6.8: PET/MR overlay for patients 5 and 8 with uptake in stents (green arrows). All images are cr-MCIR. PET images were reformatted to the MR coordinates. Left column: MR water images. Central column: overlay of both modalities. Right column: PET images only. The colour bar corresponds to the PET window in both images. In this case, the anatomic position of the uptake is obscured by the signal void generated by the stent. However, the automated inpainting of the stent voids during AC map generation mitigates potential AC artefacts for tracer uptake within stents. This figure was published in J2.

Quantitative PET Image Assessment

In the 10 patient datasets, 10 coronary plaques were identified in 9 subjects. A comprehensive overview of the computed image quality metrics is given in Tab. 6.1. The measured TBR in the 10 plaques ranged between 1.11 and 3.10 for AVG and 1.21 and 3.41 for MCIR reconstruction. The average increase in TBR was $7 \pm 7\,\%$ using cr-MCIR. The Wilcoxon signed-rank test yielded a p-value of $p<0.02$. The CBR ranged between 1.02 and 15.15 for AVG and 2.30 and 18.76 for cr-MCIR reconstructions with a mean increase of $26 \pm 38\,\%$ for cr-MCIR compared to AVG. The Wilcoxon signed-rank test applied to the data yielded a p-value of $p<0.04$. D ranged between 7.6 mm mm and 23.4 mm for the AVG and 5.47 mm and 17.91 mm for the cr-MCIR reconstructions. On average cr-MCIR decreased D by $23 \pm 18\,\%$. The Wilcoxon signed-rank test yielded a p-value of $p<0.02$.

Patient	TBR_{stat}	TBR_{dual}	ΔTBR (%)	CBR_{stat}	CBR_{dual}	ΔCBR (%)	D_{stat}	D_{dual}	ΔD (%)
$NaF_{20190307}$	1.34	1.51	13.0	2.59	4.24	64.0	12.20	11.63	-5.0
$NaF_{20190528}$	1.45	1.51	4.0	4.43	4.80	8.0	12.54	9.90	-21.0
$NaF_{20190611}$	1.30	1.21	-7.0	3.22	2.30	-29.0	7.60	7.84	3.0
$NaF_{20190625}$	3.10	3.41	10.0	15.15	18.76	24.0	23.71	17.91	-24.0
$NaF_{20191119}$	1.11	1.24	12.0	1.02	2.12	107.0	18.72	12.39	-34.0
$NaF_{20191119}$	1.58	1.68	6.0	5.46	5.93	9.0	12.57	5.47	-56.0
$NaF_{20191126}$	2.54	2.50	-2.0	9.28	8.80	-5.0	11.87	9.91	-16.0
$NaF_{20191212}$	1.37	1.48	8.0	4.49	5.72	27.0	23.39	13.09	-44.0
$NaF_{20191217}$	1.35	1.47	9.0	4.43	5.80	31.0	16.43	12.62	-23.0
$NaF_{20200114}$	1.54	1.76	14.0	5.06	6.13	21.0	10.07	9.43	-6.0
Mean			7 ± 7 *			26 ± 38 *			-23 ± 18 *

* $p < 0.05$

Table 6.1: Comprehensive overview of image quality data of motion-corrected [^{18}F]NaF PET. Plaque width D is given in mm.This table was published in J2.

6.4 Discussion

In this work, we presented an MR-based motion-compensation framework and its application to coronary [^{18}F]NaF PET/MR imaging. The data-driven patient-specific cardio-respiratory motion models enabled MCIR for both modalities which improved the assessed image quality metrics TBR, CBR, and plaque width significantly.

The quantitative assessment yielded an increase in TBR of up to 14 % and CBR of up to 107 %. Average increases were 7 ± 7 % for TBR and 26 ± 38 % for CBR. The plaque localisation in cr-MCIR reconstructions was improved compared to AVG reconstructions, reducing the diameter D on average by 23 ± 18 % with a maximum reduction of 56 %. Statistical tests on all these quantities yielded significant differences between AVG and cr-MCIR. The estimated motion model also improved the anatomical visualisation of coronary arteries in the MR cr-MCIR where both cardiac and respiratory motion contributed to an increase in the visibility and sharpness of the coronary vessels and fat surrounding the heart. The improvement of TBR and CBR depends on the location of the plaque, the motion amplitudes, and motion cycles which strongly vary between patients. Furthermore, it could be shown that a dedicated motion modelling of the AC was able to minimise mismatches between the AC map and emission data, allowing for accurate visualisation of plaque uptake in the PET images. A first improvement, as shown in Fig. 6.6, can be achieved by obtaining the AC map from a cr-MCIR MR image that is motion-corrected to the most prevalent respiratory state. Further improvement is achieved by adapting the AC map to the different motion states as part of PET MCIR (patient 7 in Fig. 6.6). In contrast to that, the AC map for vendor reconstruction was acquired during a breath-hold. Fig. 6.8 depicts an example where the breath-hold position was very different from the free-breathing position of the heart during PET data acquisition. This led to a strong AC mismatch severely impairing the visualisation of the uptake in the plaque, because AC values of the lung were used to correct PET emission data of the heart. One limitation of the presented work is that the tracer uptake could not be independently identified as coronary plaques. Furthermore, a comparison between the images achieved using motion-compensation and cardio-respiratory double-gated reconstructions were omitted. The observed TBRs using the whole dataset were too small to expect any visible plaque uptake in double-gated images. While the respiratory motion model is patient-specific, the necessity to

extrapolate the respiratory motion model from the time window of MR data acquisition onto the whole PET acquisition potentially reduces its accuracy. Large differences in respiratory amplitudes between the time during which the motion model was generated and the rest of the PET exam could make the application of the motion model less effective [70].
The number PET counts collected during the approximately 12-minute window of simultaneous MR data acquisition was not sufficient. Hence, to be able to use all acquired PET counts the motion models had to be extrapolated to the rest of the PET scan. The validity of the extrapolated motion model outside the MR window is a strong assumption that could be justified with further analysis: there have been studies [158, 218] classifying breathing patterns into three different types based on the displacement histograms measured in the diaphragm. An assessment of the probability density distribution of the surrogate signals used for data binning inside and outside the 12-minute window of the MR acquisition was performed. The distributions are displayed in Fig. 6.9 containing three different breathing patterns for three different patients. The distributions were compared visually. While we could show that there were different breathing patterns between patients, there was no intra-patient difference in the shape of the distributions inside and outside the MR acquisition window for any of the 10 acquired subjects.

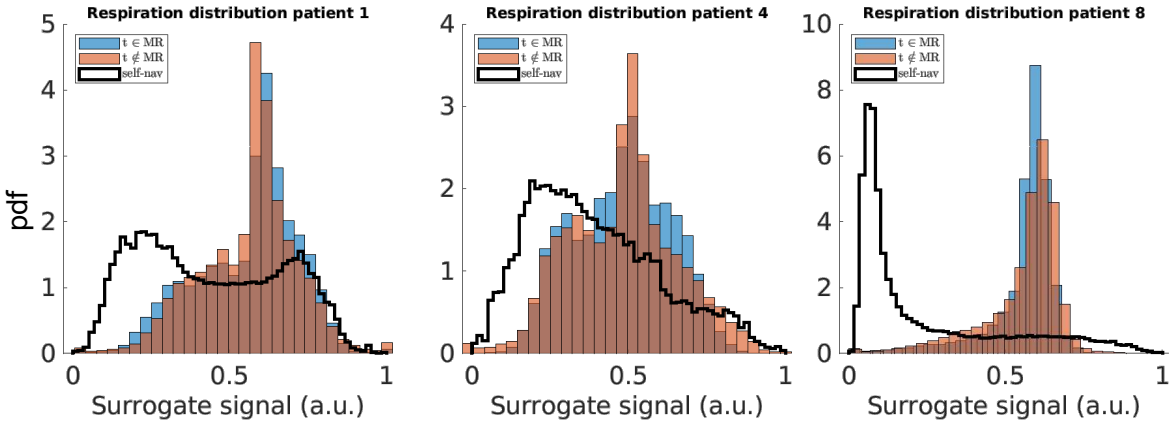

Figure 6.9: analysis of respiratory surrogate signal patterns. Blue: distribution of belt signal during the MR acquisition. Red: belt signal outside the MR window. Black: self-navigator as a direct measure of diaphragm displacement in head-foot direction. The differences in the black curves indicate different breathing patterns for the three different patients. However, there is no strong intra-patient variability in belt-signal distribution shape inside and outside the MR data acquisition window.

This extrapolation required using the respiration belt for data binning which generally is less accurate in describing motion states compared to an MR-based self-navigator[70]. This can also be deducted from Fig. 6.9 where the distributions have different shapes between the actual displacement reflected by the navigator and the belt signals. Using additional motion surrogates [219] might further improve the application of the motion model to the PET data. Also, this study neglected point-spread function modelling of the PET system. As the scatter radiation used by the vendor system is unavailable, a standard technique [220] was used for its computation. As displayed in Fig. 6.7, comparable overall image quality was achieved with a slightly higher noise level in our reconstructions than in the vendor images. Being able to use

the scatter estimation of the vendor could further improve image quality [221]. Furthermore, as depicted in Fig. 6.8 and shown by previous studies [32], the MR signal voids caused by stents do prevent the use of the MR as anatomic information to precisely locate uptake. Potential errors, however, that propagate into the MR-based AC map could be efficiently and automatically dealt with by custom morphological transformations and inpainting. The AC map outside the MR FOV was completed by the vendor-provided attenuation information of the arms. These data, if not available, could also be acquired time-efficiently in a short, free-breathing, extra scan of a few seconds before or after finalisation of the PET exam.

6.5 Conclusion

In this chapter we presented the development and application of an MR-based motion-compensation framework dedicated to [^{18}F]NaF PET/MR imaging, yielding cardio-respiratory motion-compensated high-resolution 3D MR and PET images. All required information for motion-model and AC map generation, as well as the anatomic localisation of the observed coronary uptakes, could be extracted from one comprehensive MR acquisition. Furthermore, we could show that it is feasible to use all acquired PET data in motion-corrected image reconstruction to achieve a high image quality despite the small uptakes. The presented increase in TBR and CBR, the reduction in plaque width D and mitigation of attenuation correction mismatch artefacts could enable a more reliable and reproducible plaque localisation.

7

Summary

In this thesis new methods to improve the detection and quantification of coronary atherosclerotic plaques with simultaneous [^{18}F]NaF PET/MR were developed and applied. To address the challenges posed by physiological motion a flexible data acquisition and model-based image reconstruction framework was developed able to correct motion of the heart due to both respiration and heartbeat in PET/MR. The first component consisted of a numerical simulation framework which was devised and implemented for ground truth evaluations of the accuracy of different motion estimation approaches. Furthermore, MR model-based fat water reconstruction was combined with cardio-respiratory motion correction and evaluated at 1.5 T in patients with myocardial fat infiltrations. Subsequently, these methods were extended by a motion-resolved attenuation correction and applied to successfully motion-correct simultaneous [^{18}F]NaF PET/MR in patients with known CAD. The quality of PET images necessary for a reliable and reproducible quantification of cardiovascular atherosclerosis was improved.

The generation of motion models in different applications during the project was supported by a numerical simulation framework for dynamic simultaneous PET/MR. Its functionality and implementation were described in chapter 3. The framework was designed for input based on existing scanner raw data while providing ground truth motion information. Hence, the simulations seamlessly integrate into already existing reconstruction workflows for patient data with the ability to quantify the accuracy of registrations performed on the reconstructions. It was shown that the simulation output yields realistic image data when compared to the underlying patient data, as well as realistic cardiac and respiratory motion. Its usefulness was demonstrated for simultaneous PET/MR, as well as stand-alone PET and MR applications. These included the optimisation of regularisation parameters when generating 4D cardiac and respiratory motion models for motion-resolved [^{18}F]FDG PET, as well as evaluating the registration accuracy of a non-rigid synergistic PET/MR registration algorithm. It was found that for a synergistic 4D cardiac (respiratory) motion registration weighting the PET images with 80 % (40 %) produces the optimal registration accuracy. The maximum error of the motion fields over all motion phases in the myocardium could be reduced to approximately 1 voxel.

The dynamic capabilities are not limited to motion alone, also signal and contrast changes (e.g. due to contrast agent injection) can be flexibly incorporated and combined with motion. While the framework can generate data similar to patient data with a low computational demand, it still is limited by the incomplete realism of the employed PET and MR signal models. In particular, the signal is based on analytical signal equations for steady-state sequences. Also, the presented analysis is based only on the tissue segmentation and motion of the XCAT model that was used as simulation input. Yet, as the segmentation is not bound to XCAT, every available voxelised tissue phantom with a matching motion model could be used. Hence, analysis performed with the simulation framework can be extended to multiple anatomies. Possible improvements could therefore be the use of more sophisticated signal models to extend its use to for example to T_1 -mapping or MR fingerprinting.
By incorporating the simulation into a well-maintained open-source software project, it can be accessed by a wider group of users. The final framework will be interfaced with Python and merged into SIRF as a submodule. Furthermore, in- and output of the simulation is based on open-source file formats (ISMRMRD and Interfile) which facilitates sharing the simulation output and promotes reproducible research.
An additional by-product was the implementation of the RPE encoding model into the open-source framework SIRF. This gives access to a larger variety of implemented reconstruction algorithms. It already allows reconstructing MR data with advanced methods [222] and should be able to facilitate synergistic methods for PET/MR data reconstruction in the future.

Methods that were developed aimed at an improved estimation of motion-models and attenuation correction were presented in chapter 4. On the acquisition side, a novel phase-encoding pattern for a 3D RPE k-space trajectory based on the arrangement of seeds in a sunflower was implemented. This could reduce the maximum gap sizes in k-space after retrospective motion binning and resulted in a more regular distribution of k-space points independent of whether cardiac or respiratory motion surrogates were used. On the reconstruction side, a model-based fat-water separation framework with spatial and temporal total-variation sparsity regularisation was implemented and combined with motion correction. To generate the motion models for MCIR the synergistic PET/MR registration evaluated in chapter 3 was transferred to the fat and water MR data. The motion in both fat and water image modes were jointly registered exploiting the complementary information and positive contrast the fat images provide. With the simulation framework, it was shown that using this approach both cardiac and respiratory motion can be accurately estimated from a combination of water and fat images.

Chapter 5 describes the use and validation of the novel acquisition-reconstruction methods in a stand-alone clinical cardiac MR application. It could be shown that the visualisation of adipose structures in the heart as well as coronary vessels could be improved using cardio-respiratory motion correction. Patient data for 9 subjects with suspected lipotamous metaplasia were acquired during free-breathing. The successful separation into water and fat could be validated by comparing 2D slices from a clinical routine protocol with reformatted fat images from the model-based reconstruction. The effect of motion correction on quality of the

3D fat images were subsequently assessed quantitatively using a local edge sharpness metric averaged over 4 different anatomical landmarks in different image orientations. The simulation-optimised registration parameters were largely employed when generating the patient specific motion-models for patient data in which no ground truth information to assess the registration quality is available. Respiratory MoCo yielded a $34 \pm 17\,\%$ increase in fat structure sharpness, and cardiac MoCo an additional $8 \pm 8\,\%$. In one subject myocardial fat infiltrations were detected. It was shown that the reconstruction spatially covered the entire infiltrations in 3D from one acquisition and that cardio-respiratory motion-correction strongly increased their visibility.

A clinical application of the presented model-based fat-water separated MoCo approach, however, could be hindered by long reconstruction times in the order of multiple hours. Foremost, this is because the sunflower trajectory requires non-uniform Fast Fourier Transform (NUFFT) [223] which is more computaionally costly than the standard Fast Fourier transform (FFT). Secondly, the iterative model inversion requires a more sophisticated optimisation algorithm than a standard iterative conjugate gradient. However, this has to be weighted against the advantages to the reconstruction framework over a standard Dixon approach, allowing to regularise the fat-water seperated images and exploit sparsity in the fat images. Furthermore, there was no effort in the scope of this thesis to accelerate the reconstruction times. This could potentially be achieved performing the NUFFT on a GPU instead of on a CPU.

On the acquisition side, however, a retrospective undersampling analysis to examine the effect of potential accelerations was performed. The image analysis yielded comparable image quality when reducing the theoretical acquisition time from 14 to 7 minutes where the sharpness metrics were consistent with the full dataset while the sharpness metric increased by $33 \pm 18\,\%$ for respiratory - and an additional $6 \pm 7\,\%$ for cardiac MoCo. Further reducing the acquisition time below 4 minutes was assessed in a single patient, and also able to produce high-quality motion-resolved images using preliminary machine learning image denoising. The retrospective undersampling properties of the trajectory, however, did not allow to reliably assess the fat structures. New, prospectively undersampled data would have to be acquired to assess the resulting fat image quality. The chapters 4 and 5 showed the diagnostic virtue of the presented MR sequence for cardiac fat-water applications. But also the developed methods yield image data that is able to supply a simultaneous PET/MR exam with 3D high-resolution anatomical information of the heart, motion models to employ in MCIR, as well as the basis for a fat-water resolved AC map of the thorax at the same time.

Finally, in chapter 6 all of the above methods were applied to acquire and reconstruct simultaneous [^{18}F]NaF PET/MR for 10 patients with known or suspected CAD. The generated motion-models could be successfully transferred to the PET emission and AC data. The required information to generate a motion-model and AC map generation, as well to locate the observed coronary uptakes in the anatomy could be extracted from one comprehensive MR acquisition. The motion-compensated fat-water separation was segmented into fat, soft, and lung-tissue and successfully converted into a 4-tissue AC map. Signal voids due to stents were removed using inpainting methods. Hence, the displayed simultaneous PET/MR was ideally co-registered and could provide motion, attenuation and high-resolution motion-compensated

anatomical information at once. The effect of cardio-respiratory MCIR on [^{18}F]NaF PET was analysed by evaluation the TBR, CBR and diameter of the tracer uptake in the plaques. MCIR could improve TBR by up to 14 % and CBR by up to 107 %. The average increase in was 7 ± 7 % for TBR, and 26 ± 38 % for CBR. Furthermore, cr-MCIR was able to reduce motion blurring leading to an average plaque diameter reduction of 23 ± 18 %. These quantitative measures show that MCIR was able to display the plaque uptake sharper and with higher uptake value and contrast than without motion correction.
However, there were no independent data acquired which could confirm that a PET uptake is in fact a high-risk atherosclerotic plaque. Uptakes with small TBR were, however, assumed to be atherosclerosis, if they formed a coherent structure following the vascular tree over multiple slices. A next step could be to analyse whether cr-MCIR allows the depiction of plaques that were not visible in the motion-averaged reconstructions. However, that would require the acquisition of ground truth reference data on the location of vulnerable plaques, as could, for example, be provided by intravascular OCT or US data.
The comparison with the vendor reconstruction showed that the presented STIR reconstructions feature a higher noise level. This is assumed to be due to an imperfect scatter radiation estimation. If this could improved to the quality of the vendor scatter correction a positive impact on image quality is to be expected. Also, PET MCIR with more complex optimisation algorithms [224] or the inclusion of time-of-flight (TOF) information [225–227] and partial volume corrections [143, 213] could further improve PET image quality.
Furthermore, we could show that it is feasible to use all acquired PET data in motion-corrected image reconstruction to achieve a high image quality despite the small coronary [^{18}F]NaF uptakes. To this end the motion models were generated from a time window on which the MR information was available and had to be extrapolated from this window to the duration of the entire PET scan. While the variations in the heartbeat can be assumed to be well captured during the MR acquisition time [228, 229], the respiratory patterns might not (e.g. if a subject falls asleep). However, the probability distributions for the surrogate signal amplitudes allow to draw inferences from the belt signal to the respiratory pattern [83, 218]. It was shown that there are strong inter-patient variations between these distributions, suggesting different breathing for different patients. However, there were no intra-patient differences between the distributions of the time window during which the motion models were generated and the acquisition times to which it was extrapolated for any of the ten acquired patient datasets.
The low quality of the belt signal as a respiratory surrogate generally limits the respiratory binning for the PET data. Yet, it is the only signal available during the whole scan, but is noisier and has a lower correlation with the respiratory displacement of the heart than a MR self-navigator. A navigator is, however, only available during the time window of the simultaneous data acquisition and hence unfit to extrapolate the model. Machine learning could aid in denoising the belt signal where during the period in which both surrogate signals are available for example a convolutional neural network (CNN) can be trained to determine a map between the belt signal and the self-navigator. Similar approaches have been previously used by training classical signal filters [138].

The effect of the proposed respiratory resolved AC map during cr-MCIR was compared to using an AC acquired during breath-hold as employed in the vendor reconstruction. The

Biograph mMR images occasionally contained large AC misalignments, while for cr-MCIR uptake in the vessel tree close to the heart/lung interface was better visible and was under a lower risk of superimposition of a AC-misalignment artefact. Consequentially, modelling AC motion is deemed crucial to guarantee a reproducible plaque detection independent of different respiratory amplitudes during the AC map data acquisition.

Further improvement of the developed methods could be the incorporation of a more complex surrogate signal for respiratory motion that would allow for more detailed motion models and the inclusion of hysteresis [70, 158]. Generally, the incorporation of machine learning into reconstruction has already and will play a major role in medical imaging of the future. The data acquired in the scope of this thesis would potentially benefit from an machine learning-based improved estimation of the cardiac motion [202, 230]. This highly patient-specificic and non-linear motion with large amplitudes is hard to estimate accurately from a fixed set of regularisation parameters. Also, the registration algorithms employed in this thesis are able to co-register a general images. Instead, machine learning could replace this general image registration approach by supplying algorithms dedicated to motion estimation of medical images and hence specifically optimise the motion model generation for the individiual motion types.
Machine learning can be further involved in quantiatitve applications such as parameter fitting in DCE and pharmacokinetic modelling. The presented simulation framework on the other hand could serve as a mean generate data for the validation and potentially the training of the above-mentioned machine learning applications.
The methods proposed in this thesis show that physiological motion correction can strongly improve the quality and quantification of [^{18}F]NaF -PET uptake in coronary plaques. These are an important basis for the clinical use of [^{18}F]NaF PET as a potential non-invasive way to quantify the risk of rupture posed by coronary plaques. Furthermore, with its ability to increase the TBR and CBR motion-correction could potentially reduce the time necessary for PET data acquisition. Hence, it could also help to increase patient throughput for a comprehensive clinical exam which aims at the prediction and prevention of MI with a patient-specific treatment.

8

Author's Publications

Journal Articles

J1 **Johannes Mayer**, Richard Brown, Kris Thielemans, Evgueni Ovtchinnikov, Edoardo Pasca, David Atkinson, Ashley Gillman, Paul Marsden, Matteo Ippoliti, Marcus Makowski, Tobias Schaeffter, and Christoph Kolbitsch, Flexible numerical simulation for dynamic PET-MR data, *Physics in Medicine and Biology*, 2020;65.14, p. 145003, DOI: 10.1088/1361-6560/ab7eee.

J2 **Johannes Mayer**, Thomas-Heinrich Wurster, Tobias Schaeffter, Ulf Landmesser, Andreas Morguet, Boris Bigalke, Bernd Hamm, Winfried Brenner, Marcus R. Makowski, and Christoph Kolbitsch ,Imaging Coronary Plaques Using 3D Motion-Compensated 18F-NaF PET/MR, *European Journal of Nuclear Medicine and Molecular Imaging*, 2021; DOI: 10.1007/s00259-020-05180-4

J3 **Johannes Mayer**, Yinning Jin, Thomas-Heinrich Wurster, Marcus R. Makowski, and Christoph Kolbitsch, Evaluation of synergistic image registration for motion-corrected coronary NaF-PET-MR *Phil. Trans. R. Soc. A*, 2021, 379: 20200202. 20200202., 10.1098/rsta.2020.0202

J4 Evgueni Ovtchinnikov, Richard Brown, Christoph Kolbitsch, Edoardo Pasca, Casper da Costa-Luis, Ashley G. Gillman, Benjamin A. Thomas, Nikos Efthimiou, **Johannes Mayer**, Palaw Wadhwa, Matthias J. Ehrhardt, Sam Ellis, Jakob S. Jorgensen, Julian Matthews, Claudia Prieto, Andrew J. Reader, Charalampos Tsoumpas, Martin Turner, David Atkinson, and Kris Thielemans, SIRF: Synergistic Image Reconstruction Framework, *Computer Physics Communications*, 2020;p. 107087, DOI: 10.1016/j.cpc.2019.107087.

J5 Giulio Ferrazzi, Jean Pierre Bassenge, **Johannes Mayer**, Alexander Ruh, Sebastien Roujol, Bernd Ittermann, Tobias Schaeffter, Lucilio Cordero-Grande, and Sebastian Schmitter, Autocalibrated cardiac tissue phase mapping with multiband imaging

and k-t acceleration *Magnetic Resonance in Medicine*, 2020;p. 2429-2441, DOI: 10.1002/mrm.28288.

J6 Sebastian Dietrich, Christoph Aigner, Christoph Kolbitsch, **Johannes Mayer**, Juliane Ludwig, Simon Schmidt, Tobias Schaeffter, and Sebastian Schmitter, 3D Free-breathing multichannel absolute B_1^+ Mapping in the human body at 7T, *Magnetic Resonance in Medicine*, 2020; DOI: 10.1002/mrm.28602.

J7 Richard Brown, Christoph Kolbitsch, Claire Deplancke, Evangelos Papoutsellis, **Johannes Mayer**, Evgueni Ovtchinnikov, Edoardo Pasca, Radhouene Neji, Casper da Costa-Luis, Ashley G. Gillman, Matthias J. Ehrhardt, Jamie McClelland, Bjoern Eiben, and Kris Thielemans, Motion estimation and correction for simultaneous PET/MR using SIRF and CIL, *Phil. Trans. R. Soc. A.*, 2021, 379: 20200208. 20200208., DOI: 10.1098/rsta.2020.0208

J8 **Johannes Mayer**, Edyta Blaszczyk, Alberto Cipriani, Giulio Ferrazzi, Jeanette Schulz-Menger, Tobias Schaeffter, and Christoph Kolbitsch, Cardio-Respiratory Motion-Corrected 3D Cardiac Water-Fat MRI Using Model-Based Image Reconstruction, under revision in *Magnetric Resonance in Medicine*

Conference Proceedings

C1 **Johannes Mayer**, Kris Thielemans, Richard Brown, Evgueni Ovtchinnikov, David Atkinson, Paul Marsden, Tobias Schaeffter, and Christoph Kolbitsch, Flexible Numerical Simulation Framework for Dynamic PET-MRI, *Proceedings of the 27th Annual Meeting of ISMRM*, Montreal, Canada, 2019.

C2 Christoph Kolbitsch, **Johannes Mayer**, Alberto Cipriani, Blaszczyk, Jeanette Schulz-Menger, Tobias Schaeffter, Motion-corrected 3D high-resolution fat-water imaging of the heart, *Proceedings of the 27th Annual Meeting of ISMRM*, Montreal, Canada, 2019.

C3 Sebastian Dietrich, Christoph Stefan Aigner, Juliane Ludwig, **Johannes Mayer**, Simon Schmidt, Christoph Kolbitsch, Tobias Schaeffter, and Sebastian Schmitter, Efficient high-resolution cardiac motion-corrected T1 mapping, *Proceedings of the 27th Annual Meeting of ISMRM*, Montreal, Canada, 2019.

C4 Giulio Ferrazzi, Jean Pierre Bassenge, **Johannes Mayer**, Clarissa Wink, Alexander Ruh, Michael Markl, Bernd Ittermann, Tobias Schaeffter, and Sebastian Schmitter, Autocalibrated Multiband CAIPIRINHA with k-t-acceleration: towards complete spatio-temporal coverage of the heart motion in one single breath-hold, *Proceedings of the 27th Annual Meeting of ISMRM*, Montreal, Canada, 2019.

References

[1] Shanthi Mendis, Pekka Puska, and Norrving Bo. *Global Atlas on cardiovascular disease prevention and control*. Ed. by World Health Organization. Geneva, Switzerland, 2011.

[2] Lawrence J Laslett et al. "The Worldwide Environment of Cardiovascular Disease : Prevalence , Diagnosis , Therapy , and Policy Issues A Report From the American College of Cardiology". In: 60.25 (2012).

[3] Sidney C Smith et al. "Our Time: A Call to Save Preventable Death From Cardiovascular Disease (Heart Disease and Stroke)". In: *Journal of the American College of Cardiology* 60.22 (Dec. 2012), pp. 2343–2348.

[4] Massimo F. Piepoli et al. "2016 European Guidelines on cardiovascular disease prevention in clinical practice". In: *European Heart Journal* 37.29 (2016), pp. 2315–2381.

[5] World Health Organization. *The top 10 causes of death*. URL: `https://www.who.int/news-room/fac`

[6] DGK. *Deutscher Herzbericht 2014*. 2015. URL: `https://dgk.org/daten/PA-Herzbericht.pdf`.

[7] Göran K. Hansson. "Inflammation, Atherosclerosis, and Coronary Artery Disease". In: *New England Journal of Medicine* 352.16 (2005), pp. 1685–1695.

[8] Jacob Fog Bentzon et al. "Mechanisms of Plaque Formation and Rupture". In: *Circulation Research* 114.12 (June 2014), pp. 1852–1866.

[9] Aurélien Bustin et al. "Five-minute whole-heart coronary MRA with sub-millimeter isotropic resolution, 100% respiratory scan efficiency, and 3D-PROST reconstruction". In: *Magnetic Resonance in Medicine* (2019).

[10] Marcus R Makowski et al. "Assessment of atherosclerotic plaque burden with an elastin-specific magnetic resonance contrast agent". In: *Nature Medicine* 17.3 (Mar. 2011), pp. 383–388.

[11] Marc R. Dweck et al. "Coronary arterial 18F-sodium fluoride uptake: A novel marker of plaque biology". In: *Journal of the American College of Cardiology* 59.17 (2012), pp. 1539–1548.

[12] Ryogo Minamimoto et al. "Prospective comparison of 99mTc-MDP scintigraphy, combined 18F-NaF and 18F-FDG PET/CT, and whole-body MRI in patients with breast and prostate cancer". In: *Journal of Nuclear Medicine* 56.12 (2015), pp. 1862–1868.

[13] Nikhil V. Joshi et al. "18F-fluoride positron emission tomography for identification of ruptured and high-risk coronary atherosclerotic plaques: A prospective clinical trial". In: *The Lancet* 383.9918 (2014), pp. 705–713.

[14] Jacek Kwiecinski et al. “Coronary 18F-Sodium Fluoride Uptake Predicts Outcomes in Patients With Coronary Artery Disease”. In: *Journal of the American College of Cardiology* 75.24 (2020), pp. 3061–3074.

[15] Gudrun Wagenknecht et al. “MRI for attenuation correction in PET: Methods and challenges”. In: *Magnetic Resonance Materials in Physics, Biology and Medicine* 26.1 (2013), pp. 99–113.

[16] G. Delso et al. “Performance Measurements of the Siemens mMR Integrated Whole-Body PET/MR Scanner”. In: *Journal of Nuclear Medicine* 52.12 (Dec. 2011), pp. 1914–1922.

[17] Craig S. Levin et al. “Design Features and Mutual Compatibility Studies of the Time-of-Flight PET Capable GE SIGNA PET/MR System”. In: *IEEE Transactions on Medical Imaging* 35.8 (Aug. 2016), pp. 1907–1914.

[18] Thorsten A. Bley et al. “Fat and water magnetic resonance imaging”. In: *Journal of Magnetic Resonance Imaging* 31.1 (2010), pp. 4–18.

[19] Dale Bailey et al. “Nuclear Medicine Physics.” In: *Medical Physics* 38.8 (Aug. 2011), pp. 4904–4904.

[20] Salim Yusuf et al. “Effect of potentially modifiable risk factors associated with myocardial infarction in 52 countries (the INTERHEART study): case-control study”. In: *The Lancet* 364.9438 (Sept. 2004), pp. 937–952.

[21] Jan Borén et al. “Low-density lipoproteins cause atherosclerotic cardiovascular disease: Pathophysiological, genetic, and therapeutic insights: A consensus statement from the European Atherosclerosis Society Consensus Panel”. In: *European Heart Journal* 41.24 (2020), pp. 2313–2330.

[22] Jolanda J. Wentzel et al. “Endothelial shear stress in the evolution of coronary atherosclerotic plaque and vascular remodelling: Current understanding and remaining questions”. In: *Cardiovascular Research* 96.2 (2012), pp. 234–243.

[23] *Smart Servier Medical Art*. 2020. URL: `https://smart.servier.com/`.

[24] Jason M. Tarkin et al. “Imaging Atherosclerosis”. In: *Circulation Research* 118.4 (2016), pp. 750–769.

[25] R. J.M.W. Rennenberg et al. “Vascular calcifications as a marker of increased cardiovascular risk: A meta-analysis”. In: *Vascular Health and Risk Management* 5 (2009), pp. 185–197.

[26] Marc R. Dweck et al. “Coronary Arterial 18F-Sodium Fluoride Uptake”. In: *Journal of the American College of Cardiology* 59.17 (Apr. 2012), pp. 1539–1548.

[27] Sarah Bastawrous et al. “F-NaF Skeletal PET/CT in Oncologic Practice 1”. In: *RadioGraphics* 34.1 (2014), pp. 1295–1316.

[28] Marc R. Dweck et al. “18F-Sodium Fluoride Uptake Is a Marker of Active Calcification and Disease Progression in Patients With Aortic Stenosis”. In: *Circulation: Cardiovascular Imaging* 7.2 (2014), pp. 371–378.

[29] Joo Myung Lee et al. “Clinical Relevance of 18 F-Sodium Fluoride Positron-Emission Tomography in Noninvasive Identification of High-Risk Plaque in Patients With Coronary Artery Disease”. In: *Circulation: Cardiovascular Imaging* 10.11 (Nov. 2017).

[30] Agnese Irkle et al. “Identifying active vascular microcalcification by 18F-sodium fluoride positron emission tomography”. In: *Nature Communications* 6.May (2015).

[31] Francesco Fiz et al. “18F-NaF uptake by atherosclerotic plaque on PET/CT imaging: Inverse correlation between calcification density and mineral metabolic activity”. In: *Journal of Nuclear Medicine* 56.7 (2015), pp. 1019–1023.

[32] Jack P. M. Andrews et al. “Cardiovascular 18F-fluoride positron emission tomography-magnetic resonance imaging: A comparison study”. In: *Journal of Nuclear Cardiology* (Dec. 2019).

[33] Martin Lyngby Lassen et al. “Assessment of attenuation correction for myocardial PET imaging using combined PET/MRI”. In: *Journal of Nuclear Cardiology* 26.4 (Aug. 2019), pp. 1107–1118.

[34] M. Rubeaux et al. “Motion Correction of 18F-NaF PET for Imaging Coronary Atherosclerotic Plaques”. In: *Journal of Nuclear Medicine* 57.1 (Jan. 2016), pp. 54–59.

[35] Mhairi K. Doris et al. “Optimization of reconstruction and quantification of motion-corrected coronary PET-CT”. In: *Journal of Nuclear Cardiology* 27.2 (Apr. 2020), pp. 494–504.

[36] Mhairi K. Doris et al. “Motion-corrected imaging of the aortic valve with18F-NaF PET/CT and PET/MRI: A feasibility study”. In: *Journal of Nuclear Medicine* 58.11 (2017), pp. 1811–1814.

[37] Ricella Maria Souza Da Silva and Roberto José Vieira De Mello. “Fat deposition in the left ventricle: descriptive and observacional study in autopsy”. In: *Lipids in Health and Disease* 16.1 (2017), pp. 1–7.

[38] Peter Kellman, Diego Hernando, and Andrew E. Arai. “Myocardial Fat Imaging”. In: *Current Cardiovascular Imaging Reports* 3.2 (2010), pp. 83–91.

[39] Giorgio Baroldi et al. “Lipomatous metaplasia in left ventricular scar”. In: *Canadian Journal of Cardiology* (1997).

[40] L. Pantanowitz. “Fat infiltration in the heart”. In: *Heart* 85.3 (2001), p. 253.

[41] Ify Mordi et al. “Prevalence and Prognostic Significance of Lipomatous Metaplasia in Patients with Prior Myocardial Infarction”. In: *JACC: Cardiovascular Imaging* 8.9 (2015), pp. 1111–1112.

[42] Srijita Sen-Chowdhry et al. “Cardiovascular Magnetic Resonance in Arrhythmogenic Right Ventricular Cardiomyopathy Revisited. Comparison With Task Force Criteria and Genotype”. In: *Journal of the American College of Cardiology* 48.10 (2006), pp. 2132–2140.

[43] Gaetano Nucifora et al. “Lipomatous metaplasia in ischemic cardiomyopathy: Current knowledge and clinical perspective”. In: *International Journal of Cardiology* 146.1 (2011), pp. 120–122.

[44] Jim Pouliopoulos et al. “Intramyocardial adiposity after myocardial infarction: New implications of a substrate for ventricular tachycardia”. In: *Circulation* 128.21 (2013), pp. 2296–2308.

[45] Rahul Samanta et al. “Role of adipose tissue in the pathogenesis of cardiac arrhythmias”. In: *Heart Rhythm* 13.1 (2016), pp. 311–320.

[46] Christopher X. Wong, Anand N. Ganesan, and Joseph B. Selvanayagam. “Epicardial fat and atrial fibrillation: Current evidence, potential mechanisms, clinical implications, and future directions”. In: *European Heart Journal* 38.17 (2017), pp. 1294–1302.

[47] Justus M.B. Anumonwo and Todd Herron. “Fatty infiltration of the myocardium and arrhythmogenesis: Potential cellular and molecular mechanisms”. In: *Frontiers in Physiology* 9.JAN (2018), pp. 1–7.

[48] Fumiko Kimura et al. “Myocardial fat at cardiac imaging: How can we differentiate pathologic from physiologic fatty infiltration?” In: *Radiographics* 30.6 (2010), pp. 1587–1602.

[49] James W. Goldfarb, Marguerite Roth, and Jing Han. “Myocardial fat deposition after left ventricular myocardial infarction: Assessment by using MR water-fat separation imaging”. In: *Radiology* 253.1 (2009), pp. 65–73.

[50] Minjie Lu et al. “Fat deposition in dilated cardiomyopathy assessed by CMR”. In: *JACC: Cardiovascular Imaging* 6.8 (2013), pp. 889–898.

[51] Arnold C.T. Ng et al. “Impact of Epicardial Adipose Tissue, Left Ventricular Myocardial Fat Content, and Interstitial Fibrosis on Myocardial Contractile Function”. In: *Circulation. Cardiovascular imaging* 11.8 (2018), e007372.

[52] Kristofer Nyman et al. “Cardiac steatosis and left ventricular function in men with metabolic syndrome”. In: *Journal of Cardiovascular Magnetic Resonance* 15.1 (2013), pp. 1–11.

[53] Timothy J.P. Bray et al. “Fat fraction mapping using magnetic resonance imaging: Insight into pathophysiology”. In: *British Journal of Radiology* 91.1089 (2018).

[54] Cho Kai Wu et al. “Myocardial adipose deposition and the development of heart failure with preserved ejection fraction”. In: *European Journal of Heart Failure* 22.3 (2020), pp. 445–454.

[55] Theodoros D. Karamitsos et al. “Myocardial Tissue Characterization and Fibrosis by Imaging”. In: *JACC: Cardiovascular Imaging* 13.5 (2020), pp. 1221–1234.

[56] Houchun Harry Hu et al. “ISMRM workshop on fat-water separation: Insights, applications and progress in MRI”. In: *Magnetic Resonance in Medicine* 68.2 (Aug. 2012), pp. 378–388.

[57] Maryam Nezafat et al. “Coronary MR angiography at 3T: fat suppression versus water-fat separation”. In: *Magnetic Resonance Materials in Physics, Biology and Medicine* 29.5 (2016), pp. 733–738.

[58] Christian Lücke et al. “Prevalence and functional impact of lipomatous metaplasia in scar tissue following myocardial infarction evaluated by MRI”. In: *European Radiology* 20.9 (2010), pp. 2074–2083.

[59] Young Jin Kim et al. “Lipomatous metaplasia in patients with myocardial infarction: evaluation with cardiac magnetic resonance”. In: *Journal of Cardiovascular Magnetic Resonance* 12.S1 (2010), p. 1.

[60] Jabi E. Shriki et al. “Chemical Shift Artifact on Steady-State Free Precession Cardiac Magnetic Resonance Sequences as a Result of Lipomatous Metaplasia: A Novel Finding in Chronic Myocardial Infarctions”. In: *Canadian Journal of Cardiology* 27.5 (Sept. 2011), pp. 17–664.

[61] G H Glover and E Schneider. “Three-Point Dixon Technique for True Water / Fat Decomposition with Bo Inhomogeneity Correction”. In: 383 (1991), pp. 371–383.

[62] Scott B. Reeder et al. “Multicoil Dixon Chemical Species Separation with an Iterative Least-Squares Estimation Method”. In: *Magnetic Resonance in Medicine* 51.1 (2004), pp. 35–45.

[63] Scott B. Reeder et al. “Homodyne reconstruction and IDEAL water-fat decomposition”. In: *Magnetic Resonance in Medicine* 54.3 (2005), pp. 586–593.

[64] Scott B. Reeder et al. “Cardiac CINE imaging with IDEAL water-fat separation and steady-state free precession”. In: *Journal of Magnetic Resonance Imaging* 22.1 (2005), pp. 44–52.

[65] Johan Berglund et al. “Three-point Dixon method enables whole-body water and fat imaging of obese subjects”. In: *Magnetic Resonance in Medicine* 63.6 (2010), pp. 1659–1668.

[66] Peter Kellman and Andrew E. Arai. “Cardiac imaging techniques for physicians: Late enhancement”. In: *Journal of Magnetic Resonance Imaging* 36.3 (2012), pp. 529–542.

[67] Mariya Doneva et al. “Compressed sensing for chemical shift-based water-fat separation”. In: *Magnetic Resonance in Medicine* 64.6 (2010), pp. 1749–1759.

[68] Thomas Benkert et al. “Free-breathing volumetric fat/water separation by combining radial sampling, compressed sensing, and parallel imaging”. In: *Magnetic Resonance in Medicine* 78.2 (2017), pp. 565–576.

[69] Claudio Santelli et al. “Respiratory bellows revisited for motion compensation: Preliminary experience for cardiovascular MR”. In: *Magnetic Resonance in Medicine* 65.4 (2011), pp. 1097–1102.

[70] J.R. McClelland et al. “Respiratory motion models: A review”. In: *Medical Image Analysis* 17.1 (Jan. 2013), pp. 19–42.

[71] Richard Manber et al. “Practical PET respiratory motion correction in clinical PET/MR”. In: *Journal of Nuclear Medicine* 56.6 (2015), pp. 890–896.

[72] Martin Buehrer et al. “Prospective self-gating for simultaneous compensation of cardiac and respiratory motion”. In: *Magnetic Resonance in Medicine* 60.3 (2008), pp. 683–690.

[73] Dana C. Peters et al. “2D free-breathing dual navigator-gated cardiac function validated against the 2D breath-hold acquisition”. In: *Journal of Magnetic Resonance Imaging* 28.3 (2008), pp. 773–777.

[74] Michael A. Bush et al. “Patient specific prospective respiratory motion correction for efficient, free-breathing cardiovascular MRI”. In: *Magnetic Resonance in Medicine* 81.6 (2019), pp. 3662–3674.

[75] Juliane Ludwig et al. “Pilot tone–based motion correction for prospective respiratory compensated cardiac cine MRI”. In: *Magnetic Resonance in Medicine* May (Nov. 2020), mrm.28580.

[76] Lucilio Cordero-Grande et al. “Sensitivity Encoding for Aligned Multishot Magnetic Resonance Reconstruction”. In: *IEEE Transactions on Computational Imaging* 2.3 (2016), pp. 266–280.

[77] Ashley Gillman et al. “PET motion correction in context of integrated PET/MR: Current techniques, limitations, and future projections: Current”. In: *Medical Physics* 44.12 (2017), e430–e445.

[78] D. Rueckert et al. “Nonrigid registration using free-form deformations: application to breast MR images”. In: *IEEE Transactions on Medical Imaging* 18.8 (1999), pp. 712–721.

[79] Daniel Rueckert et al. “Diffeomorphic Registration Using B-Splines”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006*. Ed. by Rasmus Larsen, Mads Nielsen, and Jon Sporring. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 702–709.

[80] Marcus Brehm et al. “Self-adapting cyclic registration for motion-compensated cone-beam CT in image-guided radiation therapy”. In: *Medical Physics* (2012).

[81] Christopher M. Rank et al. “Respiratory motion compensation for simultaneous PET/MR based on a 3D-2D registration of strongly undersampled radial MR data: a simulation study”. In: (2015), p. 941218.

[82] Christoph Kolbitsch et al. “Joint cardiac and respiratory motion estimation for motion-corrected cardiac PET-MR”. In: *Physics in Medicine & Biology* 64.1 (Dec. 2018), p. 015007.

[83] Irene Polycarpou, Charalampos Tsoumpas, and P. K. Marsden. “Analysis and comparison of two methods for motion correction in PET imaging”. In: *Medical Physics* 39.10 (Oct. 2012), pp. 6474–6483.

[84] H.M. Hudson and R.S. Larkin. “Accelerated image reconstruction using ordered subsets of projection data”. In: *IEEE Transactions on Medical Imaging* 13.4 (1994), pp. 601–609.

[85] Feng Qiao et al. “A motion-incorporated reconstruction method for gated PET studies”. In: *Physics in Medicine and Biology* 51.15 (2006), pp. 3769–3783.

[86] Joyoni Dey and Michael A. King. “Theoretical and numerical study of MLEM and OSEM reconstruction algorithms for motion correction in emission tomography”. In: *IEEE Transactions on Nuclear Science* 56.5 (2009), pp. 2739–2749.

[87] Klaas P. Pruessmann et al. “Advances in sensitivity encoding with arbitrary k -space trajectories”. In: *Magnetic Resonance in Medicine* 46.4 (Oct. 2001), pp. 638–651.

[88] Michael Lustig, David Donoho, and John M. Pauly. “Sparse MRI: The application of compressed sensing for rapid MR imaging”. In: *Magnetic Resonance in Medicine* 58.6 (2007), pp. 1182–1195.

[89] P. G. Batchelor et al. “Matrix description of general motion correction applied to multishot images”. In: *Magnetic Resonance in Medicine* 54.5 (2005), pp. 1273–1280.

[90] Manojkumar Saranathan and James Glockner. “Three-dimensional dixon fat-water separated rapid breathheld imaging of myocardial infarction”. In: *Journal of Magnetic Resonance Imaging* 38.6 (2013), pp. 1362–1368.

[91] Thanh D. Nguyen et al. “A fast navigator-gated 3D sequence for delayed enhancement MRI of the myocardium: Comparison with breathhold 2D imaging”. In: *Journal of Magnetic Resonance Imaging* 27.4 (2008), pp. 802–808.

[92] Dana C. Peters et al. “Left ventricular infarct size, peri-infarct zone, and papillary scar measurements: A comparison of high-resolution 3D and conventional 2D late gadolinium enhancement cardiac MR”. In: *Journal of Magnetic Resonance Imaging* 30.4 (2009), pp. 794–800.

[93] Aya Kino et al. “Three-dimensional phase-sensitive inversion-recovery turbo FLASH sequence for the evaluation of left ventricular myocardial scar”. In: *American Journal of Roentgenology* 193.5 (2009), pp. 381–388.

[94] Dana C. Peters et al. “Respiratory bellows-gated late gadolinium enhancement of the left atrium”. In: *Journal of Magnetic Resonance Imaging* 38.5 (2013), pp. 1210–1214.

[95] Jaime L. Shaw et al. “Left atrial late gadolinium enhancement with water-fat separation: The importance of phase-encoding order”. In: *Journal of Magnetic Resonance Imaging* 40.1 (2014), pp. 119–125.

[96] Yi Wang, Stephen J. Riederer, and Richard L. Ehman. “Respiratory Motion of the Heart: Kinematics and the Implications for the Spatial Resolution in Coronary Imaging”. In: *Magnetic Resonance in Medicine* 33.5 (1995), pp. 713–719.

[97] Li Feng et al. “XD-GRASP: Golden-angle radial MRI with reconstruction of extra motion-state dimensions using compressed sensing”. In: *Magnetic Resonance in Medicine* 75.2 (2016).

[98] Teresa Correia et al. “Optimized respiratory-resolved motion-compensated 3D Cartesian coronary MR angiography”. In: *Magnetic Resonance in Medicine* 80.6 (2018), pp. 2618–2629.

[99] Li Feng et al. “5D whole-heart sparse MRI”. In: *Magnetic Resonance in Medicine* 79.2 (2018), pp. 826–838.

[100] Davide Piccini et al. “Spiral phyllotaxis: The natural way to construct a 3D radial trajectory in MRI”. In: *Magnetic Resonance in Medicine* 66.4 (2011), pp. 1049–1056.

[101] Lorenzo Di Sopra et al. “An automated approach to fully self-gated free-running cardiac and respiratory motion-resolved 5D whole-heart MRI”. In: *Magnetic Resonance in Medicine* 82.6 (2019), pp. 2118–2132.

[102] Andrew D. Scott, Jennifer Keegan, and David N. Firmin. “Beat-to-beat respiratory motion correction with near 100% efficiency: A quantitative assessment using high-resolution coronary artery imaging”. In: *Magnetic Resonance Imaging* 29.4 (2011), pp. 568–578.

[103] Markus Henningsson et al. “Whole-heart coronary MR angiography with 2D self-navigated image reconstruction”. In: *Magnetic Resonance in Medicine* 67.2 (2012), pp. 437–445.

[104] Davide Piccini et al. “Respiratory self-navigation for whole-heart bright-blood coronary MRI: Methods for robust isolation and automatic segmentation of the blood pool”. In: *Magnetic Resonance in Medicine* 68.2 (Aug. 2012), pp. 571–579.

[105] Andrew P. Aitken et al. “100% Efficient three-dimensional coronary MR angiography with two-dimensional beat-to-beat translational and bin-to-bin affine motion correction”. In: *Magnetic Resonance in Medicine* 74.3 (2015), pp. 756–764.

[106] Konstantinos Bratis et al. “Image-navigated 3-dimensional late gadolinium enhancement cardiovascular magnetic resonance imaging: Feasibility and initial clinical results”. In: *Journal of Cardiovascular Magnetic Resonance* 19.1 (2017), pp. 1–9.

[107] Camila Munoz et al. “Motion corrected water/fat whole-heart coronary MR angiography with 100% respiratory efficiency”. In: *Magnetic Resonance in Medicine* February (2019), mrm.27732.

[108] Christoph Kolbitsch et al. “Highly efficient whole-heart imaging using radial phase encoding-phase ordering with automatic window selection”. In: *Magnetic Resonance in Medicine* 66.4 (2011), pp. 1008–1018.

[109] Johannes F.M. Schmidt et al. “Nonrigid retrospective respiratory motion correction in whole-heart coronary MRA”. In: *Magnetic Resonance in Medicine* 66.6 (Dec. 2011), pp. 1541–1549.

[110] Markus Henningsson and Rene M. Botnar. *Advanced respiratory motion compensation for coronary MR angiography*. May 2013. URL: `http://www.mdpi.com/1424-8220/13/6/6882`.

[111] Gastão Cruz et al. “Highly efficient nonrigid motion-corrected 3D whole-heart coronary vessel wall imaging”. In: *Magnetic Resonance in Medicine* 77.5 (2017), pp. 1894–1908.

[112] Aurélien Bustin et al. “3D whole-heart isotropic sub-millimeter resolution coronary magnetic resonance angiography with non-rigid motion-compensated PROST”. In: *Journal of Cardiovascular Magnetic Resonance* 5 (2020), pp. 1–16.

[113] Thomas Küstner et al. “Isotropic 3D Cartesian single breath-hold CINE MRI with multi-bin patch-based low-rank reconstruction”. In: *Magnetic Resonance in Medicine* 84.4 (2020), pp. 2018–2033.

[114] Camila Munoz et al. “Motion-corrected 3D whole-heart water-fat high-resolution late gadolinium enhancement cardiovascular magnetic resonance imaging”. In: *Journal of Cardiovascular Magnetic Resonance* 22.1 (2020), pp. 1–13.

[115] Aurélien Bustin et al. “High-dimensionality undersampled patch-based reconstruction (HD-PROST) for accelerated multi-contrast MRI”. In: *Magnetic Resonance in Medicine* 81.6 (2019), pp. 3705–3719.

[116] Haikun Qi et al. “Free-running 3D whole heart myocardial T1 mapping with isotropic spatial resolution”. In: *Magnetic Resonance in Medicine* 82.4 (2019), pp. 1331–1342.

[117] Anthony G Christodoulou et al. “Magnetic resonance multitasking for motion-resolved quantitative cardiovascular imaging”. In: *Nature Biomedical Engineering* (2018).

[118] Jaime L. Shaw et al. “Free-breathing, non-ECG, continuous myocardial T1 mapping with cardiovascular magnetic resonance multitasking”. In: *Magnetic Resonance in Medicine* 81.4 (2019), pp. 2450–2463.

[119] Yuchi Liu et al. “Cardiac Magnetic Resonance Fingerprinting: Technical Overview and Initial Results”. In: *JACC: Cardiovascular Imaging* 11.12 (2018), pp. 1837–1853.

[120] Olivier Jaubert et al. “Water–fat Dixon cardiac magnetic resonance fingerprinting”. In: *Magnetic Resonance in Medicine* 83.6 (2020), pp. 2107–2123.

[121] J.E. Bowsher et al. “Bayesian reconstruction and use of anatomical a priori information for emission tomography”. In: *IEEE Transactions on Medical Imaging* 15.5 (Oct. 1996), pp. 673–686.

[122] Matthias J. Ehrhardt et al. “Joint reconstruction of PET-MRI by exploiting structural similarity”. In: *Inverse Problems* 31.1 (Jan. 2015), p. 015001.

[123] Florian Knoll et al. “Joint MR-PET Reconstruction Using a Multi-Channel Image Regularizer”. In: *IEEE Transactions on Medical Imaging* 36.1 (Jan. 2017), pp. 1–16.

[124] Georg Schramm et al. “Evaluation of Parallel Level Sets and Bowsher’s Method as Segmentation-Free Anatomical Priors for Time-of-Flight PET Reconstruction”. In: *IEEE Transactions on Medical Imaging* 37.2 (Feb. 2018), pp. 590–603.

[125] Stefaan Vandenberghe and Paul K. Marsden. “PET-MRI: A review of challenges and solutions in the development of integrated multimodality imaging”. In: *Physics in Medicine and Biology* 60.4 (2015), R115–R154.

[126] Siemens Medical Solutions. “Biograph mCT: A wide bore and more”. In: (2018).

[127] Y. Petibon et al. “Towards coronary plaque imaging using simultaneous PET-MR: A simulation study”. In: *Physics in Medicine and Biology* 59.5 (2014), pp. 1203–1222.

[128] S. Furst et al. “Motion Correction Strategies for Integrated PET/MR”. In: *Journal of Nuclear Medicine* 56.2 (2015), pp. 261–269.

[129] Christopher M. Rank et al. “Respiratory motion compensation for simultaneous PET/MR based on highly undersampled MR data”. In: *Medical Physics* 43.12 (2016), pp. 6234–6245.

[130] Hadi Fayad et al. “4-Dimensional MRI and attenuation map generation in PET/MRI with 4-dimensional PET-derived deformation matrices: Study of feasibility for lung cancer applications”. In: *Journal of Nuclear Medicine* 58.5 (2017), pp. 833–839.

[131] Richard Manber et al. “Clinical impact of respiratory motion correction in simultaneous PET/MR, using a joint PET/MR predictive motion model”. In: *Journal of Nuclear Medicine* 59.9 (2018), pp. 1467–1473.

[132] Camila Munoz et al. *MR-Based Cardiac and Respiratory Motion-Compensation Techniques for PET-MR Imaging*. 2016.

[133] Richard Manber et al. “Joint PET-MR respiratory motion models for clinical PET motion correction”. In: *Physics in Medicine and Biology* 61.17 (2016), pp. 6515–6530.

[134] Camila Munoz et al. “Motion-corrected simultaneous cardiac positron emission tomography and coronary MR angiography with high acquisition efficiency”. In: *Magnetic Resonance in Medicine* 79.1 (Jan. 2017), pp. 339–350.

[135] Camila Munoz et al. “Respiratory- and cardiac motion-corrected simultaneous whole-heart PET and dual phase coronary MR angiography”. In: *Magnetic Resonance in Medicine* July (2018), pp. 1–14.

[136] Philip M Robson et al. “Correction of respiratory and cardiac motion in cardiac PET/MR using MR-based motion modeling”. In: *Physics in Medicine & Biology* 63.22 (Nov. 2018), p. 225011.

[137] Christoph Kolbitsch et al. “Cardiac and Respiratory Motion Correction for Simultaneous Cardiac PET/MR”. In: *Journal of Nuclear Medicine* 58.5 (May 2017), pp. 846–852.

[138] Thomas Küstner et al. “MR-based respiratory and cardiac motion correction for PET imaging”. In: *Medical Image Analysis* 42 (Dec. 2017), pp. 129–144.

[139] Christoph Kolbitsch et al. “Respiratory-resolved MR-based attenuation correction for motion-compensated cardiac PET-MR”. In: *Physics in Medicine & Biology* 63.13 (June 2018), p. 135008.

[140] Thibault Marin et al. “Motion correction for PET data using subspace-based real-time MR imaging in simultaneous PET/MR”. In: *Physics in Medicine & Biology* 65.23 (2020), p. 235022.

[141] Yoann Petibon et al. “Impact of motion and partial volume effects correction on PET myocardial perfusion imaging using simultaneous PET-MR”. In: *Physics in Medicine and Biology* 62.2 (2017), pp. 326–343.

[142] Mueez Aizaz et al. “PET/MRI of atherosclerosis”. In: *Cardiovascular Diagnosis and Therapy* 10.4 (2020), pp. 1120–1139.

[143] J. Cal-González et al. “Impact of motion compensation and partial volume correction for 18 F-NaF PET/CT imaging of coronary plaque”. In: *Physics in Medicine & Biology* 63.1 (Dec. 2017), p. 015005.

[144] Sadek A. Nehmeh and Yusuf E. Erdi. “Respiratory Motion in Positron Emission Tomography/Computed Tomography: A Review”. In: *Seminars in Nuclear Medicine* 38.3 (May 2008), pp. 167–176.

[145] Quansheng Xu, Kehong Yuan, and Datian Ye. “Respiratory motion blur identification and reduction in ungated thoracic PET imaging”. In: *Physics in Medicine and Biology* 56.14 (July 2011), pp. 4481–4498.

[146] Maxim Zaitsev, Julian Maclaren, and Michael Herbst. "Motion artifacts in MRI: A complex problem with many partial solutions". In: *Journal of Magnetic Resonance Imaging* 42.4 (2015), pp. 887–901.

[147] Beverley F. Holman et al. "The effect of respiratory induced density variations on non-TOF PET quantitation in the lung". In: *Physics in Medicine and Biology* 61.8 (2016), pp. 3148–3163.

[148] Camila Munoz et al. "MR-Based Cardiac and Respiratory Motion-Compensation Techniques for PET-MR Imaging". In: *PET Clinics* 11.2 (2016), pp. 179–191.

[149] Philip M Robson et al. "Correction of respiratory and cardiac motion in cardiac PET/MR using MR-based motion modeling". In: *Physics in Medicine & Biology* 63.22 (Nov. 2018), p. 225011.

[150] Tony Stoecker, Kaveh Vahedipour, and N Jon Shah. "HPC Simulation of Magnetic Resonance Imaging". In: *Advances in Parallel Computing* 15.May 2016 (2008), pp. 155–164.

[151] Lukas Wissmann et al. "MRXCAT: Realistic numerical phantoms for cardiovascular magnetic resonance". In: *Journal of Cardiovascular Magnetic Resonance* 16.1 (Dec. 2014), p. 63.

[152] Kelvin J. Layton et al. "Pulseq: A rapid and hardware-independent pulse sequence prototyping framework". In: *Magnetic Resonance in Medicine* 77.4 (Apr. 2017), pp. 1544–1552.

[153] Fang Liu et al. "Fast Realistic MRI Simulations Based on Generalized Multi-Pool Exchange Tissue Model". In: *IEEE Transactions on Medical Imaging* 36.2 (2017), pp. 527–537.

[154] Alexandre Fortin et al. "Flow MRI simulation in complex 3D geometries: Application to the cerebral venous network". In: *Magnetic Resonance in Medicine* 80.4 (Oct. 2018), pp. 1655–1665.

[155] Wei-Ching Lo et al. "Realistic 4D MRI abdominal phantom for the evaluation and comparison of acquisition and reconstruction techniques". In: *Magnetic Resonance in Medicine* 81.3 (Mar. 2019), pp. 1863–1875.

[156] Christos G Xanthis and Anthony H. Aletras. "coreMRI: A high-performance, publicly available MR simulation platform on the cloud". In: *PLOS ONE* 14.5 (May 2019). Ed. by Xi Chen, e0216594.

[157] C. Tsoumpas et al. "Fast generation of 4D PET-MR data from real dynamic MR acquisitions". In: *Physics in Medicine and Biology* 56.20 (2011), pp. 6597–6613.

[158] Irene Polycarpou, Georgios Soultanidis, and Charalampos Tsoumpas. "Synthesis of Realistic Simultaneous Positron Emission Tomography and Magnetic Resonance Imaging Data". In: *IEEE Transactions on Medical Imaging* 37.3 (Mar. 2018), pp. 703–711.

[159] W. P. Segars et al. "4D XCAT phantom for multimodality imaging research". In: *Medical Physics* 37.9 (2010), pp. 4902–4915.

[160] W Paul Segars et al. “Incorporation of the Living Heart Model Into the 4-D XCAT Phantom for Cardiac Imaging Research”. In: *IEEE Transactions on Radiation and Plasma Medical Sciences* 3.1 (2019), pp. 54–60.

[161] Souheil J. Inati et al. “ISMRM Raw data format: A proposed standard for MRI raw datasets”. In: *Magnetic Resonance in Medicine* 77.1 (Jan. 2017), pp. 411–421.

[162] A. Todd-Pokropek, T. D. Cradduck, and F. Deconinck. “A file format for the exchange of nuclear medicine image data”. In: *Nuclear Medicine Communications* 13.9 (Sept. 1992), pp. 673–699.

[163] Martin Uecker et al. “Berkeley Advanced Reconstruction Toolbox”. In: *Proceedings of the International Society for Magnetic Resonance in Medicine*. 2015.

[164] Kris Thielemans et al. “STIR: software for tomographic image reconstruction release 2”. In: *Physics in Medicine and Biology* 57.4 (Feb. 2012), pp. 867–883.

[165] Michael Schacht Hansen and Thomas Sangild Sørensen. “Gadgetron: An open source framework for medical image reconstruction”. In: *Magnetic Resonance in Medicine* 69.6 (June 2013), pp. 1768–1776.

[166] Evgueni Ovtchinnikov et al. “SIRF: Synergistic Image Reconstruction Framework”. In: *Computer Physics Communications* 249 (Apr. 2020), p. 107087.

[167] Robert W. Brown et al. *Magnetic Resonance Imaging*. Chichester, UK: John Wiley & Sons Ltd, Apr. 2014.

[168] Parisa Khateri et al. “Implementation of cylindrical PET scanners with block detector geometry in STIR”. In: *EJNMMI Physics* 6.1 (Dec. 2019), p. 15.

[169] Marc Modat et al. “Fast free-form deformation using graphics processing units”. In: *Computer Methods and Programs in Biomedicine* 98.3 (June 2010), pp. 278–284.

[170] Claudia Prieto et al. “3D undersampled golden-radial phase encoding for DCE-MRA using inherently regularized iterative SENSE”. In: *Magnetic Resonance in Medicine* 64.2 (2010), pp. 514–526.

[171] J.A. Schnabel et al. “Validation of nonrigid image registration using finite-element methods: application to breast MR images”. In: *IEEE Transactions on Medical Imaging* 22.2 (Feb. 2003), pp. 238–247.

[172] R. Boubertakh et al. “Whole-heart imaging using undersampled radial phase encoding (RPE) and iterative sensitivity encoding (SENSE) reconstruction”. In: *Magnetic Resonance in Medicine* 62.5 (Nov. 2009), pp. 1331–1337.

[173] Matteo Ippoliti et al. “3D nonrigid motion correction for quantitative assessment of hepatic lesions in DCE-MRI”. In: *Magnetic Resonance in Medicine* 82.5 (Nov. 2019), pp. 1753–1766.

[174] Martin Rohrer et al. “Comparison of Magnetic Properties of MRI Contrast Media Solutions at Different Magnetic Field Strengths”. In: *Investigative Radiology* 40.11 (Nov. 2005), pp. 715–724.

[175] F. Lamare et al. “Respiratory motion correction for PET oncology applications using affine transformation of list mode data”. In: *Physics in Medicine and Biology* 52.1 (2007), pp. 121–140.

[176] Torsten Rohlfing. “Image Similarity and Tissue Overlaps as Surrogates for Image Registration Accuracy: Widely Used but Unreliable”. In: *IEEE Transactions on Medical Imaging* 31.2 (Feb. 2012), pp. 153–163.

[177] K. Rohr, M. Fornefett, and H.S Stiehl. “Spline-based elastic image registration: integration of landmark errors and orientation attributes”. In: *Computer Vision and Image Understanding* 90.2 (May 2003), pp. 153–168.

[178] Alexis Roche et al. “Rigid registration of 3-D ultrasound with MR images: a new approach combining intensity and gradient information”. In: *IEEE Transactions on Medical Imaging* 20.10 (2001), pp. 1038–1049.

[179] Richard Castillo et al. “A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets”. In: *Physics in Medicine and Biology* 54.7 (Apr. 2009), pp. 1849–1870.

[180] R. Boubertakh et al. “Whole-heart imaging using undersampled radial phase encoding (RPE) and iterative sensitivity encoding (SENSE) reconstruction”. In: *Magnetic Resonance in Medicine* 62.5 (2009), pp. 1331–1337.

[181] Stefanie Winkelmann et al. “An Optimal Radial Profile Order Based on the Golden Ratio for Time-Resolved MRI”. In: *IEEE Transactions on Medical Imaging* 26.1 (Jan. 2007), pp. 68–76.

[182] Harvey E Cline and Thomas R Anthony. “Uniform K-Space Sampling with an Interleaved Fibonacci Spiral Acquisition”. In: 1995, p. 12309.

[183] Urs Gamper, Peter Boesiger, and Sebastian Kozerke. “Compressed sensing in dynamic MRI”. In: *Magnetic Resonance in Medicine* 59.2 (2008), pp. 365–373.

[184] Huanzhou Yu et al. “Multiecho water-fat separation and simultaneous R*2 estimation with multifrequency fat spectrum modeling”. In: *Magnetic Resonance in Medicine* 60.5 (2008), pp. 1122–1134.

[185] Kai Tobias Block, Martin Uecker, and Jens Frahm. “Undersampled Radial MRI with Multiple Coils . Iterative Image Reconstruction Using a Total Variation Constraint”. In: 1098.January (2007), pp. 1086–1098.

[186] Julia A. Schnabel et al. “A Generic Framework for Non-rigid Registration Based on Non-uniform Multi-level Free-Form Deformations”. In: 2001, pp. 573–581.

[187] René M. Botnar et al. “Improved coronary artery definition with T2-weighted, free-breathing, three-dimensional coronary MRA”. In: *Circulation* 99.24 (1999), pp. 3139–3148.

[188] Alex Etienne et al. “"Soap-Bubble" visualization and quantitative analysis of 3D coronary magnetic resonance angiograms”. In: *Magnetic Resonance in Medicine* 48.4 (2002), pp. 658–666.

[189] Gastao Cruz et al. “Accelerated motion corrected three-dimensional abdominal MRI using total variation regularized SENSE reconstruction”. In: *Magnetic Resonance in Medicine* 75.4 (2016), pp. 1484–1498.

[190] Christoph Kolbitsch et al. “Fully integrated 3D high-resolution multicontrast abdominal PET-MR with high scan efficiency”. In: *Magnetic Resonance in Medicine* 79.2 (2018), pp. 900–911.

[191] Johannes Mayer et al. “Flexible numerical simulation framework for dynamic PET-MR data”. In: *Physics in Medicine & Biology* 65.14 (July 2020), p. 145003.

[192] Allen P. Burke et al. “Arrhythmogenic right ventricular cardiomyopathy and fatty replacement of the right ventricular myocardium: Are they different diseases?” In: *Circulation* 97.16 (1998), pp. 1571–1580.

[193] James W. Goldfarb. “Fat-water separated delayed hyperenhanced myocardial infarct imaging”. In: *Magnetic Resonance in Medicine* 60.3 (2008), pp. 503–509.

[194] Peter Kellman et al. “Multiecho dixon fat and water separation method for detecting fibrofatty infiltration in the myocardium”. In: *Magnetic Resonance in Medicine* 61.1 (2009), pp. 215–221.

[195] Diego Hernando et al. “Robust water/fat separation in the presence of large field inhomogeneities using a graph cut algorithm”. In: *Magnetic Resonance in Medicine* 63.1 (2010), pp. 79–90.

[196] Holger Eggers and Peter Boernert. “Chemical shift encoding-based water-fat separation methods”. In: *Journal of Magnetic Resonance Imaging* 40.2 (Aug. 2014), pp. 251–268.

[197] Peter Börnert et al. “Water/fat-resolved whole-heart Dixon coronary MRA: An initial comparison”. In: *Magnetic Resonance in Medicine* 71.1 (2014), pp. 156–163.

[198] James W. Goldfarb and Meir Shinnar. “Free-breathing delayed hyperenhanced imaging of the myocardium: A clinical application of real-time navigator echo imaging”. In: *Journal of Magnetic Resonance Imaging* 24.1 (2006), pp. 66–71.

[199] Jennifer Keegan et al. “Non-model-based correction of respiratory motion using beat-to-beat 3D spiral fat-selective imaging”. In: *Journal of Magnetic Resonance Imaging* 26.3 (2007), pp. 624–629.

[200] Andrew D. Scott, Jennifer Keegan, and David N. Firmin. “High-resolution 3D coronary vessel wall imaging with near 100% respiratory efficiency using epicardial fat tracking: Reproducibility and comparison with standard methods”. In: *Journal of Magnetic Resonance Imaging* 33.1 (2011), pp. 77–86.

[201] Yabo Fu et al. “Deep learning in medical image registration: a review”. In: *Physics in Medicine & Biology* 65.20 (Oct. 2020), 20TR01.

[202] Grant Haskins, Uwe Kruger, and Pingkun Yan. “Deep learning in medical image registration: a survey”. In: *Machine Vision and Applications* 31.1 (2020), pp. 1–18.

[203] Marcus Brehm et al. “Artifact-resistant motion estimation with a patient-specific artifact model for motion-compensated cone-beam CT”. In: *Medical Physics* (2013).

[204] Hadi Wiputra et al. “Cardiac motion estimation from medical images: a regularisation framework applied on pairwise image registration displacement fields”. In: *Scientific Reports* 10.1 (Dec. 2020), p. 18510.

[205] Andreas Kofler et al. “Spatio-Temporal Deep Learning-Based Undersampling Artefact Reduction for 2D Radial Cine MRI with Limited Training Data”. In: *IEEE Transactions on Medical Imaging* 39.3 (2020), pp. 703–717.

[206] Toshiro Kitagawa et al. “18F-sodium fluoride positron emission tomography for molecular imaging of coronary atherosclerosis based on computed tomography analysis”. In: *Atherosclerosis* 263 (Aug. 2017), pp. 385–392.

[207] Mikaela L. McKenney-Drake et al. “18 F-NaF PET Imaging of Early Coronary Artery Calcification”. In: *JACC: Cardiovascular Imaging* 9.5 (May 2016), pp. 627–628.

[208] Philip M. Robson et al. “Coronary Artery PET/MR Imaging: Feasibility, Limitations, and Solutions”. In: *JACC: Cardiovascular Imaging* 10.10 (2017), pp. 1103–1112.

[209] Mika Teräs et al. “Dual-gated cardiac PET–Clinical feasibility study”. In: *European Journal of Nuclear Medicine and Molecular Imaging* 37.3 (Mar. 2010), pp. 505–516.

[210] Nicolas Karakatsanis et al. “MR-based attenuation correction in cardiovascular PET/MR imaging: challenges and practical solutions for cardiorespiratory motion and tissue class segmentation”. In: *Journal of Nuclear Medicine* 57 (2016), pp. 452–452.

[211] Hua Ai and Tinsu Pan. “Feasibility of using respiration-averaged MR images for attenuation correction of cardiac PET/MR imaging”. In: *Journal of Applied Clinical Medical Physics* 16.4 (July 2015), pp. 311–321.

[212] Martin Lyngby Lassen et al. “Triple-gated motion and blood pool clearance corrections improve reproducibility of coronary 18F-NaF PET”. In: *European Journal of Nuclear Medicine and Molecular Imaging* 46.12 (Nov. 2019), pp. 2610–2620.

[213] Jacobo Cal-Gonzalez et al. “Partial volume correction for improved PET quantification in 18F-NaF imaging of atherosclerotic plaques”. In: *Journal of Nuclear Cardiology* 25.5 (2018), pp. 1742–1756.

[214] Philip M. Robson et al. “Correction of respiratory and cardiac motion in cardiac PET/MR using MR-based motion modeling”. In: *Physics in Medicine and Biology* (2018).

[215] Michael Lustig, D L Donoho, and J M Santos. “Compressed sensing MRI”. In: *Signal Processing* March 2008 (2008), pp. 72–82.

[216] David Arthur and Sergei Vassilvitskii. “k-means ++ : The Advantages of Careful Seeding”. In: *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. New Orleans, Louisiana: Society for Industrial and Applied Mathematics, 2007, pp. 1027–1035.

[217] Nobuyuki Otsu. “A Threshold Selection Method from Gray-Level Histograms”. In: *IEEE Trans. Syst. Man Cybern* 9.1 (1979), pp. 62–66.

[218] Chi Liu et al. “The impact of respiratory motion on tumor quantification and delineation in static PET/CT imaging”. In: *Physics in Medicine and Biology* 54.24 (Dec. 2009), pp. 7345–7362.

[219] Richard Manber et al. “Practical PET Respiratory Motion Correction in Clinical PET/MR”. In: 56.6 (2015), pp. 890–897.

[220] C. Tsoumpas et al. “Evaluation of the Single Scatter Simulation Algorithm Implemented in the STIR Library”. In: *IEEE Symposium Conference Record Nuclear Science 2004.* Vol. 6. IEEE, pp. 3361–3365.

[221] Quanzheng Li et al. “Maximum a posteriori reconstruction of Biograph mMR scanner using point spread function”. In: *Journal of Nuclear Medicine* 53.supplement 1 (May 2012), pp. 2339–2339.

[222] Amir Beck and Marc Teboulle. “A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems”. In: *SIAM Journal on Imaging Sciences* 2.1 (Jan. 2009), pp. 183–202.

[223] Jeffrey A Fessler, Senior Member, and Bradley P Sutton. “Nonuniform Fast Fourier Transforms Using Min-Max Interpolation”. In: 51.2 (2003), pp. 560–574.

[224] Matthias J. Ehrhardt et al. “Faster PET reconstruction with a stochastic primal-dual hybrid gradient method”. In: *Wavelets and Sparsity XVII*. Ed. by Yue M. Lu, Manos Papadakis, and Dimitri Van De Ville. SPIE, Aug. 2017, p. 58.

[225] Maurizio Conti. “Why is TOF PET reconstruction a more robust method in the presence of inconsistent data?” In: *Physics in Medicine and Biology* 56.1 (2011), pp. 155–168.

[226] Alexander M. Grant et al. “NEMA NU 2-2012 performance studies for the SiPM-based ToF-PET component of the GE SIGNA PET/MR system”. In: *Medical Physics* 43.5 (2016), pp. 2334–2343.

[227] Edwin E.G.W. ter Voert et al. “Clinical evaluation of TOF versus non-TOF on PET artifacts in simultaneous PET/MR: a dual centre experience”. In: *European Journal of Nuclear Medicine and Molecular Imaging* 44.7 (2017), pp. 1223–1233.

[228] Arnold M. Weissler, Willard S. Harris, and Clyde D. Schoenfeld. “Systolic Time Intervals in Heart Failure in Man”. In: *Circulation* 37.2 (Feb. 1968), pp. 149–159.

[229] Richard P Lewis et al. “A critical review of the systolic time intervals.” In: *Circulation* 56.2 (Aug. 1977), pp. 146–158.

[230] Yabo Fu et al. “Deep learning in medical image registration: A review”. In: *arXiv* (2019).