Accelerating Cardiac Diffusion MRI: Concepts and Applications



Shokoufeh Golshani

University of Leeds

School of Computing and School of Medicine

Submitted in accordance with the requirements for the degree of

Doctor of Philosophy

March, 2024

Intellectual Property Statement

The candidate confirms that the work submitted is her own and that appropriate credit has been given where reference has been made to the work of others.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

The right of Shokoufeh Golshani to be identified as Author of this work has been asserted by Shokoufeh Golshani in accordance with the Copyright, Designs and Patents Act 1988.

© 2024 The University of Leeds and Shokoufeh Golshani.

Acknowledgements

The completion of this thesis has been made possible through the invaluable support, encouragement, and guidance of many individuals, and I would like to take this opportunity to extend my sincere gratitude and appreciation to each of them.

First and foremost, I wish to express my deepest gratitude to my supervisors, Professor Jurgen Schneider and Professor Alejandro Frangi, for their unwavering support, constructive discussions, and motivation throughout my PhD studies. Their expertise and mentorship have been instrumental in shaping my research endeavours, and I am incredibly grateful for the trust they placed in me to explore and pursue my projects. Their meticulous scientific approach and attention to detail have significantly enhanced my skills as a scientist and writer. Working with them has been a tremendous learning experience, not only in my PhD work but also in shaping my future career. I am particularly thankful to Professor Jurgen Schneider for his continuous support, invaluable feedback during the writing of my thesis, and extensive assistance in manuscript preparation, as well as for providing access to an excellent research facility.

I am also grateful to Dr. Irvin Teh for his valuable assistance and teaching in data acquisition, sharing his expertise in DTI data analysis, and providing insightful feedback on manuscripts. Additionally, I extend my thanks to Dr. Nishant Ravikumar for his comments and professional advice during my candidature.

Further, I wish to acknowledge the Marie Curie ITN Scholarship for their generous financial support, which has given me the privilege of being part of the B-Q Minded project and has granted me full academic freedom throughout the first three years of my research journey. I am thankful to my colleagues at the B-Q Minded project for fostering a positive and friendly atmosphere and engaging in stimulating discussions during our annual and scientific meetings. My special thanks go to the CISTIB team for their camaraderie and support.

Last but not least, my heartfelt thanks go to my family and friends for their endless spiritual support throughout my PhD journey and beyond. Their encouragement has been a constant source of strength for me. In par-

ticular, I am profoundly thankful to my husband, Yusuf ♡, for his patience and encouraging words during moments of stress and for his unwavering support, especially during the intense final stages of thesis writing.

It is with profound pride and gratitude that I dedicate this thesis to my beloved **Mom** and **Dad** for their unconditional love.

Abstract

Heart failure impacts millions of patients worldwide, with a fifty percent mortality rate within five years of diagnosis. While existing clinical diagnostic tools like magnetic resonance imaging (MRI) provide insights into cardiac morphology and global function, understanding the mechanisms driving long-term cardiac structural remodelling and functional impairment remains limited. Consequently, there has been minimal progress in early diagnosis and personalised therapy management to enhance patient outcomes.

The emergence of diffusion-weighted (DW) MRI as a unique tool for non-invasive quantification of tissue microstructure by measuring the displacement of water molecules within the tissue holds significant promise for enhancing our knowledge of structure-function relationship. Cardiac diffusion tensor imaging (cDTI) has proven to be a valuable technique for assessing the three-dimensional (3D) microstructure of the heart. However, its practical and clinical utility is limited by several inherent challenges, including the long acquisition times and low signal-to-noise ratio (SNR). Given the intrinsic trade-off in the MRI technique between the scan time, spatial resolution and SNR, current in vivo applications of cDTI have primarily sacrificed the spatial resolution and coverage to minimise the scan time and motion-related artefacts.

Nevertheless, achieving high spatial resolution and whole-heart coverage in cDTI improves the depiction of myocardial microstructure and facilitates the diagnosis of various early-stage cardiac pathologies characterised by sub-millimetric lesions, such as infarction. Traditional acquisition and reconstruction methods are inadequate in addressing the extended scan times required for high-resolution imaging. This thesis aims to address this challenge by developing fast and efficient 3D cDTI methods, with a particular focus on new sampling designs and dedicated reconstruction techniques, to enhance quantitative understanding of the heart microstructure. The methodology adopted is based on the theory of compressed sensing (CS) and leverages on the flexibility offered by the spiral acquisitions in providing shorter scan times. We first proposed a new parametrisation for the sampling density function in spiral acquisitions that is well adapted to the

diffusion-weighted signal in the heart and suitable for CS reconstructions. We then conducted a comprehensive study to optimise the proposed variable density spiral (VDS) trajectory. The trajectory conformed to the limitations of the MRI gradient hardware, ensuring its practical implementation.

Next, we proposed two extensions of the designed optimal VDS trajectory to three-dimensions based on z-variable stack of spirals and full 3D conical sampling schemes combined with an efficient 3D CS reconstruction. The proposed approach was retrospectively evaluated using a wide range of cDTI images and practical setups. The CS reconstruction utilised the group sparsity property among DW images to reduce the measurement requirements further and enhance the reconstruction performance. Correction for the phase inconsistency errors is considered in the reconstruction. Well-balanced sampling of low and high spatial frequencies in the proposed VDS design and exploiting correlated information among DW images during the reconstruction led to a higher potential for acceleration, preserved image details and diffusion features, and model robustness compared to the classical undersampling and CS reconstruction methods. Significant potential reduction in the scan time with root mean square error (RMSE) of less than 5% in diffusion parameter maps was demonstrated using around 5% of the k-space data.

Finally, we proposed a novel spiral sampling technique, referred to as SPIK (*SPIral with Keyhole*), and a sparse blind dictionary learning reconstruction to allow for further accelerations in a multislice approach. The method exploits the strong spatial and diffusion-directional correlations between DW images in both the sampling and reconstruction steps. The reconstruction technique, termed sparse blind compressed sensing (s-BCS), models the spatio-directional DW image series at each pixel as a sparse linear combination of bases in a dictionary. The dictionary and the sparse coefficients were then jointly estimated from undersampled data. Quantitative analysis demonstrated the superior performance of the proposed method compared to the low-rank and conventional CS models. Furthermore, the method exhibited robustness to patient- and sequence-specific motion, resulting in minimal degradation of parameter maps in the simu-

lated free-breathing states. Clinically relevant cardiac microstructural parameter maps derived from ex vivo cDTI data and realistic numerical in vivo human heart phantoms demonstrated comparable quality to fully sampled reference images within an acceptable error range.

In conclusion, significant theoretical advances have been achieved in accelerating 3D cDTI, offering potential improvements in quantification of cardiac microstructure within physiologically acceptable scan times and while maintaining image quality. However, experimental validation of these advancements is left for future studies. These advancements could pave the way for new diagnostic tools to deepen our understanding of the causes and effects of heart failure and other cardiac conditions. Moreover, these improvements may provide essential information for personalising future in-silico cardiac analyses and interventions.

Contents

1 Introduction **1**

1.1 Significance . . . 2
1.2 Structure of the Heart . . . 3
1.2.1 Gross Morphology . . . 3
1.2.2 Microstructure . . . 6
1.2.3 Laminar Structure . . . 6
1.3 Heart Structure and Function Relationship . . . 7
1.3.1 Heart Pathologies . . . 9
1.4 Myocardium Microstructure Characterisation . . . 11
1.5 Motivation . . . 14
1.6 Research Aim and Objectives . . . 16
1.7 Thesis Outline . . . 17

2 Principles of MRI and Diffusion MRI **19**

2.1 Nuclear Magnetic Resonance Physics . . . 20
2.1.1 Precession . . . 20
2.1.2 Excitation . . . 22
2.1.3 Relaxation . . . 22
2.1.4 NMR signal . . . 23
2.2 MR Imaging . . . 24
2.2.1 Signal Localisation . . . 24
2.2.2 MRI Pulse Sequence . . . 27
2.2.3 Spatial Encoding and k-Space . . . 29
2.2.4 k-Space Trajectories . . . 31

2.2.5 Sampling of k-Space Trajectories . . . 32
2.2.6 Image Reconstruction . . . 36
2.3 Diffusion MRI . . . 39
2.3.1 Diffusion Principles . . . 39
2.3.2 Diffusion-weighted MR Signal . . . 41
2.3.3 Diffusion Tensor Imaging (DTI) . . . 44
2.4 Summary . . . 50

3 State-of-the-Art Techniques in Cardiac DTI 51
3.1 Methodological Challenges in cDTI . . . 52
3.1.1 Motion Artefacts . . . 52
3.1.2 Low Signal-to-Noise Ratio . . . 53
3.1.3 Long Scan Time . . . 53
3.1.4 Limited Spatial Coverage and Resolution . . . 54
3.1.5 Gradient Field Imperfections . . . 54
3.2 Solutions . . . 55
3.2.1 Diffusion Encoding Sequence Development . . . 55
3.2.2 Fast Data Acquisition . . . 59
3.2.3 Incomplete k-Space Sampling Techniques . . . 62
3.3 Compressed Sensing Theory . . . 65
3.3.1 Sensing . . . 66
3.3.2 Recovery . . . 67
3.4 Compressed Sensing in MRI . . . 67
3.4.1 Compressed Sending in cDTI . . . 68
3.5 Summary . . . 71

4 Optimal Design of the Sampling Density Function for Spiral-CS cDTI 73
4.1 Introduction . . . 74
4.2 Sampling Design for CS cDTI . . . 75
4.2.1 A Priori Knowledge . . . 75
4.2.2 Variable Density Sampling . . . 75
4.2.3 CS Limitations in MRI . . . 76
4.3 Empirical and Quantitative Study of Optimal VDS Design . . . 80
4.3.1 Design Considerations . . . 80

4.3.2 Spiral Trajectory Design . . . 81
4.3.3 A Generic Method . . . 84
4.3.4 Experimental Validation with Real cDTI Data . . . 88
4.3.5 The Studied Pipeline . . . 90
4.3.6 Computational Setup . . . 92
4.3.7 Image Quality Metrics . . . 92
4.4 Results . . . 93
4.4.1 Effect of the Sampling Density Function . . . 93
4.4.2 Effect of the Trajectory Configuration . . . 97
4.4.3 Effect of the DW Direction . . . 100
4.4.4 Performance of the VDS Trajectory . . . 101
4.5 Discussion . . . 103
4.6 Summary . . . 107

5 Efficient High-resolution 3D Spiral-CS cDTI 109
5.1 Introduction . . . 110
5.2 Materials and Methods . . . 110
5.2.1 3D Imaging . . . 110
5.2.2 Fully 3D VDS: VD-Cones Sampling . . . 113
5.2.3 $k-q$ Ordering Scheme . . . 114
5.2.4 Image Reconstruction with Phase Error Compensation . . . 115
5.2.5 Optimisation Algorithm . . . 119
5.2.6 Parameter Setting . . . 120
5.3 Experiment Design . . . 121
5.3.1 Data Acquisition . . . 122
5.3.2 Sampling Schemes Comparison . . . 123
5.3.3 Image Reconstruction . . . 124
5.3.4 Quantitative Analysis . . . 125
5.4 Results . . . 126
5.5 Discussion . . . 133
5.5.1 Scan Time Reduction . . . 136
5.6 Summary . . . 137

6 Accelerated cDTI with Spiral Keyhole (SPIK) and Blind Dictionary Learning **138**

6.1 Introduction . . . 139
6.2 Materials and Methods . . . 139
6.2.1 Spiral Keyhole Imaging . . . 139
6.2.2 Image Reconstruction . . . 142
6.2.3 Diffusion-weighted Signal Representation . . . 143
6.2.4 Reconstruction Formulation . . . 144
6.2.5 Optimisation Algorithm . . . 145
6.3 Experiments . . . 148
6.3.1 Datasets . . . 148
6.3.2 Sampling Schemes . . . 149
6.3.3 Comparison of Reconstruction Methods . . . 150
6.3.4 Analysis Methods . . . 151
6.4 Results . . . 151
6.4.1 Influence of *Keyhole* Radius . . . 151
6.4.2 Algorithm Convergence . . . 152
6.4.3 Comparison of Reconstruction Schemes . . . 153
6.4.4 Comparison of Sampling Schemes . . . 157
6.4.5 Dictionary Size . . . 158
6.4.6 Comparison on the MRXCAT cDTI phantom . . . 158
6.5 Discussion . . . 164
6.6 Summary . . . 168

7 Conclusion And Future Work **169**

7.1 Contributions . . . 170
7.2 Limitations and Future Work . . . 173

A Spiral Data Simulation (Simulation Principles) **176**

References **179**

List of Figures

1.1 Heart anatomical structure 4
1.2 Schematic drawing of the trilayered structure of the heart wall 5
1.3 The correlation between left ventricular pressure and the electrocardiogram (ECG) 6
1.4 Basic cardiac segments and planes 7
1.5 Microstructure of the myocardium 7
1.6 A schematic diagram of myocardial laminar architecture 8
1.7 Cardiac structure remodelling in various pathological conditions 10
1.8 Torrent-Guasp myocardial band model 11
1.9 Confocal imaging of cardiomyocytes 12
1.10 Representative example of 3D DTI tractography of cardiomyocyte orientations 13
1.11 Inherent trade-off in MRI 15

2.1 Proton spin and precession 21
2.2 RF excitation and signal detection 23
2.3 Plot illustrating relaxation times 24
2.4 Slice selection 26
2.5 Multiple slice selection 27
2.6 Phase encoding 27
2.7 Frequency encoding 28
2.8 Spin-echo sequence 29
2.9 Correspondence between k-space (left) and image space (right) via the Fourier transform (FT) 31

2.10 Examples of 2D k-space trajectories . . . 33
2.11 Illustration of the Cartesian and non-Cartesian sampling schemes . . . 35
2.12 Basic regridding idea. . . . 38
2.13 A schematic of the molecular diffusion trajectory of a single water molecule 39
2.14 Complex diffusion in a tissue. . . . 41
2.15 Schematic illustration of the pulsed-gradient spin-echo (PGSE) sequence introduced by Stejskal and Tanner . . . 42
2.16 Schematic representation of dephasing for spins in a single voxel under the influence of a standard Stejskal-Tanner sequence . . . 43
2.17 Schematic representation of the diffusion tensor . . . 46
2.18 Schematic diagram of the correspondence of the diffusion tensor eigenvectors to the cardiomyocyte and sheetlets orientations . . . 48
2.19 Helix angle definition . . . 49

3.1 Schematic diagrams of common DW sequences used in the heart . . . 56
3.2 Flowchart describing the off-resonance artefacts correction in EPI. . . . 61
3.3 Schematic of multi-shot EPI and spiral imaging . . . 63
3.4 k-Space undersampling . . . 64

4.1 The b_0 and diffusion-weighted images of an ex vivo human heart (a-c) and their corresponding k-space image (d-f) . . . 76
4.2 Illustration of the impact of nulling different regions of k-space on a reconstructed cardiac image . . . 77
4.3 A human heart MRI image (a) and its wavelet representation (b) . . . 78
4.4 Proposed generic sampling density function . . . 84
4.5 Resulting k-space sampling trajectory . . . 86
4.6 Flowchart illustrating the proposed pipeline . . . 87
4.7 Examples of the studied DW images of a human heart . . . 89
4.8 Examples of the studied sampling density functions . . . 91
4.9 Performance analysis of all the examined sampling density functions in the proposed pipeline . . . 95
4.10 Quantitative analysis of all the examined sampling density functions in the proposed pipeline . . . 96

4.11 Illustration of the residual artefacts resulting from four arbitrary examined density functions in the proposed pipeline 97
4.12 SSIM scores plotted for all the examined density functions in the proposed pipeline 97
4.13 Quantitative analysis of the combination of sampling density function parameters and the spiral readout time 98
4.14 Quantitative analysis of the combination of the spiral readout time and the number of shots for a fixed sampling density function 99
4.15 SSIM results from three density functions in all examined readout durations and number of shots 100
4.16 SSIM scores for eight randomly selected sampling density functions . . . 101
4.17 Performance comparison of uniform density and variable density spiral trajectories 102
4.18 3D and 2D views of the point spread function of the UDS and VDS trajectories 103
4.19 PSF cross-section 104
4.20 The resulting VDS trajectory 104
4.21 Gradients and slew rates waveforms in the resulting VDS trajectory . . 105

5.1 Different Extensions of 2D spiral trajectories to 3D 112
5.2 Sampling patterns in the $k_y - k_z$ plane 113
5.3 Sampling patters in the $k_x - k_y$ plane 114
5.4 3D Cones sampling 115
5.5 Proposed incoherent $k - q$ undersampling pattern 116
5.6 Illustration of $k - q$ sampling schemes used for comparison 124
5.7 Reconstructed long-axis mid-ventricular images with the BJS method using various undersampling schemes 126
5.8 Reconstructed FA and MD maps with the BJS method using various undersampling schemes 127
5.9 Performance comparison of the SNAILS and the proposed z-variable VD-SOS schemes 128
5.10 RMSE of FA and MD maps from different reconstruction methods . . . 129
5.11 Reconstruction performance of the BJS method and its adaptive version against various ISNR values 130

5.12 Reconstruction performance of the BJS method and its adaptive version against reduced number of diffusion directions . . . 131
5.13 RMSE of FA, MD, and HA maps comparing the proposed z-variable VD SOS and cones sampling schemes . . . 131
5.14 Reconstructed DTI parameter maps . . . 132
5.15 Performance of phase correction in the adaptive BJS reconstruction . . . 133

6.1 Schematic representation of the SPIK acquisition scheme . . . 140
6.2 The proposed SPIK sampling . . . 141
6.3 Schematic representation of the k-space trajectory of three-shot SPIK . . . 142
6.4 The MRXCAT DW images used to validate the proposed approach . . . 149
6.5 Representative free-breathing MRXCAT DW images . . . 150
6.6 Influence of *keyhole* radius (α) . . . 152
6.7 Convergence of the proposed continuation scheme as a function of β_1 and β_2 . . . 153
6.8 Performance Comparison of different reconstruction methods . . . 154
6.9 Intensity profiles of different reconstruction methods . . . 155
6.10 SER scores as a function of number of iterations comparing the performance of different reconstruction methods . . . 155
6.11 FA maps obtained from different reconstruction methods . . . 156
6.12 MD maps obtained from different reconstruction methods . . . 156
6.13 Comparison of the SPIK sampling and the optimised multi-shot spiral sampling . . . 157
6.14 Evaluation of the performance of the proposed s-BCS method using different numbers of dictionary atoms . . . 159
6.15 Choice of regularisation parameters τ_1 and τ_2 . . . 159
6.16 Reconstructed MRXCAT cDTI data simulated using the proposed sparse BCS (case I) . . . 160
6.17 Reconstructed MRXCAT cDTI data using the proposed sparse BCS (case II) reconstruction . . . 161
6.18 Reconstruction performance of the proposed approach in DTI parameters using noisy free-breathing MRXCAT cDTI data . . . 163
6.19 Reconstruction performance of the proposed approach against the direct NUFT reconstruction using free-breathing MRXCAT cDTI data . . . 163

6.20 Reconstruction performance of the proposed approach in DTI maps using free-breathing MRXCAT cDTI data 164
6.21 Reconstruction performance of the proposed approach in DTI maps using noisy free-breathing MRXCAT cDTI data 164

List of Tables

4.1 Notations used in the proposed sampling function parametrisation . . . 85
4.2 Correspondence between SSIM index and Mean Opinion Scores (MOS) . 93
4.3 Concluded optimal variables for the proposed VDS design 106

5.1 Regularisation parameters at different undersampling rates 121
5.2 Quantitative comparison of undersampling schemes 127
5.3 Performance comparison at a high undersampling factor 128

6.1 Parameters used in MRXCAT-CDTI simulations 148
6.2 Quantitative comparison of the stack of SPIK sampling and the optimised multi-shot spiral sampling 158
6.3 Quantitative analysis of the proposed approach using breath-hold MRXCAT cDTI data 162
6.4 Quantitative analysis of the proposed approach in DTI parameter maps using breath-hold MRXCAT cDTI data 162

Abbreviations

General

m	meter
mm	millimeter
μm	micrometer
T	Tesla
mT	millitesla
s	second
ms	millisecond
min	minute
2D	two-dimensional
3D	three-dimensional
rad	radian
MHz	megahertz

Technical

CVD	cardiovascular diseases
RV	right ventricle
LV	left ventricle
ECG	electrocardiogram
DCM	dilated cardiomyopathy
HCM	hypertrophic cardiomyopathy
MRI	magnetic resonance imaging
dMRI	diffusion MRI
DTI	diffusion tensor imaging
MD	mean diffusivity
FA	fractional anisotropy
HA	helix angle
SNR	signal-to-noise ratio
cDTI	cardiac DTI
CS	compressed sensing
DW	diffusion-weighted
SPIK	Spiral with Keyhole
NMR	nuclear magnetic resonance
RF	radiofrequency
FID	free induction decay
EPI	echo planar imaging
FOV	field of view
SE	spin-echo
FFT	fast Fourier transform
PDF	probability density function
PGSE	pulsed-gradient spin-echo
ADC	apparent diffusion coefficient
STEAM	stimulated echo acquisition mode
bpm	beats per minute
CPR	conjugate phase reconstruction
PI	Parallel imaging
DWI	diffusion-weighted image
VDS	variable density spiral
UDS	uniform density spiral
SSIM	structural similarity
RMSE	root mean square error
PSNR	peak signal-to-noise ratio
PSF	point spread function
SOS	stack of spirals
BJS	bidomain joint sparsity
BCS	blind compressed sensing

FISTA	fast iterative shrinkage-thresholding algorithm	ADMM	alternating direction method of multipliers

Chapter 1

Introduction

In this chapter, we explore the importance of measuring the microstructure of the heart. We then introduce the key methods available for this purpose and discuss their advantages and disadvantages. Finally, we elaborate on the goals of this research and the thesis outline in the subsequent chapters.

1.1 Significance

The heart is undeniably one of the most vital organs in the human body, continuously pumping blood rich in oxygen and nutrients to sustain life. Heart diseases, collectively referred to as cardiovascular diseases (CVD), stand as the leading cause of mortality worldwide and in Europe, accounting for an estimated 32% (19 million deaths) of all global deaths [1, 2]. As life expectancy increases, the proportion of global deaths attributed to CVD is expected to continue rising. During the pandemic, patients with CVD faced some of the highest risks of mortality and complications when contracting COVID-19. The longer-term implications for cardiovascular health in individuals who have recovered from COVID-19 are becoming more apparent and acknowledged, leading to evolving research in this field. According to the British Heart Foundation (2022), there are approximately 7.6 million individuals living with CVD in the UK, which is roughly double the combined count of those living with cancer and Alzheimer's [3]. Moreover, CVD causes approximately one-quarter of all deaths in the UK, resulting in over 160,000 fatalities annually. This equates to about one death every three minutes. Beyond the human toll, CVD exerts significant financial burdens, with an estimated annual cost of €282 billion to the European Union economy in 2021 [4].

It is worth noting that over 80% of CVD-related deaths are attributed to heart attacks and strokes, and a significant one-third of these fatalities occur prematurely in individuals under the age of 70. Early detection of CVD is thus of paramount importance. Identifying individuals at risk of CVD and ensuring they receive appropriate and timely treatment can help prevent premature deaths. However, CVD diagnosis often occurs at late symptomatic stages, leading to delayed interventions and reduced effectiveness of medical care. Many cardiac diseases potentially leading to heart dysfunction impact the mechanical properties of the cardiac muscle at the very early stages. The optimal mechanical functioning of the heart strongly relies on the geometric structure and arrangement of cardiac muscle cells, or cardiomyocytes. Therefore, gaining a thorough understanding of the architecture of the myocardium in both healthy and pathological

states helps enhance our interpretations and predictions of structural changes in the heart due to cardiac diseases and explore effective therapeutic approaches. Such knowledge offers insights into developing early-stage biomarkers and identifying key risk factors before CVDs fully develop. Techniques that allow for quantitative assessment of both normal and abnormal myocardial structure and function support preventive care and early cardiovascular treatment. They also aid in guiding future treatment strategies for timely and effective delivery of treatment to affected patients, ultimately helping to reduce patient morbidity and mortality.

1.2 Structure of the Heart

The general organisation of cardiac tissue has been a subject of research for centuries. The heart is a muscular organ characterised by varying levels of complexity. Its fundamental contractile units, or building blocks, are the *cardiomyocytes*. The geometric arrangement of these cardiomyocytes is believed to play a crucial role in efficient contraction. While there have been many observations regarding cardiomyocyte architecture, we still lack a comprehensive understanding of how they are precisely organised, interconnected, and collectively contribute to the heart's overall function.

1.2.1 Gross Morphology

The primary function of the heart is to pump blood throughout the body, and its macroscopic structure is well-designed for optimal performance. It comprises four chambers: the right and left ventricles and the right and left atria. The right atrium receives the deoxygenated blood from the systemic circulation and pumps it into the right ventricle (RV) through the Tricuspid valve. The RV is responsible for sending deoxygenated blood from the body to the lungs through the Pulmonary valve. Conversely, oxygenated blood from the respiratory system is collected in the left atrium, from where it flows through the Mitral valve into the left ventricle (LV). The blood in the LV is then ejected through the Aortic valve into the systemic circulation. The anatomical structure of the heart is depicted in Figure 1.1. Heart valves play a crucial role in ensuring that blood flows through the heart in one direction only. As we may expect, the LV wall is much thicker than the RV wall because the systemic circulation is larger and offers more resistance to the blood outflow.

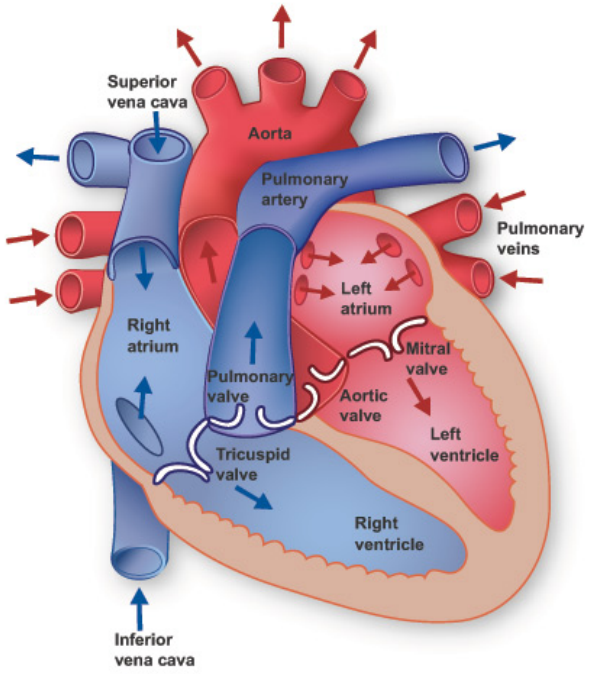


Figure 1.1: Heart anatomical structure. The right heart (in blue) ejects deoxygenated blood to the lungs. The left heart (in red) pumps oxygen-saturated blood to the vascular system. (Adapted from [5])

The heart wall, enclosing its chambers, is composed of three layers, as depicted in Figure 1.2. Each layer has specialised structure and function that collectively ensure the proper functioning of the heart and the efficient pumping of blood to all organs in the body. The outermost layer is the epicardium, which constitutes part of the protective pericardium layer (visceral pericardium). It is a serous membrane lining the exterior of the heart wall and produces pericardial fluid, which acts as a lubricant, facilitating smooth motion between the inner and outer layers of the pericardium. The inner layer is called the endocardium, consisting of thin and smooth epithelial tissue that covers the inner surface of all heart chambers and valves. It connects to the endothelium of the vessels entering and leaving the heart and plays a crucial role in prolonging myocardial contraction. Between these layers is the myocardium, the thick muscular layer responsible for contraction. It primarily consists of basic units of cardiac muscle cells called cardiomyocytes. The precise organisation of these cardiomyocytes within the myocardium is crucial for the optimal heart function.

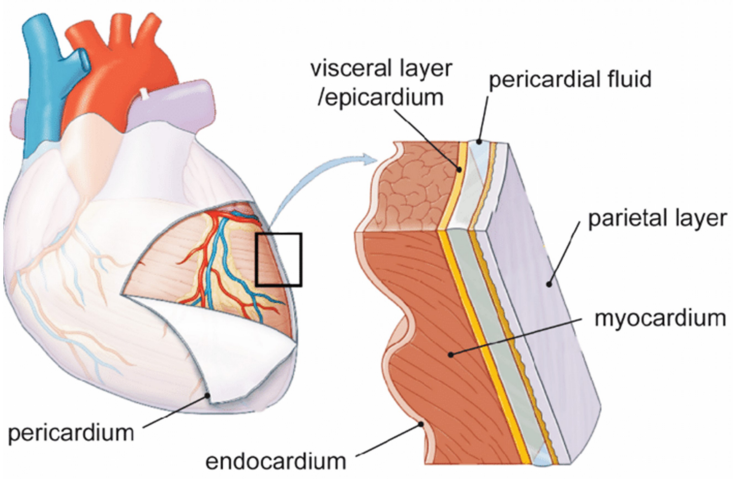


Figure 1.2: Schematic drawing of the trilayered structure of the heart wall. (Adapted from [6])

The coordination of the cardiac muscle contractions, which propel blood to various parts of the body, is facilitated by electrical impulses propagating throughout the heart. The electrical activity of the heart can be detected using electrodes placed on the chest to record an electrocardiogram (ECG), as illustrated in Figure 1.3. Each heartbeat constitutes a cardiac cycle, divided into two phases: systole (contraction) and diastole (relaxation). The systolic phase begins with the R-wave on the ECG, during which ventricular pressure increases, causing the aortic valve to open and allowing blood to be rapidly ejected into the vascular system. Following contraction at the end of systole, the aortic valve closes, and ventricular repolarisation occurs, leading to a decrease in pressure. Shortly after, the mitral valve opens, allowing the ventricle to fill with blood from the pulmonary system during the diastolic phase.

The two ventricles are separated by the septum. The plane where the ventricles meet the atria is called the basal plane. The tip of the heart, the farthest point from the base, is called the apex. The heart's long axis extends from the apex to the centre of the base. Cross-sectional slices perpendicular to this long axis are referred to as axial, short-axis, or transverse slices, as shown in Figure 1.4.

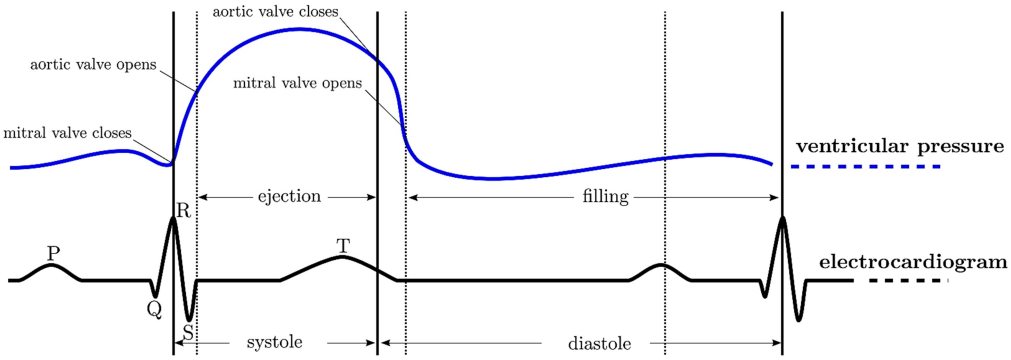


Figure 1.3: The correlation between left ventricular pressure and the electrocardiogram (ECG). The ECG measures the electrical activity of the heart, which triggers the contraction and relaxation of cardiomyocytes, thus explaining ventricular function.

1.2.2 Microstructure

The myocardium comprises the fundamental muscle tissue of the heart and is distinguished by its striated nature. It is a complex matrix composed of collagen, muscle cells or myocytes, and various proteins, as depicted in Figure 1.5. Cardiac muscle cells are elongated cigar-shaped cells, typically measuring between 50 to 150 μm in length and 10 to 20 μm in diameter [8]. These cells are unique to the heart and are not found in other types of tissue. Cardiomyocytes are interconnected to neighbouring cells through structures called intercalated discs, and at their ends, cardiomyocytes often share boundaries with multiple cells, forming a mesh-like structure. The intercalated disc is a characteristic feature of cardiac muscle tissue. These discs allow for robust mechanical connections between cells and facilitate rapid transmission of electrical impulses from one cell to another.

1.2.3 Laminar Structure

The ventricular myocardium exhibits a complex laminar structure, where myocytes are organised into branching layers that are typically 3 or 4 cells thick, forming structures called myolaminae or sheets, as shown in Figure 1.6. These sheets are interconnected by an extracellular collagen network. This intricate collagen matrix enables tight coupling between cardiomyocytes within the sheets while allowing for some degree of flexibility between adjacent sheets, where direct myocyte-to-myocyte coupling is absent. This

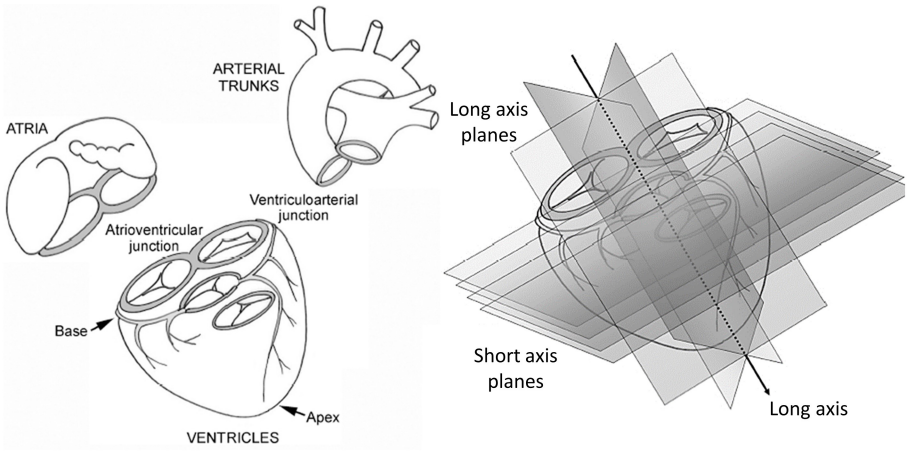


Figure 1.4: Basic cardiac segments and planes. (Reproduced from [7])

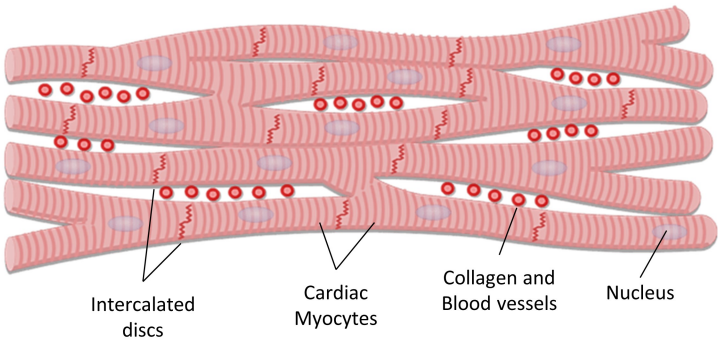


Figure 1.5: Microstructure of the myocardium. (Reproduced from [9])

loose coupling between sheets facilitates the sliding movement of sheets against each other during the cardiac cycle, contributing to the heart's efficient functioning [10, 11]. Additionally, the geometric arrangement of sheets within the heart wall demonstrates a high degree of regularity, with orientations changing smoothly from approximately -60° at the endocardium to about +60° at the epicardium, exhibiting a characteristic transmural orientation.

1.3 Heart Structure and Function Relationship

The function of the heart can be observed at different levels. As the heart performs its pumping action, it moves within the pericardial sac. This mechanical movement can be visualised at a macroscopic level using conventional medical imaging techniques. The

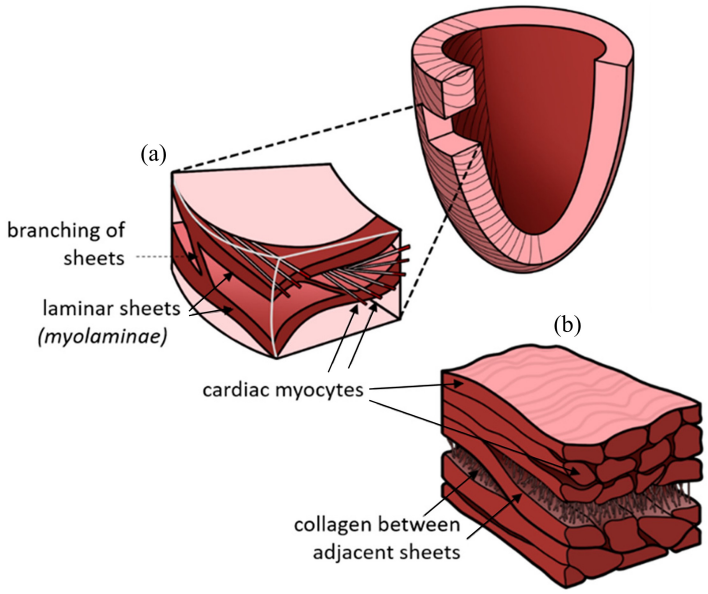


Figure 1.6: A schematic diagram of myocardial laminar architecture. (a) Transmural arrangement of cardiomyocytes into laminar sheets throughout the myocardial wall. (b) A closer examination of the sheet structure, where cardiomyocytes are stacked in groups of four and held together by collagen struts. (Reproduced from [12])

base of the heart is restricted from rotating in the basal plane due to the presence of vessels, and the apex is constrained from moving up or down. The remaining motion is characterised by the twisting of the LV, which is more complex than a simple squeeze and release. Twist refers to the rate of change in the rotation angle within the LV, from the apex to the base. During ventricular contraction, there is transmural wall thickening, which serves a vital role in overcoming resistance and ejecting blood into the body's circulation. This behaviour has been documented by numerous studies and has been shown to be altered in damaged hearts [13, 14]. Additionally, variations in twist throughout the cardiac cycle have been found to change with age.

At a microscopic level, the contraction of individual myocardial cells drives the heart movement. Although cardiomyocytes typically shorten by about 14%, resulting in only an approximate 8% increase in cardiomyocyte diameter, this alone cannot account for the observed 28-50% increase in average wall thickness during contraction. This substantial wall thickening is best explained by considering secondary changes in cellular orientation. Particularly, the sliding of sheets relative to each other is con-

sidered a primary mechanism for myocardial wall thickening [15]. The laminar model also significantly influences the spread of activation within the myocardium as it has been shown that the electrical impulse triggering cell contraction propagates rapidly along the direction of the myocytes, with a velocity approximately three times faster than in the perpendicular plane.

Hence, the spatial arrangement of cardiomyocytes in sheets is believed to play a critical role in explaining the complex structural changes related to the heart function. To delve further into the significance of the correlation between heart structure and function, let us explore some common pathologies linked to structural changes in the heart.

1.3.1 Heart Pathologies

Each type of heart disease involves a complex interplay of multiple factors and has various stages with symptoms like breathlessness and fatigue, but often, irreversible damage occurs before these symptoms. According to American Heart Association guidelines, in stage B, patients do not show symptoms but exhibit structural remodelling of the myocardium [16], making early diagnosis challenging yet crucial. Diseases that directly damage the heart muscle are referred to as cardiomyopathy. In cardiomyopathy, the heart may become abnormally enlarged, thickened, or stiffened, affecting all its chambers to develop different mechanical properties from healthy tissue, thereby impacting the heart function. There are several types of cardiomyopathy, with the three main ones being dilated cardiomyopathy, hypertrophic cardiomyopathy, and restrictive cardiomyopathy. Figure 1.7 illustrates the structural remodelling associated with these heart pathologies.

Dilated cardiomyopathy (DCM), also known as congestive cardiomyopathy, is the most common type of heart muscle disease. This condition is characterised by the enlargement of the heart and weakening of the heart muscle wall [17]. This stretching of the cardiomyocytes can affect the ventricles' ability to pump blood efficiently and may cause the heart valves to leak. Consequently, the heart fails to pump enough blood to the body, resulting in congested lungs and heart failure. DCM can also result in abnormal heart rhythms.

Hypertrophic cardiomyopathy (HCM), also known as asymmetric septal hypertrophy, is the second most common type of cardiomyopathy. HCM is characterised

by the disordered growth of heart muscle cells, leading to the thickening of the heart wall [18]. This thickening reduces the flexibility of the heart wall and can potentially block blood flow out of the heart (known as obstructive HCM), ultimately leading to replacement fibrosis and left ventricular remodelling [19]. HCM is a major cause of heart failure, sudden death, and abnormal heart rhythms. It is also the leading cause of sudden cardiac death among seemingly healthy athletes and can have a genetic component.

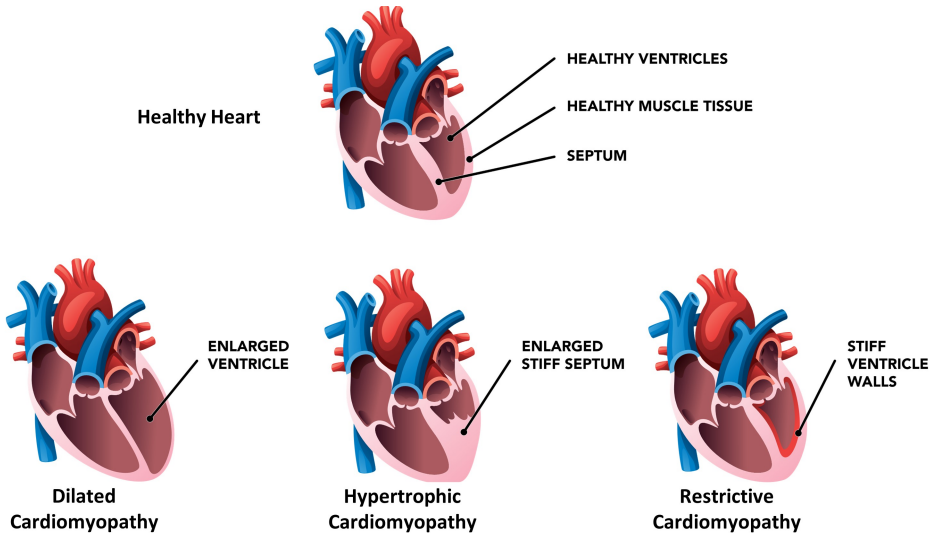


Figure 1.7: Cardiac structure remodelling in various pathological conditions. (Reproduced from [20])

Early detection of heart disease is widely recognised as a crucial step in reducing mortality rates among heart failure patients. While risk factors like high blood pressure and high cholesterol are well-known, there is no current technology for detecting the initial stages of heart failure due to a limited understanding of the underlying mechanisms. Obtaining a comprehensive three-dimensional (3D) view of the cardiac architecture in vivo offers a novel perspective on understanding the dynamic changes in healthy cardiac microstructure due to its function and the pathophysiological alterations due to diseases. This, in turn, helps in unraveling the causes of heart pathologies based on how they impact the underlying myocardial architecture prior to affecting the macro functional responses such as cardiac output.

1.4 Myocardium Microstructure Characterisation

Insights into cardiac cellular architecture initially stemmed from traditional techniques like dissection and histological analyses of specific tissue regions [21]. As far back as the 1600s, Lower described the heart as having a helical arrangement [22]. In 1969, using advanced imaging technology of the time, Streeter described the complex and highly organised structure of cardiomyocytes throughout the LV. He discovered that the angles change linearly from the inner layer (endocardium) to the outer layer (epicardium) [23]. This led to the proposal that the myocardium consists of sheets of helical cardiomyocytes with orientation gradually shifting from the epicardium to the endocardium [24].

In the early 1970s, another histological study by Torrent-Guasp aimed at isolating cardiac muscle cells from the collagen network. This study revealed that the muscle cells are intertwined to form a single band, which winds around both ventricles in a double helical pattern [25]. The presence of an oblique fold in the center of the band was exposed, which acts as a demarcation between the basal and apical loops. Photographs of a dissected heart displaying this arrangement can be seen in Figure 1.8.

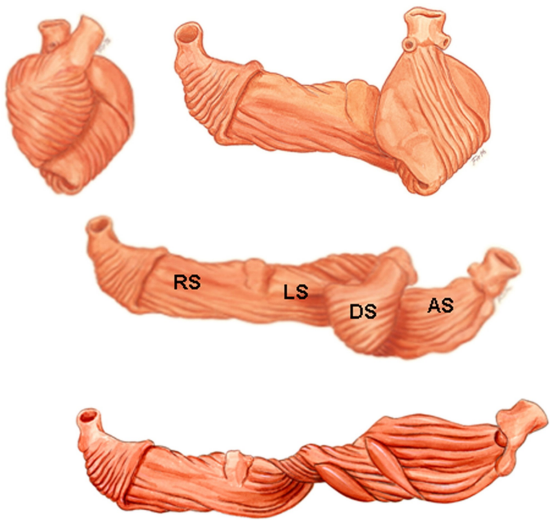


Figure 1.8: Torrent-Guasp myocardial band model. RS: right segment; LS: left segment; DS: descending segment; AS: ascending segment. (Reproduced from [26])

However, these methods have several limitations: they can be subjective, making them challenging to reproduce, often qualitative rather than quantitative, and are de-

structive, time-consuming, and inherently two-dimensional (2D).

Later, quantitative polarized light microscopy was utilised to investigate the three-dimensional (3D) orientations of cardiomyocytes within the ventricular mass [27, 28]. In confocal microscopy, following fixation, the heart is sliced, and the direction of the cardiomyocytes is physically measured in each cross-section, allowing for the 3D reconstruction of cellular orientations with an impressive spatial resolution of 0.1 mm $\times$ 0.1 mm $\times$ 0.5 mm [29, 30]. A noteworthy advantage of this technique is its objectivity, as it removes the need for human intervention during the measurement process; however, it is inherently destructive and requires tissue freezing before imaging, which can damage the tissue and potentially introduce bias in the resulting orientation quantification. Specifically, the procedures used in the fixation and sectioning of the specimen can affect the organisation of cardiomyocytes, potentially leading to differences in the measured angles. To address this, some studies repeat the experiments on different samples or sections and report the average angle measurements [29]. Additionally, it is only suitable for examining limited tissue sample volumes. An example of confocal imaging of cardiomyocytes is shown in Figure 1.9.

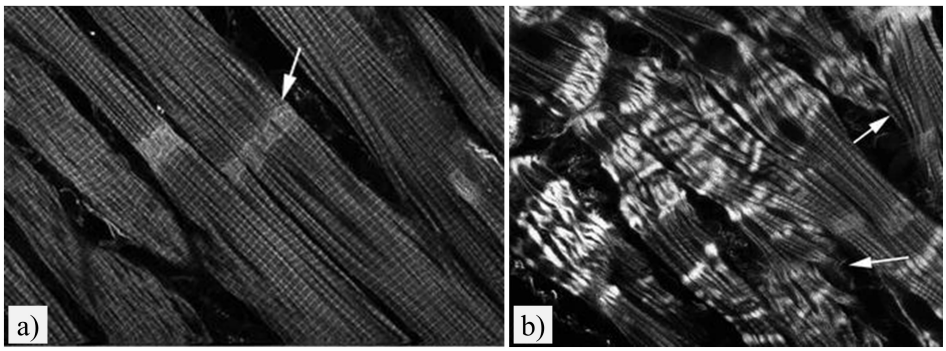


Figure 1.9: Confocal imaging of cardiomyocytes. Single optical sections of the left ventricle from (a) a healthy heart and (b) a hypertrophied heart are shown. Arrows point to intercalated disks. (Reproduced from [31])

Currently, histological sectioning combined with confocal microscopy is considered the gold standard for mapping cardiac microstructure and assessing laminar orientation. Yet, it is important to note that both techniques are invasive and thus restricted to ex vivo heart examinations and cannot be used for in vivo examination of the heart.

More recently, diffusion magnetic resonance imaging (dMRI) has emerged as a new

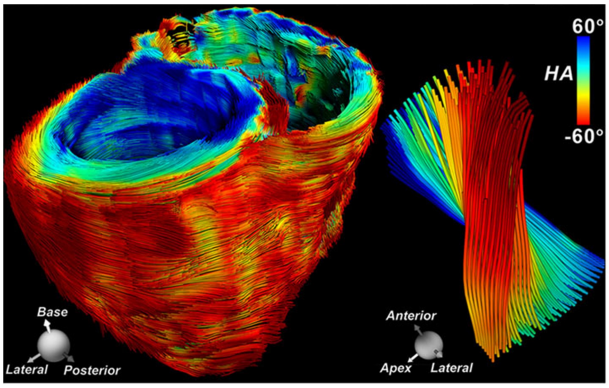


Figure 1.10: Representative example of 3D DTI tractography of cardiomyocyte orientations; HA: helix angle. (Reproduced from [32])

and promising technique for depicting the 3D microstructure of tissues [33, 34]. This technique does not require the use of any contrast agent, as the contrast is generated by the natural movement of water molecules already present in the body. It is entirely non-invasive and has thus enabled the quantification of the tissue microstructure in vivo. The most commonly used dMRI technique in the heart today is diffusion tensor imaging (DTI) [35–37], which leverages the inherent anisotropy of cellular structures to determine their orientation, as illustrated in Figure 1.10. It stands as a powerful tool and is currently the only non-invasive method for investigating the laminar architecture within the in vivo human heart [38–44]. Cardiac DTI (cDTI) studies have revealed that the cellular architecture in the heart is more intricate than a simple combination of sheets [45], offering significant new insights over traditional methods. Extensive studies have also shown that the organisation of cardiomyocytes has a significant impact on the heart function [46, 47]. Specifically, it has been observed that the torsion and local stress of the ventricle are highly sensitive to the local distribution of cellular orientations. Subsequent research in healthy human and animals [48] has highlighted the significant importance of the cardiomyocytes orientations in cardiac output and in cardiac remodelling following infarction. DTI-derived biomarkers, such as the mean diffusivity (MD), fractional anisotropy (FA), and helix angle (HA), serve as useful indicators to determine major microstructural properties and have led to exciting theoretical and computational developments. These metrics can serve as early-stage biomarkers in various cardiac pathologies, including myocardial infarction. For instance, previous

ex vivo dMRI studies conducted in chronic myocardial infarction porcine models have demonstrated an increase in the MD of infarcted regions. This increase has been linked to alterations in the ventricle shape and disarray of the cellular orientations around the infarcted zone, as identified through histological analysis. While the DTI technique has facilitated a comprehensive and quantitative understanding of the myocardium microstructure and has offered insights into complex processes associated with different cardiac diseases, it still faces several limitations that have yet to be addressed.

1.5 Motivation

Robust quantification of tissue microstructure using diffusion MRI necessitates: (i) multiple measurements ($\geq$ 6) along distinct diffusion-encoding directions to ensure accurate capturing of the orientation information [49], and (ii) scan repetitions ($\approx$ 4-6) for signal averaging to enhance the inherently low signal-to-noise ratio (SNR) of DTI experiments [50]. The resulting prohibitively lengthy scan time is thus one of the main challenges complicating the routine application of cDTI. Specifically, prolonged data acquisition can cause substantial problems if the patient or subject is unable to remain motionless during the measurements and, in particular, when imaging body sections with unavoidable movements such as the heart.

It is also important to note the intrinsic trade-off in MRI between acquisition time, SNR, and spatial resolution (Figure 1.11), meaning that adjusting one of these factors will inevitably impact the other(s). In other words, improving one aspect may have adverse effects on the others unless additional adjustments are implemented. Therefore, research efforts aimed at addressing challenges in in vivo cDTI applications and its clinical translation have mainly resulted in compromises in spatial resolution and myocardium coverage. Consequently, most in vivo human cDTI studies have been limited to imaging only a few short-axis slices (three to five) with an in-plane resolution of $\geq$ 2 mm [35].

As discussed, the cardiac tissue has a highly complex multi-scale architecture, which can be decomposed into three main scales: the cell scale (10-100 μm), the sheetlet scale (100-1000 μm), and the tissue scale (1-10 mm). While the orientation and helical arrangement of cardiomyocytes at the tissue scale have been extensively studied, the precise understanding of how sheetlets are organised throughout the myocardial wall remains unclear. Some studies support the dominance of a single laminar sheetlet pop-

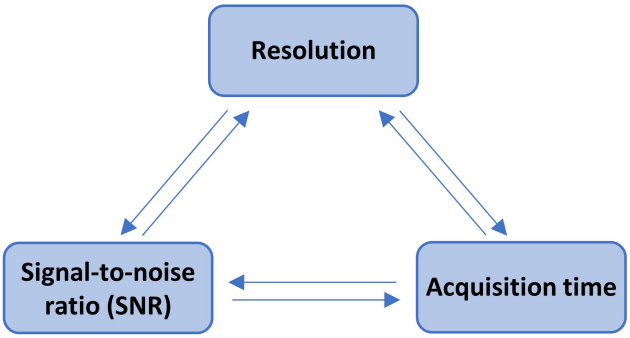


Figure 1.11: Inherent trade-off in MRI

ulation, while others suggest the coexistence of two local myocardial sheet populations [51]. Acquiring high-resolution isotropic voxels across the myocardium syncytium enables detailed and robust mapping of the heart microstructure in three dimensions and allows for resolving small groups of cardiomyocytes in sheet arrangement. Additionally, high-resolution DTI holds promise for enhancing the understanding and diagnosis of various early-stage cardiac pathologies characterised by sub-millimetric lesions or morphological changes, such as infarction [52, 53].

Improving spatial resolution with full coverage of the LV myocardium would enhance the accuracy of the derived DTI metrics and, subsequently, the quality of coherent tractography results, better aligning with histological studies. Although high-magnetic-field and advanced gradient systems can provide sufficient SNR to improve the spatial resolution and despite the advancements in correcting for bulk and physiological motion [54], extended acquisition times still remain a significant obstacle to widespread use of high-resolution 3D cDTI.

In ex vivo settings, cDTI studies have achieved isotropic resolutions at orders of 50-100 μm in rodent hearts and sub-mm/mm in large animal models (including human hearts) using conventional 2D multi-slice or 3D Cartesian-based sequences. Acquisition times for these studies typically range from several hours to multiple hours, far exceeding clinically acceptable limits. With current sequences, implementing 3D cDTI in vivo remains impractical due to challenges associated with imaging a moving heart and the increased susceptibility to motion-induced artefacts. To expand its applicability, substantial improvements in both the speed and robustness of whole-heart cDTI are necessary. This underscores the need for new acceleration strategies capable of

addressing the extended continuous acquisition time in high-resolution 3D imaging to fully unlocking the vast potential of DTI in the heart.

1.6 Research Aim and Objectives

The primary goal of this thesis is to facilitate high spatial resolution and whole-heart coverage cDTI while maintaining scan times comparable to 2D scans and without compromising SNR or image quality. Specifically, the objectives are to achieve a high spatial resolution of 1 mm, 3D myocardial coverage with 30 slices, and an error of $\leq$ 4% in DTI parameter maps, all within a scan time of $\leq$ 10 minutes. To achieve this, we explore efficient novel sampling schemes combined with dedicated reconstruction techniques capable of achieving high accelerations necessary for 3D cDTI within clinically acceptable scan times.

In recent decades, the emerging theory of compressed sensing (CS) [55–57] has shown significant potential for reducing MRI scan times. CS has offered a fresh perspective on both data sampling and signal reconstruction from significantly undersampled measurements. This approach departs from the traditional Shannon-Nyquist sampling theorem [58, 59] and leverages the inherent sparsity or compressibility of signals in specific domains (such as the pixel or wavelet transform domain) to enable artefact-free reconstructions through specialised nonlinear algorithms. While CS has been successfully applied in cDTI [60–63] to shorten the scan time, its implementation has primarily relied on basic undersampling patterns with points arranged on equidistant or Cartesian grids, which have limitations in achieving very high acceleration factors (≥ 8).

Taking advantage of alternative sampling schemes known as non-Cartesian MRI, we first revisit the design of a variable-density spiral trajectory and introduce a new parametrisation for an optimal sampling density function. This function conforms to MR hardware constraints, is well-adapted to the diffusion-weighted signal, and is particularly suitable for CS reconstructions, enabling high acceleration factors.

We then extend the proposed optimal sampling scheme to sparse 3D trajectories based on stack-of-spirals and conical schemes. These trajectories efficiently cover the combined 4D frequency and diffusion directional domain (k, q-space) within a short time interval and offer high spatial and directional incoherency, making them well-suited for CS application. Furthermore, we adopt a constrained reconstruction framework

to recover the 3D non-Cartesian data, effectively utilising prior information on the joint sparsity of diffusion-weighted (DW) images acquired at different directions. The incorporation of two regularisation terms provides complementary prior information, thereby improving reconstruction accuracy. We adapt an efficient fast algorithm to solve the large-scale composite optimisation problem and introduce a phase correction procedure to account for shot-to-shot inconsistencies.

Finally, we propose a novel sampling approach called SPIK (SPIral with Keyhole) and a sparse blind dictionary learning method, which leverage the high structural similarities between DW images while efficiently capturing diffusion-related information with high resolution. The method utilises the proposed optimised variable-density spiral sampling patterns. The reconstruction exploits a data-adaptive dictionary which allows for sparser representation of the diffusion signal and effectively eliminates blurring artefacts through learning correlated information among DW images. By imposing sparsity constraints, the reconstruction quality is further improved. We demonstrate that this approach can yield reasonable images within a significantly reduced measurement time, an achievement not attainable with conventional sampling schemes.

We use retrospective evaluations on a broad range of cDTI datasets, leaving prospective evaluations for future studies. We thoroughly discuss and address practical concerns, including the compatibility of the proposed trajectories with the gradient system, the competitiveness of the proposed method compared to other sampling trajectories, and the sensitivity of the trajectories to imaging and system imperfections. This work has the potential to pave the road towards the wide adoption of high-resolution 3D cDTI.

1.7 Thesis Outline

The research work is organised as follows:

Chapter 2 gives a brief introduction to the principles of MRI, including MR image formation and image reconstruction. Particular emphasis is placed on the k-space formalism, a domain that serves as the Fourier conjugate to the conventional spatial domain containing the magnetisation of the imaged object. The chapter proceeds to introduce the principles of the diffusion MRI (dMRI), and the diffusion tensor model. Finally, we will discuss diffusion tensor imaging (DTI) in the heart, which concerns the

main application of this work.

Chapter 3 discusses the methodological challenges and the state-of-the-art techniques in cDTI, reviewing their respective advantages and disadvantages. It also delves into the theory of compressed sensing and its application to cDTI.

Chapter 4 proposes a new parametrisation of the sampling density function for a spiral trajectory, well adapted to the diffusion-weighted signal in the heart and suitable for CS reconstructions. It then presents an empirical and quantitative study to determine the optimal parameter values for efficient, high-resolution sampling of diffusion-weighted images (DWIs) with a given input SNR while maintaining a desired level of image quality. This spiral-CS analysis was undertaken because a quantitative guide on the optimal spiral sampling in a given acquisition scenario was lacking, resulting in a trial-and-error approach in selecting the appropriate undersampling parameters in MRI compressed sensing applications.

Chapter 5 introduces two extensions of the optimal spiral-CS to three dimensions, enabling efficient whole-heart cDTI. The CS reconstruction is also extended to utilise the shared sparse patterns among DWIs in both the total variation and wavelet domains to allow for reduced amount of measurements. This method called bidomain joint sparsity (BJS) also accounted for the data inconsistencies in multi-shot imaging approaches.

Chapter 6 introduces an accelerated sampling and reconstruction technique for efficient whole-heart cDTI, achieved through spatio-directional undersampling. The proposed method utilises a novel sampling called spiral with keyhole (SPIK) and a reconstruction strategy based on blind dictionary learning. The approach leverages the spatial and directional similarities or correlations between DWIs to efficiently sample the unique diffusion-related information in each DW direction while sampling the correlated information in a segmented manner to reduce the scan time.

Chapter 7 summarises the main achievements of the thesis and offers an outlook for future work.

Chapter 2

Principles of MRI and Diffusion MRI

The discovery of nuclear magnetic resonance (NMR) brought about a revolution in the field of medical imaging. Today, MRI has become the leading modality for diagnosing various common diseases, in particular cancer and stroke. It has significantly contributed to advancing our understanding of the human body, offering unprecedented levels of structural and functional insights. MRI boasts exceptional attributes for both preclinical and clinical applications, including excellent soft-tissue contrast. The breakthrough of MRI was further supported by its remarkable ability to provide diverse types of contrast, ranging from proton density, diffusion, flow, and temperature, up to complex quantities like perfusion or brain activation. This chapter provides an overview of MRI principles, focusing on a signal processing perspective. The discussion is limited to key aspects that pertain to the subsequent chapters of the thesis. For more detailed information on this technique, the reader is referred to related textbooks [64, 65]. Additionally, we delve into the diffusion process within biological tissues and how it is measured using MRI to quantify the underlying tissue structure.

2.1 Nuclear Magnetic Resonance Physics

To understand how MR images of body tissues are obtained, it is necessary to consider the classical description of NMR physics. Thus, we begin by exploring some interesting phenomena known as precession, excitation, and relaxation.

2.1.1 Precession

Human body tissues contain a significant amount of water, and within each water molecule are hydrogen nuclei, or protons (^{1}H). Each proton, which carries a positive electrical charge, continuously spins around an axis, as shown in Figure 2.1(a), acting like a tiny magnet. The cumulative effect of these small magnetic fields generated by proton spins is called *net magnetisation*. In the human body, these miniature magnetic moments are oriented randomly, resulting in a net magnetisation that cancels out and effectively becomes zero.

However, when these protons are exposed to a strong external magnetic field B_0, a phenomenon known as precession occurs. In addition to "spinning" around their own axis, the nuclei will "precess" or rotate around the axis of B_0 due to their magnetic moment, as shown in Figure 2.1(b). Consequently, some of the protons will precess

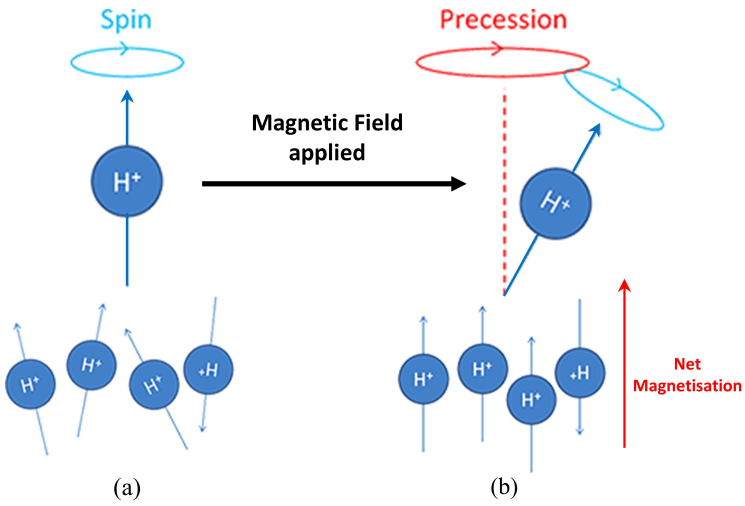


Figure 2.1: (a) Proton spin around its own axis. (b) Proton precession around the axis of the external magnetic field. (Reproduced from [66])

aligned with the field (the low-energy state), while others will precess in the opposite direction (the high-energy state), resulting in a non-zero net magnetisation M_0 in the direction of the applied magnetic field. According to the Boltzmann distribution, a slightly higher proportion of protons will occupy the lower energy state, as systems tend to favour lower energy configurations at thermal equilibrium. This imbalance between the two states is what gives rise to the measurable net magnetisation.

The frequency at which spins' precession occurs is often referred to as the *Larmor frequency*, denoted by ω_0. This frequency is determined by the strength of the primary magnetic field B_0 and the gyromagnetic ratio γ of the proton ^{1}H:

$$\omega_0 = \gamma B_0 \tag{2.1}$$

The gyromagnetic ratio γ, which has been determined through experimentation, relates the angular momentum to the magnetic moment for a specific type of nucleus. For hydrogen nuclei, for instance, the gyromagnetic ratio is 267.5×10^6 rad/s/T or 42.57 MHz/T, corresponding to a Larmor frequency of 64 MHz in a 1.5 T magnetic field.

2.1.2 Excitation

In MRI, B_0 is conventionally oriented along the $z-$axis, corresponding to the long axis of the patient, and therefore, a net magnetisation, referred to as the longitudinal magnetisation M_z, in the same direction as B_0 is generated. M_z is related to the proton density, a key parameter of interest in MRI experiments. To generate an MR signal, another magnetic field, known as the B_1 field or radiofrequency (RF) pulse, is applied to temporarily deflect the alignment of the spins with B_0, causing the net magnetisation to flip into the transverse $(x-y)$ plane. This process is called *RF excitation*, and the field is called an RF pulse since it uses a rotating magnetic field at a high (radio wave) frequency. The angle at which M_z deviates from its alignment with B_0 is called the *flip angle*, determined by the amplitude and duration of the RF pulse. Typically, the flip angle is set to 90° to ensure that the longitudinal magnetisation M_z is fully transferred into the transverse magnetisation M_{xy}, as illustrated in Figure 2.2(a). If the frequency of the RF pulse matches the Larmor frequency of the protons, *resonance* occurs, resulting in all proton spins being in phase or exhibiting complete phase coherence. The B_1 magnetic field is much weaker, typically a few Gauss, compared to the 1.5 T strength of the B_0 field (1 T = 10^4 Guass), and it is applied transiently for only a few milliseconds.

2.1.3 Relaxation

After the B_1 field is switched off, the only magnetic field affecting the magnetisation is B_0. At this point, the net magnetisation of the proton spins, which lies in the transverse plane, begins to return to its equilibrium state, aligning once again with the primary magnetic field B_0 (Figure 2.2(b)). To accomplish this realignment, the magnetisation must release the energy it gained from the RF pulse, a process known as *relaxation*. Early NMR experiments revealed the occurrence of two distinct relaxation mechanisms [67], as illustrated in Figure 2.3. The first mechanism, known as *spin-lattice* or *longitudinal relaxation*, describes the exponential recovery of the longitudinal magnetisation (M_z) to the equilibrium state (M_0) along the B_0 axis. This phenomenon arises from energy exchange between the protons and their environment, causing the excited system to return to its thermal equilibrium state. It is characterised by a time constant called T_1 relaxation time, which defines how long it takes for the longitudinal magnetisation to reach 63% of its maximum value. The second mechanism, known as *spin-spin*

or *transverse relaxation*, refers to the dephasing of the transverse magnetisation due to frequency fluctuations caused by proton interactions. This process is characterised by a time constant called T_2 relaxation time, which indicates the rate at which the transverse magnetisation has lost approximately 63% of its initial value.

These relaxation times, T_1 and T_2, are inherent properties of the individual sample material and are thus of significant importance in MRI because they affect both the signal intensity and the contrast in the resulting images.

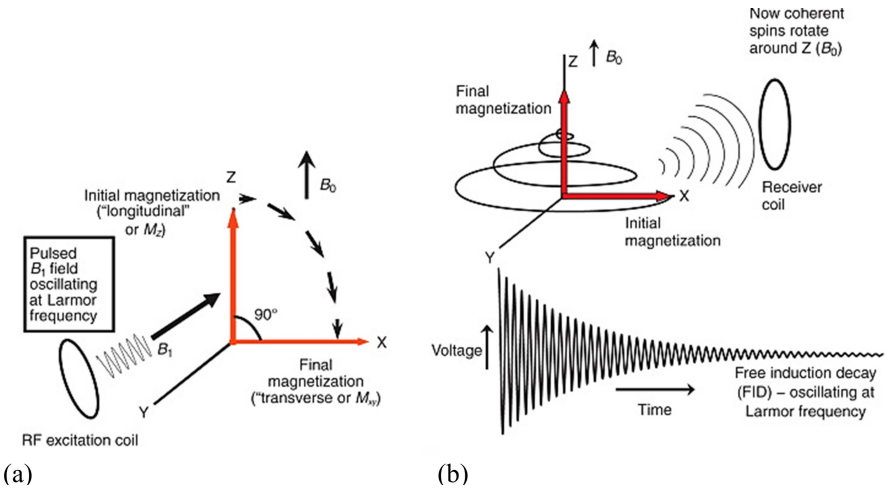


Figure 2.2: RF excitation and signal detection: (a) During the application of a 90° RF pulse, longitudinal magnetisation (M_z) transitions into transverse magnetisation (M_{xy}), (b) Once the RF pulse is turned off, the longitudinal magnetisation continues to precess around the z-axis at the Larmor frequency, while the transverse magnetisation decays in the $x - y$ plane, as equilibrium is restored. This results in a coherent oscillating magnetic field that can be detected by a receiver coil as the free induction decay (FID). The FID is a decaying sinusoidal voltage signal with a frequency equal to the Larmor frequency of the spins. RF: Radiofrequency. (Adapted from [68])

2.1.4 NMR signal

Once the RF pulse is deactivated, the relaxation process begins. If a receiver coil is positioned in the transverse plane and perpendicular to the transverse magnetisation M_{xy}, the net magnetisation can be detected as an oscillating signal, as illustrated in Figure 2.2(b). The fluctuations in the magnetisation will generate corresponding magnetic

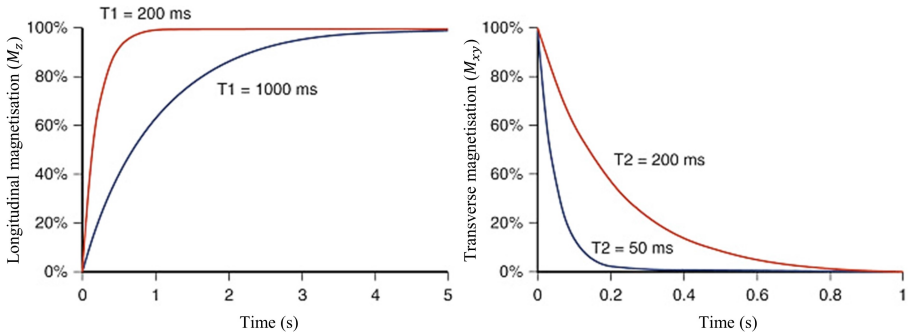


Figure 2.3: Plots illustrating T_1 and T_2 relaxations for two different tissues after a 90° RF pulse. (Adapted from [68])

field fluctuations within the receiver coil, which induces an electric current/voltage. This decaying sinusoidal signal, referred to as the free induction decay (FID), with a frequency equal to the Larmor frequency of the spins constitutes the signal measured in MRI. The variation in magnetisation is mathematically described by the Bloch equation, first formulated by Felix Bloch in 1946, as follows:

$$\frac{\partial \vec{M}}{\partial t} = \vec{M} \times \gamma \vec{B}_0 + \begin{pmatrix} \frac{M_x}{T_2} \\ \frac{M_y}{T_2} \\ \frac{M_0 - M_z}{T_1} \end{pmatrix} \tag{2.2}$$

2.2 MR Imaging

2.2.1 Signal Localisation

The NMR experiment allows the reception of a signal from the sample, but it does not provide spatial assignment of the resonating protons. By incorporating additional magnetic *gradient* fields that vary locally, we can acquire the spatial information. These gradient fields at a specific location, denoted by $\vec{r} = (x, y, z)$, can be independently switched in the x, y, and z directions and generate a spatially variant magnetic field as:

$$\vec{B}_G(\vec{r}, t) = (\vec{G}(t) \cdot \vec{r})\ \vec{e}_z = (G_x(t) \cdot x + G_y(t) \cdot y + G_z(t) \cdot z)\ \vec{e}_z \tag{2.3}$$

Here, G_x indicates the gradient field slope in the x-direction. These gradients change the magnetic field strength, causing the precession frequency of the spins to differ depending on their location. In other words, the magnetic field strength can be

expressed as the sum of these components:

$$B(x, y, z, t) = B_0 + G_x(t)x + G_y(t)y + G_z(t)z \tag{2.4}$$

This fundamental concept was introduced by Lauterbur in 1973 [69]. To spatially localise the MR signal from different voxels (three-dimensional volume elements containing protons), three essential steps are employed to encode spatial information:

Slice selection: The initial step in spatial encoding is to apply a magnetic field gradient for selecting the slice plane, as depicted in Figure 2.4. This process is known as slice selection, and it involves activating a slice-select gradient (G_{ss}) perpendicular to the desired slice plane. As previously mentioned, spin excitation occurs only when the RF pulse is applied at the Larmor frequency of those spins. In the presence of a gradient, the Larmor frequency varies along the direction of the gradient, allowing only spins located within a specific thin slice perpendicular to the gradient direction to be in resonance and emit a signal. The RF pulse associated with the slice-select gradient and the adjusted resonance frequency is referred to as the *selective* pulse. The thickness of the selected slice (ST) depends on both the frequency bandwidth of the selective RF pulse ($\Delta\omega$) and the strength of the slice-select gradient (G_{ss}):

$$ST = \frac{\Delta\omega}{\gamma G_{ss}} \tag{2.5}$$

By altering the RF pulse frequency, multiple slices at different positions in the body can be acquired, as illustrated in Figure 2.5. Once a slice has been selected, the MRI signal emitted from it must be localised along both axes within that slice. This localisation is achieved through phase and frequency encoding.

Phase encoding: The first technique for locating the MR signal within a slice is phase encoding, which determines one axis in the xy-plane of the slice by manipulating the phase of the precessing spins. This can be achieved by briefly activating a gradient along that axis, then deactivating it. Prior to the application of the phase encoding gradient (G_{PE}), all spins within the selected slice precess with the same frequency and are in phase. However, when the phase encoding gradient is activated, the magnetic field strength, and thus the precession frequency of spins along the axis of the phase gradient, is altered. As a result, the speed of precession of the spins varies from their

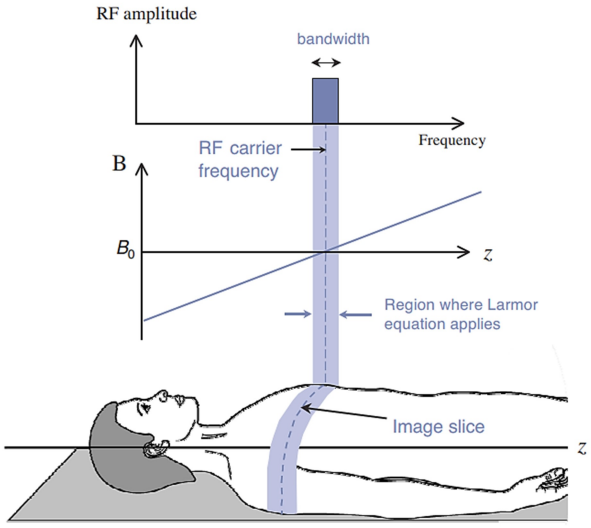


Figure 2.4: Slice selection: Selective excitation of an image slice by simultaneously applying a shaped RF pulse and a field gradient. (Reproduced from [70])

neighbours. When G_{PE} is switched off, the spins within the slice return to precessing at the same Larmor frequency but are now at different phases. This phase discrepancy among the spins is harnessed to encode the positions of the spins along the phase encoding direction, as illustrated in Figure 2.6.

Frequency encoding: Localisation along the other axis of the slice is performed by frequency encoding. Similar to slice selection, the presence of a magnetic field gradient results in a corresponding gradient in the precession frequencies of protons along that direction. By employing the slice-select gradient, we have already limited the protons being excited to a specific slice. To effectively encode the spatial information in the x-axis (for example), a frequency encoding gradient (G_{FE}) is activated along that axis, resulting in an alteration of the magnetic field strength. Consequently, the precession frequency of the spins within the slice vary slightly due to this gradient, enabling the identification of the signal based on its frequency, as illustrated in Figure 2.7. The frequency encoding gradient is typically activated during signal reception, and it is often referred to as the *readout* gradient.

It is important to note that different assignments of gradient directions are possible, such as slice-selection in the y-direction and readout gradient in the z-direction. In fact, any linear combination of gradients can be used, enabling the acquisition of arbitrary

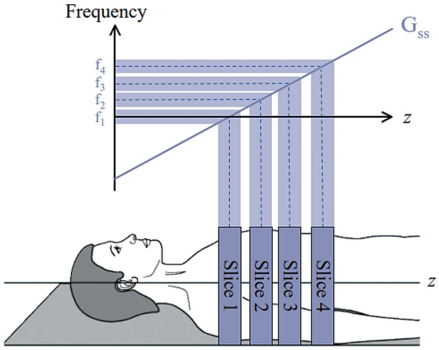


Figure 2.5: Multiple slices at different locations can be selected by adjusting the RF pulse centre frequency. (Reproduced from [70])

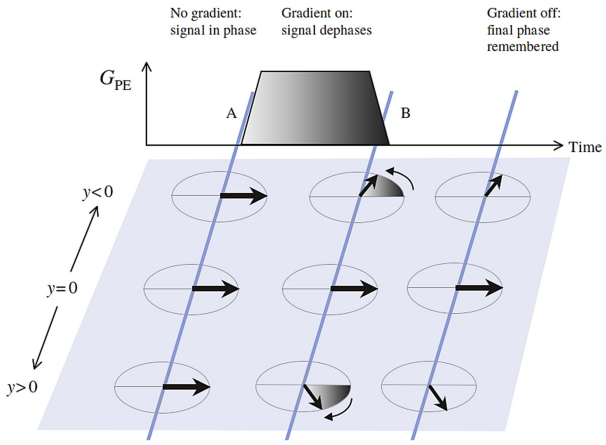


Figure 2.6: Phase encoding results in the signal returning to the Larmor frequency but with position-dependent phase shifts. (Reproduced from [70])

slice orientations.

2.2.2 MRI Pulse Sequence

An MRI pulse sequence is a predetermined set of defined RF and gradient pulses, typically repeated multiple times during a scan. The time interval between pulses, as well as the amplitude and shape of the gradient waveforms, control the reception of the NMR signal and influence the characteristics of the MR images (e.g., their contrast). Over the years, numerous imaging methods have been introduced, with spin-echo (SE) pulse sequences being among the earliest developed and still widely used forms of

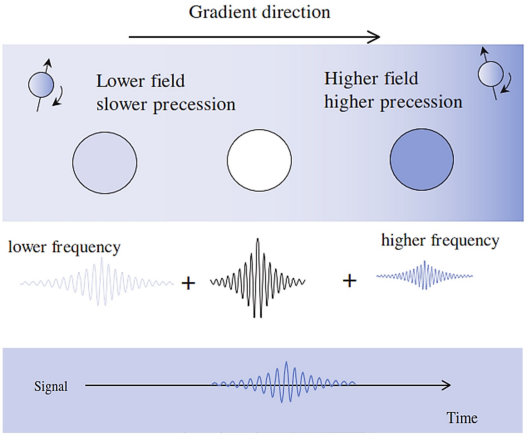


Figure 2.7: Frequency encoding results in position-dependent frequency changes. (Reproduced from [70])

MRI pulse sequences. These sequences offer versatility, allowing for the acquisition of T_1-weighted, proton density, and T_2-weighted images by adjusting the pulse sequence timing. SE sequences utilise an excitation pulse that is normally 90°, followed by a 180° refocusing pulse to reverse the effects of field inhomogeneities, as illustrated in Figure 2.8. The duration between successive excitations is termed the *repetition time* (TR), while the time around which the signal is acquired is called the *echo time* (TE).

Briefly, in a SE sequence, the process begins with the application of a selective 90° RF pulse, causing all proton spins to precess in phase and the M_{xy} signal to reach its maximum. Following this pulse, the M_{xy} signal rapidly decays due to free induction decay. Some proton spins are fast and lead, while others are slow and lag, leading to dephasing. After a time interval ($\frac{TE}{2}$), a 180° RF pulse is applied, which is twice as long as the 90° pulse in the transverse plane. This causes all proton spins to rotate by 180°. Another time interval ($\frac{TE}{2}$) passes, bringing the proton vectors back into phase, resulting in the M_{xy} signal reaching its peak again. This point is known as the echo, which is the signal that is measured. The decay in the signal from the original 90° pulse to the echo is attributed to the tissue characteristic T_2 effect, with the influence of magnetic field inhomogeneities minimised. These steps also involve the application of slice selection, frequency and phase encoding gradients to limit the measured signal to particular spatial locations, as discussed earlier.

To effectively decode the spatial information in the acquired signal, we need to per-

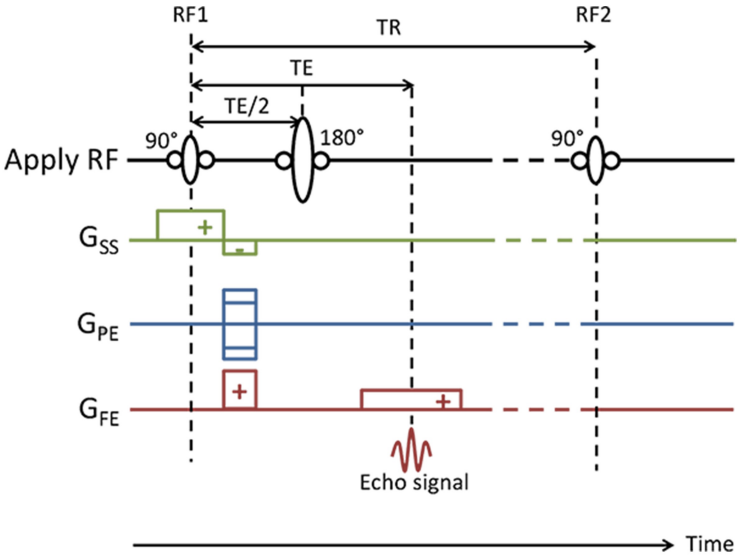


Figure 2.8: Schematic of a spin-echo pulse sequence. RF: radio frequency pulse; TE: echo time; TR: repetition time; G_{SS}: slice select gradient; G_{PE}: phase encoding gradient; G_{FE}: frequency encoding gradient.

form multiple separate phase encodings for each row of pixels in the the G_{PE} direction. This means repeating an entire MR pulse sequence for each different phase-encoding gradient step. For example, in an image of $N \times N$ pixels, N phase-encoding steps are required. This results in long acquisition times in MRI compared to other imaging modalities.

2.2.3 Spatial Encoding and k-Space

Thus far, we have discussed how to locate the net magnetisation within the body. After a 90° RF pulse, the rotation of the transverse magnetisation can be described in the complex plane with relaxation effects being ignored for simplicity as:

$$M_{xy}(\vec{r},t) = M_{xy}(\vec{r},0) \cdot e^{-t/T_2(\vec{r})} \cdot e^{-i\omega_0 t} \approx M_{xy}(\vec{r},0) \cdot e^{-i\omega_0 t} \tag{2.6}$$

Applied magnetic field gradients alter the field strength experienced by protons linearly with position:

$$\Delta B_z = \vec{G} \cdot \vec{r} \tag{2.7}$$

leading to spatially linear variation in the resonance frequency:

$$\omega = \gamma \vec{G} \cdot \vec{r} \tag{2.8}$$

This concept forms the basis for spatial encoding. With time-varying gradients, phase dispersion that depends on both time and location is induced in the magnetisation. Thus, the variation in frequency can be represented by an additive phase term (disregarding relaxation):

$$M_{xy}(\vec{r}, t) = M_{xy}(\vec{r}, 0) \cdot e^{-i\left(\omega_0 t + \ \phi(\vec{r},\ t)\right)} \tag{2.9}$$

The MRI signal detected through receiver coils originates from the precession of the transverse magnetisation, hence it is directly proportional to the spin density $m(\vec{r})$. By demodulating the received signal at a central frequency of $\omega_0 = \gamma B_0$, the resulting signal at a specific voxel v (disregarding relaxation effects) will become:

$$s_v(\vec{r}, t) = m(\vec{r})\ e^{-i\phi(\vec{r},\ t)} \tag{2.10}$$

At a given time point, the received signal is the summation of all signals originating from individual voxels:

$$S(t) = \int_{volume} m(\vec{r})\ e^{-i\phi(\vec{r},\ t)}\ d\vec{r} \tag{2.11}$$

The accumulated phase at an arbitrary point in space $\vec{r} = (x, y, z)$ can be derived from the cumulative gradient area, or the time integral of the activated gradients:

$$\phi(\vec{r}, t) = 2\pi \int_0^t \frac{\gamma}{2\pi} \vec{G}(\tau) \cdot \vec{r}\ d\tau \tag{2.12}$$

This implies that we can activate and deactivate the gradients to initiate and halt the phase accumulation, and also reverse the gradient polarity to negate any previous phase accumulation. By substituting

$$\vec{k}(t) = \frac{\gamma}{2\pi} \int_0^t \vec{G}(\tau)\ d\tau \tag{2.13}$$

we introduce the concept of k-space, which is a simplified representation of the phase accumulation resulting from magnetic field gradients [71]. The detected signal can be rewritten as:

$$S(t) = \int_{\vec{r}} m(\vec{r})\ e^{-i2\pi \vec{k}(t) \cdot \vec{r}}\ d\vec{r} \tag{2.14}$$

Interestingly, Equation 2.14 demonstrates that the imaged object and the acquired MR signal have a Fourier transform relation. This underscores the significance of the k-space representation, as it describes how MRI samples data in the Fourier transform domain, or the spatial frequency domain, of the object's net magnetisation. The relationship between image domain and k-space is analogous to that of time and temporal frequency domains in signal processing. In other words, k-space is essentially the Fourier domain counterpart of image space. By acquiring samples at a sufficient number of locations within k-space, according to Nyquist theory, the recovery of the imaged object can be achieved through an inverse Fourier transformation, as illustrated in Figure 2.9.

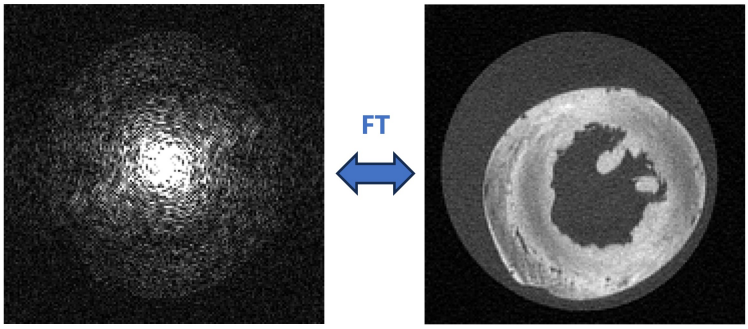


Figure 2.9: Correspondence between k-space (left) and image space (right) via the Fourier transform (FT).

2.2.4 k-Space Trajectories

We now define:

$$\begin{aligned} k_x(t) &= \frac{\gamma}{2\pi}\int_0^t G_x(\tau)d\tau \\ k_y(t) &= \frac{\gamma}{2\pi}\int_0^t G_y(\tau)d\tau \\ k_z(t) &= \frac{\gamma}{2\pi}\int_0^t G_z(\tau)d\tau \end{aligned} \tag{2.15}$$

$k_x(t)$, $k_y(t)$, and $k_z(t)$ describe a three-dimensional (3D) path, referred to as the k-space *trajectory* along which the signal is sampled. The speed and acceleration of k-space

traversal are controlled by the respective gradients:

$$\begin{aligned} \dot{\vec{k}} &= \frac{\gamma}{2\pi}\vec{G} \\ \ddot{\vec{k}} &= \frac{\gamma}{2\pi}\dot{\vec{G}} \end{aligned} \tag{2.16}$$

In MRI, the derivative of the gradient waveforms $\dot{\vec{G}}$ is called the *slew rate*. Owing to hardware and physiological constraints (e.g., peripheral nerve stimulation), the encoding gradients are bounded to a maximum amplitude and a maximum slew rate. These constraints pose challenges in designing trajectories since they need to be smooth with limited speed and acceleration. This is probably why the most common sampling patterns are based on simple geometrical shapes such as lines. Additionally, the rapid decay of the MR signal ($\simeq 50$ ms) often prevents the simultaneous measurement of all required data. Consequently, k-space trajectories typically consist of multiple segments, referred to as *shots*, which sequentially fill the k-space grid in a segmented acquisition approach.

The most widely used k-space trajectory is a Cartesian raster [72], where parallel and evenly spaced lines of the k-space grid are acquired during the frequency-encoded readout. Each Cartesian line corresponds to a specific value of the phase-encoding gradient. One drawback of the Cartesian raster trajectory is relatively long scan times [1] because each line necessitates a separate RF excitation pulse (or multiple pulses, if signal averaging is used). Another Cartesian-based k-space trajectory is echo planar imaging (EPI) [73]. EPI was developed to speed up the Cartesian raster by collecting multiple k-space lines in a single excitation in a zigzag pattern through the k-space. Numerous non-Cartesian trajectories have also been developed for MRI, mostly based on geometric patterns that were gradually improved or optimised. Figure 2.10 represents some of the most widespread sampling patterns used for 2D MRI scans.

2.2.5 Sampling of k-Space Trajectories

Shannon-Nyquist Theory

Most physical signals in the real world are analog, but modern signal processing systems predominantly operate in the digital domain. To effectively handle these analog signals in digital systems, a fundamental process known as “sampling” is required.

[1] Unless echo train spin refocusing is used.

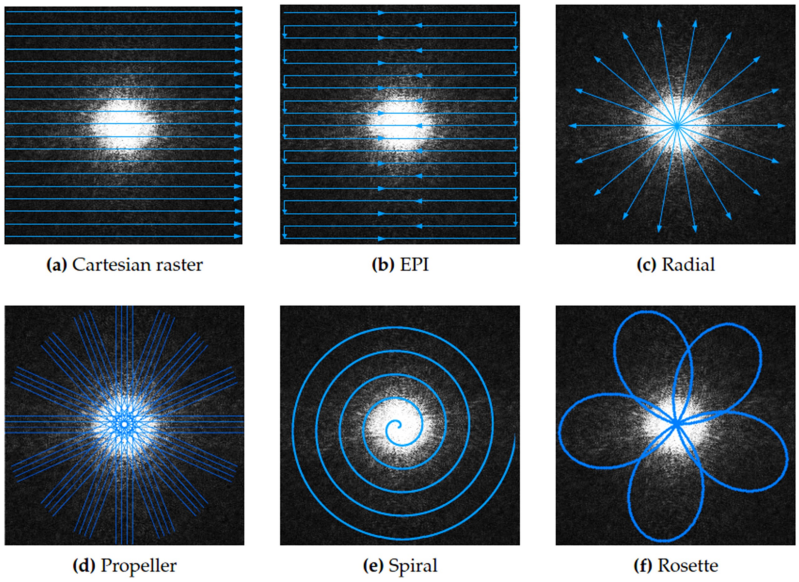


Figure 2.10: Examples of 2D k-space trajectories. The blue lines represent the traversal paths of the trajectories, while the gray-scale background depicts the k-space. (Adapted from [74])

Sampling involves converting an analog signal into a digital format by capturing discrete measurements at specific time intervals. However, it is generally impossible to perfectly reconstruct a signal from a series of sampled measurements without prior knowledge or certain assumptions about the signal. Over time, engineers have refined their understanding of practical assumptions and how they can be generalised.

The foundation of this transformation in signal processing lies in the groundbreaking work of several pioneers, including Shannon [58] and Nyquist [59]. The pivotal concept is known as the *Shannon-Nyquist Sampling Theorem*. According to this theorem, for continuous-time signals with a limited frequency range (band-limited signals), if the sampling rate (also known as the Nyquist rate), is more than twice the highest frequency component of the signal, then the original signal can be accurately reconstructed from its discrete samples. This theorem provides the sufficient (but not necessary) conditions under which a signal can be sampled and reconstructed without any loss of essential information.

It is important to note that, in many important and emerging applications, the conventional Nyquist sampling rate might be too high to implement, or the signal

bandwidth might be unknown in advance. Therefore, in practice, data compression is often employed immediately after sensing to address the logistical and computational challenges associated with high-dimensional data.

In MRI, during data acquisition, the continuous Fourier transform of the object is discretely sampled at a specific rate. To prevent image aliasing, also known as wrap-around artefacts, the Shannon-Nyquist criterion should be applied to k-space sampling. Let us first consider the case of 2D Cartesian sampling. It is convenient to introduce the field of view (FOV), which defines the extent of the object in the spatial domain to be imaged. With an imaging FOV of $\text{FOV}_x \times \text{FOV}_y$ and a matrix size of $N_x \times N_y$, we achieve a resolution of $\Delta x = \frac{\text{FOV}_x}{N_x}$ and $\Delta y = \frac{\text{FOV}_y}{N_y}$.

To comply with the Shannon-Nyquist criterion and attain the desired resolution, the k-space sampling should satisfy the following conditions:

$$\Delta k_x \leq \frac{1}{\text{FOV}_x} \quad \text{and} \quad \Delta k_y \leq \frac{1}{\text{FOV}_y} \tag{2.17}$$

$$k_{x_{max}} \leq \frac{1}{2\Delta_x} \quad \text{and} \quad k_{y_{max}} \leq \frac{1}{2\Delta_y} \tag{2.18}$$

where k_{max} denotes the maximum spatial frequency sampled in each k-space axis and Δk represents the spacing between consecutive spatial frequencies. From Equation 2.18, it can be seen that the k_{max} relates to the spatial resolution and determines how good details are captured in each spatial dimension. Assuming sufficient SNR, the higher the k_{max} or the farther frequencies we capture in the k-space, the better the details and fine structures will be resolved in the reconstructed image.

In the case of non-Cartesian sampling, such as radial or spiral schemes, due to the irregular distribution of samples in k-space, as illustrated in Figure 2.11, it is not straightforward to derive a direct relationship between k-space sampling requirements and the Shannon-Nyquist criterion; however, two criteria can be introduced: an azimuthal sampling interval Δk_θ along a radial line or spiral interleaf and a radial sampling interval Δk_r. For example, in radial sampling, the sample spacing additionally depends on the number of radial lines or *spokes* (n_s), which can be freely chosen as another sequence parameter. However, assuming a matrix size of $N \times N$, it is commonly recommended in the literature to choose the number of spokes according to [75]:

$$n_s \geq \frac{\pi}{2} N \tag{2.19}$$

This selection ensures that the distance between samples on neighbouring spokes satisfies the Shannon-Nyquist criteria, and an isotropic spatial resolution of $\Delta x = FOV/N$ is guaranteed. Therefore, Nyquist-sampled 2D radial trajectories require more encoding steps than a comparable Cartesian acquisition; approximately 1.57 times more for spokes spanning over π rad.

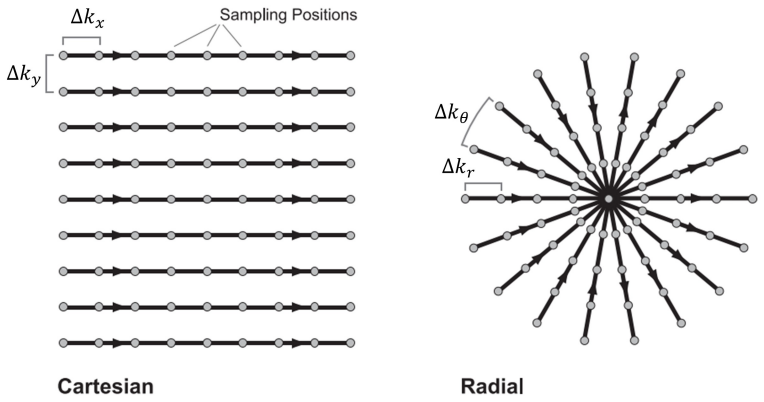


Figure 2.11: Illustration of the Cartesian and non-Cartesian sampling schemes.

One approach to adjust the distance between the measured samples is to modify the temporal sampling rate of the receiver. A higher sampling rate results in a smaller sample distance in k-space. Additionally, from Equation 2.16, it is evident that the speed of movement in k-space is dependent on the amplitude of the gradient fields. Therefore, the sampling distance can be reduced either by increasing the sampling rate or by reducing the gradient strength. While the latter extends the total duration of the sampling, it provides an enhanced SNR. This is because lowering the gradient strength reduces the spectrum of Larmor frequencies used for spatial encoding, while the noise floor can be assumed to be uniformly distributed in the frequency space. Consequently, it is useful to introduce the terms "readout bandwidth" or "receiver bandwidth" as:

$$BW = \frac{1}{\Delta t} \tag{2.20}$$

Here, Δt denotes the sampling interval or dwell time of the receiver. Frequently, this value is normalised by the number of acquired samples N, resulting in the bandwidth per pixel:

$$BW/pixel = \frac{1}{\Delta t \cdot N} = \frac{1}{T_s} \tag{2.21}$$

where T_s represents the total duration of the sampling period. With a constant number of samples, a high value of the bandwidth corresponds to a short acquisition time with low SNR, while a low value indicates a longer readout time with higher SNR.

In practice, the operator typically specifies the FOV, the number of samples, and the bandwidth per pixel, while the other parameters are derived from these inputs. The required dwell time of the receiver is thus given by:

$$\Delta t = \frac{1}{BW/pixel \cdot N} \tag{2.22}$$

and the amplitude of the readout gradient can be computed by:

$$G_{ro} = \frac{1}{\gamma \cdot FOV \cdot \Delta t} \tag{2.23}$$

2.2.6 Image Reconstruction

To ensure the feasibility of a later reconstruction of the object, it is required to collect sufficient Fourier samples along the trajectory, and an inverse Fourier transform can then be employed to recover the image:

$$I(\vec{r}) = \int_{\vec{k}} S(\vec{k}(t))\ e^{i\vec{k}(t)\cdot\vec{r}}\ d\vec{k} \tag{2.24}$$

It should be noted that the received signal is complex (see Equation 2.11), and thus the resulting image matrix has complex entries, but typically, the magnitude values are presented (Figure 2.9).

In the Cartesian sampling scheme, all k-space samples lie on an equidistant grid. This allows for efficient reconstruction through a fast Fourier transformation (FFT). The samples are arranged on a raw data matrix according to their respective k-space position, and then an FFT is applied in each dimension:

$$I(x,y) = \sum_{k_x} \sum_{k_y} S(k_x,k_y) e^{-i2\pi(xk_x+yk_y)} \tag{2.25}$$

Reconstruction of Irregular Data

When k-space is sampled along non-equidistant points, e.g., in radial/spiral sampling, direct calculation of Equation 2.25 becomes time-consuming and practically infeasible, even with modern high-performance computer systems. Adding to the complexity, the non-uniform sampling density in non-Cartesian trajectories, i.e., higher density in the

centre of k-space than in the outer regions, yields significant blurring in the reconstructed image. This issue occurs because the discrete Fourier transform in Equation 2.25 does not account for the locally varying distances between sample points. Therefore, the signal from densely sampled regions accumulates, leading to an overestimation of low spatial frequencies.

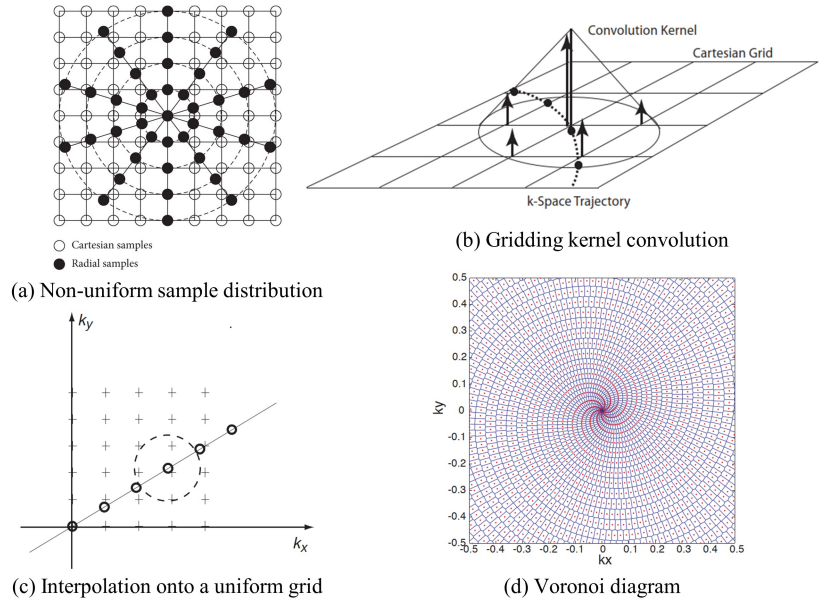
One approach to reconstruct images from non-Cartesian data is to resample the data onto a Cartesian grid before applying an FFT. This technique is called *regridding*, which has been introduced to medical imaging by O'Sullivan [76]. The idea involves interpolating the measured samples onto a uniform grid, as illustrated in Figure 2.12. Specifically, each sample is convolved with a small kernel, wide enough to cover the neighbouring grid points. The convolution introduces modulation effects in the image space, highlighting the importance of selecting an appropriate interpolation technique to mitigate these effects. This approach has the advantage of exploiting all the data, making it "SNR efficient". However, a density compensation is required to account for the possibility that the samples may be concentrated in specific areas of k-space, resulting in an increased contribution to surrounding grid points. Another consideration is the selection of the reconstruction grid density (i.e., the grid oversampling factor α). The basic gridding process follows these straightforward steps:

- Convolve the k-space data with a gridding kernel.
- Resample the convolved data onto evenly spaced grid points.
- Apply the inverse FFT (IFFT) to the resampled data.
- Multiply the result by an apodisation correction function ($\frac{1}{c(x,y)}$) to account for the convolution effects.

The reconstructed image $\hat{m}(x, y)$ after the entire gridding process is given by [77]:

$$\hat{m}(x,y) = \mathrm{IFFT}\left\{ \Big[[S(k_x,k_y) \times W(k_x,k_y)] * C(k_x,k_y) \Big] \times \Pi\Big(\frac{k_x}{\frac{\Delta k_x}{\alpha}}, \frac{k_y}{\frac{\Delta k_y}{\alpha}} \Big) \right\}(x,y) \times \frac{1}{c(x,y)} \tag{2.26}$$

here, $*$ represents the convolution operator. $S(k_x, k_y)$ is the acquired MR signal, $W(k_x, k_y)$ is a weighting function for density compensation, $C(k_x, k_y)$ is the convolution kernel, and Π is the comb function associated with the Cartesian grid. $c(x, y)$ is the reciprocal of the inverse Fourier transform of the kernel $C(k_x, k_y)$ and is determined completely by the choice of kernel.



(a) Non-uniform sample distribution

(b) Gridding kernel convolution

(c) Interpolation onto a uniform grid

(d) Voronoi diagram

Figure 2.12: Basic regridding idea: Data samples lie on some non-uniform grid in k-space. Each data point is conceptually convolved with a gridding kernel, and that convolution is evaluated at the adjacent grid points. The Voronoi plot shows k-space region associated with each sample. The density corresponds to the inverse of the area of these regions.

The Kaiser-Bessel kernel is commonly used for $C(k_x, k_y)$ because it offers a good trade-off between accuracy and computational efficiency [77]:

$$K_{KB}(d) = \begin{cases} \frac{1}{L} I_0(\beta\sqrt{1-(2d/L)^2}) & |d| \leq \frac{L}{2} \\ 0 & |d| > \frac{L}{2} \end{cases} \tag{2.27}$$

Here, $d = |\mathbf{k} - \mathbf{k}_j|$ ($\mathbf{k} = (k_x, k_y)$) represents the distance between the jth sample and the interpolated position, L denotes the desired kernel width, $I_0(d)$ is the zero-order modified Bessel function of the first kind, and β is a shape parameter determined by an equation provided by Beatty et al. [78].

Selecting the appropriate kernel width is crucial for minimising artefacts and achieving accurate reconstruction. This is typically determined through empirical analysis of the point spread function (PSF) in image space. In practice, kernel widths ranging from 1.5 and 4 grid points are commonly used, with an oversampling factor often set between 1.5 and 2. Additionally, a trade-off between interpolation accuracy and computational

efficiency is considered when choosing these parameters. The density compensation $W(k_x, k_y)$ can be calculated analytically using geometrical considerations in certain cases (e.g., radial sampling) or by employing a Voronoi diagram (Figure 2.12(d)).

2.3 Diffusion MRI

Diffusion magnetic resonance imaging (dMRI) is an emerging and promising non-invasive technique for assessing the 3D microstructure of tissues [79]. This technique has gained widespread use in studying the human brain but has seen limited application in the heart [80, 81]. DMRI is utilised to quantify the diffusion properties of water molecules within biological tissues by estimating the average displacement of water molecules. This provides valuable insights into the microstructure of the diffusion environment within tissues.

2.3.1 Diffusion Principles

Free Diffusion

"Diffusion" in biological tissues typically refers to the random, Brownian motion of water molecules driven by thermal energy; molecules collide with each other with no net displacement in any preferential direction, as illustrated in Figure 2.13.

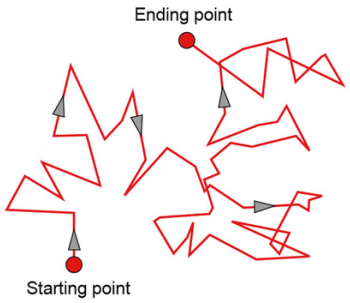


Figure 2.13: A schematic of the molecular diffusion trajectory of a single water molecule

Since diffusion is a random process, meaning that molecules starting from the same initial position and diffusing for the same amount of time are likely to end up at different locations, the concept of the diffusion displacement distribution, also referred

to as the displacement probability density function (PDF), was introduced [82]. This PDF, denoted as $P(\mathbf{r}'|\mathbf{r}, t)$, represents the possibility for a given water molecule to move from its initial position $\mathbf{r}$ to another position $\mathbf{r}'$ within a specified diffusion time (t). In a perfectly uniform and homogeneous medium, diffusion is equally random and isotropic, meaning there is an equal probability of movement in all directions. Thus, the PDF takes on a Gaussian shape. This is determined through analytical solutions of the diffusion equation without the imposition of geometric constraints (i.e. diffusion barriers) and can be expressed as:

$$P(\mathbf{r}|\mathbf{r}', t) = \frac{1}{\sqrt{(4\pi Dt)^3}} \exp\Big(- \frac{|\mathbf{r}' - \mathbf{r}|^2}{4Dt} \Big) \tag{2.28}$$

where D is called the diffusion coefficient.

Hindered and Restricted Diffusion

As mentioned, in free diffusion, water molecules can move freely in all spatial directions, such as in cerebrospinal fluid. However, in biological tissues, the movement of water molecules encounters varying degrees of obstacles in different directions like vascular structures, cell membranes, etc. This leads to "restricted" or "hindered" diffusion, where movements of water molecules are constrained in certain spatial directions depending on their environment, as illustrated in Figure 2.14, highlighting the difference between the two. In such environments, we anticipate anisotropic diffusivity behavior, and additional considerations come into play in describing the PDF. These include incorporating boundary conditions to account for the effects of permeable and semi-permeable barriers, as well as considering the geometric microstructural properties of the tissue being studied. The PDF for restricted and hindered diffusion depends on factors such as diffusion time, diffusion coefficient, and, most importantly, the structures present in the diffusion environment. Different tissues in the human body have unique cellular architectures and specific proportions of intra- and extracellular compartments. Consequently, they possess characteristic diffusion properties. Variations in the PDF can thus provide valuable insights into tissue structure and alterations in physical properties, making it a valuable tool for understanding and diagnosing diseases.

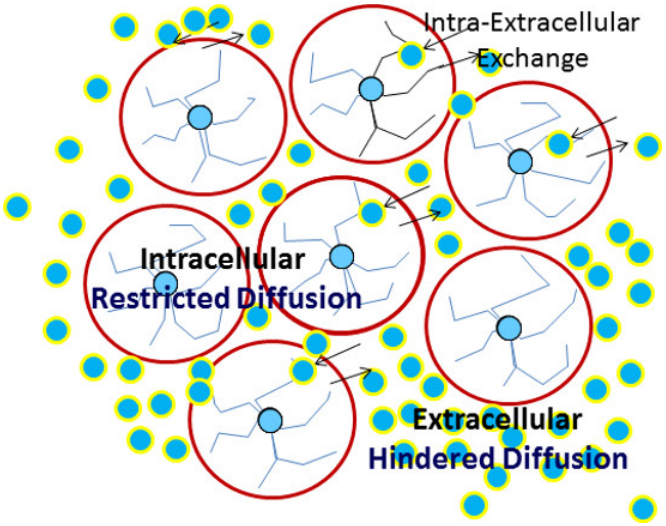


Figure 2.14: Complex diffusion in a tissue: Schematic representation of water diffusion in a cellular system showing the movement of water molecules in the extracellular compartment (green dots) and intracellular compartment (blue dots), with arrows indicating water exchange between these two compartments. (Adapted from [83])

2.3.2 Diffusion-weighted MR Signal

The effect of diffusion on the MRI signal was first observed by Hahn in 1950 [84]. A few years later, Torrey [85] introduced an extended version of the Bloch equations that accounted for molecular diffusion. He demonstrated that in the presence of molecular diffusion, the variation in transverse magnetisation can be described as:

$$\frac{\partial \vec{M}}{\partial t} = \gamma \vec{M} \times \vec{B}_0 + \begin{pmatrix} \frac{M_x}{T_2} \\ \frac{M_y}{T_2} \\ \frac{M_0 - M_z}{T_1} \end{pmatrix} + D\nabla^2 \vec{M} \tag{2.29}$$

where D is the diffusion coefficient and ∇^2 denotes the Laplacian operator. The first two terms on the right-hand side of Equation 2.29 correspond to the original Bloch equation, while the third term represents the contribution of molecular diffusion. This Equation is referred to as the "Bloch-Torrey" equation. To capture molecular diffusion information, modifications were made to the traditional MRI pulse sequence, resulting in the development of a new imaging technique known as diffusion-weighted magnetic resonance imaging (DW-MRI) [86].

Stejskal-Tanner Sequence

The initial diffusion-weighted sequence was pioneered by Stejskal and Tanner in 1965 [87]. This sequence was based on a conventional spin-echo (SE) sequence, but with the addition of a symmetric pair of diffusion encoding gradients on either side of the refocusing 180° RF pulse in each of the three encoding directions, as shown in Figure 2.15. These gradients have equal strength (G) and duration (δ) and are symmetrically positioned around the 180° pulse, with a temporal separation between their starting points denoted as Δ. This sequence effectively labeled the magnetic properties of spin transverse magnetisation in relation to their spatial positions and is commonly referred to as the pulsed-gradient spin-echo (PGSE) sequence.

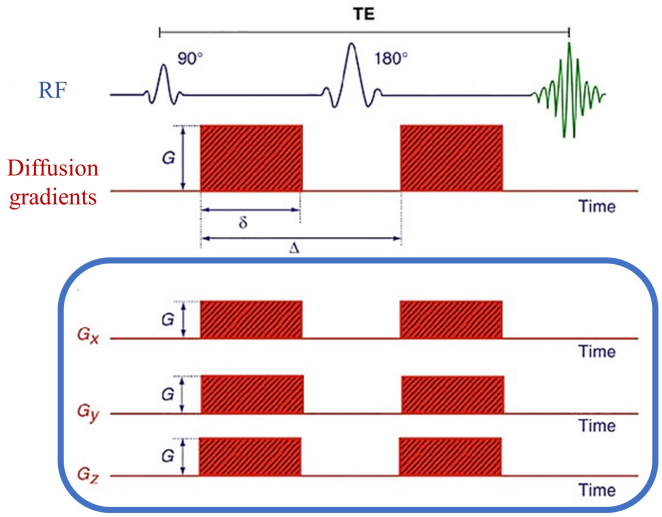


Figure 2.15: Schematic illustration of the Pulsed Gradient Spin Echo (PGSE) sequence introduced by Stejskal and Tanner. The sequence includes 90° excitation and 180° refocusing RF pulses to generate a spin echo, with two strong gradient pulses of magnitude (G) and duration (δ), separated by time interval (Δ), applied for diffusion weighting. (Reproduced from [88])

Following the initial 90° RF pulse, the spins precess in the xy-plane at the Larmor frequency. The first diffusion gradient offsets the phase of the spins by an amount proportionate to their average position during the application of the gradient (δ). In essence, the first diffusion encoding gradient serves to label the spins based on their positions along its direction. The second diffusion gradient, due to the 180° refocusing

pulse, leads to equal but opposite dephasing, if the spins have not moved during the time interval (Δ) between the two gradients. In other words, the second gradient would reverse the phase shift induced by the first gradient, ensuring complete rephasing of the spins if there was no diffusion. However, as the spins move randomly due to diffusion, the magnetic field experienced during the application of the first gradient differs from that experienced during the second gradient, leading to incomplete rephasing and phase dispersion. This dispersion causes signal loss in the resulting diffusion-weighted MR image, as illustrated in Figure 2.16. The farther the spins have diffused during the diffusion time, the less perfect the rephasing, resulting in a smaller amplitude of the final signal. Therefore, more diffusion is indicated by a more attenuated signal.

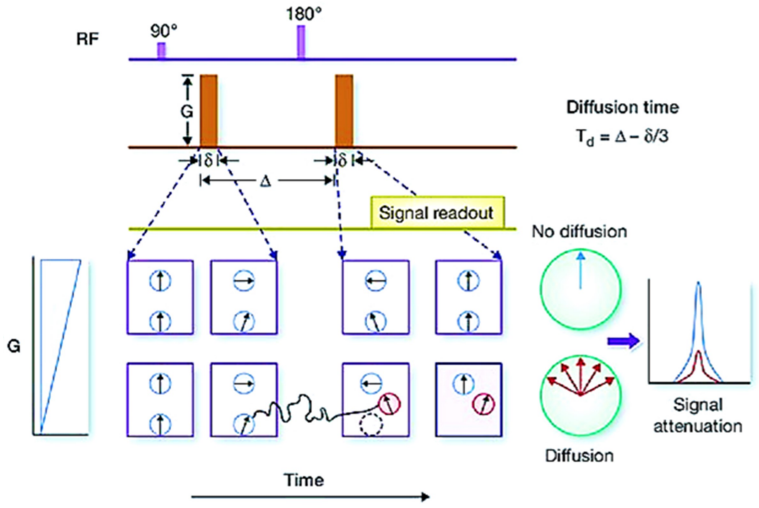


Figure 2.16: Schematic representation of the dephasing process for spins within a single voxel influenced by a standard Stejskal-Tanner sequence. (Reproduced from [89])

The diffusion signal detected by this sequence is given by:

$$\frac{S(b)}{S(0)} = \exp\Big(-\gamma^2 \delta^2 G^2 (\Delta - \frac{\delta}{3}) D \Big) \tag{2.30}$$

where γ denotes the gyromagnetic ratio, and $S(0) = M_0 \exp(-t/T_2)$ represents the signal obtained when no diffusion gradient is applied. A comprehensive derivation of Equation 2.30 can be found in [90]. This formula is commonly expressed as:

$$\frac{S(b)}{S(0)} = \exp(-bD) \tag{2.31}$$

where $b = \gamma^2\delta^2G^2(\Delta - \delta/3)$ is called the b-value (measured in s/mm^2), which depends on the characteristics of the protocol employed for acquiring the diffusion-weighted MR signal (pulse sequence shown in Figure 2.15). The concept of this diffusion weighting factor was initially introduced by Le Bihan in 1986 [91]. Equation 2.31 remains valid for an arbitrary diffusion-sensitising gradient pulse waveform $\vec{G}(t)$ applied during a total time T, with the b-value calculated according to:

$$b = \gamma^2 \int_0^T dt \left[\int_0^t d\tau \vec{G}(\tau) \right]^2 \tag{2.32}$$

It should be noted that in the original derivation of the Equation 2.30, rectangular-shaped pulses were considered. However, this shape cannot be realised in real MRI scanners because the gradient pulse cannot instantaneously transition from zero to its maximum value. Therefore, it was later corrected by [92] to accommodate the non-zero pulse ramp time ζ as:

$$S(b) = \exp\left(-\gamma^2\delta^2G^2\left[(\Delta - \frac{\delta}{3}) + \frac{\zeta^3}{30} - \frac{\delta\zeta^2}{6}\right]D\right) \tag{2.33}$$

It is also worth noting that with this sequence, diffusion information can be measured along only one direction, which is defined by the orientation of the applied pair of diffusion encoding gradients. In practice, the sequence is often used with different gradient orientations and/or amplitudes to enable the measurement of diffusion in three dimensions.

2.3.3 Diffusion Tensor Imaging (DTI)

In traditional DW-MRI, the diffusion process is simply described by a single scalar parameter, referred to as the apparent diffusion coefficient (ADC), using a single exponential decay model:

$$ADC = \frac{1}{b} \ln\left(\frac{S(b)}{S(0)}\right) \tag{2.34}$$

While studying diffusion along a single direction can offer valuable insights into many pathological conditions, it lacks the ability to provide information about the directional characteristics of diffusion. As mentioned earlier, we expect diffusion to display anisotropy in biological tissues. To explore this anisotropy, the concept of a 3D approach known as the diffusion tensor was introduced by Basser et al. [33].

Diffusion Tensor Model

In an anisotropic scenario, it is necessary to consider the measurement direction. Therefore, the scalar diffusion term D in Equation 2.31 is replaced by a rank-2 tensor. Consequently, Equation 2.31 transforms into:

$$\frac{S(b)}{S(0)} = \exp(-b\mathbf{g}^T\mathbf{D}\mathbf{g}) \tag{2.35}$$

Here, $\mathbf{D}$ is the so-called Diffusion Tensor (DT), which is a 3×3 matrix of the form:

$$\begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{pmatrix}$$

The diagonal components D_{xx}, D_{yy} and D_{zz} represent the diffusion coefficients measured along three mutually orthogonal directions corresponding to the three axes of the scanner magnet. The six off-diagonal terms represent the correlation between each pair of these three orthogonal directions. In the case of isotropic diffusion, the diagonal elements are all the same ($D_{xx} = D_{yy} = D_{zz} = ADC$), and the tensor takes on a diagonal form. However, for anisotropic diffusion, the diagonal components are generally unequal, and non-zero off-diagonal terms exist. According to the thermodynamic laws, diffusivities in mirror-image directions should be the same; i.e., $D_{ij} = D_{ji}$, with $i, j = x, y, z$. Therefore, the diffusion tensor is a 3×3 symmetric, positive-definite matrix, with only six unique elements.

For this positive-definite tensor, it is possible to establish an orthogonal coordinate system, with the three unit normal vectors referred to as *eigenvectors*. In this new system, the tensor transforms into a diagonal form with three positive entries, known as *eigenvalues*:

$$\mathbf{DE} = \mathbf{EA} \quad \text{where} \quad \mathbf{E} = (e_1|e_2|e_3); \quad \mathbf{A} = \begin{vmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{vmatrix} \tag{2.36}$$

where e_1, e_2, and e_3 denote the eigenvectors and λ_1, λ_2, and λ_3 are the eigenvalues. By convention, eigenvalues are labeled in descending order of magnitude ($\lambda_1 \geq \lambda_2 \geq \lambda_3$). The major eigenvector of the diffusion tensor corresponding to λ_1 indicates the primary diffusion direction, i.e., the direction of fastest diffusion. In anisotropic fibrous tissues,

such as the brain, the major eigenvector determines the fibre tract axis of the tissue. However, this interpretation holds true only in regions where fibre tracts do not cross or branch.

A simple method to visualise the diffusion tensor is by using an ellipsoid, which may be more or less elongated depending on the extent of anisotropy, as shown in Figure 2.17. The shape and size of the ellipsoid in a voxel give a visual representation of how likely water is to move in different directions within that voxel. In other words, this ellipsoid represents an iso-surface of diffusion probability: its axes align with the eigenvectors representing the diffusion directions, and their lengths correspond to the respective eigenvalues representing the diffusivity along each eigenvector direction.

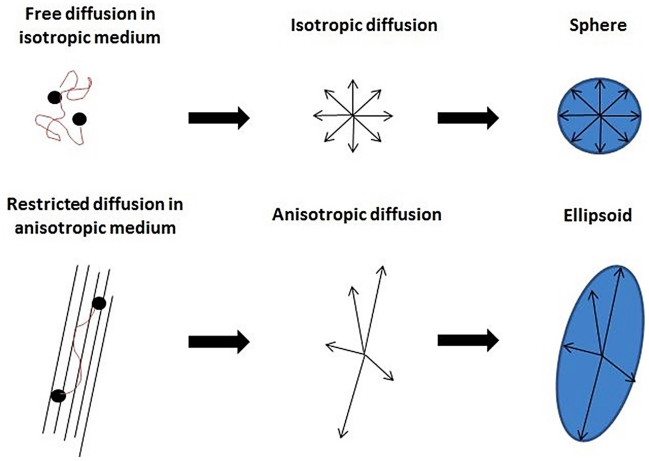


Figure 2.17: Schematic representation of the diffusion tensor. Arrows show the diffusion pattern in three dimensions.

Tensor Estimation

Considering the symmetry of $\mathbf{D}$, the logarithm of the diffusion signal in Equation 2.35 in each voxel can be written as:

$$\ln S = \ln S_0 - b_{xx}D_{xx} - b_{yy}D_{yy} - b_{zz}D_{zz} - 2b_{xy}D_{xy} - 2b_{xz}D_{xz} - 2b_{yz}D_{yz} \tag{2.37}$$

For simplicity, the right-hand side of Equation 2.37 can be written as:

$$\ln S = BX \tag{2.38}$$

with

$$\begin{aligned} B &= (-b_{xx}, -b_{yy}, -b_{zz}, -2b_{xy}, -2b_{xz}, -2b_{yz}, 1) \\ X &= [D_{xx}, D_{yy}, D_{zz}, D_{xy}, D_{xz}, D_{yz}, \ln S_0]^T \end{aligned} \tag{2.39}$$

To determine the six unique components of the diffusion tensor, a minimum of seven measurements are required. This set of measurements includes one non-weighted measurement (S_0) and at least six weighted measurements (S_i, where $i = 1, ..., N$, with $N \geq 6$) with diffusion gradients $g_i = (x_i, y_i, z_i)$ applied at different non-collinear directions. We then need to solve the following set of linear equations:

$$\ln S_i = B_i X \tag{2.40}$$

When $N = 6$, the diffusion tensor can be directly calculated using:

$$X = B_i^{-1} \ln S_i \tag{2.41}$$

However, when the number of diffusion encoding directions is more than six, there is no exact inverse B_i^{-1}. In this case, Equation 2.40 can still be solved by calculating a pseudo-inverse matrix $B^\dagger$:

$$X = (B_i^T B_i)^{-1} B_i^T \ln S_i = B^\dagger \ln S_i \tag{2.42}$$

Various methods exist for calculating the so-called pseudo-inverse matrix, which allows for fitting the acquired data to a rank-two tensor with six degrees of freedom. The most commonly used methods include linear least squares (LLS), weighted linear least squares (WLLS), and non-linear least squares (NLLS), along with their constrained counterparts that enforce the diffusion tensor to be positive definite. A comprehensive description of these fitting methods is beyond the scope of this thesis, and additional information on this topic can be found in [79].

Diffusion Parameters

Once the diffusion tensor is calculated, several diffusion characteristics are typically extracted to describe the underlying structures of the diffusion environment. Some key parameters commonly used for cardiac tissue characterisation are:

Eigenvectors: The direction of the primary eigenvector e_1 of the diffusion tensor, corresponding to the largest eigenvalue, is considered to represent the local tissue orientation. However, defining a primary cellular direction in the heart, given the complex arrangement of cardiomyocytes, is not as straightforward as in other muscle fibers. Therefore, it is common practice to determine a local principal direction aligned with the long axis of the cylindrical-shaped cardiomyocyte at each location within the heart [93]. Consequently, the primary eigenvector e_1 is believed to indicate the mean orientation of the cardiomyocytes in the heart [42]. Additionally, evidence suggests that the second eigenvector of the diffusion tensor is linked to the direction of the sheetlets in the myocardium [40, 44]. And the third eigenvector aligns with the direction of the sheetlet normal (perpendicular to the sheetlets). Figure 2.18 illustrates the relationship between the diffusion tensor eigenvectors and the orientations of cardiomyocytes and sheetlets.

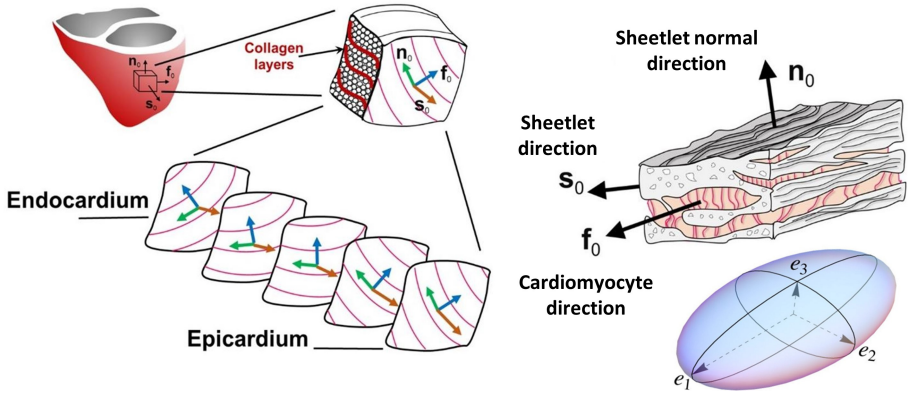


Figure 2.18: Schematic diagram of the correspondence of the diffusion tensor eigenvectors (e_1, e_2, e_3) to the cardiomyocyte and sheetlet orientations. (Reproduced from [94])

Mean Diffusivity: Mean diffusivity (MD) provides information about the average diffusion within a voxel and is typically expressed in units of mm^2/s, with higher values indicating greater diffusion. MD is calculated as the average of the three apparent diffusion coefficients (ADCs) measured along the three axes of the diffusion tensor. Assuming that the tensor axes align with the x, y, and z directions, the MD can be

expressed as:

$$\mathrm{MD} = \frac{\mathrm{trace}\mathbf{D}}{3} = \frac{D_{xx} + D_{yy} + D_{zz}}{3} = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3} \tag{2.43}$$

Fractional Anisotropy: Fractional anisotropy (FA) quantifies the extent to which the diffusion profile differs from the isotropic case or how much the diffusion ellipsoid deviates from a spherical shape. Higher values of FA indicate the presence of a higher level of oriented structures. This measure is dimensionless, and its value falls within the range of 0 (indicating isotropy) to 1 (representing complete anisotropy). It is defined as:

$$\mathrm{FA} = \sqrt{\frac{3}{2}}\sqrt{\frac{(\lambda_1 - \mathrm{MD})^2 + (\lambda_2 - \mathrm{MD})^2 + (\lambda_3 - \mathrm{MD})^2}{(\lambda_1)^2 + (\lambda_2)^2 + (\lambda_3)^2}} \tag{2.44}$$

Helix Angle: Helix Angle (HA) is typically defined as the angle between the short-axis plane (circumferential direction) and the primary eigenvector, which is assumed to align with myocytes, projected onto the local endocardial tangent plane. A schematic representation of the helix angle and examples of both positive and negative helix angles are illustrated in Figure 2.19.

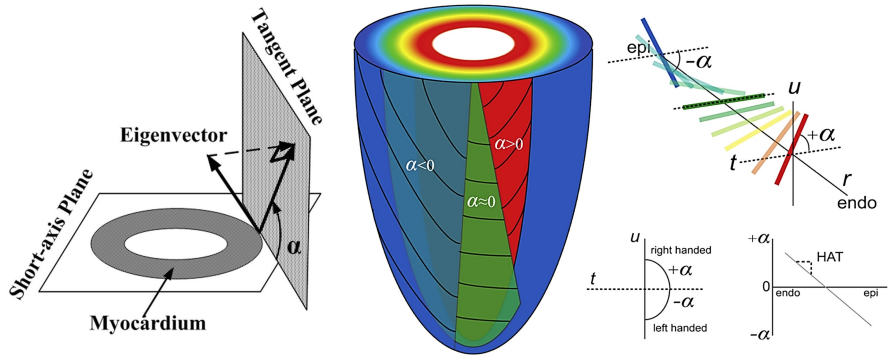


Figure 2.19: Helix angle definition. HAT stands for helix angle transmurality, which refers to the slope of the helix angle versus the transmural depth of the myocardium. This parameter describes how the orientation of cardiomyocytes changes from the epicardium (epi) to the endocardium (endo). (Reproduced from [95])

The orientation of cardiomyocytes within the heart wall changes smoothly from approximately -60° at the endocardium to about +60° at the epicardium. When viewed from the apex, the sub-epicardial cardiomyocytes follow a left-handed helical path,

while the sub-endocardial cardiomyocytes follow a right-handed helical path. In the mid-wall region of the heart, the cellular pathways are predominantly circumferential. This arrangement of cardiomyocytes, which is consistent across different species, has been shown to evenly distribute stress and strain across the ventricular wall, allowing for optimal mechanical functioning of the heart [96].

2.4 Summary

In this chapter, we first introduced the principles of magnetic resonance imaging (MRI), covering fundamental concepts of proton magnetisation, precession, excitation, and relaxation. Additionally, we explained spatial encoding techniques, signal detection, and the k-space representation, aiming to provide a clear understanding of how MRI functions and how MR images are generated. Following this, we outlined the physical process of molecular diffusion, including both free diffusion and diffusion within biological tissues. Lastly, we discussed how MRI captures diffusion-sensitive signals and its unique ability to non-invasively characterise the tissue microstructure by examining the diffusion of water molecules within the tissue. Changes in the tissue microstructure due to pathologies can affect the diffusion of water molecules, thereby making them detectable by diffusion MRI.

Furthermore, we explained the most commonly utilised dMRI technique in the heart, namely diffusion tensor imaging (DTI). The myocardium displays high anisotropy, making it well-suited for characterisation using key parameter metrics derived from the diffusion tensor.

Chapter 3

State-of-the-Art Techniques in Cardiac DTI

DTI has proven to be a valuable and efficient tool for the non-destructive analysis of the 3D structure of the myocardium in both healthy and diseased hearts. It has been utilised to investigate the remodelling of infarcted myocardium in various species, including pigs [97], mice [53], humans, sheep, and rats [98], often in ex vivo studies. Myocardial remodelling involves structural and functional alterations and has been observed in the surrounding adjacent and remote zones in infarcted myocardium. Despite the exceptional attributes, diffusion MRI techniques, including DTI, face several inherent limitations, some of which pose significant challenges in cardiac DTI applications and have thus hindered their routine clinical use. In this chapter, we will discuss the methodological challenges in cardiac DTI (cDTI) and explore the developments and strategies employed to mitigate them.

3.1 Methodological Challenges in cDTI

3.1.1 Motion Artefacts

Conventional dMRI techniques are highly sensitive to bulk motion, due to the strong magnetic field gradients used to achieve diffusion weighting. In the absence of bulk motion and under on-resonance conditions, a typical Stejskal-Tanner diffusion experiment should result in completely refocused transverse magnetisation, with no net residual phase (ϕ) in the transverse plane. However, in the presence of bulk motion, diffusion encoding gradients generate a non-zero net residual phase in the transverse plane ($\phi \neq 0$), resulting in an additional signal attenuation component:

$$S(b) = S(0)e^{-bD}e^{-i\phi} \tag{3.1}$$

Here, $\phi = \gamma \int \vec{G} \cdot \vec{r}\, dt$ represents the phase shift caused by the displacement $\vec{r}$ of proton spins during the application of the gradient pulse $\vec{G}$. Unlike the diffusion attenuation term, e^{-bD}, caused by the molecular displacements due to water self-diffusion, the attenuation due to the tissue motion, $e^{-i\phi}$, is spatially varying and several orders of magnitude larger. Consequently, it results in a substantial loss of signal.

Therefore, the primary technical challenge in cardiac diffusion imaging is related to the heart motion, which arises from two main sources: the rhythmic contraction of the heart due to its pumping action and the displacement of the heart due to breathing (respiratory motion). While techniques such as breath-holding or respiratory navigation

have been effective in mitigating the respiratory motion, compensating for the effects of cardiac contraction is more complex and demanding. It involves using ECG signals to trigger the acquisition at specific time points, referred to as "sweet spots". Nonetheless, variations between heartbeats, especially in cases of arrhythmias or ectopic heartbeats, can introduce artefacts. Additionally, in clinical settings, requiring patients to hold their breath can pose challenges, particularly for those who are unstable or have impaired respiratory function. Failure to account for the impact of cardiac motion can lead to severe blurring or signal distortions in DW MR images [35, 99], compromising the reliability and interpretability of diffusion MRI data. Addressing and mitigating the motion artefacts are thus essential for ensuring the clinical utility and accuracy of diffusion MRI in cardiac tissue characterisation and disease diagnosis.

3.1.2 Low Signal-to-Noise Ratio

Another challenge is posed by the inherently low SNR of diffusion-weighted (DW) images. As discussed in the previous chapter, the DW signal relies on the attenuation of the classical magnetic resonance signal by the diffusion process. In addition, the considerably shorter T_2 relaxation time of the myocardium compared to brain white matter imposes limitations on the echo time, contributing to lower signal amplitudes in the heart [37]. To enhance SNR, averaging of multiple acquisition repetitions can be performed. However, before averaging, registration techniques are essential to mitigate smoothing and partial volume artefacts resulting from differences in the position of the heart during each acquisition

3.1.3 Long Scan Time

Estimating the diffusion tensor theoretically necessitates the use of a minimum of six linearly independent diffusion-encoding directions along with a non-weighted, or b_0, measurement, which already leads to extended scan times. However, practical research has shown that over-specifying the tensor with more diffusion encoding steps (up to 30 gradient directions) is desirable for accurate reconstruction of the diffusion tensor and robust derivation of the eigenvectors and diffusion-related parameter maps. Added to this, often acquiring 2 to 6 repetitions of each acquisition to achieve sufficient SNR and obtaining multiple slices spanning the entire myocardium for 3D microstructure mapping is required. Consequently, this results in prolonged scan durations exceeding

20 minutes, which is far beyond the clinical utility.

3.1.4 Limited Spatial Coverage and Resolution

As discussed in section 1.5, MRI faces an inherent limitation that imposes a trade-off between scan time, SNR and spatial resolution. This trade-off, particularly in in vivo cDTI, has resulted in relatively low in-plane spatial resolution, typically around 2 mm, and limited spatial coverage. In clinical practice, to maintain imaging times within a manageable range of 10 minutes, spatial coverage is often reduced to a few slices (typically 3 to 5), each with a slice thickness of 6-8 mm [100]. Consequently, comprehensive characterisation of microstructure and lesions such as infarct volume may be compromised, potentially resulting in an inaccurate estimation of the true lesion volume.

3.1.5 Gradient Field Imperfections

In practical SE-based diffusion imaging, technical challenges arise from imperfections in diffusion-weighting gradient fields, including the effects of eddy currents and concomitant fields. Under quasi-stationary conditions, the behavior of magnetic flux can be described as follows:

$$\begin{aligned} \vec{\nabla} \cdot \vec{B} &= 0 \\ \vec{\nabla} \times \vec{B} &= 0 \end{aligned} \tag{3.2}$$

In a dMRI sequence, the strong diffusion gradient pair that must be switched within a considerable time period, induces time-varying eddy currents within the magnet bore. These eddy currents consist of (i) a spatially invariant term, (ii) linear field gradient terms along the x, y, and z axes, often with similar characteristics to the inducing gradient pulses, and (iii) higher-order eddy current terms, which are typically negligible and can be disregarded [101]. The effective magnetic flux due to these eddy currents can be represented by [75]:

$$\vec{B}(t) = \vec{\varepsilon}_0(t) + \vec{G}_{Eddy}(t) \cdot \vec{r} + \vec{B}_c(t) \tag{3.3}$$

where $\vec{G}_{Eddy} = (\varepsilon_x, \varepsilon_y, \varepsilon_z)$, with its components representing the eddy currents induced along each spatial axis, whose magnitude is determined by the amplitude, duration, and separation of the diffusion pulses. $\vec{r} = (x, y, z)$ is a position vector referenced to the

centre of the magnet, and B_c, known as the concomitant field component, is expressed up to the second spatial order as:

$$B_c = \frac{1}{2B_0}\left[\frac{\varepsilon_z^2}{4}(x^2+y^2) + (\varepsilon_x^2+\varepsilon_y^2)z^2 - \varepsilon_x\varepsilon_y xy - \varepsilon_y\varepsilon_z yz\right] \tag{3.4}$$

It is worth noting that a specific eddy current can originate from its corresponding gradient term (e.g., G_x leading to ε_x), or from a cross-term (e.g., G_x leading to ε_y). The symmetric arrangement of the diffusion-encoding gradients around the refocusing pulse eliminates concomitant gradients, and these terms are often ignored; however, their quadratic dependence on the gradients may become particularly problematic in ex vivo imaging where very strong gradients are employed.

Residual eddy current effects can introduce errors in estimating tensor orientation and DTI parameters if left uncorrected. Specifically, when Cartesian readouts like EPI are used, they can cause distortion or susceptibility artefacts. These effects can be mitigated prospectively using bipolar diffusion gradients[1] with modified sequence timings [102] or measured and corrected for during image reconstruction [103].

3.2 Solutions

We discussed the methodological challenges associated with cDTI. In the following sections, we will review the techniques proposed in the literature to address these limitations, particularly focusing on reducing the scan time and improving spatial resolution and coverage. These topics are relevant to our proposed methodology in the subsequent chapters.

3.2.1 Diffusion Encoding Sequence Development

Current in vivo cDTI techniques that aim to mitigate the inherent sensitivity to motion in diffusion imaging can be categorised into two methods, both focusing on significantly reducing the diffusion encoding time:

(A) The first method utilises a stimulated-echo acquisition mode (STEAM) sequence, splitting the diffusion encoding over two consecutive heartbeats [35, 104].

[1] diffusion-weighting gradients with opposite polarities

(B) The second method employs a spin echo (SE) sequence similar to diffusion MRI sequences used in neuroimaging but with specialised research hardware capable of ultra-high gradient strength [105] or minimal gradient moment nulling [106].

Figure 3.1 shows a schematic of these two commonly used DW sequences in the heart.

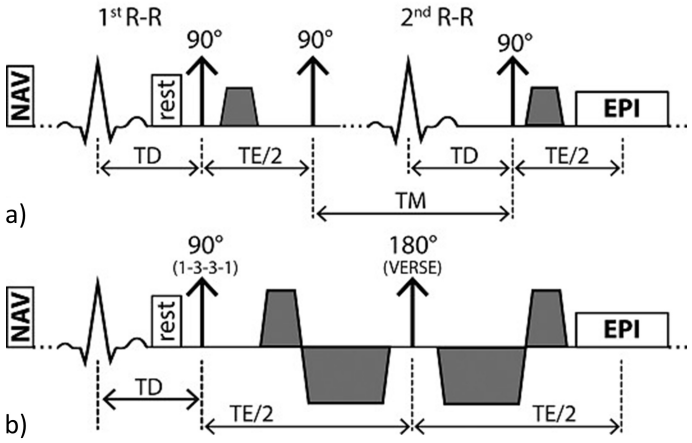


Figure 3.1: Schematic diagrams of common DW sequences used in the heart: (a) Diffusion-weighted stimulated echo (DW STEAM) sequence spanning over two consecutive heartbeats, (b) Spin echo acquisition with second-order motion compensated diffusion encoding gradients (M2-SE). T_D: trigger delay from R-wave to RF pulse; T_E: echo time; T_M: mixing time; NAV: respiratory navigator. (Adapted from [107])

STEAM

Initial in vivo cardiac diffusion MRI studies utilised a stimulated echo approach to spread the diffusion encoding period over two neighbouring R-R intervals, assuming periodicity for cardiac motion [108]. The majority of the diffusion encoding occurs during the so-called "mixing time" of a single R-R cycle, which limits sensitivity to bulk motion. The mixing time (T_M) is relatively long ($T_M = RR - T_E/2$), and only a small and very brief diffusion-encoding gradient is needed to achieve an adequate b-value with a comparatively short T_E. The DW signal is typically acquired during mid-systole using single-shot EPI, either under breath-hold conditions [104] or using navigator-based free-breathing techniques [35]. Because the stimulated echo forms over two consecutive heartbeats, the b-value for DW STEAM sequence depends on the heart

rate. The calculation of the b-value is given by:

$$b = \gamma^2 G^2 \left[\delta^2 \frac{60}{HR} - \frac{\delta^3}{3} + \frac{\zeta^3}{30} - \frac{\delta\zeta^2}{6} \right] \tag{3.5}$$

with δ representing the diffusion gradient pulse duration (its plateau), ζ denoting the duration of the gradient pulse slope, and HR representing the heart rate in beats per minute (bpm).

One drawback of this strategy is the assumption that bulk motion is precisely periodic during two consecutive heartbeats, which is generally true in healthy individuals for specific cardiac phases but may not hold in patients with arrhythmia or variable heart rates. If there are variations in heart rate, the apparent diffusion coefficient can be expressed as follows:

$$ADC_{actual} = \frac{b_{assumed}}{b_{actual}} ADC_{assumed} \tag{3.6}$$

For instance, assuming a heart rate of 60 bpm and a b-value of 500 s/mm^2, an increase of 10 bpm in the heart rate results in an underestimation of the diffusion coefficient by 14%, whereas a decrease in heart rate by 10 bpm leads to an overestimation of ADC by 20%. When the current heart rate is known for each acquisition, it is possible to correct the diffusion tensor by expanding the b-matrix into multiple b-values for each acquisition [107].

Another drawback of this sequence arises from the fact that the stimulated echo produces only half of the signal compared to a spin echo, resulting in an undesirable penalty on the already attenuated DW signal. Also, STEAM leads to double the scan time of an SE sequence, as it uses two heartbeats to produce each DW signal. Therefore, it has a lower SNR efficiency [107]. The measured diffusion is additionally influenced by the strain history of the myocardium between the encoding gradients [37], although the extent of this impact remains a subject of debate [109].

Motion-compensated Spin Echo

As discussed, the standard Stejskal-Tanner diffusion encoding experiment during the contraction of the heart leads to the deterioration of the DW signal. This issue can be understood by examining the phase accumulation of the magnetisation while a magnetic field gradient is applied:

$$\phi_m = \gamma \int_0^T \vec{G}(t) \cdot \vec{x}(t) \; dt \tag{3.7}$$

where T represents the duration of the diffusion encoding gradient, $\vec{G}(t)$ is the gradient waveform, and $\vec{x}(t)$ is the spatial trajectory of material points. Using the Taylor approximation of the trajectory:

$$x(t) \mid_0 = \sum_n \left(\frac{\partial}{\partial t}\right)^n x(t) \mid_{t=0} \frac{1}{n!} t^n \tag{3.8}$$

we get:

$$\phi_m = \gamma \int_0^T |\vec{G}(t)| \; \vec{x}_{t=0} \; dt + \gamma \int_0^T |\vec{G}(t)| \; \dot{\vec{x}}_{t=0} \; t \; dt + \frac{1}{2}\gamma \int_0^T |\vec{G}(t)| \; \ddot{\vec{x}}_{t=0} \; t^2 \; dt + \cdots \tag{3.9}$$

thus,

$$\phi_m = M_0 + M_1 v + M_2 a + \cdots \tag{3.10}$$

From Equation 3.9, it is evident that the accumulated phase can be computed based on the initial spatial coordinates of the spins, their temporal derivatives, and the gradient moments $\vec{M}_n$:

$$\begin{aligned} M_0 &= \int_0^T |\vec{G}(t)| dt \\ M_1 &= \int_0^T t |\vec{G}(t)| dt \\ M_2 &= \int_0^T t^2 |\vec{G}(t)| dt \\ &\vdots \\ M_n &= \int_0^T t^n |\vec{G}(t)| dt \end{aligned} \tag{3.11}$$

The gradient configuration in a PGSE sequence leads to significant deviations of the gradient moments from zero, rendering the sequence sensitive to bulk (large-scale) motion. In the heart, it is essential to account for non-rigid deformation and displacement of the myocardium, and it has been shown that compensating for tissue velocity and acceleration is crucial for preserving diffusion information integrity.

A straightforward strategy to make the PGSE sequence insensitive to motion involves transforming the monopolar diffusion-encoding gradients in the standard sequence into bipolar gradients with zero net area [110]. This gradient setup was demonstrated to effectively null up to the first gradient moment (velocity compensation, M_1) [111]. Therefore, it is best employed when the tissue is either stationary or undergoing rigid displacement at a relatively constant velocity. However, this condition is

typically met for only very brief time intervals during the peak of the cardiac cycle (mid-systole), which lasts for approximately 10-20 ms. Achieving sufficient b-values under these conditions necessitates specialised high-performance gradient systems, which are not commonly available in clinical settings (2 to 10 times larger than what clinical systems typically offer). To compensate for acceleration (M_2) in the PGSE sequence, a more complicated gradient pulse series needs to be incorporated, as shown in Figure 3.1(b). This gradient configuration extends the echo time required for the same b-value. However, recent findings indicate that the advantages of employing such configurations outweigh the penalty of the longer T_E [112, 113].

In motion-compensated SE sequences, the gradient capability significantly affects the achievable echo time, and hence the SNR. This method has been demonstrated to be feasible on clinical systems with high gradient capabilities ($G_{max} \approx 80$ mT/m). Importantly, this method is applicable across a wide range of heart rates found in cardiac patients, making it a promising approach for examining cardiac microstructure in clinical settings.

3.2.2 Fast Data Acquisition

To mitigate artefacts resulting from cardiac and respiratory motion, it is also crucial to keep the time required for acquisition of each DW image as short as possible. This can be achieved through employing fast imaging techniques. Currently, the most widely used method in cDTI is a single-shot echo planar imaging (EPI), where the complete k-space of a 2D DW image (one slice) is acquired with one RF excitation.

Single-shot EPI, however, comes with some challenges that may counteract its advantages, including degradation of image quality and limited spatial resolution. Due to its longer acquisition window compared to the conventional Cartesian scheme, which acquires only one k-space line during each acquisition window, single-shot EPI is highly susceptible to artefacts arising from off-resonance effects. Off-resonance spins are characterised as spins that do not resonate at the expected Larmor frequency, mainly due to imperfections that cause deviations of the actually experienced field strength from the assumed field strength. For these particular spins, Equation 2.14 is modified by an additional exponential term [65], which leads to a blur in the reconstructed image:

$$S(t) = \int m(\vec{r}) e^{-i2\pi\vec{k}(t)\cdot\vec{r}} e^{-i\Delta\omega(\vec{r})t} \, d\vec{r} \tag{3.12}$$

where, $\Delta\omega(\vec{r})$ represents the off-resonance frequency at location $\vec{r}$, caused by different factors such as magnetic field inhomogeneity, eddy currents, tissue susceptibility, and chemical shift. Importantly, these factors may vary across the image area in distinct manners, leading to spatially local and time-dependent phase evolution in the affected regions of the object. The appearance and strength of the resulting artefacts depend on the trajectory shape and acquisition duration.

In Cartesian schemes, this non-zero residual phase will interfere with the stepwise phase encoding, leading to issues such as geometrical distortions and signal dropout in EPI images. This will in turn result in direction-dependent distortions in DW image series and deterioration of DTI-derived parameter maps if left uncorrected. In contrast, non-Cartesian trajectories (e.g., spiral or radial) do not have a distinct phase encoding direction and sample k-space continuously and in multiple directions. This characteristic helps spread off-resonance effects across the image, thereby minimising the localised distortions seen in Cartesian phase encoding. Spiral acquisition, in particular, is a promising fast imaging method comparable to EPI. Generally, spiral trajectories provide several advantages including shorter echo times (T_E), improved SNR, and inherent robustness against motion and flow artefacts. However, these methods have not yet been widely adopted in cDTI applications. Notably, off-resonance effects in spiral schemes typically manifest as artefacts perceived as less visually disruptive, i.e., blurring rather than geometrical distortions. This more uniform distribution of error energy across the image plane retains a higher visibility of diagnostic information.

Off-resonance Correction

Due to the spatial dependency of the off-resonance frequency, it becomes challenging to determine a singular inverse solution for Equation 3.12. Extensive research has been conducted to address the correction of the resulting images through post-processing. Noll et al. [114, 115] introduced a technique known as conjugate phase reconstruction (CPR), which, while effective, was computationally demanding. CPR involves creating a field map that maps off-resonance frequency $\Delta\omega$ to spatial locations $\vec{r}$, which is then used to demodulate the image at the correct frequency for each spatial pixel. Time-segmented and frequency-segmented CPR are approximations to the original method that reduce the number of demodulation processes but are still beyond practical utility. Figure 3.2 outlines the steps involved in the frequency-segmented CPR method. Altern-

ative methods have thus been explored to improve frequency-segmented CPR, in terms of computation time, image quality, and theoretical accuracy. One such method is the multi-frequency interpolation technique [116], which expresses arbitrary demodulation frequencies as a linear combination of a predefined set of demodulation frequencies as:

$$e^{i\Delta\omega t_k} = \sum_{i=0}^{L} C_i(\Delta\omega) e^{i\Delta\omega_i t_k} \tag{3.13}$$

This generalisation allows pixels to be demodulated at a frequency closer to the exact value by using a linear combination of already demodulated and reconstructed pixels. Also, it provides three to four times faster computations while maintaining equal image quality. The coefficients C_i can be calculated through a least squares approximation and added to the demodulated images after the inverse Fourier transformation step in Figure 3.2. The final image is constructed by combining each pixel from different demodulated images.

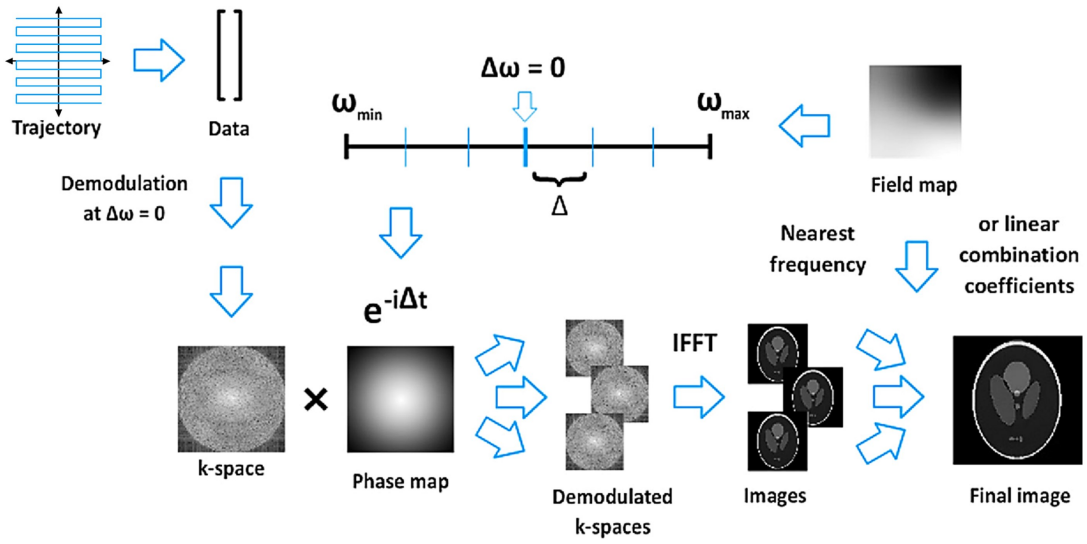


Figure 3.2: Flowchart describing the off-resonance artefacts correction in EPI. (Reproduced from [117])

Multi-shot Acquisition

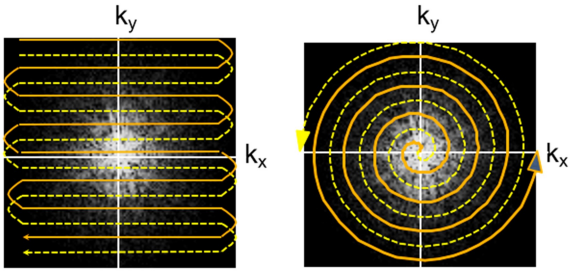
Multi-shot (also known as interleaved) imaging technique (Figure 3.3), which has a shorter acquisition window and fills the k-space in multiple segments, can help reduce

the off-resonance artefacts and improve the spatial resolution of single-shot acquisitions, but may result in an increased scan time. While motion or phase offsets within a single readout interval of several milliseconds is often negligible, a simple delay in the gradient response or even minor motion between different k-space segments can cause signal phase inconsistencies, resulting in inconsistent data being merged during the reconstruction procedure and thus leading to sever image distortions. In Cartesian sampling schemes, including EPI, this leads to signal modulation along neighbouring phase encoding rows, resulting in the prominent occurrence of shifted copies of the object in the corresponding direction, known as ghosting artefacts. However, the spiral sampling geometry effectively circumvents the impact of phase inconsistencies seen in multi-shot schemes. Since each spiral shot begins sampling from the k-space centre, any inconsistencies primarily result in misalignment of this information, rather than creating the same disruptive ghosting artefact. Considering the properties of the Fourier transform, this shift in k-space encoding translates to a phase shift in the image space. The variable phase modulation of the spatial information encoded in each shot causes signal cancellation and object blurring when the shots are combined. Due to the oversampling of the k-space centre, spiral shots contain redundant information, resulting in an averaging effect that compensates for errors in individual shots. This feature is another reason for the enhanced motion robustness of spiral sampling, which, combined with the absence of ghosting artefacts, holds significant promise for imaging moving organs or uncooperative patients.

It is essential to implement a compensation mechanism for mitigating the factors contributing to these inconsistencies. Navigator-based phase correction can help eliminate these shot-to-shot phase inconsistencies [118]. Some techniques have incorporated the gradient impulse response function (GIRF) into image reconstruction to correct for errors arising from k-space trajectory distortion [119].

3.2.3 Incomplete k-Space Sampling Techniques

Imaging speed is a crucial factor in MRI, and in particular, in diffusion MRI applications. Speeding up image acquisition while maintaining diagnostic quality is crucial in many respects. It directly reduces the costs of MR imaging, alleviates patient discomfort, reduces the risk of motion artefacts, and potentially enhances image quality. Specifically in dMRI applications, the time saved can be used to increase the spatial



(a) Multi-shot EPI (b) Multi-shot Spiral

Figure 3.3: Schematic of multi-shot EPI and spiral imaging.

resolution or to perform additional scans for improving the SNR. However, MRI scanners often operate at the limits of potential imaging speed due to constraints such as hardware performance (e.g., gradient amplitudes and slew rates), physiological factors, and the available SNR per unit time. Consequently, achieving faster imaging is typically accomplished by reducing the amount of acquired data, which may involve skipping some of the data that would normally be acquired. However, when data is acquired at sub-Nyquist rates, it can result in images corrupted with undersampling artefacts, as illustrated in Figure 3.4. The field of accelerated MRI addresses how to mitigate these aliasing artefacts and produce artefact-free images, either through dedicated hardware or advanced reconstruction techniques. In this section, we briefly discuss the major image acceleration techniques that have been successfully applied in the cDTI.

Hardware-based Approaches

Parallel imaging (PI) has emerged as the most widely adopted technique in clinical practice [121] to reduce the scan time. Using an array of localised receiver coils that obtain complementary information, PI acquires a partial amount of k-space data through reduced phase encoding steps. The undersampled data result in aliased images, as illustrated in Figure 3.4(b). To reconstruct artefact-free images, various PI algorithms have been developed, including sensitivity encoding (SENSE) [122], generalized autocalibrating partially parallel acquisitions (GRAPPA) [123], and iterative self-consistent parallel imaging reconstruction (SPIRiT) [124]. Coil sensitivity information is required for these techniques, which can be obtained either through a pre-scan or by acquiring additional k-space lines (auto-calibration).

However, PI techniques are typically constrained by the Nyquist sampling rate,

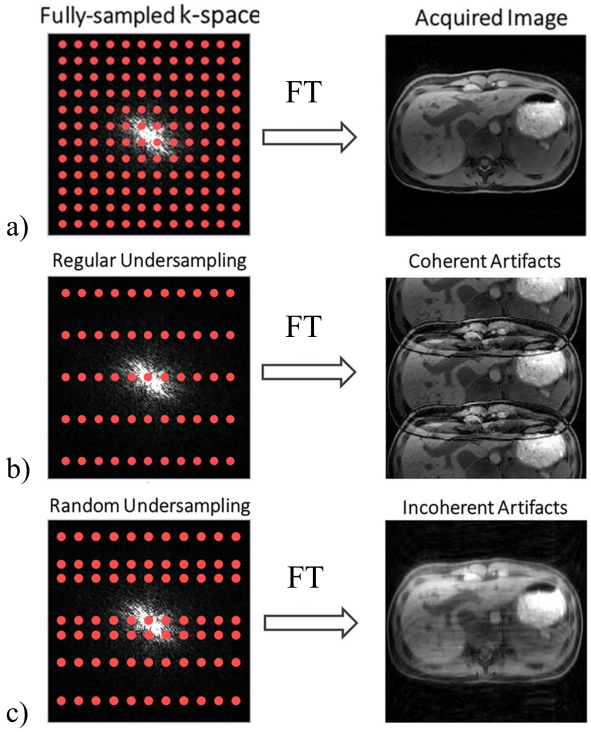


Figure 3.4: k-Space undersampling: (a) A fully sampled k-space results in an artefact-free image through Fourier transform (FT) reconstruction, (b) Regular undersampling produces coherent replicas of the signal structure, (c) Random undersampling generates incoherent artefacts resembling added noise. (Adapted from [120])

and the achieved acceleration is primarily limited to low to moderate factor values [125]. Higher acceleration can be obtained at the expense of reduced SNR, which is particularly problematic for SNR-starved diffusion MRI.

Software-based Approaches

In recent years, compressed sensing (CS) theory [55–57] has gained growing interest within the MR community. CS aims to reduce the number of measurements and acquisition time by exploiting the compressibility of MR images. With CS, data can be significantly undersampled compared to fully-sampled (sampled at the Nyquist rate) acquisition, while still enabling optimal image recovery. Various CS-based image reconstructions have been successfully applied in cDTI [60–63, 126–128].

In the following section, we will delve into the theory of compressed sensing and ex-

plain its fundamental concepts. Understanding these concepts will provide a foundation for comprehending the primary work undertaken in this thesis.

3.3 Compressed Sensing Theory

Two fundamental principles in the development of the CS theory are sparsity (or compressibility) and incoherence:

Sparsity: Sparsity requires the signal to have a sparse representation in some domain. This concept has been widely utilised in signal processing and approximation theory for tasks like compression and denoising. Many natural signals can be represented sparsely or well-approximated using only a few elements in a suitable basis or dictionary. In mathematical terms, a discrete signal $x \in \mathbb{R}^n$ is considered k-sparse if it has at most k non-zero entries, i.e., $\|x\|_0 \leq k$, $k \ll n$.

Signals in general are not inherently sparse, but they can have a sparse representation in a specific basis, commonly referred to as the transform domain. This means that these signals can be represented as linear combinations of a few basis functions from a "sparsifying" transform or dictionary. Generally, two main types of sparsifying transforms can be considered: 1) predefined transforms, which are constructed based on a mathematical model of the data, such as Fourier or wavelet transforms, and 2) data-adaptive transforms, which are typically constructed through learning a dictionary that best represents the data using a training set. In either case, we can still describe the signal x as k-sparse, with the understanding that it can be expressed as:

$$x = \Psi c = \sum_{i=1}^{n} c_i \psi_i \tag{3.14}$$

where c_i are the transform coefficients, Ψ is the basis or dictionary with vectors $\{\psi_i\}$ as columns called atoms, and $\|c\|_0 \leq k$.

Incoherence: In addition to sparsity, incoherence is another fundamental premise underlying CS. For the sake of simplicity in notation, here we consider a pair of orthobases of $\mathbb{R}^n(\Phi, \Psi)$. However, it is worth noting that this specific choice of bases is not a strict requirement. The basis Φ is used for signal sensing, while the basis Ψ is used for sparsely representing the signal x. The coherence, also referred to as *mutual* coherence,

between the sensing basis Φ and the sparsifying basis Ψ is defined as follows [129]:

$$\mu(\Phi, \Psi) = \max_{1 \leq k,j \leq n} | \langle \phi_k, \psi_j \rangle | \tag{3.15}$$

with ϕ_k and ψ_j denoting the kth and jth column of Φ and Ψ, respectively.

The coherence $\mu(\Phi, \Psi)$ measures the largest correlation between any two elements of the sensing basis Φ and the sparsifying basis Ψ, and for orthonormal bases (such as spikes and sinusoids), it falls within the range $\lceil 1/\sqrt{n}, 1 \rceil$ [129]. In CS, we are primarily interested in low-coherence pairs, where the two bases exhibit greater independence and are referred to as *incoherent*. Random matrices with independently and identically distributed (i.i.d.) entries, such as Gaussian or ± 1 binary entries, typically exhibit very low coherence with any fixed Ψ.

3.3.1 Sensing

Let $x \in \mathbb{R}^n$ represent the signal and ϕ_k denote the sensing waveforms. The sensing process can be expressed as:

$$y_k = \langle x, \phi_k \rangle, \quad k = 1, \cdots, m \tag{3.16}$$

If the sensing waveforms are Dirac delta functions (spikes), then y is a vector of sampled values of x in the time or space domain. If the sensing waveforms are sinusoids, then y is a vector of Fourier coefficients; this is the sensing process used in MRI, which can be represented in matrix form as:

$$y = \Phi x \tag{3.17}$$

where $\Phi \in \mathbb{R}^{n \times m}$ is the sensing matrix and $y \in \mathbb{R}^m$ is the obtained signal.

Under the conventional sensing paradigm, m must be at least equal to n (when the sensing waveforms are Dirac delta functions-spikes). However, in CS, the matrix Φ maps $\mathbb{R}^n$ into $\mathbb{R}^m$, where $m \ll n$. This results in a highly under-determined problem, with too few equations and too many unknowns. Thus, two main theoretical questions arise. First, how to choose the type of sensing matrix to ensure that it preserves the information in the signal. Second, how many measurements should be collected to ensure that these measurements will be sufficient to recover the original signal. The robustness of the CS relies on a notion known as the *restricted isometry property (RIP)* [55]. If a matrix Φ satisfies the RIP, then this condition alone is adequate for various algorithms to successfully recover a sparse signal from noisy measurements:

$$(1 - \delta_k)\|x\|_2^2 \leq \|\Phi x\|_2^2 \leq (1 + \delta_k)\|x\|_2^2 \quad \forall \delta_k \in (0, 1) \tag{3.18}$$

3.3.2 Recovery

The problem of recovering a sparse signal x from an undersampled observation vector y is typically addressed through an optimisation of the form:

$$\min_{x\in\mathbb{R}^n} \|x\|_0 \quad \text{s.t.} \quad y = \Phi x \tag{3.19}$$

where $\|x\|_0 = \#(i|_{x_i\neq 0})$ denotes the ℓ_0-norm, which counts the number of non-zero elements in a vector. This problem requires combinatorial optimisation, making it computationally challenging. It has been demonstrated that when x is sufficiently sparse [56], the ℓ_0-norm can be effectively replaced by the ℓ_1-norm:

$$\min_{x\in\mathbb{R}^n} \|x\|_1 \quad \text{s.t.} \quad y = \Phi x \tag{3.20}$$

with $\|x\|_1 = \sum_{i=1}^{n} |x|$ representing the sum of absolute values of the elements in a vector. This Equation can be reformulated as an unconstrained problem and with a more general regularisation penalty as:

$$\hat{x} = \arg\min_{x\in\mathbb{R}^n} \|y - \Phi x\|_2^2 + \mathcal{R}(x) \tag{3.21}$$

Here, $\|y-\Phi x\|_2^2$ is called the data consistency term and $\mathcal{R}(x)$ is a regularisation term that makes the under-determined problem well-posed. Various algorithms can be utilised to solve these Equations and recover the original signal from its compressed measurements. Examples include matching pursuit [130], basis pursuit [56, 131], iterative thresholding [132], proximal gradient methods [133], and many others.

Proximal gradient methods are particularly well-suited for solving inverse problems with large-scale datasets. A notable algorithm in this category is the fast iterative shrinkage-thresholding algorithm (FISTA), developed by Beck and Teboulle [134], which builds upon early work by Nesterov [135] for general convex problems. FISTA incorporates an effective acceleration scheme, making it particularly efficient in recovery of high dimensional data. Additionally, the composite splitting algorithm (CSA) and its accelerated version (FCSA) [136, 137] were proposed to solve composite regularisation problems, offering further advancements in signal recovery techniques.

3.4 Compressed Sensing in MRI

By leveraging signal sparsity in a transform domain and undersampling k-space in a manner that generates non-coherent artefacts, CS enables accurate recovery of MR

images using appropriate non-linear reconstruction techniques, potentially leading to significant reductions in the scan time. Regular undersampling schemes, commonly employed in parallel imaging, yield artefacts that manifest as coherent replicas of the signal structure (Figure 3.4(b)). Theoretically, incoherent aliasing artefacts, resembling noise-like patterns in the sparse domain, are generated by random selection of k-space locations. However, in practice, achieving purely random sampling faces hardware constraints. As an alternative, pseudo-random undersampling is widely implemented in Cartesian acquisition by skipping some phase encoding lines in the k-space, as illustrated in Figure 3.4(c).

The choice of a sparsifying transform depends on the specific clinical application. There are two commonly used sparsifying transforms in cardiac MRI: 1) the total variation (TV) or spatial finite difference, and 2) discrete wavelet transform (DWT). For a 3D image (I), the first-order forward TV operators along the three spatial coordinates (i, j, k) are given by:

$$\begin{aligned}(\nabla_1 I)_{ijk} &= I_{i+1,j,k} - I_{i,j,k} \\ (\nabla_2 I)_{ijk} &= I_{i,j+1,k} - I_{i,j,k} \\ (\nabla_3 I)_{ijk} &= I_{i,j,k+1} - I_{i,j,k}\end{aligned} \tag{3.22}$$

The TV model tends to favor images with piece-wise constant structures. While it effectively preserves edge details, it can result in image smoothing and blurring. To mitigate this inherent limitation of the TV model, several extensions have been introduced, including nonlocal total variation (NLTV) [138, 139] and total generalised variation (TGV) [140–142]. Traditional wavelet transforms commonly used in CS MRI [57, 137, 143] are effective in representing point-like singularities, but they often exhibit limited performance when dealing with singularities in higher dimensions, such as edges, contours, or regular textures [144]. To address this limitation, a multi-scale geometric analysis approach such as contourlets [145] or shearlets [141] can be used to achieve a more sparse representation of piece-wise smooth images containing rich geometric information like edges and curves. Combining wavelet transform with finite differences often provides better performance, especially in accelerated static imaging.

3.4.1 Compressed Sending in cDTI

As discussed, CS MRI aims to significantly reduce the number of measurements required for image reconstruction, thereby accelerating overall imaging speed. Numerous studies

have focused on applying CS to accelerate cDTI. Examples include exploiting additional TV constraint in the diffusion directional dimension besides the conventional spatial TV [61], and utilising the location and orientation information of the edges from a reference image through directional TV [63]. While classical CS theory suggests a minimum bound for the number of measurements required for accurate reconstruction [55, 56], it was later shown that the number of measurements can be reduced by utilising additional priors besides the fundamental transform-sparsity constraint in the CS formulation or the underlying structures in the signal [146, 147].

DTI involves acquisition of a series of diffusion-weighted images $\{x(\mathbf{r}, q)\}_{q=1}^{N}$ at N diffusion encoding directions and/or diffusion weighting factors. Because these images are obtained from the same anatomical cross-section with the same tissue boundaries, they are highly correlated. Yet, the intensity values in corresponding pixels may differ. In other words, the angular (diffusion-directional) profiles of the voxels, denoted as the n-dimensional vectors:

$$\mathbf{S}_i = [x(\mathbf{r}_i, q_1), x(\mathbf{r}_i, q_2), \cdots, x(\mathbf{r}_i, q_N)]^T; \quad i = 1, \cdots, M \tag{3.23}$$

are correlated or linearly dependent. Here, M is the number of voxels and q denotes the diffusion parametric dimension, referred to as the q-space. Motivated by multichannel systems, we can concatenate the diffusion signal in the combined spatial ($\mathbf{r}$) and angular (q) domain ($k-q$ space) in a matrix form to exploit the correlations [148]:

$$\mathbf{X} = \begin{bmatrix} x(\mathbf{r}_1, q_1) & \cdots & x(\mathbf{r}_1, q_N) \\ \vdots & & \vdots \\ x(\mathbf{r}_M, q_1) & \cdots & x(\mathbf{r}_M, q_N) \end{bmatrix} \tag{3.24}$$

The rows of $\mathbf{X}$ correspond to the voxels and the columns represent the diffusion angular samples (q-points). Using CS, the unknown image described in Equation 3.21 transforms into a spatial-angular ($k-q$ space) image sequence matrix $\mathbf{x} \in \mathbb{C}^{M \times N}$, and the sensing operator extends to $\Phi : \mathbb{C}^{M \times N} \rightarrow \mathbb{C}^{N_k}$, where N_k represents the total number of k-space samples across the entire diffusion-weighted series.

Since the sparsity of MR images in transformation bases is a key factor influencing CS performance, sparser representations can either yield the same performance with fewer samples or achieve better reconstruction quality with a given number of samples. The extensive correlations among DWIs can be utilised to allow for sparser representations, thereby making higher accelerations feasible.

One approach is through utilising the group sparsity property, inspired by distributed compressed sensing. This property implies that each individual image $x(\mathbf{r}, q_i)$ in the DWI series exhibits sparsity in a specific transform domain, determined by the matrix $\mathbf{\Psi}$, and shares similar sparsity patterns along the diffusion-encoding dimension with other images, such that the zero and non-zero coefficients occur at the same locations across DWIs. The group sparsity prior of DW images in the wavelet transform domain has been exploited in cDTI to further reduce the number of required measurements [60]. The regularisation penalty promoting group sparsity is characterised by an $\ell_{1,2}$-norm with an explicit expression of:

$$\|\mathbf{X}\|_{1,2} = \sum_{i=1}^{n} \left\|\mathbf{X}^{(i)}\right\|_2 \tag{3.25}$$

where $\mathbf{X}^{(i)}$ is the ith group in $\mathbf{X}$. In DTI, the sparse coefficients at one voxel along the diffusion encoding dimension can be considered as one group

Another approach involves leveraging the underlying data structure. Since the rows of $\mathbf{X}$ are linearly dependant, its rank is defined as $L < \min(M, N)$, making it a low-rank (LR) matrix. This low-rank property allows for the approximation of $\mathbf{X}$ through matrix factorisation, where it can be explicitly represented as the product of two (or more) lower-dimensional matrices:

$$\mathbf{X} \approx \underbrace{\mathbf{U}}_{M\times L}\underbrace{\mathbf{V}^T}_{L\times N} \tag{3.26}$$

In this representation, the matrix $\mathbf{V} \in \{\mathbb{C}^{L\times N} : V_{jk} = v_j(k)\}$ contains "diffusion parametric basis functions" that span the low-dimensional diffusion subspace. $\mathbf{U} \in \{\mathbb{C}^{M\times L} : U_{jk} = u_k(r_j)\}$ represents the corresponding spatial coefficients that span the spatial subspace. The low-rank structure provides significant potential for undersampling, as the matrix $\mathbf{X}$ has a total of MN complex elements (equivalent to $2MN$ real values), while it possesses only $2(M+N-L)L$ degrees of freedom. Image reconstruction can be carried out by finding a low-rank matrix $\mathbf{X}$ that fits the data consistently [148], or by estimating $\mathbf{U}$ and $\mathbf{V}$ separately in distinct steps [149]. Combination of LR modelling and transform-sparsity was used in [126], and complementary techniques that embedded phase correction components into the reconstruction framework [127] and accounted for the local LR regularisation [62] were then developed to improve the image quality and reduce the reconstruction errors.

Data-adaptive sparsifying transforms or dictionaries have shown to yield superior image sparsity compared to the predefined transforms [150]. Dictionary learning methods involve learning a dictionary as a sparse basis from elemental patches of specific image instances or training datasets. This approach was used to improve the representation of the diffusion signals without making assumptions about the underlying diffusion process and could effectively reduce the acquisition time of cDTI [128]. However, some of these methods rely on large dictionaries generated by simulations of all conceivable parameter combinations. One drawback of this approach is its vulnerability to subject/sequence specific inconsistencies such as motion and eddy current induced artefacts, primarily due to the inability of the dictionary basis functions to account for signal variations between the training dataset and the actual data. Another issue with directly employing this method for 3D cDTI is the rapid expansion of the dictionary size with an increasing number of parameters, thus requiring a large number of training datasets leading to increased complexity in the nonlinear recovery algorithm.

Previous works addressing the challenge of prolonged data acquisition in cDTI have been primarily focused on improving the sparse approximation and accurate recovery of DW images from undersampled measurements. They have utilised commonly used Cartesian undersampling schemes [151–153] to induce incoherent aliasing artefacts, without any application-specific optimisation on the sampling parameters.

3.5 Summary

In this chapter, we outlined the technical challenges that complicate the utilisation of DTI in the heart, explored the solutions and approaches that have been devised to tackle these challenges, as well as their advantages and disadvantages. In summary, DTI is known for its lengthy acquisition times required to obtain high-quality data (both in terms of high spatial resolution and high SNR), significantly limiting its practical and clinical utility for imaging the human heart [40, 41, 104, 154]. Currently, a two-dimensional (2D) multislice EPI acquisition has been the most commonly used method in both clinical and preclinical settings. With this approach, imaging of only a few short-axis slices (three to five) has been realised with an in-plane resolution of $\approx$ 2 mm, for limited number of diffusion directions in in vivo human cDTI studies [35, 127, 155]. Efforts have been recently made towards achieving whole left ventricular coverage by enabling acquisition of up to 12 slices, but still compromising the spatial

resolution in the long-axis direction (slice thickness = 8 mm) [156, 157]. While 3D acquisition techniques have enabled whole-heart coverage in ex vivo setting with isotropic resolution at the orders of μm in rodent hearts [158, 159] and mm in large animal models (including the human heart) [40], they currently suffer from a prolonged scan time far beyond the clinical range. The extended scan time in 3D imaging can lead to increased susceptibility to cardiac motion-induced artefacts in in vivo applications. Therefore, a significant endeavor in this field is the development of techniques that can enable higher-resolution whole-coverage cDTI within clinically feasible scan times. These advancements hold promise for furthering our understanding of cardiovascular diseases and improving patient care. By employing acceleration techniques, it becomes possible to significantly reduce scan times while still achieving high spatial resolution data. This allows for the accurate assessment of cardiovascular status within shorter acquisition times. The research presented in this thesis focuses combining compressed sensing and optimal spiral sampling in the field of cDTI, drawing inspiration from and building upon the topics discussed in this chapter.

Chapter 4

Optimal Design of the Sampling Density Function for Spiral-CS cDTI

4.1 Introduction

In the previous chapter, we explained that achieving whole-heart coverage with high spatial resolution ($\leq$ 1 mm) and minimal artefacts is an ongoing challenge in cardiac diffusion MRI, particularly in in vivo applications. We delved into the cDTI literature that have applied the CS theory [160] to accelerate acquisition time while preserving image quality. However, these methods rely on common Cartesian undersampling techniques, which are inefficient in at least one direction and allow for limited acceleration factors that are insufficient for a 3D scan to achieve comparable scan times to 2D scans. Therefore, there is a need for a new strategy capable of providing high acceleration factors without compromising image quality or SNR. The inherent incoherence in the artefacts, manifested as diffuse aliasing patterns, produced by non-Cartesian k-space trajectories, such as spiral trajectories, in multiple dimensions, makes them a more appealing alternative. Furthermore, spiral trajectories offer potential for shorter scan times and are less susceptible to bulk motion and flow effects [161], which improves capturing diffusion information from the heart.

The aim of this chapter is to design an optimal acquisition strategy for clinically-feasible whole-heart cDTI that can enable acquiring at least 8-10 slices covering the entire myocardium within a single R-R interval (in less than a minute), through combination of a novel spiral sampling technique and an efficient CS reconstruction approach. To achieve this objective, specific goals regarding resolution, timing, and SNR must be established. Additionally, parameters in the spiral design, such as readout duration and sampling density function, and the reconstruction strategy need to be carefully considered and defined.

We have thus divided our proposed methodology into two parts. The first part, presented in this chapter, focuses on optimal parametrisation of the sampling density function in a spiral k-space trajectory given a particular experimental setup. Once the optimal parameters are identified, the subsequent chapter extends the sampling method to 3D imaging and enhances the CS reconstruction performance by incorporating the joint sparsity property among the diffusion-weighted images (DWIs). In chapter 6, we introduce a novel data-adaptive sampling and reconstruction technique to further improve acquisition efficiency and reconstruction quality.

4.2 Sampling Design for CS cDTI

In this section, we highlight a few considerations and open questions to provide context for our experimental study.

4.2.1 A Priori Knowledge

Figure 4.1 illustrates the normalised b_0 and two arbitrary diffusion-weighted images of a human heart with their corresponding k-space representations. From this, we can draw two observations: (i) most of the energy of the signal is concentrated near the centre of the k-space, and (ii) the distribution of signal energy is altered due to the diffusion encoding. Specifically, we observed that the amplitude of the captured diffusion-weighted signal (k-space data) drops to 1% of its maximum within $\approx$ 10% of the maximum k radius in all DW directions. Thus, when designing an optimal sampling strategy for a cDTI scan, the energy distribution of the signal should be taken into account, and it may not be appropriate to choose the same sampling function for both non-weighted and weighted scans.

It is thus reasonable to measure low spatial frequencies more precisely while allocating fewer sample points to areas with less information. However, in cDTI, one important point to consider is whether any high-frequency information related to the diffusion process might be compromised if higher frequencies are not adequately acquired. This information could potentially be encoded in the phase of the signal. In this study, we primarily focus on the diffusion-weighted data and aim to design sampling functions that sufficiently capture diffusion-related information for near-perfect signal recovery. We will also explore if there should be a diffusion direction-dependent factor consideration when designing the sampling density.

Figure 4.2 illustrates the effects of removing different k-space regions in the reconstructed image. Low spatial frequencies contribute to the signal contrast while higher spatial frequencies contain the information related to the fine image details and edges.

4.2.2 Variable Density Sampling

In MRI, the acquisition time is directly proportional to the number of points sampled in k-space, given a specific gradient system constraint and acquisition bandwidth. Tra-

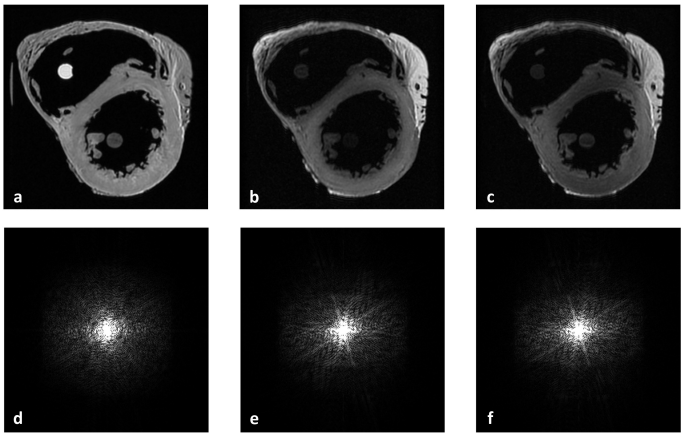


Figure 4.1: The b_0 and diffusion-weighted images of an ex vivo human heart (a-c) and their corresponding k-space image (d-f).

ditionally, the number of sampled points is determined by the Nyquist sampling rate, considering the field of view (FOV) and spatial resolution requirements. However, this approach imposes a minimum limit on scan time, which can be restrictive, particularly in applications requiring 3D imaging. To overcome the Nyquist-imposed scan time limitation, variable density k-space sampling has been successfully implemented [162]. Variable density sampling allows the density of sampling, or the sampling interval (Δk), to vary as a function of the k-space location. This flexibility enables the allocation of scan time based on signal characteristics, scan time constraints, and the acceptable image quality for a given application. Leveraging the knowledge that most of the energy of an image is concentrated near the k-space origin, more scan time can be strategically allocated to sampling low spatial frequencies while higher spatial frequencies are undersampled. By modifying the gradient waveforms within the constraints of maintaining a smooth trajectory, sampling points in k-space can be selectively chosen in a highly flexible manner.

4.2.3 CS Limitations in MRI

Early theoretical findings in the original CS theory relied on incoherent properties for the sampling and sparsity bases. However, MRI signals are highly structured and can be modelled much more accurately than arbitrary k-sparse signals. An effective model to characterise this structure is the *sparsity by levels* in wavelet bases. As illustrated in

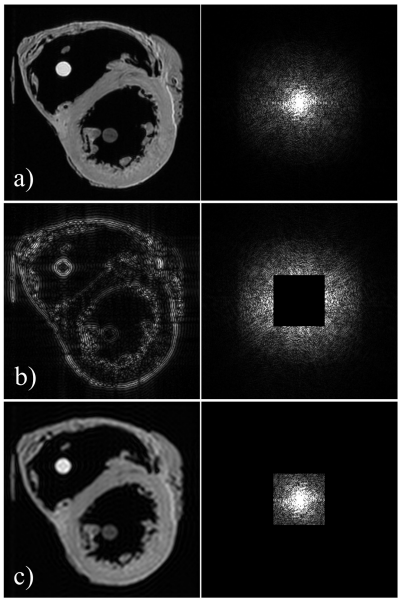
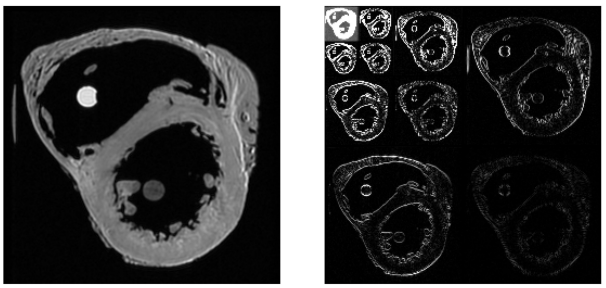


Figure 4.2: Illustration of the impact of nulling different regions of k-space on a reconstructed cardiac image: (a) image obtained from the full k-space data, (b) image obtained from the peripheral k-space data, exhibiting only edges and fine details, and (c) image obtained from the central k-space data, showing the coarse structure with missing details.

Figure 4.3, each wavelet sub-band of the image contains a specific number of non-zero coefficients, constrained by a known quantity at each scale. Moreover, the conventional method of acquiring data in MRI differs significantly from independent Gaussian measurements, as Fourier transform values of the image are probed along continuous trajectories. In this challenging context, the sampling theory is notably less comprehensive compared to the case of Gaussian measurements. Recently, advances have been proposed by various researchers [153, 163, 164]. Without delving into the theorems, let us outline the main conclusions and limitations of these studies. First, it has been suggested that the use of incoherent sampling and reconstruction bases may not be necessary if the sparsity pattern is adapted to the sampling scheme. This adaptation is particularly relevant for MRI, where most images are compressible in the wavelet domain. Second, an important finding indicates that sampling should be distributed along variable density patterns that consider the signal structure [165]. In MRI, this can be translated into low frequencies should be sampled more frequently than high

frequencies, not simply because there is more information in the low frequencies, but rather due to more subtle reasons.



(a) Human heart image **(b)** Wavelet representation

Figure 4.3: (a) A human heart MRI image and (b) its wavelet representation for a Haar wavelet basis and 4 levels of decomposition.

Therefore, to fully understand the CS reconstruction limits in a specific MRI application, it is essential to precisely describe the relationships between wavelet bases, Fourier bases, and sparse signals at different levels. The current major limitations can be summarised as follows [166]:

- The constants featured in the theorems are typically too large to be practically useful. This is primarily due to inaccurate proof techniques, and it is currently unknown to which extent these constants can be reduced. Establishing phase transition results for structured sampling schemes, similar to those in Gaussian measurements, still requires significant mathematical advances.

- Designing optimal sampling densities remains an open problem. While theorems provide quantitative upper bounds on the required number of measurements for perfect reconstruction, proving optimality necessitates tight lower bounds. In practice, designing effective sampling densities is therefore heuristic.

- Most results are restricted to orthogonal wavelet transforms and ℓ_1-minimisation reconstructions. However, it is now recognised that more effective reconstruction results can be achieved in practice by employing redundant transforms, learned dictionaries, or non-convex regularisers. The understanding of CS in this context is still incomplete.

- The denoising capabilities of ℓ_1 regularisations are not correctly accounted for. In

practice, it is often observed that, in addition to allowing the recovery of missing information, ℓ_1 minimisation also serves as a regulariser capable of denoising the data.

- Limited knowledge exists regarding measurements collected along curves, such as spiral paths while a full analysis of parallel lines as proposed in Lustig's seminal work [57] was provided in [164].

Overall, existing theoretical results provide a good qualitative analysis of CS, allowing for designing effective sampling and reconstruction schemes. However, to date, theorems are unable to provide quantitative conclusions on the optimal sampling considerations and the number of needed measurements in CS MRI. In this study, we are particularly interested in exploring the sampling design requirements for variable density spiral (VDS) trajectories. VDS acquisitions have proven useful in CS MRI due to their ability to generate aliasing artefacts with lower coherence compared to those produced by standard spiral trajectories. We aim to shed light on the following question:

In Spiral-CS cDTI, how should an optimal VDS trajectory be parameterised as a function of the FOV, resolution, and the available SNR, in order to maintain a desired image quality while minimising the scan time?

Previous studies combining CS with undersampled spiral trajectories have provided qualitative comparisons of various heuristic VDS designs [167]. In particular, in diffusion MRI, the VDS design proposed by Kim et al. [168] has been widely adopted, yet there exists no generic method to relate data/sequence specifications to sampling design parameters. As the choice of target density significantly affects the quality of CS reconstruction, our approach aims to find an optimal density function that best adapts to the spatial resolution and scan time constraints.

Here, we introduce a quantitative analysis of practical parametrisation for undersampled spiral trajectories. Such research could provide valuable insights into designing application-specific MRI sequences, allowing for the selection of appropriate parameters for each scan without relying on a trial-and-error approach. However, this poses a challenging evaluation due to various factors influencing CS reconstruction quality, such as the resolution and the SNR for a specific acquisition [125]. Therefore, we suggest undertaking an empirical analysis of the aforementioned question, specifically focusing

on cDTI data. By employing a generic method based on retrospective undersampling, we present the results using a conventional ℓ_1-based CS reconstruction from various DW images of the heart. This analysis allows for quantitative assessment of image quality variations, leading to the identification of optimal threshold values.

4.3 Empirical and Quantitative Study of Optimal VDS Design

We have put forward the following methodology to empirically investigate the optimal parametrisation of a VDS sampling for CS cDTI. The findings may offer valuable practical and quantitative insights into undersampled spiral trajectory designs, particularly for cDTI.

4.3.1 Design Considerations

To achieve whole-heart coverage in cDTI, it is essential to account for various timing factors. This includes considering the total time available for imaging all slices at a given heart rate, the time needed to image each individual slice, and the total number of diffusion encoding directions. Therefore, the total imaging time of a multislice SE-based spiral sequence can be determined by:

$$T_{total} = \Big((T_D + T_{ro} \times N_s) \times N_{slc}\Big) \times N_q \tag{4.1}$$

where T_D denotes the duration of the two diffusion weighting gradients, T_{ro} is the duration of a spiral shot, N_s is the number of spiral shots required to fill the entire k-space, N_{slc} is the total number of slices, and N_q denotes the number of diffusion encoding directions. It should be noted that for simplicity, we excluded considerations such as the echo time and the time required for cardiac gating, fat saturation, and the suppression of other body parts, which are typically necessary for every cardiac MRI acquisition.

As indicated in Equation 4.1, the overall imaging time (T_{total}) can be shortened by reducing the number of diffusion-weighted directions, the number of spiral shots, the readout duration per shot, or the diffusion time. While the T_D can be reduced, this comes at the expense of SNR and reduced diffusion weighting. $N_q \geq 6$ also depends on the available SNR and cannot be arbitrarily reduced, as it will result in biased and

inaccurate diffusion estimates. In the case of a fully sampled spiral with a fixed spatial resolution, the product of N_s and T_{ro} remains approximately constant. Increasing the number of interleaves may result in phase inconsistencies between shots, particularly in the presence of pronounced eddy currents. Conversely, decreasing the number of interleaves while extending the readout time per spiral shot yields more susceptibility to field inhomogeneity and off-resonance effects.

Therefore, to maintain a fixed spatial resolution and consistent diffusion weighting, the most straightforward approach to reducing T_{total} is to decrease the total number of shots by reducing the total number of samples. This can be achieved through efficient k-space undersampling by incorporating an optimal variable density sampling function. Here, our goal is to enable the acquisition of each slice during a single heartbeat. Typically, at rest, heart rates range from 60 to 90 beats per minute (bpm), corresponding to R-R intervals of 667-1000 ms. In practice, within each repetition time (T_R), multiple slices can be acquired by adopting an interleaved approach [157]. Thus, the total imaging time can be expressed as:

$$T_{total} = \left(T_D + (T_{ro} \times N_s) \times \frac{N_{slc}}{n_{prep}}\right) \times N_q = T_R \times n_{prep} \times N_q \tag{4.2}$$

where n_{prep} is the total number of PGSE pulses, and $\frac{N_{slc}}{n_{prep}}$ determines the number of slices acquired followed by each PGSE preparation. Acquiring only two slices per preparation achieves a 50% reduction in the time needed to capture each image. Additionally, this minimises susceptibility to artefacts caused by motion.

4.3.2 Spiral Trajectory Design

In this section, we proceed with the optimal design of a base 2D spiral that lies in a $k_x - k_y$ plane, and the extension to three dimensions will be discussed in the next Chapter.

The standard uniform density spiral trajectory, also known as the Archimedian spiral, for a specified radial field-of-view FOV_r and number of shots (N_s) can be described in the polar coordinate system by the following relation [169]:

$$k_r = \frac{N_s}{2\pi FOV_r}\theta \tag{4.3}$$

where $k_r = \sqrt{k_x^2 + k_y^2}$ denotes the distance from the k_z axis.

For a spiral trajectory with a variable sampling rate, the sampling density becomes a *design* variable. Therefore, we need an accurate definition of the sampling density. Previous works described the relative sampling density compared to the Nyquist rate as a function of time [162] or the turning angle θ [170], or by using a simple analytic method [168, 171]; however, it is more comprehensive to consider the sampling density as a function of the k locations. Here, we define the sampling density as the inverse of the sampling interval Δk. With this definition, the sampling density becomes equivalent to the FOV in the case of uniform density sampling (Equation 2.17). Therefore, the FOV becomes a *variable* and is termed the *effective* FOV. It is important to highlight that the *effective* FOV differs from the nominal FOV, which remains constant and defines the region of interest in the image domain. We use this *effective* FOV as the design parameter for the optimal VDS trajectory.

As mentioned earlier, the purpose of variable density k-space sampling is to undersample in the low-energy, high-spatial-frequency region. Without any prior knowledge about the object being imaged, it seems reasonable to consider a radially symmetric density function. The rationale behind this choice is our intention to create sampling patterns that are rotation and translation invariant and adaptable to any organ. This means the patterns can reconstruct images independently of arbitrary rigid transformations. Therefore, we consider the effective FOV as a function of the distance from the k_z axis ($\sqrt{k_x^2 + k_y^2}$):

$$FOV_r(k_r) = FOV_{xy}\ \rho(k_r) \tag{4.4}$$

where $\rho(k_r)$ represents a radially decreasing function ranging between 1 to 0, and FOV_{xy} is the nominal FOV in the $x - y$ plane.

This effective FOV will determine the rate of increase in the radial direction compared to the rate of turning through the following relationship [162, 170]:

$$\frac{\Delta k_r}{\Delta \theta} = \frac{N_s}{2\pi FOV_r(k_r)} = \frac{N_s}{2\pi} \frac{1}{FOV_{xy}\ \rho(k_r)} \tag{4.5}$$

The VDS trajectory can now be determined by specifying key parameters including the desired number of shots, sampling time duration, number of points per shot, and the nominal FOV as:

$$C = k_r(t)\ \exp\Big(i \frac{\alpha FOV_{xy}\ \rho(k_r)}{N_s} \int_0^{T_{ro}} k_r(t) dt\Big) \tag{4.6}$$

where α is a constant that determines the number of spiral turns and helps maintain the targeted readout time. According to Equation 2.18, $k_{r_{max}} \leq \frac{1}{2\Delta r}$, indicating the

relationship between the maximum sampled k radius and the spatial resolution in the radial direction.

It is essential to note that while k-space trajectories follow a continuous path, the signal is sampled only at discrete intervals $\Delta t[i]_{1\leq i\leq N} \in [0,\ T_{ro}]$ along this path. Due to the inherent band-limiting filters in MRI receiver hardware, the measured information is essentially averaged over segments of trajectories between two sampled points [65]. In order for the designed theoretical trajectories to be implementable on an actual MRI scanner, hardware limitations should also be considered. The gradient system is limited by rise time (or slew rate) during the initial phase of the acquisition and will then maintain a constant amplitude for the remainder of the sampling duration. In addition, to fully exploit the flexibility of the spiral trajectory and maximise sampling efficiency, it is logical to sample at a rate at least equal to the gradient receiver bandwidth rBW [172, 173]. Hence, to minimise the filtering effects of the receiver hardware, the path in k-space between two consecutive samples should be smaller than the size of a k-space pixel ($\Delta k = 1/FOV$). Therefore, the constraints imposed by the hardware and the effective FOV include:

$$\begin{aligned} \|G\| &\leq \min\Big(\frac{2\pi}{\gamma}\frac{1}{FOV_r(k_r)}\frac{1}{T_s}, G_{max}\Big) \\ \|S\| &\leq S_{max} \end{aligned} \tag{4.7}$$

where G is the gradient waveform, S is the slew rate, and $T_s = 1/rBW$ is the sampling interval or the dwell time. These hardware constraints impose limitations on the trajectory speed and acceleration and can be formulated as inequality range limitations for each of the time points in the discrete waveform $\mathbf{k}[i]_{1\leq i\leq N}$:

$$\begin{aligned} \|\dot{\mathbf{k}}[i]\| &= \left\|\frac{\mathbf{k}[i]-\mathbf{k}[i-1]}{\delta t}\right\| < \min\Big(\frac{\gamma}{2\pi}G_{max}, \frac{1}{FOV_r(k_r)}\frac{1}{T_s}\Big) \\ \|\ddot{\mathbf{k}}[i]\| &= \left\|\frac{\mathbf{k}[i-1]-2\mathbf{k}[i]+\mathbf{k}[i+1]}{\delta t^2}\right\| < \frac{\gamma}{2\pi}S_{max} \end{aligned} \tag{4.8}$$

Here, δt represents the gradient sampling period, defined as the raster time. The trajectory meeting the above constraints can be determined numerically. We used the optimum gradient design Toolbox (OptGrad V0.2) [174, 175] to compute the practical gradient waveforms and the corresponding actual trajectory in k-space. Given the designed theoretical k-space trajectory, the toolbox calculates the actual trajectory that satisfies the specified gradient constraints by solving a time-optimal control problem

to get from one point to the other in minimum time. We opted for the rotationally invariant option to ensure that all spiral shots adhere to the same constraints. This approach will guarantee that an MRI gradient system is capable of playing out such waveforms. The resulting actual k-space trajectories were then used for all our evaluations. Multi-shot trajectory was obtained by rotating the spiral interleave around the k-space centre with a uniform angle spacing: $\Delta\theta = 2\pi/N_s$.

4.3.3 A Generic Method

We propose a generic function ρ of the radius k_r, which describes the relative sampling density compared to the Nyquist rate, as illustrated in Figure 4.4.

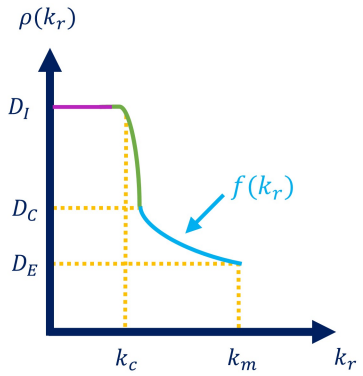


Figure 4.4: Proposed generic sampling density function, colour-coded into three sections corresponding to those expressed in Equation 4.9. Refer to Table 4.1 for definitions of all notations used.

The proposed function is parameterised into two primary domains, with a smooth transition region that enables multilevel sampling, consistent with recent findings in CS MRI [163], as discussed earlier. Specifically, the first domain supports sampling at the Nyquist rate up to a predefined radius, k_c, while the second domain allows for undersampling to achieve a shorter acquisition time. The transition between these regions is governed by a Fermi-Dirac relation, which adjusts the sampling density, D_C, to account for the decrease in signal energy. This function is mathematically described

Table 4.1: Notations used in the proposed sampling function parametrisation

Notations	Meaning
D_I	Starting density[1]
D_C	Central density
D_E	Ending density
k_m	Maximum spatial frequency (cm^{-1})
k_c	Fully sampled k-space radius (percentage of k_m)
τ	Steepness parameter[2]
$f(k_r)$	Function governing the density variation in the under-sampling region[3]
$T_{ro} \in \mathcal{R}$	Readout duration (ms)
$N_s \in \mathcal{S}$	Number of spiral shots
$T \in \mathcal{T}$	Total acquisition time (ms)
G_{max}	Maximum gradient amplitude (mT/m)
S_{max}	Maximum gradient slew rate (mT/m/ms)
T_s	Dwell time (μs)
FOV_{xy}	Nominal FOV in $x-y$ plane (cm)
$\sigma \in \Sigma$	Noise level
$\text{SNR} \in \mathcal{SNR}$	Signal-to-noise ratio
$\mathbf{x}_{ref} \in \mathbb{C}^{N^2}$	Reference Image
$\mathbf{x}_{\text{SNR}}$	Noisy image characterised by its input SNR
$\hat{\mathbf{x}}_{\rho_{km}}$	Reconstructed image from retrospective VDS data of sampling density function ρ_{k_m}
$Q; Q_0$	Image quality metric; image quality threshold
$\mathbf{z} \in \mathbb{C}^{N^2}$	Wavelet coefficients of image $\mathbf{x}$
$\mathbf{y} \in \mathbb{C}^m$	Measured Fourier data
$\mathcal{F}_\Omega$	Fourier Transform over the support $\Omega \subseteq \{1, ..., N^2\}$
$\mathbf{\Psi}$	q-level discrete wavelet transform ($2^q \leq N$)

[1] $D_I > 1$ corresponds to oversampling near the k-space centre, $D_I = 1$ satisfies the Nyquist criterion, and $D_I < 1$ means undersampling in central parts.

[2] determines the sharpness of the transition region

[3] Four different functions were examined.

in Equation 4.9:

$$\rho(k_r) = \begin{cases} D_I & \text{if } k_r \leq k_c - \epsilon \\ D_I - \frac{D_I - D_C}{1+\exp(-\tau(k_r - k_c))} & \text{if } k_c - \epsilon < k_r \leq k_c + \epsilon \\ f(k_r) & \text{if } k_r \leq k_m \end{cases} \tag{4.9}$$

where ϵ is a small margin, and τ determines the sharpness or steepness of the transition region (higher τ values correspond to sharper transition). All the notations are listed in Table 4.1. We divided all the parameters into three categories: (i) sampling design parameters, (ii) hardware system parameters, and (iii) data-specific parameters. The representation of the fully sampled radius k_c and the maximum radius k_m in k-space, as well as the resulting variable-density sampling trajectory—featuring a densely sampled centre and an undersampled periphery—compared to a uniform-density spiral, is shown in Figure 4.5.

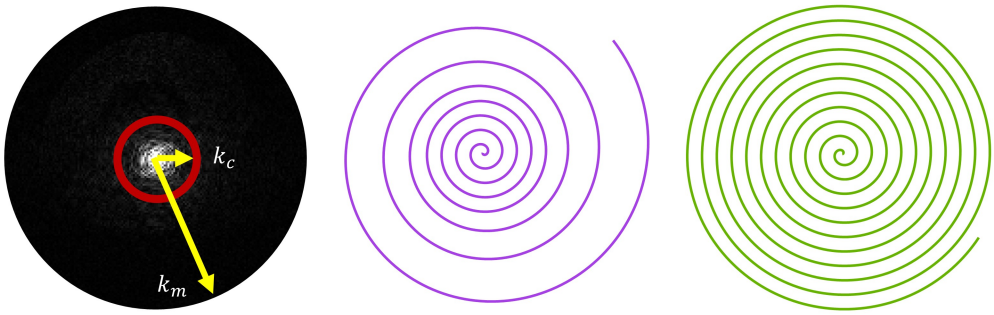


Figure 4.5: (a) Representation of the k_c and k_m parameters in the k-Space of a DW image, (b) the resulting variable-density spiral trajectory, which densely samples the centre while undersampling the periphery, compared to (c) a uniform-density sampling, where all samples are acquired at a constant rate and the entire k-space is fully sampled.

First, we constructed a set of sampling density functions, denoted by $\rho_{k_m} \in P$, where k_m represents the range of studied spatial resolutions RES ($k_m = \frac{1}{RES}$). VDS trajectories with different acquisition times ($T \in \mathcal{T}$), denoted by C^T_{spi}, were obtained by varying the readout time $T_{ro} \in \mathcal{R}$ and the number of shots $N_s \in \mathcal{S}$. Here, $\mathcal{R}$ and $\mathcal{S}$ represent the range of studied readout times and number of interleaves, respectively. Spiral data is then obtained by applying forward non-uniform fast Fourier transform

(NUFT) [176] to the spatial frequency locations on the specified k-space trajectories using a broad range of short-axis midventricular DW images (see Appendix A for more details).

Complex-valued zero-mean Gaussian noise with varying standard deviations $\sigma \in \Sigma$ were added to the complex k-space of the reference images ($\mathbf{x}_{ref}$). The resulting noisy k-space data was then Fourier transformed to obtain images with various input SNR levels ($\mathbf{x}_{SNR}$). Here, Σ represents the range of studied noise levels. Each noise level was then expressed as its corresponding input SNR $\in \mathcal{SNR}$, a metric commonly used in the MRI community. Finally, the simulated VDS data were reconstructed using a proper CS reconstruction. The proposed pipeline is outlined in Algorithm 1 and illustrated through a flowchart in Figure 4.6.

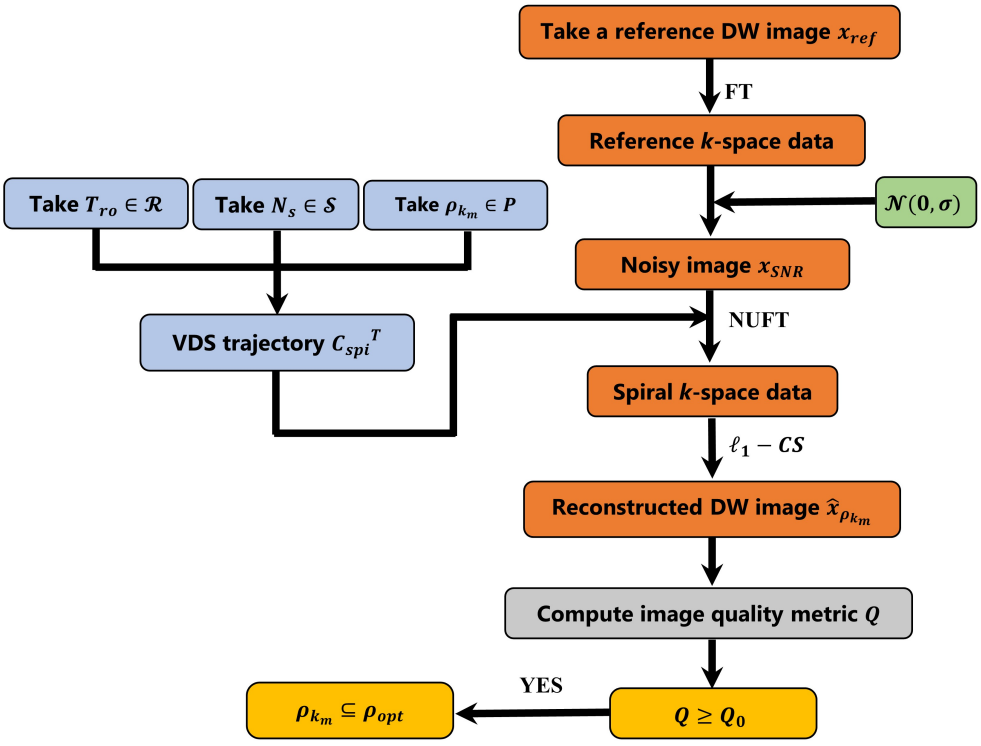


Figure 4.6: Flowchart illustrating the proposed pipeline.

We now introduce the concept of an optimal sampling density function $\rho_{opt}(Q_0, T, \text{SNR})$ as a function of (i) the targeted image quality metric Q expressed in terms of a quality

threshold Q_0, (ii) the acquisition time of the VDS trajectory, and (iii) the input SNR:

$$\begin{aligned} \rho_{opt}(Q_0, T, \mathrm{SNR}) = & \max_{\rho_{km}} Q(\hat{\mathbf{x}}_{\rho_{km}}, \mathbf{x}_{SNR}) \\ & \text{s.t.} \quad Q(\hat{\mathbf{x}}_{\rho_{km}}, \mathbf{x}_{SNR}) \geq Q_0 \end{aligned} \tag{4.10}$$

Here, $\mathbf{x}_{SNR}$ represents the reference image with a given SNR, $\hat{\mathbf{x}}_{\rho_{km}}$ denotes the corresponding reconstruction from a retrospectively spirally sampled data using a sampling function ρ_{km}, which is defined by three free parameters k_c, D_C, and D_E, and Q_0 is a quality threshold.

To isolate and examine the individual effects of either the sampling function or the noise level, one of these parameters was kept constant while the other was varied. This framework enables a systematic exploration of the impact of the proposed sampling parametrisation on the quality of reconstructed images from VDS data.

Algorithm 1 Find $\rho_{opt}(Q_0, T, \mathrm{SNR})$

Input: Image quality metric Q with a threshold of Q_0

for a spiral trajectory C^T_{spi} of sampling duration T **do**
  Take a reference DW image $\mathbf{x}_{ref}$
  for σ in Σ **do**
    $\mathbf{x}_{SNR} \leftarrow \mathbf{FFT^{-1}}\Big(\mathbf{FFT}(\mathbf{x}_{ref}) + \mathcal{N}(0, \sigma)\Big)$
    for ρ_{k_m} in P **do**
      $Spiral\ data \leftarrow sample(\mathbf{x}_{SNR}, \rho_{k_m})$
      $\hat{\mathbf{x}}_\rho \leftarrow reconstruct(Spiral\ data)$
      $Q(\hat{\mathbf{x}}_\rho, \mathbf{x}_{SNR})$
    end for
  end for
end for

Output: $\rho_{opt}(Q_0, T, \mathrm{SNR}) \quad \forall T \in \mathcal{T},\ \mathrm{SNR} \in \mathcal{SNR}$

4.3.4 Experimental Validation with Real cDTI Data

We evaluated the performance of the proposed generic method using real cDTI images. Our experiments focused on ex vivo cDTI data, which enabled assessing the performance of the proposed strategy independently from motion and physiological considerations.

DW images of an ex vivo normal human heart[1] [40, 48] were utilised for retrospective evaluations. These images were acquired with a 3D fast spin echo (FSE) sequence using a four-element phased array coil on a 1.5 T GE CV/I MRI scanner (GE Medical System, Wausheka, WI, USA) equipped with an enhanced gradient system (maximum amplitude = 40 mT/m and slew rate = 150 T/m/s). Phase encoding was applied in both the y and z directions to generate 3D image data. Because of the long duration of the diffusion gradients and the short T_2 of fixed myocardium ($T_2 \approx 65$ ms), only two echoes were used. The corresponding imaging parameters were as follows: image size = $256 \times 256 \times 134$ mm, resolution = $0.43 \times 0.43 \times 1.0$ mm^3, a total of 21 diffusion gradient directions (including 2 non-weighted and 19 DW images), and a b-value of 1500 s/mm^2. The total scan time was approximately 70 hours with a TR of 700 ms. We used three short-axis mid-ventricular slices from the dataset for analysis[2]. The examined images are illustrated in Figure 4.7.

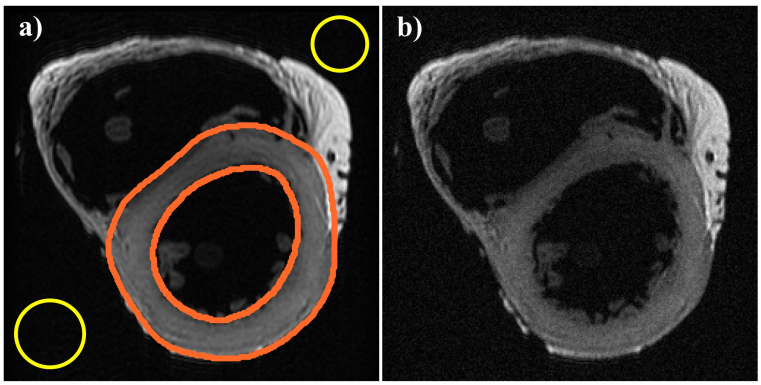


Figure 4.7: DW images of a human heart with the highest studied SNR > 30 (a) and the lowest input SNR of 12 (b). The orange contours indicate the region-of-interest (ROI) in the myocardium, and the small yellow circles represent the area in the background used to calculate the input SNR.

[1] Available at https://www.cvrgrid.org/data/ex-vivo.html

[2] Since only two echoes were employed, the correction for the b-value due to the echo train was minimal and thus not considered.

4.3.5 The Studied Pipeline

Input SNR

We investigated a range of input SNR values from 12 to 35 for the DW images to resemble the SNR range typically found in clinical examinations. The input SNR was computed as the ratio of the mean signal in a region-of-interest (ROI) in the LV myocardium, indicated by orange contours in Figure 4.7(a), to the averaged standard deviation of the signal in the background (yellow circles in Figure 4.7(a)). It should be noted that this approach for calculating the SNR in MRI images is valid when the images were acquired with a single receiver coil. However, since information on the actual noise was not available, this approach could provide a reasonable approximation. The noisy image corresponding to the most extreme SNR value is displayed in Figure 4.7(b).

Sampling Function Parameters

We considered the following setting:

- Gradient hardware specification of a clinical scanner [177, 178]: $G_{max} = 43$ mT/m, $S_{max} = 180$ mT/m/ms, $T_s = 2.8$ μs
- Isotropic spatial resolution = 1 mm, corresponding to $k_m = 5$ cm^{-1}
- Nominal FOV$_{xy}$ = 12 cm
- $D_I = 1$ to allow sampling at the Nyquist rate near the k-space centre

We then examined the proposed VDS trajectory by varying the sampling density parameters as follows: the ratio of central k-space radius to the maximum k-space radius (k_c/k_m) ranged from 1% to 20%, D_C ranged from 0.8 to 0.5, D_E ranged from 0.5 to 0.2, and $f(k_r)$ was selected from the set {uniform, linear, quadratic, Gaussian}. The number of shots (N_s) decreased from 8 to 1, while the readout time (T_{ro}) increased from 5 to 15 ms. To address off-resonance effects, the readout duration was limited to 15 ms, as a prior study [177] demonstrated that no correction was necessary at this spiral sampling duration. Finally, τ was set to 12, as it resulted in a satisfactory transition region. Examples of the studied sampling density functions $\rho(k_r)$ are illustrated in Figure 4.8.

We also implemented a uniform density spiral (UDS) trajectory with the same setting for comparison. An undersampling or acceleration factor is defined as the ratio of the total acquisition time T_{acq}^{UDS} for the UDS trajectory compared to the VDS trajectory T_{acq}^{VDS}: $R = \frac{T_{acq}^{UDS}}{T_{acq}^{VDS}}$. As the k-space samples are optimally spread along a variable density adapted to compressed sensing, the proposed method maximises the amount of information measured per shot, reducing the number of shots needed for high-resolution imaging. Consequently, very high acceleration factors (AF) can be achieved.

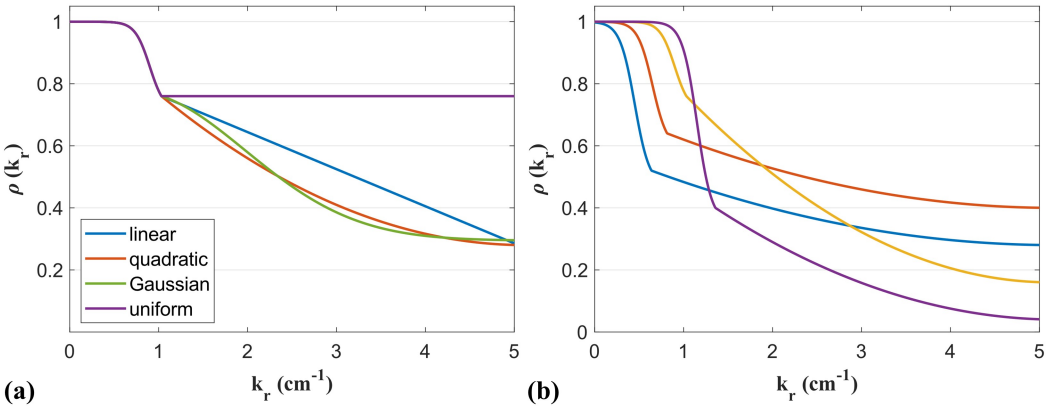


Figure 4.8: Examples of the studied sampling density functions: (a) Four examined functions governing the undersampling region ($f(k_r)$), (b) Variations in D_C and k_c in the proposed density function.

Image Reconstruction

The VDS DW images were reconstructed using a non-Cartesian CS reconstruction by solving an optimisation problem of the form:

$$\hat{\mathbf{z}} = \underset{\mathbf{z}}{\operatorname{argmin}}\{\frac{1}{2}\|\mathbf{y} - \mathcal{F}_\Omega \mathbf{\Psi}\mathbf{z}\|_2^2 + \lambda_1\|(\mathbf{z})\|_1 + \lambda_2 TV(\mathbf{\Psi}\mathbf{z})\} \tag{4.11}$$

Here, $\mathbf{y} \in \mathbb{C}^m$ represents the spiral Fourier data, $\hat{\mathbf{x}} = \mathbf{\Psi}\mathbf{z} \in \mathbb{C}^{N^2}$ denotes the reconstructed image, and $\mathbf{z}$ represents the wavelet coefficients. $\mathcal{F}_\Omega$ denotes the NUFT operator, and $\{\lambda_i\}_{i=1}^2 > 0$ are the regularisation parameters. The NUFT was implemented with the Kaiser-Bessel Kernel [78] utilising Voronoi density compensation [179]. Additionally, *Symmlet* 8 with a 4-level decomposition was used as the sparsifying transform ($\mathbf{\Psi}$). The reconstruction problem was solved using the conjugate gradient

method [57]. The initial reconstruction obtained from regridding served as the input for the CG algorithm, and no preconditioner was applied. Convergence was defined as the point at which the rate of change in the cost function (Equation 4.11) dropped below a threshold of 1×10^{-5}. The regularisation parameters were tuned over a range of $[10^{-6}; 10^{-2}]$, and the values that led to the best reconstruction quality based on the criterion described below were selected.

4.3.6 Computational Setup

All reconstructions detailed in this study were performed using MATLAB (Release 2021a, MathWorks, Natick, Massachusetts, USA) on a PC equipped with a 2.54 GHz 8-core processor and 128 GB RAM.

4.3.7 Image Quality Metrics

To quantitatively evaluate the image quality, we utilised three reference-based metrics: the structural similarity (SSIM) index [180], the normalised root mean square error (NRMSE in %), and the peak SNR (PSNR). The SSIM index is designed to model the human visual system (HVS) and has become increasingly popular in the MRI community for assessing image quality [125]. NRMSE and PSNR are widely used intensity-based metrics, calculated as:

$$\begin{aligned} \mathrm{NRMSE}(\hat{\mathbf{x}}, \mathbf{x}) &= \frac{\|\hat{\mathbf{x}} - \mathbf{x}\|_2}{\|\mathbf{x}\|_2} \times 100 \\ \mathrm{PSNR}(\hat{\mathbf{x}}, \mathbf{x}) &= 10 \log \left(\frac{\max(|\mathbf{x}|)^2}{MSE} \right) \end{aligned} \tag{4.12}$$

where $\hat{\mathbf{x}}$ is the reconstructed image and $\mathbf{x}$ is the reference image. The SSIM index ranges between 0 (indicating no correspondence with the reference) and 1 (indicating a perfect match with the reference), and the NRMSE varies between 0% and 100%, representing the error percentage. SSIM captures structural relationships and effectively identifies distortions that impact the visual perception of an image (e.g., blurring or contrast changes), whereas NRMSE tends to focus on numerical differences and is more sensitive to minor pixel discrepancies. Utilising both metrics together can provide a comprehensive evaluation of image quality, addressing both perceptual differences and numerical accuracy.

For each reconstruction, the corresponding reference image was selected as the Cartesian image. Additionally, all metrics were computed over the myocardium ROI.

Table 4.2: Correspondence between SSIM index and Mean Opinion Scores (MOS) between an image and a reference [181]

SSIM	MOS	Quality	Impairment
≥ 0.99	5	Excellent	Imperceptible
$[0.95, 0.99]$	4	Good	Perceptible but not annoying
$[0.88, 0.95)$	3	Fair	Slightly annoying
$[0.5, 0.88)$	2	Poor	Annoying
≤ 0.5	1	Bad	Very annoying

Since the SSIM and PSNR metrics showed similar trends, PSNR performance is only presented in the first part of the results, with NRMSE and SSIM indices being used for the remainder of the study. Moreover, to quantitatively express the desired image quality, we chose a specific SSIM threshold, denoted as Q_0, with a set value of 0.88, and a NRMSE threshold of 4%. This value is regarded as indicative of fair image quality [181]. A correlation between SSIM scores and mean opinion scores (MOS) is presented in Table 4.2.

The NRMSE and SSIM values from all experiments ($\approx$ 200,000) were analysed using a blocked factorial N-way analysis of variance (ANOVA). By employing factorial ANOVA, we can evaluate main effects, which reflect the influence of individual factors (such as central density or ending density) on a particular response variable. We can also determine whether interactions exist between factors, i.e., whether the effect of a specific factor is influenced by one or more other factors. This analysis aimed to identify the optimal combination of the sampling density parameters and the spiral trajectory specifications. The model also incorporated the interactions between trajectory and sampling density function parameters, with significance determined at a $P-\text{value} < 0.05$.

4.4 Results

4.4.1 Effect of the Sampling Density Function

We observed that all the examined $f(k_r)$ relations in the sampling density function $\rho(k_r)$ exhibited similar trends as the parameters were varied. Among them, the "quad-

ratic" function consistently yielded higher SSIM and PSNR scores and lower NRMSE. Therefore, in this section, we are only presenting the analysis of this function in the undersampling region.

Figure 4.9 shows the SSIM and PSNR scores averaged over all studied DW images as a function of the percentage of the Nyquist sampled radius to the maximum k radius (k_c/k_m) for all combinations of the sampling density functions $\rho(k_r)$. The results shown correspond to the VDS trajectory comprising of $N_s = 8$ and $T_{ro} = 15$ ms. Figure 4.10 illustrates the corresponding NRMSE values.

It can be observed that for this spiral trajectory, almost all the sampling density functions have reasonably acceptable performance. For $D_C \geq 0.7$, the averaged SSIM values remain relatively constant, irrespective of the increase in the Nyquist sampled radius or the variations in the ending density D_E. However, for $D_C < 0.7$, the SSIM typically decreases with increasing k_c, with the lowest quality observed when $D_E = 0.5$. Similar trends are observed for PSNR values, and NRMSE also shows comparable but reversed trends (Figure 4.10). For all evaluated DW images, the main effects of sampling density function parameters and the interaction between each pair of these parameters were all significant ($P < 0.001$), except for the interaction between k_c and D_E, which was insignificant ($P = 0.2133$). When controlling for the D_C and D_E, and interactions between parameters, $k_c/k_m = 1\%$ had the highest estimate for SSIM ($P < 0.05$). However, PSNR scores showed the highest scores at $k_c/k_m = 20\%$. This is expected as more signal is captured with increased Nyquist sampled radius. When controlling for the k_c and D_E, and interactions between parameters, $D_C = 0.8$ yielded the highest SSIM estimate ($P < 0.05$). And the highest SSIM score was estimated with no significant difference for $D_E = 0.3$ and $D_E = 0.2$, when controlling for the other two parameters.

Figure 4.11 shows the effect of the density function variations in terms of the residual artefacts in the reconstructed images after a few iterations, compared to the reference image. The absolute error images are scaled 15 times for better visualisation. We observed that smaller ending density (D_E) leads to increased undersampling in the outer k-space regions and produces more noise-like artefact patterns. Conversely, reducing the central density D_C resulted in more structural artefacts. The significant effect of D_C on the artefacts is also manifested in the SSIM scores plotted in Figure 4.12 for all the examined density functions. Reducing the D_C degraded the image quality, leading

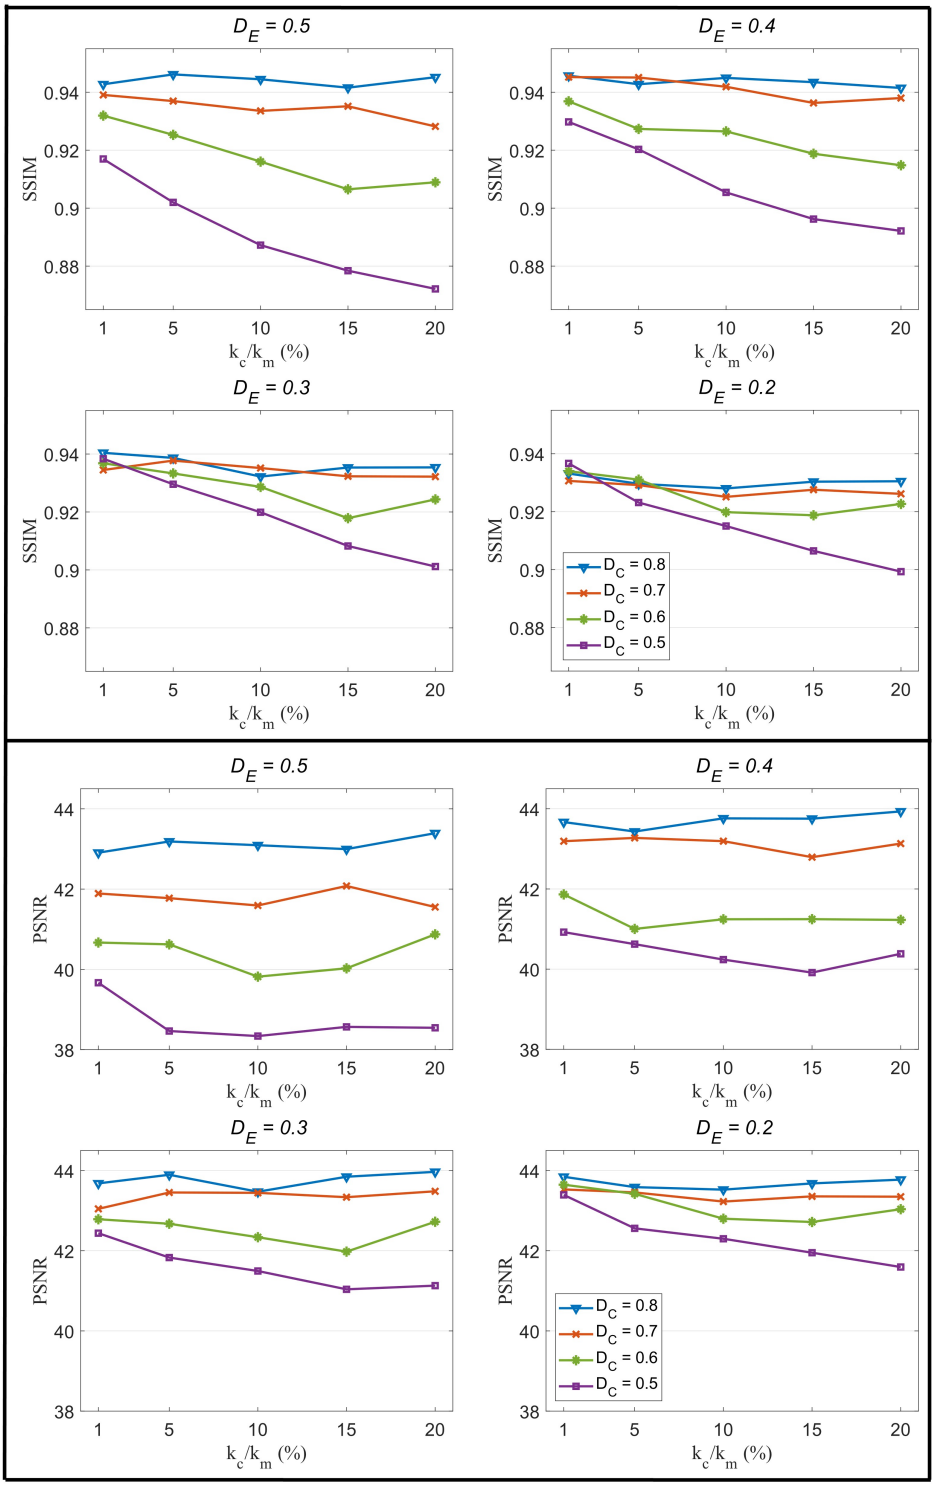


Figure 4.9: (Caption on next page.)

Figure 4.9: (Continued from previous page).) Performance analysis of all the examined sampling density functions in the proposed pipeline. SSIM and PSNR scores averaged over all DW images are plotted as a function of the ratio of the central radius to the maximum k radius (k_c/k_m). Different plots illustrate the effect of variations in D_C and D_E.

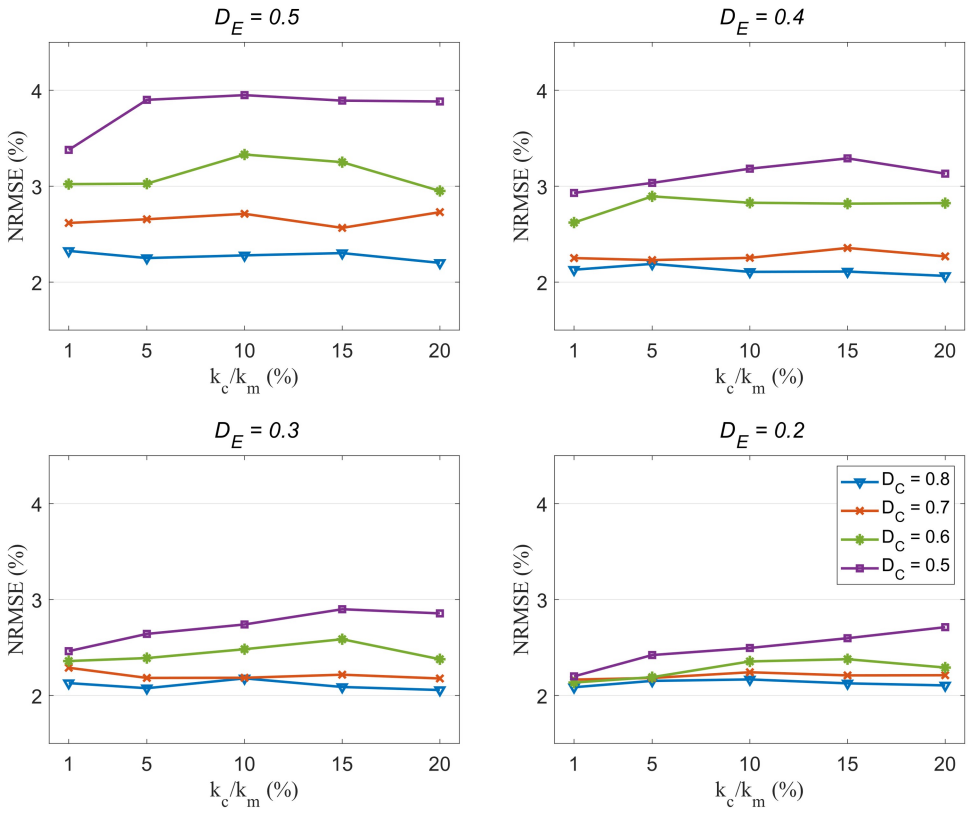


Figure 4.10: Quantitative analysis of all the examined sampling density functions in the proposed pipeline. NRMSE scores (in %) averaged over all DW images are plotted as a function of the ratio of the central radius to the maximum k radius (k_c/k_m). Different plots illustrate the effect of variations in D_C and D_E.

to low SSIM scores falling below the threshold. Additionally, the SSIM scores were significantly different from each other for all D_C values ($P < 0.05$), indicating that the value of D_C has a significant impact on the quality of the reconstructed images.

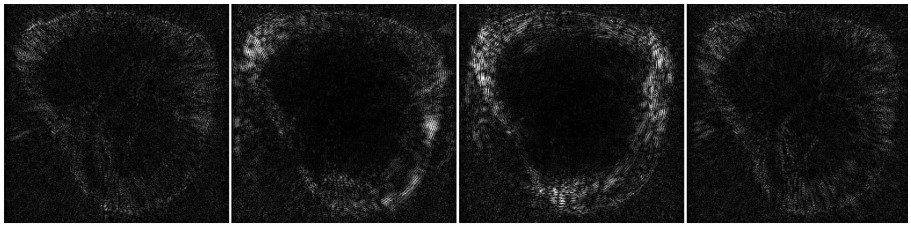


Figure 4.11: Illustration of the residual artefacts resulting from four arbitrary examined density functions in the proposed pipeline. The Absolute error images between the reconstruction and the reference are displayed scaled by 10X for better visualisation ($|\hat{x}_{\rho_{km}} - x_{ref}| \times 10$).

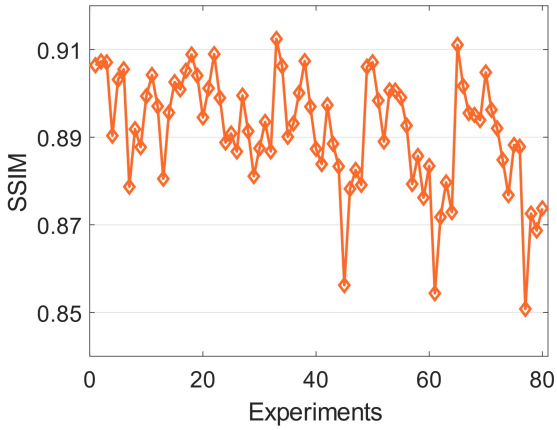


Figure 4.12: SSIM scores obtained from one DW image plotted for all the examined density functions in the proposed pipeline.

4.4.2 Effect of the Trajectory Configuration

Readout Duration

Figure 4.13 compares the quantitative performance of the examined VDS trajectories with varying readout durations and a fixed number of shots $N_s = 8$. The results are shown for the sampling function with $k_c/k_m = 20\%$ and $D_C = 0.8$ across the entire range of D_E. The averaged SSIM and NRMSE scores over all DW images are displayed. It was observed that as the readout time decreased, the reconstruction quality degraded, and large aliasing artefacts appeared when $T_{ro} = 5$ ms, resulting in the worst reconstruction quality. This is clearly manifested in Figure 4.13, with the

lowest SSIM scores and the highest NRMSE values in VDS trajectories with $T_{ro} = 5$ ms among all examined VDS trajectories.

Again, the main effects of sampling density function parameters, the readout duration, and the interaction between parameters were all significant ($P < 0.05$), except for the interaction between k_c and D_E, which was insignificant ($P = 0.2083$). We found that as the readout time decreased, the optimal value for k_c, where $\|Q(\hat{\mathbf{x}}_{k_c}) - Q_0\|_2^2$ is minimised, shifted towards higher values ($k_c/k_m = 20\%$), while the optimal value for D_E was pushed towards the smallest studied value of 0.2. The substantial difference between NRMSE values at $D_E \geq 0.3$ for $T_{ro} = 5$ ms implies that for shorter readout times, it is desirable to increase the undersampling ratio towards the outer k-space parts. Spiral readout durations ≥ 10 ms exhibited acceptable performance for this high-resolution imaging setting, and no significance difference was observed within the group with $T_{ro} \geq 10$ and $D_E = 0.3, 0.2$ for both SSIM and NRMSE metrics.

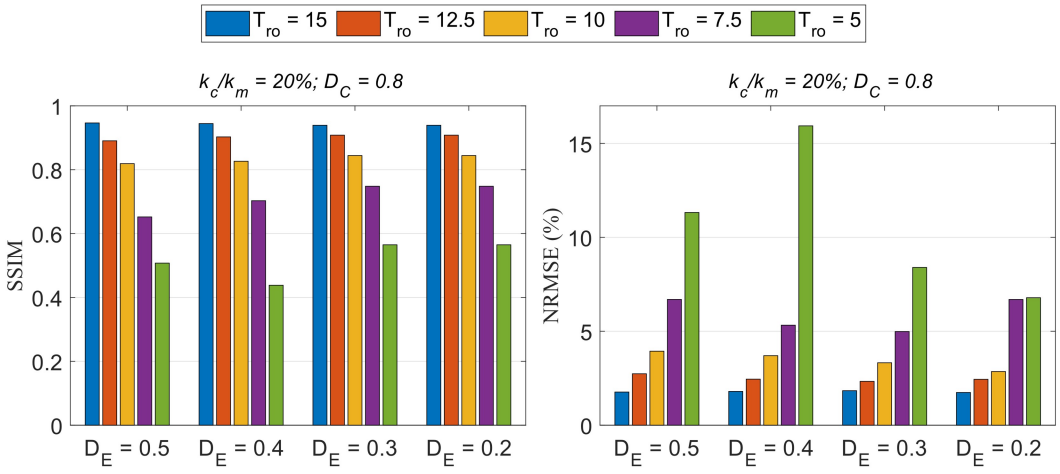


Figure 4.13: Quantitative analysis of the combination of sampling density function parameters and the spiral readout time in the proposed pipeline. The SSIM scores (left) and NRMSE scores (right) averaged over all DW images are displayed for varying D_E and T_{ro} values. It is evident that the image quality degrades with decreasing T_{ro}, and VDS trajectories with $T_{ro} \geq 10$ fall within the acceptable quality threshold.

Number of Shots

Figure 4.14 compares the quantitative performance of the examined VDS trajectories with varying readout durations and number of shots for a fixed sampling density function ($k_c/k_m = 20\%$, $D_C = 0.8$, and $D_E = 0.2$). The averaged SSIM and NRMSE scores over all DW images are displayed. Reducing the number of shots led to pronounced artefacts, particularly when $T_{ro} < 10$ ms. For $T_{ro} \geq 10$ ms, reducing the number of shots to 5 still resulted in quality metrics above the acceptable threshold, while reducing the number of shots further, regardless of the readout duration, led to high NRMSE and low SSIM scores beyond the acceptable range. In the ANOVA analysis, when assessing the effect of the individual parameters, the SSIM scores showed significant differences with the variations in the readout duration ($P < 0.001$), unlike the variations in the number of shots, which were insignificant ($P > 0.5$). The interaction between these two parameters for VDS trajectories with a fixed sampling density function was also found to be insignificant ($P > 0.5$). This suggests that for a fixed sampling density function, the readout duration has a substantial impact on the quality of the resulting images.

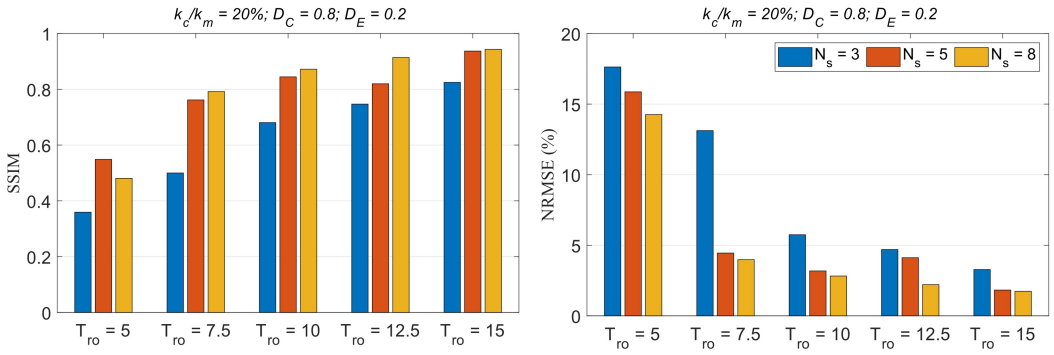


Figure 4.14: Quantitative analysis of the combination of the spiral readout time and the number of shots for a fixed sampling density function. The SSIM scores (left) and NRMSE scores (right) averaged over all DW images are displayed. The quality falls beyond the acceptable range when $T_{ro} < 10$ ms, regardless of the number of shots.

Figure 4.15 presents the combined effect of the number of shots (N_s) and readout duration (T_{ro}) and summarises the VDS trajectories that lead to acceptable quality metrics. Figure 4.15(a) displays the mean SSIM scores for three different sampling density functions. In Figure 4.15(b), the SSIM scores are plotted against the total

acquisition time, $T_{acq} = T_{ro} \times N_s$, using the most frequently reported parameter group in the analyses, where no significant differences were observed: $k_c/k_m = 20\%$, $D_C = 0.8$, $D_E = 0.2$.

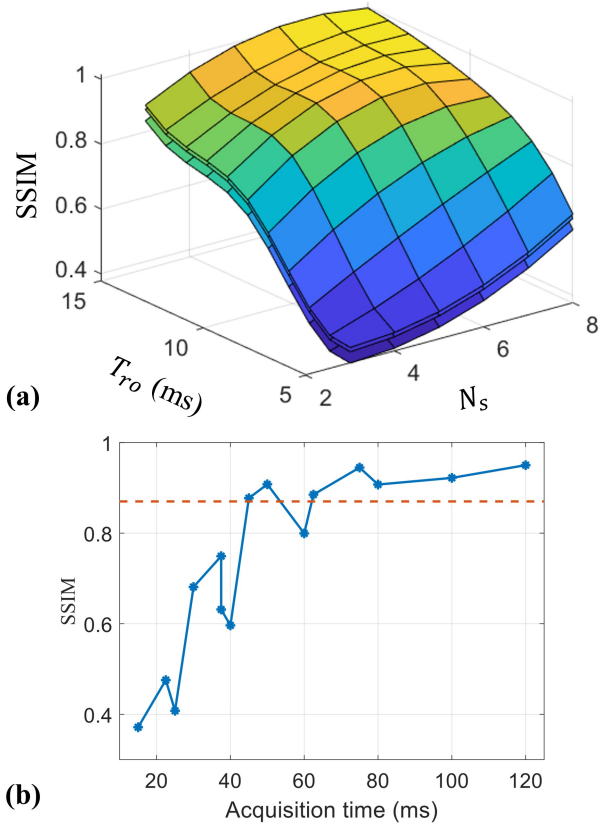


Figure 4.15: (a) A 3D view showing the mean SSIM (structural similarity index measure) scores as a function of readout duration (T_{ro}) and number of spiral shots (N_s) in the proposed VDS trajectory. The three curves correspond to three density functions. (b) SSIM scores as a function of the total acquisition time $T_{acq} = T_{ro} \times N_s$. The horizontal line indicates the target image quality, set at $Q_0 = 0.88$.

4.4.3 Effect of the DW Direction

Figure 4.16 displays the SSIM scores obtained for six arbitrary DW images at eight randomly selected sampling density functions from the proposed pipeline ($S_i : \ i = 1, \cdots, 8$. The results shown correspond to the VDS trajectory with $T_{ro} = 15$ ms and $N_s = 8$. For a fixed sampling density function, variations in the reconstructed

image quality and subsequently the SSIM scores are observed between different DW directions, with some sampling density functions yielding more differences than others. The ANOVA analysis also showed that the main effect of the DW direction and the sampling density, as well as the interactions between them, were significant ($P < 0.001$).

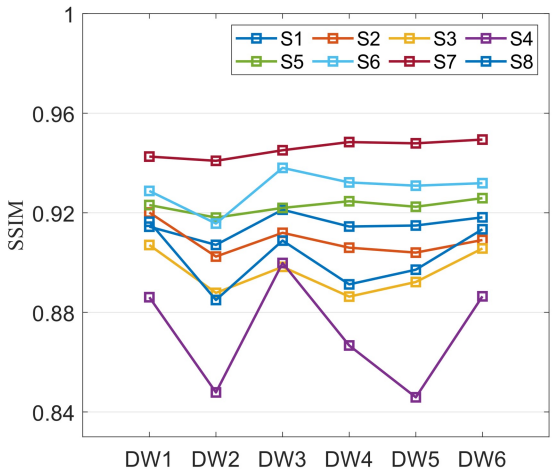


Figure 4.16: SSIM scores obtained from six arbitrary DW images are plotted for eight randomly selected sampling density functions $\{S_i\}_{i=1}^{8}$ from the proposed pipeline. Large variations between the reconstruction quality of DW images at some sampling density functions can be interpreted as sub-optimality of those density functions in fully capturing the diffusion-related information in all DW directions.

4.4.4 Performance of the VDS Trajectory

The performance of the proposed VDS trajectory is compared with corresponding fully sampled and undersampled UDS trajectories in Figure 4.17. All trajectories were simulated with a readout duration (T_{ro}) of 15 ms. The fully sampled UDS trajectory consisted of 80 spiral shots, while both the undersampled UDS and the VDS trajectories had 8 shots. The DW images were obtained using a regridding reconstruction technique. It should be noted that the Nyquist rate for a fully sampled spiral imposes $N_s = \left\lfloor \frac{FOV_{xy}}{RES_{xy}} \right\rfloor \approx \frac{\text{matrix size}}{2}$, which corresponds to $N_s \approx 120$ in our case. The resulting acceleration factors achieved for the UDS and VDS trajectories with respect to the Nyquist sampled spiral are $R = 1.5$ and $R = 15$, respectively. As observed

from the error images in Figure 4.17(a), standard spiral trajectories could effectively tolerate mild undersampling factors, with streaking artefacts appearing outside the myocardium ROI. However, increasing the undersampling factor results in coherent aliasing artefacts with pronounced inaccuracies. In contrast, the VDS design yields non-coherent noise-like artefacts suitable for CS reconstructions.

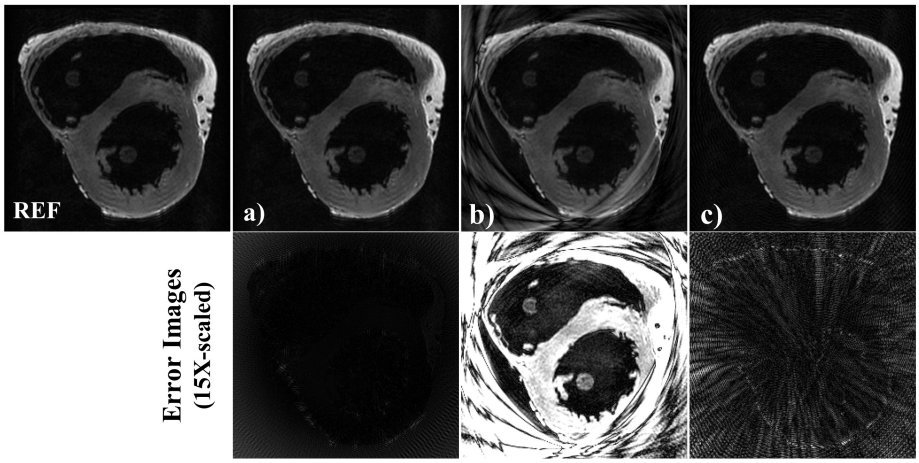


Figure 4.17: Performance comparison of (a) fully sampled and (b) undersampled uniform density spiral (UDS) trajectories with (c) the proposed VDS trajectory. DW images are obtained with the gridding method. The absolute difference maps, scaled by a factor of 15 for better visualisation of artefacts ($|\hat{x} - x_{ref}| \times 15$), are shown in comparison to the reference.

Another useful method for analysing the performance of a sampling design is through the point spread function (PSF). The PSF is defined as the modulus of the Fourier transform of the sampling pattern viewed as a set of Dirac impulses. Figure 4.18 displays the 2D and 3D views of the PSFs corresponding to a spiral trajectory with $T_{ro} = 10$ ms and $N_s = 5$, comparing the VDS and UDS schemes. The cross-sections of the functions are plotted in Figure 4.19 for better visualisation of the energy distribution of the sampling schemes. It can be seen that the undersampling in the UDS scheme leads to elevated side peaks, resulting in severe aliasing artefacts.

Figure 4.20 illustrates how the designed trajectory adjusts to comply with the gradient hardware constraints, demonstrating the preservation of the target sampling density in the proposed radial density function. The corresponding gradient and slew rate

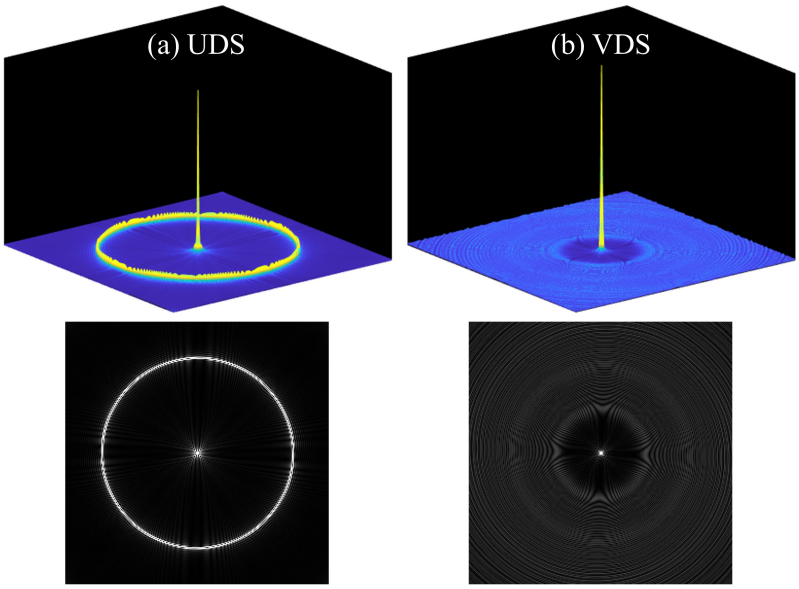


Figure 4.18: 3D (top) and 2D (bottom) views of the point spread function of the UDS and VDS trajectories with the same readout and number of shots.

waveforms in the x and y axis are illustrated in Figure 4.21.

4.5 Discussion

In cDTI, the spatial resolution is often sacrificed to acquire a few slices covering the entire myocardium in the available scan time. Here, by leveraging the motion robustness of spiral acquisitions and their flexibility in achieving high spatial resolutions, we explored an optimal variable density spiral (VDS) trajectory to enable spatial resolutions as high as 1 mm while efficiently shortening the acquisition time and preserving a targeted image quality. However, this approach necessitates undersampling in the k-space beyond the Nyquist rate, which subsequently leads to artefacts in the reconstructed images. Fortunately, compressed sensing can help recovering the images with minimal artefacts. But it has some requirements and lots of factors influencing the resulting image quality, which are not as yet mature in the case of spiral sampling and structured MRI images. Therefore, in this study, we proposed a novel density function parametrisation for high-resolution spiral trajectories, well-adapted to the k-space data, to efficiently shorten the acquisition time.

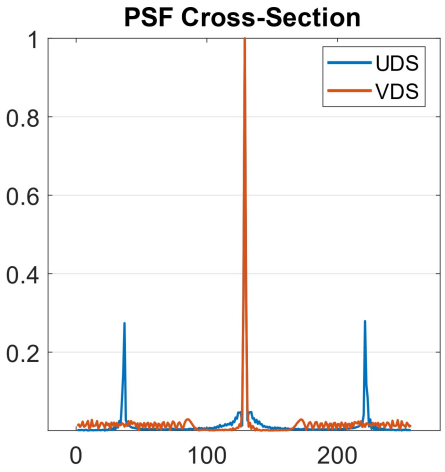


Figure 4.19: The PSF cross-section, comparing the undersampling performance of VDS and UDS trajectories.

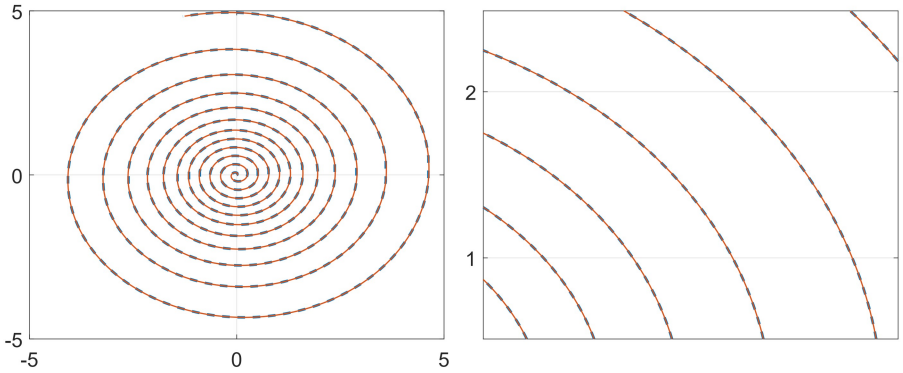


Figure 4.20: The resulting VDS trajectory: the blue line represents the designed trajectory and the dashed red line shows the practical time-optimal trajectory using the OptGrad toolbox.

An empirical study was performed to quantitatively determine the optimal parameters for the proposed density function using various trajectories. The experiments were focused on a wide range of real human heart DW images to ensure the capture of diffusion-related information. We investigated the individual and interactional effects of all the parameters involved in the density function—namely, the Nyquist sampling radius, central and ending densities, and the function governing the undersampling region—as well as spiral trajectory specifications, including readout duration and num-

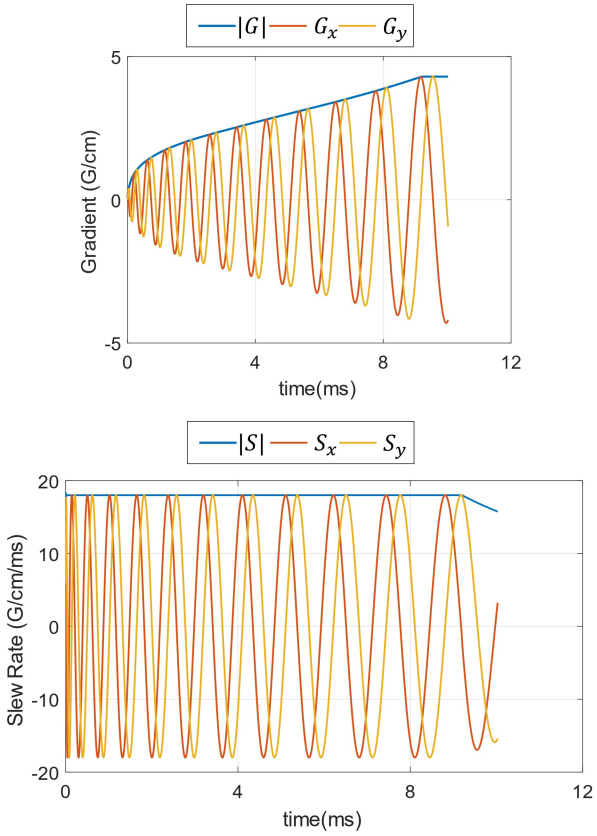


Figure 4.21: Gradients and slew rates waveforms in the resulting VDS trajectory.

ber of shots. Our analysis quantitatively confirmed the advantages of the proposed density function in efficient capturing of the required information and improving the CS reconstruction. The proposed sampling density function demonstrated consistent performance across a wide range of examined parameter values for a fixed trajectory. However, changes in the readout and number of shots could potentially compromise its performance, with the readout time having a greater influence.

Overall, the SSIM quality metric displayed greater variability compared to the intensity-based metrics, namely PSNR and NRMSE. In other words, the SSIM scores were more restrictive, whereas the NRMSE and PSNR scores did not exhibit significant differences across a broader range of parameters. Therefore, SSIM may serve as a more specific quality measure for assessing the performance of the methods during optimisation. Nevertheless, we also tracked NRMSE to ensure that the reconstruction error remained below a set threshold and used PSNR to monitor noise suppression effects,

which becomes particularly relevant when higher undersampling rates are applied.

The readout duration is constrained by off-resonance artefacts, as longer durations can cause pronounced blurring and may require extensive post-processing or additional acquisition time, which can be costly. Hence, minimising these artefacts upfront is preferable. Although the extent of these artefacts depends on field homogeneity, studies have shown that in cardiac imaging, $T_{ro} \leq 15$ ms does not necessitate correction [177]. In theory, maintaining the same sampling density for shorter readout durations can be achieved by increasing the number of shots. However, in practice, especially in diffusion MRI applications, eddy currents induced by diffusion gradients can lead to magnetisation variations between shots. This can cause phase inconsistencies between the data acquired during each spiral shot, resulting in artefacts when the data are combined to produce the final image. Therefore, minimising the number of shots is more desirable.

Based on the statistical analyses and considering groups where no significant differences were found in SSIM scores, we conclude that the values reported in Table 4.3 represent the optimal parameters for the proposed VDS trajectory in a multislice approach. (For $T_{ro} = 15$ ms, $N_s = 3$ also resulted in acceptable quality). Theses values consistently had the smallest point estimates for NRMSE, the highest SSIM and PSNR, and a larger fully sampled k-space centre suitable for calibrations that may be required in prospective acquisitions. Therefore, they were chosen for our methodology in the next chapter. Given the short T_2 relaxation time of the myocardium (≈ 54 ms), multislice approaches face the challenge of acquiring more than one slice per TR (where TR equals one heartbeat). To fulfil our primary objective of acquiring 6-8 slices in less than a minute, the condition $N_{slc} \times N_s \times TR \leq 60$ s must be met. This requires that $N_s \leq 8.5$ and the number of slices (N_{slc}) be around 7-12. Based on this, the parameter pairs of $T_{ro} = 10$ ms with $N_s = 5$ or $T_{ro} = 15$ ms with $N_s = 3$ can be selected as the optimal settings for multislice spiral imaging with the proposed density function parameters.

Table 4.3: Concluded optimal variables for the proposed VDS design

Density function parameters	Spiral trajectory specifications
$k_c/k_m = 20\%$, $D_C = 0.8$, $D_E = 0.2$	$T_{ro} \geq 10$ ms, $N_s \geq 5$

With the proposed density function, larger acceleration factors can indeed be em-

ployed, allowing for the recovery of high-resolution details. However, the enhanced performance at high resolution is practically limited by the available SNR in the acquisition. We identified a minimum SNR threshold necessary to achieve the expected image quality. In situations where the intrinsic SNR is limiting, it may be more advantageous to use a smaller acceleration factor on a lower resolution scan to achieve optimal results. Moreover, compressed sensing is expected to enhance the visually apparent SNR of the image beyond what is anticipated based on these theoretical considerations. Better reconstruction quality could be generally achieved by tuning the regularisation parameters more carefully for each individual experiment. The relationship between the encoding direction and the information distribution should also be considered when optimising the sampling density. As shown, in our analyses, for some density functions, variations between the quality of the reconstructed DW images were observed that can be attributed to the limitation of those density functions to capture the required information for that particular DW direction. We did not account for this in the present study, but this highlights the need for DW direction-specific sampling designs. For simplicity, we considered the two-step approach, however, once the desired density is obtained, the SPARKLING method can be used [182] to simultaneously optimise the k-space and the gradient waveforms corresponding to the target density. Here, we considered Nyquist sampling near the k-space centre with $D_I = 1$. Oversampling of the k-space centre can be achieved by setting $D_I > 1$. This approach has been explicitly used in non-Cartesian acquisitions for post-processing corrections. For instance, low-resolution artefact-free images can be computed using densely sampled areas near the k-space centre to correct for motion artefact. However, this oversampling comes at the expense of increasing the readout time, or reducing the spatial resolution for the same readout duration.

4.6 Summary

In this chapter, we introduced a new parametrisation for sampling density function in VDS trajectories. For a given MR protocol with defined imaging parameters (TR, TE, and readout duration), this method proved to be effective in improving sampling performance. The versatility of this approach was demonstrated across various spiral setups and diffusion-weighted images. With the proposed VDS sampling, the acquisition time for acquiring 10 slices in each DW image with 1 mm resolution and con-

sidering the same T_R of 700 ms potentially reduce from 14 min to 21 s. Our findings based on real cDTI images provide quantitative guidelines for selecting VDS trajectory parameters for high-resolution imaging within the compressed sensing framework. The potential of the proposed VDS design was showcased, and we anticipate that employing more advanced compressed sensing reconstruction techniques capable of fully utilising the correlations between DW images would further enhance its performance. These advancements will be explored and presented in the next chapter.

Chapter 5

Efficient High-resolution 3D Spiral-CS cDTI

5.1 Introduction

In the previous chapter, we demonstrated a generic method to define an optimal sampling density function for a tailored VDS trajectory. The proposed VDS method can be readily extended to design optimal 3D trajectories, thus enabling 3D imaging which favors higher SNR compared to 2D imaging. In addition to providing general advantages of no gaps or "cross-talk" between different slices and retrospective reformats at arbitrary planes, a 3D imaging approach helps eliminate misregistration artefacts when computing 3D tensor parameters, enhance the accuracy of the derived DTI metrics, and improve the quality of tractography results, particularly if diagnostic information is to be extracted [98]. It also gives more degrees of freedom in trading off SNR and scan time, which can be notably beneficial to cDTI applications.

In this chapter, we extend the proposed VDS design to three dimensions directly using stacks of 2D VDS trajectories, similar to the standard stack-of-spirals (SOS) approach, and a full 3D conical sampling pattern [183]. Full 3D trajectories have the potential to incorporate a true 3D variable density, presumably necessary for optimal compressed sensing in 3D. Additionally, they are not constrained by limitations in the undersampling factor along the partition (stacking) direction, which in practice rarely exceeds 4. We conduct retrospective implementation on a series of real ex vivo cDTI data in both clinical and preclinical settings and compare the results with extensions of previously used sampling schemes in CS cDTI studies.

5.2 Materials and Methods

5.2.1 3D Imaging

Stack-of-spirals sampling

The stack-of-spirals trajectory comprises spirals in the $k_x - k_y$ plane with phase encoding in the k_z direction. Due to the flexibility offered by both the spiral trajectory and phase encoding for arbitrary density selection, this trajectory is well-suited for variable density sampling. The rapid coverage of k-space in the $k_x - k_y$ plane using spirals makes it particularly suitable for fast imaging applications.

In the typical configuration of 3D stack of spirals (SOS), identical spiral trajectories are generally employed for each $k_x - k_y$ plane. In this scenario, the point spread

function (PSF) mirrors that of the simpler 2D case. The entire energy of the PSF is concentrated in the slice corresponding to $z = 0$, with no energy dispersion into other slices. Examples of the standard stack of uniform density spirals (UD-SOS) and variable density spirals (VD-SOS) are depicted in Figure 5.1(a,b). Here, the k_z steps are identical and constant, equal to the inverse of the nominal FOV in the z direction. The samples in the k_z direction lie on a uniform grid, therefore, these trajectories can be reconstructed with 2D gridding followed by a Fourier transform along the third direction of the k-space, thereby reducing the computation cost of CS reconstruction [184]. However, for a nominal FOV_z of 12 cm and spatial resolution of 1 mm, we need to acquire 120 slices, according to the Nyquist criterion. One approach to reduce the scan time of the standard SOS trajectories is to reduce the number of interleaves in the outer k-space planes by reducing the target resolution or maximum k radius, and therefore, to cover a spherical k-space instead of a cylinder. An example of such trajectories, referred to as spherical stack-of-spirals (S-SOS) is shown in Figure 5.1(c).

Reducing the required number of slices would further reduce the scan time. We can achieve this by utilising the variable density also along the k_z direction. To best approximate radial symmetry and a true variable density across all directions, we employed spirals that are unevenly spaced in k_z and exhibit varying density within each $k_x - k_y$ plane. The phase-encoding locations in the k_z direction are established by specifying the spacing between phase encodes as $1/FOV_z(k_z)$, where $FOV_z(k_z)$ is the effective FOV defined by:

$$FOV_z(k_z) = FOV_{zn}\ \rho(k_r, k_z) \tag{5.1}$$

Here, FOV_{zn} is the nominal FOV in the z direction. We used the same parameters for the sampling function in the k_z direction as the optimal values defined in the k_r direction. By utilising the k_z dependence in the sampling density function $\rho(k_r, k_z)$, the spirals become sparser for phase-encode locations that are farther away from the k-space origin. We propose to modify the density function in Equation 4.5 as follows [184]:

$$\rho(k_r, k_z) : k_c^j = k_c(k_{zj}) = k_c^0\ \exp\Big(-\frac{k_{zj}^2}{k_{zc}^2}\Big) \tag{5.2}$$

where k_c^0 is the fully sampled center radius in the central slice, k_c^j is the corresponding radius for the jth slice, and k_{zc} is a parameter that regulates the amount of under-

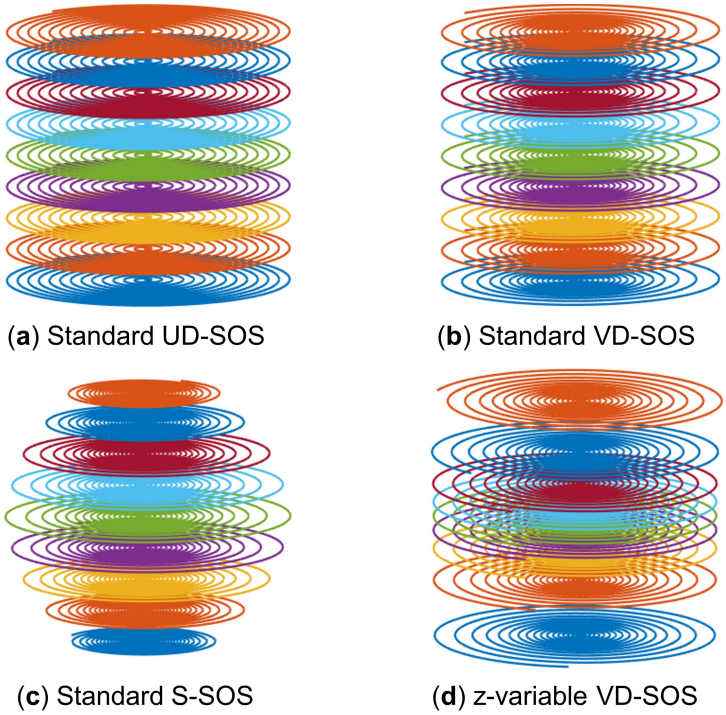


Figure 5.1: Different Extensions of 2D spiral trajectories to 3D. Nine consecutive slices are shown with the same colour scheme. UD: uniform density; VD: variable density; S: spherical; SOS: stack of spirals.

sampling.

With this design, the readout time for outer planes in k-space becomes shorter than for the centre, if keeping the target resolution and number of shots in each plane constant. However, it is advisable to maintain similar readout times for all planes to avoid different off-resonance effects along the k_z direction. Furthermore, the scan time is more related to the total number of the interleaves than to the total readout time. Therefore, we reduced the number of interleaves in the outer planes to achieve approximately equal readout times per interleaf, resulting in further reduction of the scan time. After selecting the number of shots in the central stack N_s^0, the mass of each plane can be tailored to the plane density by reducing the number of shots as:

$$N_s^j = N_s^0 \frac{\int \rho(k_r, k_z)}{\int \rho(k_r, 0)} \tag{5.3}$$

where N_s^j represents the number of shots in the jth slice of altitude k_z. Figure 5.1(d) illustrates such a stack, we hereafter refer to as the z-variable VD-SOS. Figures 5.2 and

5.3 respectively illustrate the k_z and k_{xy} views of these sampling patterns.

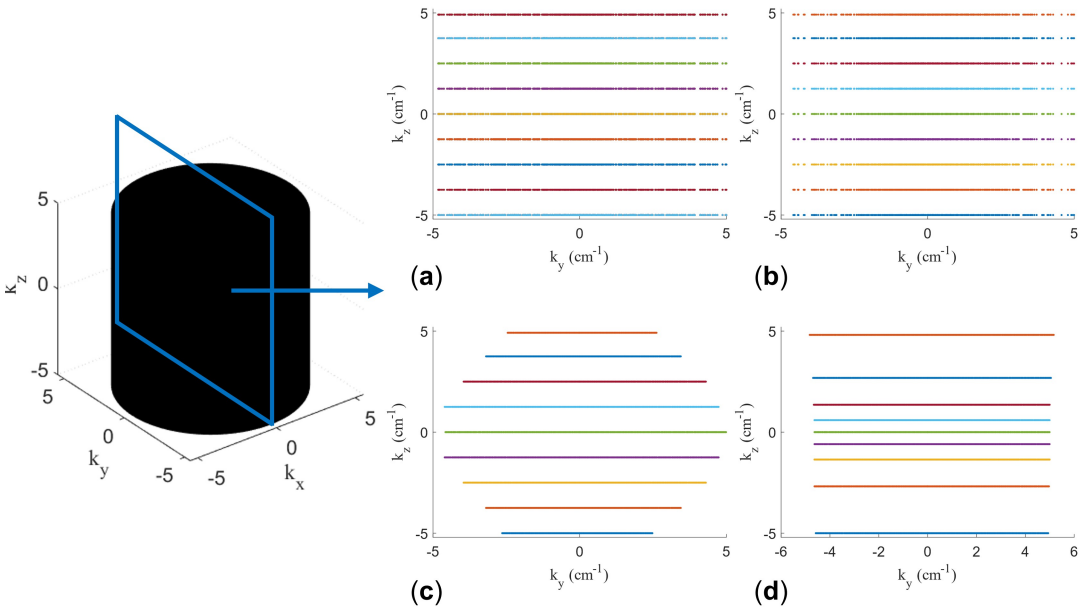


Figure 5.2: Sampling patterns in the $k_y - k_z$ plane corresponding to the trajectories depicted in Figure 5.1.

5.2.2 Fully 3D VDS: VD-Cones Sampling

The 3D cones trajectory combines elements of radial and 2D spiral trajectories to extend the advantages of time-efficient, motion-robust non-Cartesian trajectories into three dimensions. This trajectory comprises conical surfaces onto which interleaved spiral trajectories are projected. Key design parameters for the 3D cones trajectory include the transverse and longitudinal field of views (i.e., FOV_{xy} and FOV_z) and spatial resolutions (RES_{xy} and RES_z). The initial step in designing the 3D cones trajectory involves selecting a set of elevation angles for concentric conic surfaces:

$$\phi^j = \sin^{-1}\left(\frac{k_z^j}{k_m}\right) \tag{5.4}$$

where $k_m = 5/RES_z$, is the maximum spatial frequency in the k_z direction, and k_z^j is the latitude corresponding to the jth transverse plane. The maximum spatial frequency for the projected 2D spiral is then computed as $k_r = k_m \times cos(\phi^j)$, and the resulting VD spiral is designed. Again, we used the same optimal sampling density function in

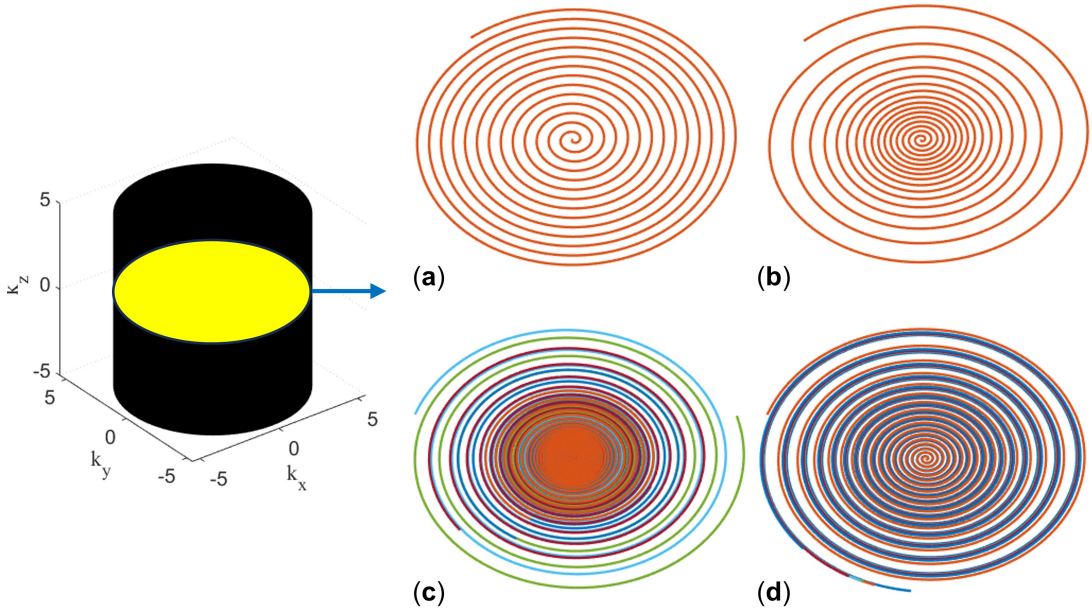


Figure 5.3: Sampling patters in the $k_x - k_y$ plane corresponding to the trajectories depicted in Figure 5.1.

both the k_r and k_z directions. The gradients were then optimised using the OptGradV2 toolbox, and the resulting optimised trajectory was used for the subsequent evaluations. Figure 5.4 compares a standard uniform cones trajectory and the proposed z-variable VD cones trajectory.

5.2.3 $k - q$ Ordering Scheme

The multi-shot trajectories are formed by rotating the designed spiral/cone interleaf about the k_z axis for each $k_x - k_y$ plane/conic surface. In CS-cDTI, the $k - q$ sampling pattern is another crucial consideration, as we want to maximise the incoherence between DW images. We evaluated the golden angle ordering scheme $GA - q$. The spiral interleaves are distributed uniformly within a DW image but are rotated by a golden angle ($\Delta\theta = (\sqrt{5} - 1)/2 \times 180° \approx 111.25°$) [185] between DWIs, as shown in Figure 5.5.

The golden-angle ordering scheme ensures that successive sampling patterns are distinct and non-redundant. With each new acquisition in q-space, the sampling trajectory is rotated by an irrational fraction of a full circle (approximately 111.25°). This approach prevents overlapping of k-space samples across DW directions, resulting in

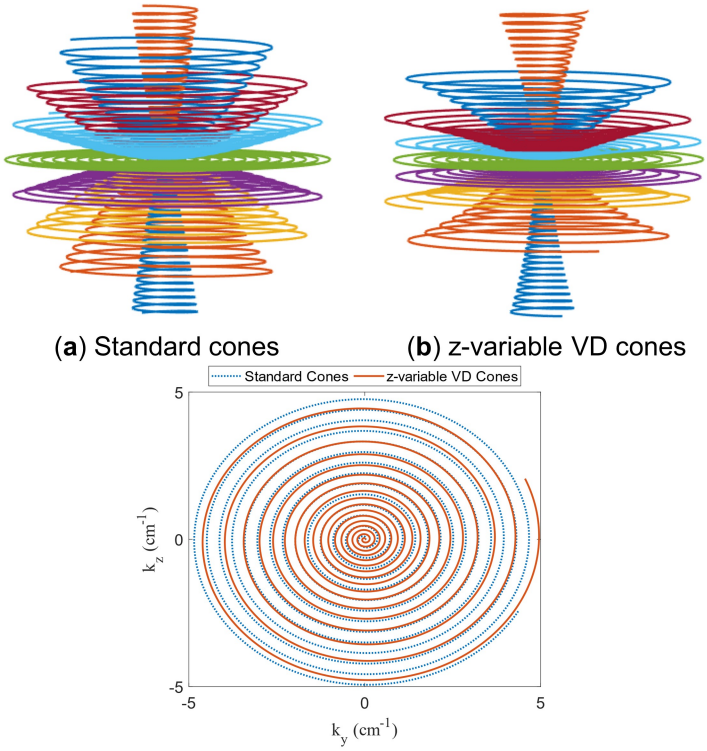


Figure 5.4: 3D Cones sampling.

pseudo-random coverage over multiple acquisitions and enhancing incoherence, which is particularly advantageous for the proposed CS reconstruction.

To distribute the energy across slices/conic surfaces, we introduced a "random rotation" ($d\theta$) to the entire spiral trajectory within each $k_x - k_y$ slice encoding plane/cone. The same rotation was applied consistently to all interleaves within the same slice encoding plane/cone encoding surface. This approach allowed for through-plane randomisation without altering the in-plane sampling pattern. This way, uniform and incoherent sampling of the $k - q$ space is obtained, contributing to a more accurate CS reconstruction.

5.2.4 Image Reconstruction with Phase Error Compensation

As mentioned in Section 3.4.1, group sparse modelling can lead to a lower sampling requirement, or alternatively with the same number of measurements, to a more accurate reconstruction [60]. Here, we propose to incorporate the group sparsity among DWIs

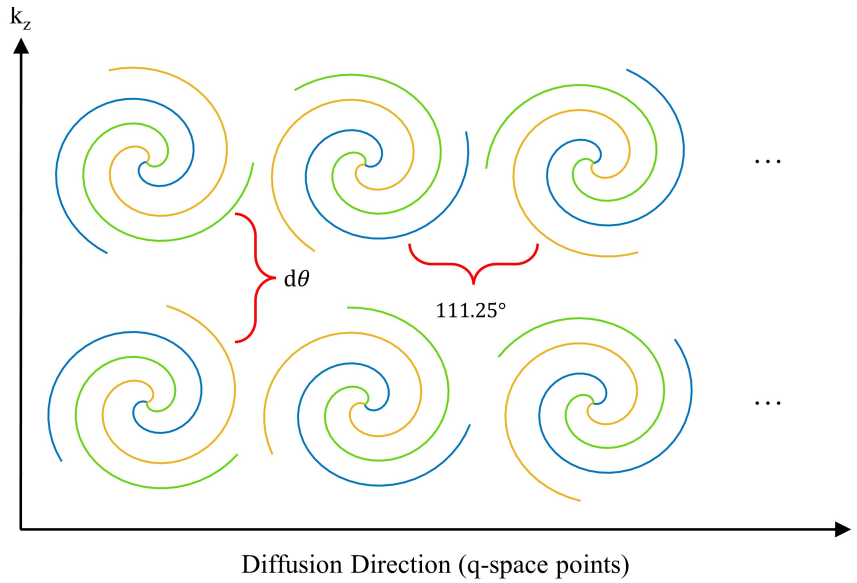


Figure 5.5: Proposed incoherent $k-q$ undersampling pattern: $d\theta$ represents the small angle rotations between slices and 111.25° corresponds to the golden angle rotations between diffusion-weighted directions.

in both the wavelet and TV domains into a unified CS reconstruction formulation to enhance the reconstruction performance and simultaneously recover all images.

Let $\mathbf{X} = [x_1, x_2, \cdots, x_N] \in \mathbb{C}^{M\times N}$ represent a matrix formed by stacking all 3D images in a DWI series as column vectors, $\mathbf{Y} = [y_1, y_2, \cdots, y_N] \in \mathbb{C}^{N_k\times N}$ be the respective undersampled k-space data, and $\mathbf{E} \in \mathbb{C}^{N_k N\times MN}$ denote the 3D spatial encoding operator, the bidomain joint sparsity (BJS) reconstruction is formulated as:

$$\hat{\mathbf{X}} = \underset{\mathbf{X}}{\operatorname{argmin}} \left\{ \frac{1}{2} \sum_{q=1}^{N} \left\| \mathbf{Y}(:,q) - \mathbf{E}_q \mathbf{X}(:,q) \right\|_2^2 + \lambda_1 \|\mathbf{\Psi X}\|_{1,2} + \lambda_2 JTV(\mathbf{X}) \right\} \tag{5.5}$$

where $M = n_x \times n_y \times n_z$ represents the total number of voxels, N is the number of DW images, $\mathbf{\Psi}$ is a 3D wavelet transform, and JTV refers to the joint total variation constraint.

The $\ell_{1,2}$-norm constraint characterises the joint wavelet sparsity (JWS) regularisation, which has been previously used to accelerate cDTI [60]. This constraint is equivalent to applying the ℓ_2-norm on each row, considering the sparse coefficients at each voxel along the diffusion encoding dimension as one group, and then performing the ℓ_1-norm on the resulting vector to promote spatial sparsity. This can be explicitly

expressed as:

$$\|\mathbf{\Psi X}\|_{2,1} = \sum_{\nu=1}^{M} \sqrt{\sum_{q=1}^{N} \left([\mathbf{\Psi X}]_{\nu q}\right)^2} \tag{5.6}$$

Rooted in the observation that the variation of image gradients should be similar for the same spatial position across multiple DW directions; the joint TV (JTV) regularisation can be defined as either the isotropic or anisotropic multidimensional (colour) TV semi-norms [186]:

$$\begin{aligned} JTV_a(\mathbf{X}) &= \sum_{\nu=1}^{M} \sum_{q=1}^{N} \Big(|[\nabla_1 \mathbf{X}]_{\nu q}| + |[\nabla_2 \mathbf{X}]_{\nu q}| + |[\nabla_3 \mathbf{X}]_{\nu q}| \Big) \\ JTV_{iso}(\mathbf{X}) &= \sum_{\nu=1}^{M} \sqrt{\sum_{q=1}^{N} \Big(([\nabla_1 \mathbf{X}]_{iq})^2 + ([\nabla_2 \mathbf{X}]_{iq})^2 + ([\nabla_3 \mathbf{X}]_{iq})^2 \Big)} \end{aligned} \tag{5.7}$$

where ∇_1, ∇_2, ∇_3 respectively denote the forward finite-difference operators along the three spatial coordinates (i, j, k) centered in the pixel ν. It is worth mentioning that these two constraints provide complementary prior information, thus, further improving the reconstruction accuracy.

In practical scenarios, a major challenge in reconstruction of DTI data, particularly those obtained from multishot acquisitions, is the uncorrelated phase variations between different shots/interleaves and across different diffusion encoding directions, arising from eddy currents and motion. These phase inconsistencies could negatively impact the effectiveness of group sparsity in the BJS reconstruction, leading to signal cancellation and image degradation in the reconstructed DWIs and subsequent errors in the DTI metrics if not accounted for. The phase errors corresponding to η-th shot and q-th diffusion-encoding direction can be modelled as a unit-magnitude phase map:

$$p_{\eta,q}(\mathbf{r}) = e^{j\phi_{\eta,q}(\mathbf{r})} \tag{5.8}$$

Previous multi-shot spiral DTI techniques employed either a variable-density spiral trajectory with oversampling in the central k-space [171, 187] or a single-shot spiral navigator echo [177]. These approaches aimed to create a low-resolution estimation of motion-induced phase errors for each shot and subsequently correct for such artefacts. However, these methods come with a drawback of a longer scan time, for instance, up to twice as long for a variable-density spiral acquisition with a typical oversampling factor of 4. This extended scan time may pose challenges in achieving both high

spatial and angular resolutions, especially in clinical applications. In this work, we adapted an iterative phase correction method [188] to the BJS implementation that can inherently correct for phase errors without requiring additional data acquisition or increasing the scan time. Specifically, a *phase* image from the k-space data of each shot was reconstructed. As motion-induced phase errors typically exhibit slow spatial variations, only the central k-space data of each shot corresponding to the fully sampled k_c radius were used and zero-padded to the full spatial frequency range to construct a low-resolution (LR) image: $\hat{x}_\eta^{LR} = \mathcal{NF}^{-1}(y_\eta)$. The phase errors were then estimated as $\phi_\eta(\mathbf{r}) = -\measuredangle\, \hat{x}_\eta^{LR}$, where $\measuredangle$ represents the angle (or phase) of the complex-valued image $\hat{x}_\eta^{LR}$. Additional steps were implemented to enhance phase estimation, especially for acquisitions with a high number of shots or a high spatial resolution or b-value (resulting in a low SNR). These steps include phase unwrapping and smoothing with a 9×9 median filter.

Once phase errors are known, these phase maps can be expressed as an additive term into the encoding operator [187]. To derive an expression for $\mathbf{E}_q$, we first define $\mathbf{F}$ as the Fourier encoding matrix and $\mathbf{P}$ as the phase error in the matrix form:

$$\mathbf{F} = \begin{bmatrix} e^{-j\mathbf{k}_{1,1}\cdot\mathbf{r}_1} & e^{-j\mathbf{k}_{1,1}\cdot\mathbf{r}_2} & \cdots & e^{-j\mathbf{k}_{1,1}\cdot\mathbf{r}_{N^2}} \\ e^{-j\mathbf{k}_{2,1}\cdot\mathbf{r}_1} & e^{-j\mathbf{k}_{2,1}\cdot\mathbf{r}_2} & \cdots & e^{-j\mathbf{k}_{2,1}\cdot\mathbf{r}_{N^2}} \\ \vdots & \vdots & e^{-j\mathbf{k}_{k,\eta}\cdot\mathbf{r}_\rho} & \vdots \\ e^{-j\mathbf{k}_{n_k,n_\eta}\cdot\mathbf{r}_1} & e^{-j\mathbf{k}_{n_k,n_\eta}\cdot\mathbf{r}_2} & \cdots & e^{-j\mathbf{k}_{n_k,n_\eta}\cdot\mathbf{r}_{N^2}} \end{bmatrix} \tag{5.9}$$

$$\mathbf{P} = \begin{bmatrix} e^{j\phi_1(\mathbf{r}_1)} & 0 & \cdots & 0 \\ 0 & e^{j\phi_1(\mathbf{r}_2)} & \cdots & 0 \\ \vdots & \vdots & e^{j\phi_\eta(\mathbf{r}_\rho)} & \vdots \\ 0 & 0 & \cdots & e^{j\phi_{n_\eta}(\mathbf{r}_{N^2})} \end{bmatrix} \tag{5.10}$$

Here, $\mathbf{k}_{k,\eta}$ represents the k-th sampling point of the η-th shot, and $\mathbf{r}_\rho$ is the ρ-th pixel of a DW image. Both $\mathbf{r}$ and $\mathbf{k}$ are expressed in vector form to emphasize their representation in 3D image space and k-space, respectively. $\phi_\eta(\mathbf{r}_\rho)$ is the phase error at position $\mathbf{r}_\rho$ during the η-th shot, and it can be linear or nonlinear. With these definitions, the encoding matrix $\mathbf{E}$ can be expressed as:

$$\mathbf{E} = \mathbf{\Omega G F P} = \mathbf{\Omega} \mathcal{N} \mathcal{F} \mathbf{P} \in \mathbb{C}^{(n_k n_\eta)\times M} \tag{5.11}$$

where $\mathbf{\Omega}$ denotes the sampling mask matrix for different DW directions, containing ones for the interested frequency locations and zeros for those needed to be skipped,

G denotes the gridding and sampling density correction procedure necessary to handle the non-Cartesian k-space samples; we define the **GF** as the 3D non-uniform Fourier transform (NUFT) operator $\mathcal{NF}$, and **P** is a stack of diagonal matrices corresponding to DWIs with entries equal to the estimated phase error at each voxel and for each shot. n_k is the number of k-space sampling points per shot, n_η is the number of spiral interleaves/shots, and M is the number of voxels in the image domain.

5.2.5 Optimisation Algorithm

The BJS problem is a large-scale problem, and therefore, it is crucial to develop fast algorithms to efficiently solve it. Here, we adapted a fast composite splitting algorithm (FCSA) [137] and extended the implementation to comply with non-uniform complex data and to three dimensions. Briefly, the non-smooth convex problem (Equation 5.5) was decomposed into two simpler regularisation subproblems by: (a) splitting variable $\mathbf{X}$ into two variables $\{\mathbf{X}_i\}_{i=1,2}$; (b) performing operator splitting to independently minimise the JWS subproblem over $\mathbf{X}_1$ and JTV subproblem over $\mathbf{X}_2$, and (c) linearly combining the resulting optimal solutions. Let $f(X) = \frac{1}{2}\sum_{q=1}^{N} \|\mathbf{Y}_q - \mathbf{E}_q\mathbf{X}_q\|_2^2$ be a convex and smooth function with the gradient ∇f and Lipschitz constant L_f, and $g_1(\mathbf{X}) = \lambda_1\|\mathbf{\Psi X}\|_{1,2}$ be a convex but non-smooth function, and $g_2(\mathbf{X}) = \lambda_2 JTV(\mathbf{X})$. Algorithm 2 outlines the pseudo-code for solving the BJS problem. In brief, the solution begins by reducing the value of f through a step in the negative gradient direction. Then, the value of each subproblem is reduced by applying a proximal mapping operation [134], and the process is repeated. The proximal map (or proximal operator) is widely used in optimisation algorithms, especially when dealing with non-differentiable functions. The proximal operator of a continuous convex function $g(x)$ (with a given scalar $t > 0$) is defined as:

$$prox_t(g)(u) \triangleq \arg\min_x \left(g(x) + \frac{1}{2t}\|x - u\|_2^2\right) \tag{5.12}$$

The solution of the proximal map can be interpreted as finding a point x that balances minimising $g(x)$ with staying close to the input point u. The parameter t serves as a step size that controls the trade-off between the two objectives by penalising larger deviations from u.

The first subproblem has a closed-form solution by soft thresholding:

$$(\mathbf{X}_2)_i = \mathbf{\Psi}^T\left(\max\left(1 - \frac{2\lambda_1/L_f}{\|(\mathbf{\Psi X})_i\|_2}, 0\right)(\mathbf{\Psi X})_i\right) \tag{5.13}$$

Algorithm 2 BJS Algorithm

Input: $L_f, \lambda_1, \lambda_2 > 0,\ t^1 = 1,\ \mathbf{X}^0 = \mathbf{Z}^1 = 0,\ k = 0$

for $k = 1$ **to** K **do**

$\mathbf{U} = \mathbf{Z}^k - \frac{1}{L_f}\nabla f(\mathbf{Z}^k)$

$\mathbf{X}_1 = prox_{1/L_f}(g_1)(\mathbf{U})$

$\mathbf{X}_2 = prox_{1/L_f}(g_2)(\mathbf{U})$

$\mathbf{Q}^k = \frac{\mathbf{X}_1+\mathbf{X}_2}{2}$

$\mathbf{X}^k = \arg min\{C(\mathbf{X}) : \mathbf{X} = \mathbf{Q}^k, \mathbf{X}^{k-1}\}$

$t^{k+1} = \frac{1+\sqrt{1+4(t^k)^2}}{2}$

$\mathbf{Z}^{k+1} = \mathbf{X}^k + \frac{t^k}{t^{k+1}}(\mathbf{Q}^k - \mathbf{X}^k) + \frac{t^k-1}{t^{k+1}}(\mathbf{X}^k - \mathbf{X}^{k-1})$

Stop: $\frac{|C^k - C^{k-1}|}{|C^k|} < 10^{-5}$

end for

where $(.)_i$ denotes the i-th row of the matrix. There is no closed-form solution for the second subproblem; thus, we approximated it using a modified fast joint-gradient projection (FJGP) method [189] that is extended to handle 3D complex data. The final reconstructed image matrix was obtained in an iterative framework, accelerated by a FISTA scheme [134]. FISTA provides fast convergence performance suitable to tackle large-scale optimisation problems. Because calculation of the Lipschitz constant for this large-scale problem was not straightforward, we used the monotone version of the FISTA to ensure a decreasing cost function. The iterations terminated when the relative change in the cost function C was less than an error tolerance of 10^{-5}.

5.2.6 Parameter Setting

The selection of regularisation parameters $\{\lambda_i\}_{i=1}^2$ is particularly important in the thresholding and gradient projection steps to provide a good balance between the denoising purpose and maintaining fine image features. We empirically found that increasing λ_1 and λ_2 would respectively result in pronounced checkerboard-like artefacts and oversmoothing of the images, whereas too small values led to insufficient suppression of the undersampling artefacts. We determined the optimal values based on an analysis of the relative error compared to a fully sampled reference data, and

the values reported in Table 5.1 were fixed for all datasets. In the absence of fully sampled ground truth, the regularisation parameters are typically determined using the L-curve strategy. However, finding the appropriate (optimal) regularisation parameters for high-dimensional problems is a tedious task. We, therefore, introduced an adaptive noise variance estimation step utilising the robust median estimator [190] to adjust the threshold values in both subproblems at every iteration based on the current image update as:

$$\hat{\lambda}_i^k = \frac{Median(|\mathbf{U}(:)|)}{0.6745}/N \tag{5.14}$$

This equation is used to obtain a robust yet adaptive estimate of the noise level in the data, setting an appropriate threshold for the denoising of wavelet or total variation coefficients[1]. We refer to this implementation as *adaptive* BJS. This way, the dependence of the algorithm on the choice of hyperparameters was effectively eliminated. This is especially helpful in large-scale and high-dimensional optimisation problems and in practice where the reference data are not known. Improved performance could be achieved by performing adaptive thresholding for each sub-band separately in the wavelet domain.

Table 5.1: Regularisation parameters at different undersampling rates

	λ_1	λ_2
$\mathrm{N_{int}} = 32, 16$	0.005	0.0006
$\mathrm{N_{int}} = 11, 8$	0.012	0.0015
$\mathrm{N_{int}} = 6, 5$	0.034	0.0023

5.3 Experiment Design

The performance of the proposed approach was evaluated using three ex vivo cDTI datasets. Due to the unavailability of the spiral pulse sequence, acquisitions were fully-sampled on a Cartesian grid, consisting n_z slices and n_y lines collecting n_x Fourier

[1] The constant $\frac{1}{0.6745}$ is used to convert the median absolute deviation into an estimate of the standard deviation σ in a Gaussian distribution. The scale of the noise estimate is adjusted by normalising with the total number of DW images N.

samples without oversampling for an image size of $M = n_x \times n_y \times n_z$. The non-uniform k-space data on a desired trajectory were then retrospectively generated by applying forward 3D NUFT [179] on the complex-valued fully sampled Cartesian images. All animal studies adhered to the UK Home Office for animal care and the NIH Guide for the Care and Use of Laboratory Animals.

5.3.1 Data Acquisition

The first dataset was acquired from an isolated rat heart using a standard pulsed gradient spin-echo sequence with a conventional 3D Cartesian imaging. The experiment was conducted on a 7T preclinical scanner (BioSpec 70/20 USR, Bruker Biospin, Ettlingen, Germany) with a microscopy gradient system (Resonance Research Inc., Billerica, US, maximum gradient strength = 1.5 T/m, rise time = 100 μs). Acquisition parameters were as follows: FOV = $17.6 \times 17.6 \times 21.6$ mm^3, isotropic resolution = 100 μm^3, TE/TR = 13.37/250 ms, 3 non-weighted (b_{ref}) images, and 27 DW images with a b-value of 1000 s/mm^2. The total scan time was 64.5 h. This dataset was used to tailor the sampling pattern properties, determine the optimal regularisation parameters, and compare various methods.

The second dataset consisted of five excised rat hearts acquired in a previous study [158] with a 3D DW fast spin-echo sequence. Briefly, imaging was performed on a 9.4T preclinical MRI scanner (Agilent, Santa Clara, CA, USA) with a shielded gradient system (max gradient strength = 1 T/m, rise time = 130 μs), and transmit/receive birdcage coil (inner diameter = 20mm; Rapid Biomedical, Rimpar, Germany). Acquisition parameters were: TE/TR = 9.3/250 ms, echo spacing = 4.9 ms, echo train length = 8, FOV = $16 \times 16 \times 20$ mm^3, resolution = 100 μm^3 isotropic, number of non-DW images = 8, number of DW directions = 61, b = 1000 s/mm^2, and total scan time = 15 h 20 min. A pixel-wise T_2 map was estimated using additional 3D multiecho spin echo data. DTI data were then corrected for the T_2 weighting prior to further processing. Only one b_0 and the first 12 DW images were used in our evaluations. This dataset was used to demonstrate the generalisability and robustness of the proposed sampling and reconstruction framework.

The third dataset corresponded to an ex vivo normal human heart [40, 48] acquired with a 3D fast spin echo (FSE) sequence using a four-element phased array coil on a 1.5 T GE CV/I MRI scanner (GE Medical System, Wausheka, WI, USA). The

following imaging parameters were as follows: image size = $256 \times 256 \times 134$, resolution = $0.43 \times 0.43 \times 1.0$ mm^3, and 21 diffusion gradient directions (2 non-DW and 19 DW images at $b = 1500$ s/mm^2). We used one b_{ref} image and the first 12 DW images of this dataset to evaluate the applicability of the proposed method to clinical setting and the effectiveness of the phase correction procedure in the reconstruction.

5.3.2 Sampling Schemes Comparison

We compared the proposed z-variable VD SOS trajectory to the 3D extensions of previous sampling schemes used in CS cDTI studies:

(i). Standard stack of SNAILS (self-navigated interleaved spirals) [171, 177] with the oversampling ratio $\alpha = 4$ and the same readout duration and gradient specifications as our VDS trajectory (Figure 5.6(b)). Randomised angular offsets were added to all uniformly-space interleaves for each DW direction.

(ii). Stack of 1D-Random sampling along the phase encoding direction (k_y) with a Gaussian density distribution [127] (Figure 5.6(c)). For each slice and each diffusion direction, different random patterns were generated with additional four lines at the k-space centre for every diffusion direction.

(iii). Poisson disk strategy with a Gaussian density distribution applied along two phase-encoding directions [63] (Figure 5.6(d)). Different random patterns were used for each q-space point.

We define the undersampling factor R as the ratio of the acquired phase encoding lines in the fully sampled Cartesian data to the maximum number of examined readouts (for the spiral/cones trajectory, examined readouts = number of interleaves in each $k_x - k_y$ plane/conic surface $\times$ number of k_z steps). To ensure a fair comparison across strategies, all trajectories were designed to support the same isotropic spatial resolution.

We also compared the performance of both z-variable VD-SOS and cones trajectories with their standard uniform-density counterparts, adequately sampled, serving as a baseline for assessing the image quality while accounting for any interpolation errors in the NUFT operator during the spiral data generation. The duration of each spiral arm was kept constant to allow for similar off-resonance performance for the different trajectories.

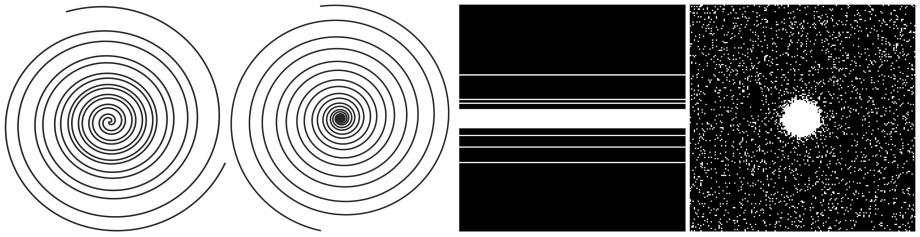


Figure 5.6: Illustration of $k-q$ sampling schemes used for comparison. (a) Proposed z-variable VD, (b) SNAILS, (c) 1D-RAND, (d) Poisson disk (2D-RAND). For simplicity, trajectories corresponding to one representative k_z plane and q-point are shown.

5.3.3 Image Reconstruction

We performed retrospective undersampling of the DTI datasets with varying acceleration factors (R). All DW images were simultaneously reconstructed using the BJS method (Equation 5.5). The non-diffusion weighted measurements (b_{ref}) were only mildly undersampled with a factor of two ($R = 2$) and reconstructed separately with a simple gridding method. The sparse domain $\mathbf{\Psi}$ was chosen as the 3D $Daubechies - 4$ wavelet transform with 4 level decomposition, and the anisotropic joint total variation was used. The algorithm generally took 10-20 iterations to converge. To effectively remove the burden of computing many NUFTs per iteration, i.e., two for each DWI, and minimise memory usage, we pre-calculated the Fourier exponentials corresponding to the frequency locations on a dense trajectory with variable density compensation weights using Beatty's Kaiser-Bessel kernel method [78] and stored them as a look-up matrix.

We also demonstrated the advantage of using the bidomain joint constraints over using only wavelet group sparsity constraint (JWS, setting $\lambda_2 = 0$ in Equation 5.5). The reconstruction performance of the BJS method with optimal regularisation parameters was evaluated against its *adaptive* version. For a given undersampling factor, identical masks were used in different reconstructions.

The effect of noise on the performance of the proposed approach was studied by adding zero-mean complex Gaussian noise to the raw k-space data with varying standard deviations σ_n to reduce the SNR, from around 35 to 15 (approximating the SNR of in vivo cases). The robustness of the proposed method against decreasing number

of diffusion directions was also investigated by performing uniform subsampling in the q-space.

To demonstrate the effectiveness of the phase compensation strategy in practical scenarios, we simulated the phase errors in the range comparable to those measured experimentally and calculated for the zth slice as:

$$\phi_z(x, y) = \phi_0 + ax + by + cx^2 + dy^2 + exy \tag{5.15}$$

with $|\phi_0| \leq \pi$, $|a|, |b| \leq \pi/l$, $|c|, |d|, |e| \leq \pi/(2l^2)$, and $l \times l$ is the in-plane image matrix size. We evaluated the performance of the phase-compensated reconstruction with both linear and nonlinear phase errors.

First, we reconstruct complex images from k-space data of individual interleaves. Then, a random spatially nonlinear phase (Equation 5.15) was added to each image prior to transforming back to the k-space. The resulting corrupted data were then reconstructed with the proposed phase-compensated BJS method.

5.3.4 Quantitative Analysis

The magnitude of the reconstructed images were fitted to a mono-exponential model on a pixel by pixel basis using an in-house weighted linear least-squares algorithm [81]. After computing the 3D diffusion tensor, conventional metrics, viz. mean diffusivity (MD), fractional anisotropy (FA), and helix angle (HA) maps were calculated and compared between the reference and the reconstructed images. In all comparisons, the 3D myocardium was chosen as the ROI, and the corresponding DWIs and DTI maps from the fully sampled Cartesian datasets were served as the reference, unless otherwise stated. Image quality metrics, including the NRMSE (%), peak SNR (PSNR), and the structural similarity index (SSIM) [180] were calculated, and the mean values (averaged over all DWIs) were reported.

For the diffusion parameter maps, we compared the RMSE and the normalised bias. We calculated the global MD and global FA (i.e., the average value in the ROI) to measure the normalised bias as: $\beta = |(GP_{rec} - GP_{ref})/GP_{ref}|$, with GP_{ref} and GP_{rec} respectively denoting the global parameter map of the reference dataset and the reconstructed dataset.

5.4 Results

Comparison of Sampling Schemes

Figure 5.7 shows a representative DW image reconstructed from the first dataset with the BJS method using different sampling schemes at $R \simeq 9$. The absolute difference maps with respect to the reference image are also included. The two spiral sampling schemes outperformed their Cartesian counterparts, yielding incoherent artefacts, which could be effectively removed during the reconstruction. The proposed z-variable VD-SOS showed improved reconstruction quality compared with the stack of SNAILS sampling. Cartesian 1D-Random sampling resulted in the most degraded image quality with visible aliasing artefacts. A similar performance is observed in the resulting FA and MD maps, as shown in Figure 5.8. Table 5.2 summarises the quantitative results.

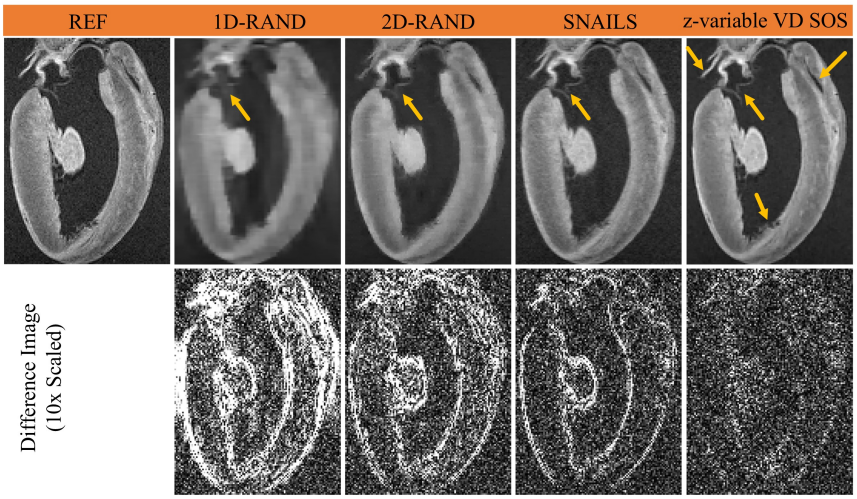


Figure 5.7: Long-axis mid-ventricular images of one representative diffusion direction and the corresponding absolute difference maps ($|\mathbf{X}_{rec} - \mathbf{X}_{ref}| \times 10$) obtained from the first dataset with the BJS method using various undersampling schemes. The arrows point out the difference in the recovery of fine image details.

Figure 5.9 and Table 5.3 compare the proposed z-variable VD SOS and the SNAILS sampling at a higher undersampling rate ($R = 22$). The proposed scheme shows significant improvement over the SNAILS. The oversampling of the centre and coarser

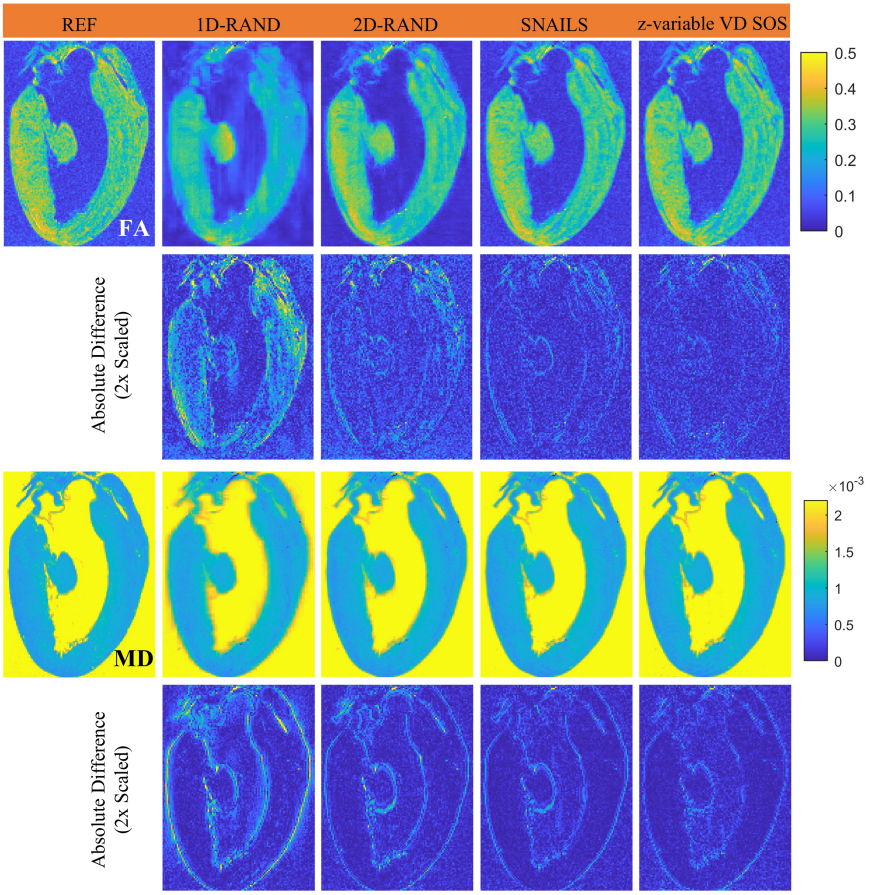


Figure 5.8: FA and MD (mm^2/s) maps reconstructed from the first dataset with the BJS method using various undersampling schemes. The absolute difference maps are also included for better evaluation of the sampling performance ($|\mathbf{I}_{rec} - \mathbf{I}_{ref}| \times 2$).

Table 5.2: Quantitative comparison of undersampling schemes

	SSIM	PSNR (dB)	DW_{RMSE}	FA_{RMSE}	MD_{RMSE} (mm^2/s)
1D-RAND	0.858	23.2	0.056	0.081	15.0×10^{-5}
2D-RAND	0.966	25.8	0.044	0.060	11.5×10^{-5}
SNAILS	0.971	28.2	0.033	0.042	6.8×10^{-5}
z-variable VD-SOS	**0.987**	**31.5**	**0.024**	**0.038**	$\mathbf{6.1}\times10^{-5}$

sampling of the outer k-space parts in the SNAILS scheme yields more severe streak artefacts and pronounced blurring in the reconstructed images with increased undersampling, while the proposed scheme offers a better trade-off between sampling the low and high spatial frequencies, allowing for better restoration of image details and reduced artefacts. The BJS reconstruction yields a careful balance between denoising and preserving edges. The pre-calculated NUFT coefficients significantly reduced the BJS algorithm runtime from several hours ($\sim$ 8 h) to about 93.7 min to recover all 27 DWIs

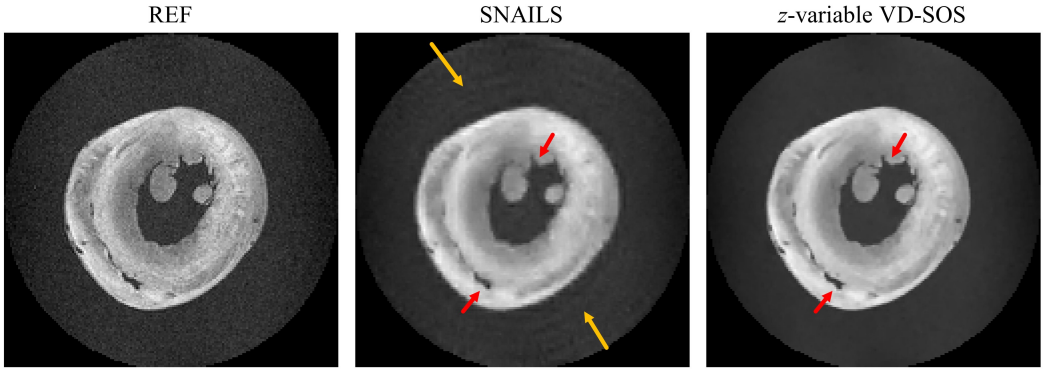


Figure 5.9: Performance comparison of the SNAILS and the proposed z-variable VD-SOS schemes at higher undersampling rates using BJS reconstructions. Arrows point to the residual undersampling artefacts (yellow arrows) and blurred image edges (red arrows).

Table 5.3: Performance comparison at a high undersampling factor

	PSNR (dB)	DW_{RMSE}	FA_{RMSE}	MD_{RMSE} (mm^2/s)
SNAILS	25.3	0.046	0.063	13.1×10^{-5}
z-variable VD	**29.3**	**0.031**	**0.043**	$\mathbf{7.6}\times10^{-5}$

Comparison of Reconstruction Schemes

The performance of different reconstruction methods using the proposed z-variable VD-SOS scheme at various undersampling rates (represented as the number of examined

interleaves) is demonstrated in Figure 5.10. The BJS method consistently yields better reconstruction quality with reduced number of interleaves compared to the JWS reconstruction. This is reflected in the lower RMSE values in the FA and MD maps, particularly at higher undersampling rates ($N_s \leq 8$). The adaptive BJS resulted in similar reconstruction performance with the BJS.

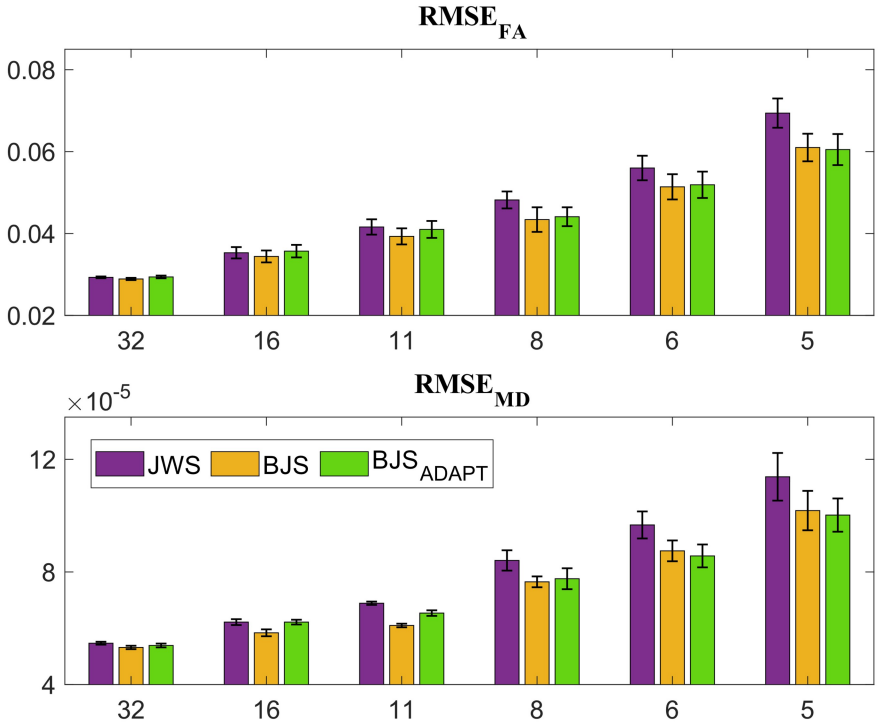


Figure 5.10: RMSE of FA and MD maps obtained from different reconstruction methods at varied undersampling rates indicated by the number of interleaves examined at the central k_z plane. (MD has a unit of mm^2/s)

Robustness Against Noise and DW Directions

The effects of the number of DWIs and the ISNR on the performance of the proposed BJS method and its adaptive version at various undersampling factors are presented in Figures 5.11 and 5.12, respectively. Undersampling is performed using the proposed z-variable VD-SOS scheme on the second dataset. Plots in Figure 5.11 demonstrate that the RMSE in the DTI indices increases with increased values of input noise, more

substantially when ISNR = 10. For a given undersampling factor, the BJS can effectively suppress the undersampling artefacts even when as few as 5 interleaves are used provided that the noise is at low to moderate levels (ISNR $\geq$ 20), whereas higher noise levels adversely affect the effectiveness of artefact suppression, leading to larger errors in the parameter maps.

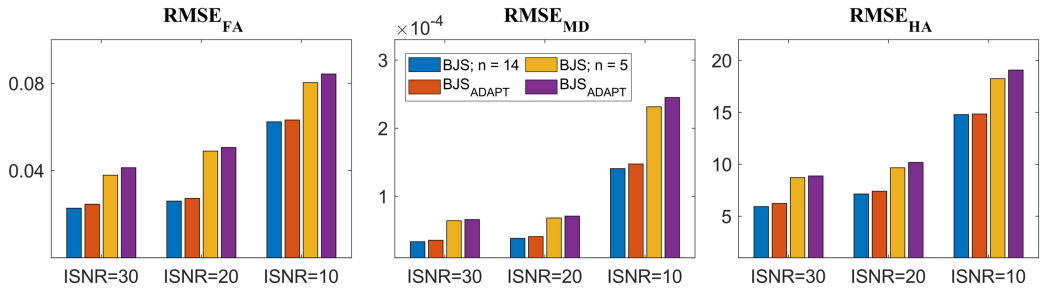


Figure 5.11: Reconstruction performance of the BJS method and its adaptive version against various ISNR values. RMSE of FA, MD (mm^2/s), and HA (°) maps of one representative heart from the second dataset are reported at two undersampling rates.

Plots in Figure 5.12 show the robustness of the proposed approach against reduced number of diffusion directions. With increased undersampling factor, the effect of decreasing number of DWIs on the accuracy of the estimated DTI maps becomes more noticeable, with highest RMSEs measured when N = 6 directions were used. This implies that to achieve the best reconstruction results and increase the benefit of joint sparsity in the BJS method, a balance between the undersampling rate and the number of DWIs should be met.

Comparison of the z-variable VD SOS and cones Trajectories

The RMSE plots (mean ± SD across hearts) of the DTI parameter maps obtained with the BJS reconstruction and its adaptive version using the proposed z-variable VD SOS and cones schemes at various undersampling factors are shown in Figure 5.13. With the same design parameters, the SOS is observed to yield better reconstruction quality compared to the cones scheme. This can be attributed to the inherent characteristic of the oversampling of the k-space centre in the cones scheme, which leads to a low-pass filter-like behaviour, resulting in the blurring of the reconstructed images and higher errors. This can be particularly noticed from Figure 5.14 representing the

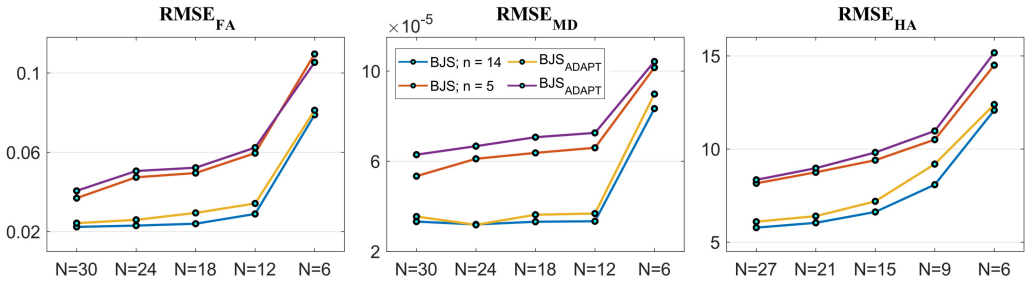


Figure 5.12: Reconstruction performance of the BJS method and its adaptive version against reduced number of diffusion directions. RMSE of FA, MD (mm^2/s), and HA (°) maps of one representative heart from the second dataset are reported at two undersampling rates.

reconstruction results of one illustrative heart obtained with the BJS method using five interleaves in the proposed z-variable VD SOS and cones sampling schemes. The adaptive BJS provides comparable reconstruction quality with similar trends at various undersampling factors.

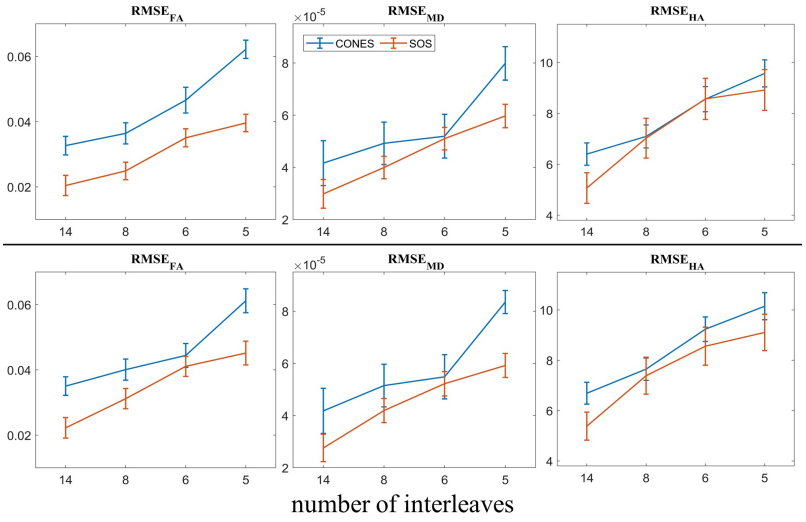


Figure 5.13: RMSE of FA, MD (mm^2/s), and HA (°) maps comparing the proposed z-variable VD SOS and cones sampling schemes obtained with the BJS reconstruction (top panel) and its adaptive version (bottom panel) at various undersampling factors indicated by the number of interleaves examined at the central k_z plane/cone.

Reconstructed DTI parameter maps, including MD (mm^2/s), FA, and HA (°), of

a representative heart from the second dataset, are compared using the proposed z-variable VD SOS and cones sampling schemes. The comparison also includes difference maps, scaled by a factor of 4 for better visualisation.

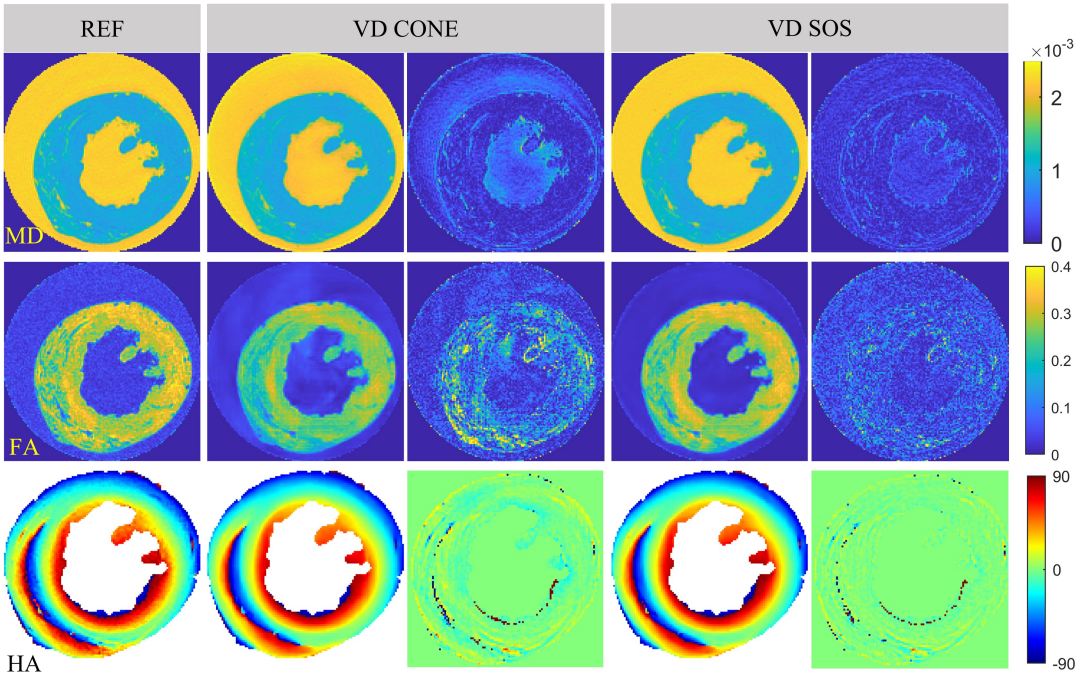


Figure 5.14: Reconstructed DTI parameter maps (MD (mm^2/s), FA, HA (°)) of a representative heart in the second dataset comparing the proposed z-variable VD SOS and cones sampling schemes. The difference maps, scaled by a factor of 4 for better visualisation ($|\mathbf{I}_{rec} - \mathbf{I}_{ref}| \times 4$), are also included.

Phase Error Compensation

Figure 5.15 compares one representative DW image and the FA and MD maps reconstructed from the third dataset with and without phase error compensation using z-variable VD-SOS. The results shown are obtained with the *adaptive* BJS method. It is apparent that images reconstructed without any phase correction are severely distorted. With the phase compensation, the image quality is significantly improved, particularly in the myocardium ROI. This improvement is more clearly evident in the absolute difference maps scaled by a factor of 4. The RMSE reduced from 0.133 to 0.053 in the FA maps and from 15.3×10^{-5} to 4.8×10^{-5} (mm^2/s) in the MD maps

when estimated phase error maps from the 20% of the central k-space were used in the reconstruction to compensate for the phase inconsistencies, further indicating the effectiveness of the phase correction in the BJS reconstruction.

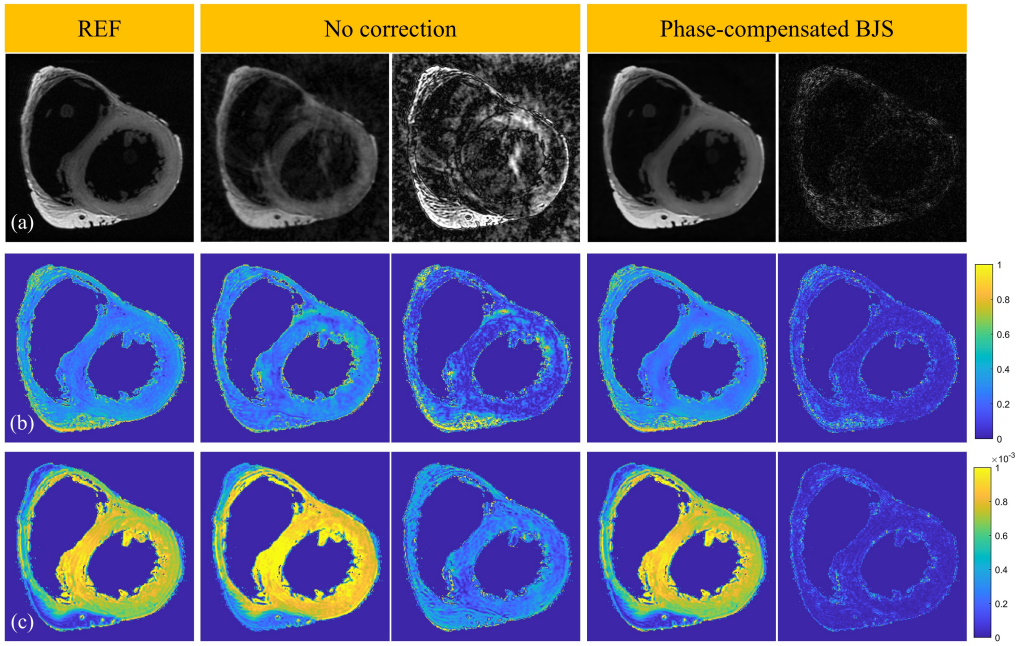


Figure 5.15: Performance of the phase correction in the adaptive BJS reconstruction using z-variable VD-SOS sampling. One representative DW image and the FA and MD (mm^2/s) maps obtained from the third dataset are shown. Random nonlinear phase errors were added to each shot at 10 randomly selected k_z planes. The corrupted data is then reconstructed with and without the phase correction. The absolute difference maps, scaled by a factor of 4 ($|\mathbf{X}_{rec} - \mathbf{X}_{ref}| \times 4$), are also included for improved visualisation and performance assessment.

5.5 Discussion

This work presents a theoretical framework to evaluate a novel sampling/reconstruction approach using variable-density 3D spiral trajectories combined with a bidomain joint CS reconstruction. Our experiments aimed to show that the proposed approach can enable whole-heart cDTI at a high nominal resolution while ensuring accurate quantitative assessment of the microstructural indices, which can benefit the investigation of

fine structures reported in histology. The proposed approach relies on improvements to three aspects of high-resolution 3D imaging: sampling scheme, image reconstruction, and phase inconsistency containment, all of which need to be addressed prior to prospective implementation on an MR system.

The optimal sampling of the image information, concentrated at the centre of k-space, with the proposed z-variable VD stack-of-spirals scheme provided a flexible way of reducing acquired data while leading to all reconstructions exhibiting less visible artefacts, particularly in the myocardium ROI, even at high undersampling factors. The undersampling is well tolerated because of the diffuse aliasing properties of the VD sampling scheme. The proposed BJS reconstruction exploited the group sparsity characteristic of DWIs and enabled simultaneous recovery of all images, improving the computational speed. Both qualitative and quantitative results suggested improved overall image quality in the reconstructions leveraging the group sparsity property, with the proposed BJS reconstruction outperforming the use of a single constraint alone, i.e., JWS reconstruction. The JTV regularisation was particularly suitable for sharpening large edges (results not included here), while the JWS constraint was more effective at removing noise-like background artefacts and re-introducing fine image details. These features were combined in the proposed BJS reconstruction, successfully removing undersampling artefacts and restoring the image structure information. Lower reconstruction errors, better preserved structural similarity to the reference, and improved noise suppression were consistently obtained through the BJS reconstruction at all undersampling factors. The superiority of the proposed method over other reconstruction techniques was more prominent at higher undersampling factors $\mathrm{N_{int}} \geq 8$.

We evaluated the performance of the proposed approach for varied number of DWIs and ISNR values. Better estimation of DTI indices was obtained from data with higher ISNR at a given undersampling factor. For moderate undersampling factors (e.g., $\mathrm{N_{int}} = 14$), the reconstruction could effectively suppress the input noise, and no significant changes were observed in RMS errors at ISNR ≥ 20. For extreme undersampling factors (e.g., $\mathrm{N_{int}} = 5$), the noise contamination besides the appearance of undersampling artefacts led to more variations in the measured error with reduced ISNR. The reconstruction showed robustness against the number of DW images when adequate number of diffusion directions were used ($N \geq 12$), but reducing N further led to the amplified errors even at a moderate undersampling rate ($\mathrm{N_{int}} = 14$). This is

partly due to the reduced reconstruction accuracy resulting from insufficient correlated information and partly due to the error in the tensor estimation. Therefore, with strong undersampling, adequate number of diffusion directions are required to ensure high reconstruction fidelity and more reliable DTI parameter maps. For a given undersampling factor, MD showed the least dependence on the number of diffusion directions, while HA and FA showed more observable changes.

For quantitative MRI applications, such as DTI, where many measurements are typically required, both the accuracy and the speed of the reconstruction become essential factors for the efficiency of the whole imaging process. The proposed BJS algorithm has a considerably fast convergence rate, which makes it specifically desirable for tackling 3D non-Cartesian data reconstruction. Since the two subproblems are decoupled, they can be solved in parallel to improve the reconstruction speed. Using a Graphical Processing Unit (GPU) would also contribute to further reduction of the computational time and alleviates the large memory requirement for storing NUFT coefficients, thus facilitating in-line reconstructions.

Selecting the optimal regularisation parameter in iterative algorithms is usually a tedious task, especially for large-scale optimisation problems, and requires many heuristic trials and/or relies on the availability of a reference dataset. We addressed this concern by introducing an adaptive version of the BJS algorithm, which still provides better reconstruction quality compared to the classical CS methods but with slightly higher RMSEs than the BJS method with optimised hyperparameters at moderate undersampling factors. The two methods, however, lead to similar results at extreme undersampling rates. Improved methods of adaptive noise estimation, automatic parameter setting, and variable-density compensation kernels (required for the NUFT operation) may help in improving the reconstruction accuracy.

An inherent phase map estimation and correction is also introduced. A phase map could be effectively estimated from the fully sampled central k-space of each interleave and incorporated into the encoding matrix in the forward model in the BJS reconstruction. The effect of noise and number of shots in the estimated phase map were beyond the scope of this study, and yet remained to be investigated in the future. The estimation of the phase map can be possibly made more accurate using an iterative scheme to recalculate the phase error of each shot from the updated image at every iteration at the expense of increased reconstruction time [191].

5.5.1 Scan Time Reduction

The proposed BJS reconstruction combined with the incoherent k, q VD undersampling strategy could result in considerable reduction of the scan time with minimal loss in image quality and accuracy of diffusion metrics. In this work, the SOS and conical schemes are demonstrated for comparison and the extent of undersampling achieved through each trajectory is evaluated. With the same design parameters, both trajectories provided fairly comparable reconstruction quality with similar trends at moderate undersampling factors, while the SOS outperforming the cones at aggressive undersampling rates ($N_{int} = 5$). The proposed VD sampling enables reduced coherent aliasing artefacts to appear along the stacking direction (k_z), facilitating higher undersampling factors. Specifically, our results show that reconstruction with $\sim 5\%$ error is feasible using $\leq 5\%$ of the k-space data per direction. Considering the same acquisition parameters as in the preclinical setting, this can be translated into reducing the acquisition time of each 3D DW image with a high nominal resolution of 100 μm^3 from 2.15 h in full sampling to approximately $4 - 6$ min. For the clinical setting with $T_R = 1000$ ms, this corresponds to an acquisition time of $10 - 12$ min for whole-heart imaging at full FOV with a spatial resolution of 1 mm^3 isotropic. Additional acceleration can be achieved by employing parallel imaging along the k_z direction.

We evaluated different multishot schemes using preclinical and clinical settings. Since, the actual number of interleaves required depends on the readout time and gradient hardware constraints, there should be considered a trade-off between the number of interleaves and the length of each interleave (or equivalently the readout time). In preclinical settings, with high gradient capabilities and no strict motion-induced artefacts, within a shorter readout time, sufficient sampling of the k-space until larger k_{max} (corresponding to higher spatial resolution) can be achieved while phase-induced artefacts arising, for example, from B_0 inhomogeneities will be minimised. With limiting the readout time to 3 ms, we expect the phase accrual during the readout could be possibly minimised, thus rendering the acquisition insensitive to off-resonance artefacts. This is also consistent with the analysis presented in [192], indicating that the off-resonance effect is not severe as long as $\Delta f < 1/(2T_{ro})$, where T_{ro} is the readout duration. However, in clinical settings, where smaller gradient amplitudes and slew rates are available and motion-induced phase errors are more severe, longer readouts with smaller number of shots may be helpful to effectively avoid the shot-to-shot phase

inconsistencies but may in turn necessitate some additional steps to correct for these artefacts [193, 194]. Similar to the shot-shot phase correction, the k-space trajectory errors can be modeled into the image encoding matrix and jointly estimated during the reconstruction [195].

The gain in the acquisition time achieved in this study can be potentially utilised for more complex microstructure models, such as multishell b-value and diffusion spectrum imaging (DSI) schemes, where a higher number of diffusion-encoding directions and weightings (q-points) are required, and acceleration techniques are more needed.

5.6 Summary

In summary, this chapter introduced a novel 3D sampling combined with an efficient CS reconstruction designed MRI to generate high spatial resolution DW images of the heart within significantly reduced acquisition times. The technique utilises an optimal variable density spiral trajectory, resulting in undersampled, incoherent k, q-space filling, followed by compressed sensing (CS) reconstruction. High accelerations of 25X were achieved with maintained image quality, allowing for accurate determination of cardiac microstructural parameters. Compared to the multislice sampling strategies, 3D volume imaging would be more advantageous as it allows for a higher degree of acquisition undersampling and much higher flexibility in incoherence of the undersampling artefacts, allowing for better compressed sensing (CS) reconstruction. Moreover, T_E could be further reduced due to the absence of a slice selection refocusing gradient, and thereby leading to enhanced SNR.

Chapter 6

Accelerated cDTI with Spiral Keyhole (SPIK) and Blind Dictionary Learning

6.1 Introduction

High-resolution cDTI holds promise for enhancing the understanding and diagnosis of various early stage pathologies characterised by sub-millimetric lesions or morphological changes, such as early detection of HCM [19]. While providing sufficient signal-to-noise ratio (SNR) to improve spatial resolution is required for accurate characterisation of diffusion-related parameter maps and tractography results, challenges persist due to lengthy scan times and susceptibility to motion artefacts. Despite advancements in correcting for bulk and physiological motion, extended acquisition times remain a significant impediment to achieving high SNR and resolution imaging, particularly in clinical applications.

In Chapters 4 and 5, we relied on an extensive study on the optimal sampling parameters using an efficient CS reconstruction method. In this chapter, we focus on introducing a novel sampling scheme, referred to as Spiral Keyhole ("SPIK") combined with a blind dictionary learning reconstruction to further enhance the DW image quality and allow for shorter scan times, which can favor more complex diffusion MRI techniques including the multi-shell dMRI in the heart. Our focus is to minimise the phase inconsistencies arise from motion and eddy current induced artefacts and to allow for improved correction for any residual artefacts.

6.2 Materials and Methods

6.2.1 Spiral Keyhole Imaging

In this study, we propose a novel sampling strategy inspired by the Echo-planar imaging with keyhole [196, 197], which was originally developed in functional MRI (fMRI) to achieve a high temporal resolution while eliminating the phase errors and image distortions attributed to long acquisition window in traditional EPI. Our proposed sampling scheme, referred to as SPIral Keyhole (SPIK), combines the acquisition of a central part of k-space with a continuous and interleaved update of high spatial frequencies in the peripheral parts. The fundamental idea behind SPIK is based on utilising the extensive correlations between diffusion-weighted images to share portions of the k-space, taking advantage of the fact that the central part of k-space, referred to as the "keyhole", contains most of the useful image information and defines the coarse appearance and contrast of an image. SPIK optimally samples this central region according to the

Nyquist criterion for every DW image, ensuring that image contrast and angular (diffusion directional) SNR maintains efficiently at a level comparable to that of single-shot spiral. Meanwhile, the outer regions of the k-space, where information about the image details and edge definition lies, are shared among DWIs and sparsely sampled through fewer shots.

In other words, only one low-resolution spiral interleave is employed to encode k-space centre for every DW image while a single shot of a high-resolution multi-shot spiral encodes the k-space periphery. As a result, SPIK achieves faster scan time and fewer image distortions compared to conventional multi-shot spiral imaging, yet reaches comparable high spatial resolution, making it a valuable sampling for cDTI studies.

Figure 6.1 illustrates the schematic representation of the SPIK acquisition scheme. Each DW k-space is divided into two regions: a *keyhole* region (KK) and a *sparse* region (KS). A high-resolution three-shot spiral encodes the KS region for three consecutive measurements, while a low-resolution single-shot spiral samples the KK region for every measurement.

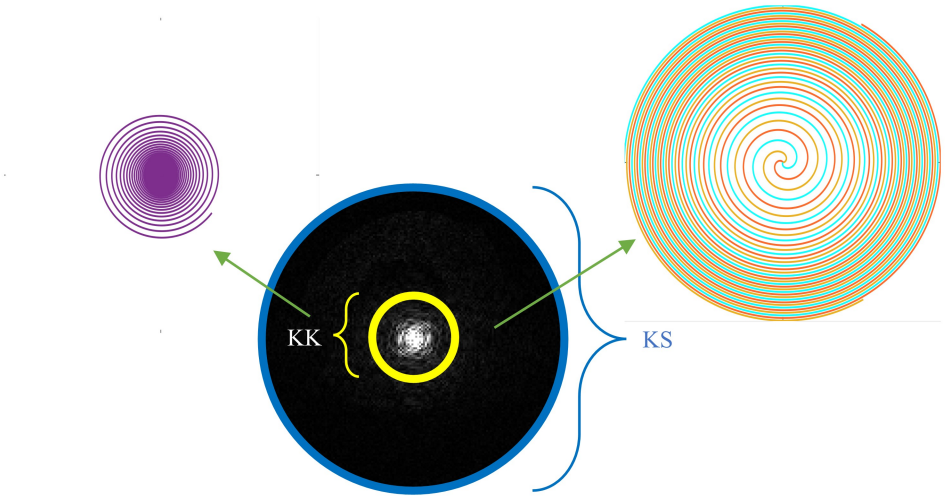


Figure 6.1: Schematic representation of the SPIK acquisition scheme. Each k-space is divided into two distinct regions: a keyhole region (KK) and a sparse region (KS). A high-resolution three-shot spiral encodes the KS region for 1st, 2nd, and 3rd measurements. A low-resolution single-shot samples the KK region for every measurement.

The *keyhole* spiral was designed using the optimal sampling density function/effective FOV introduced in chapter 4. The same sampling density function vertically mirrored

such that the higher sampling density is assigned to the farthest k radius and lower density to the lower k radius was used to design the peripheral multi-shot spiral, as depicted in Figure 6.2. For the *keyhole* spiral, the maximum spatial frequency was determined as a fraction of the maximum spatial frequency k_m corresponding to the desired high spatial resolution; i.e., $k_m^{KK} = \alpha\ k_m$, with $0 < \alpha < 1$.

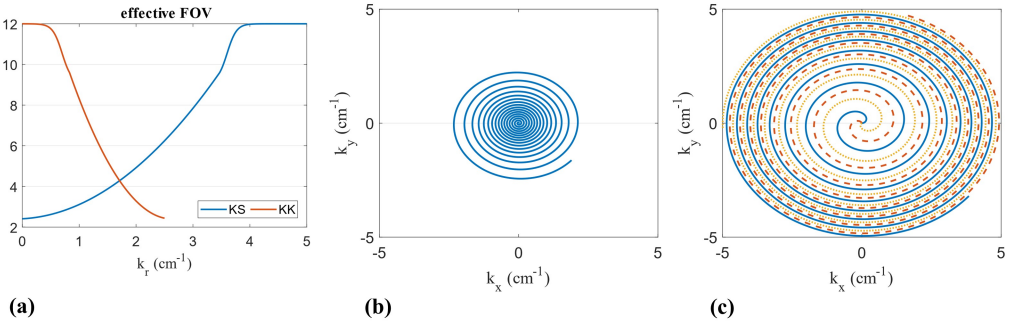


Figure 6.2: (a) The effective FOV corresponding to the sampling density functions used for VDS trajectories in the KK and KS regions. (b) The resulting Keyhole spiral trajectory. (c) The multi-shot spiral trajectory for sparsely sampling the k-space periphery. The spiral shot numbers used for sampling the 1st, 2nd, and 3rd scans are denoted with solid, dashed, and fine-dashed lines in (c).

By utilising the correlations in the sparse data region from a few consecutive scans and updating the *keyhole* region with every measurement, complete k-space data can be obtained for each DW scan. The k-space trajectory of a three-shot SPIK used for sampling three consecutive DW directions is depicted in Figure 6.3. To introduce incoherency, *keyhole* trajectories were rotated with the golden angle ordering.

With this scheme, the central *keyhole* region, consisted of unique DW signal, is sufficiently sampled with the optimal VDS trajectory, displayed in red in Figure 6.3. The peripheral k-space is acquired sparsely in an interleaved manner, displayed in blue in Figure 6.3, such that after every three scan, the sparse k-space region will be densely acquired. Effectively, KK is updated after every scan, while KS is updated every fourth scan. Consequently, by only doubling the scan time compared to a single-shot acquisition, high spatial resolutions can be achieved.

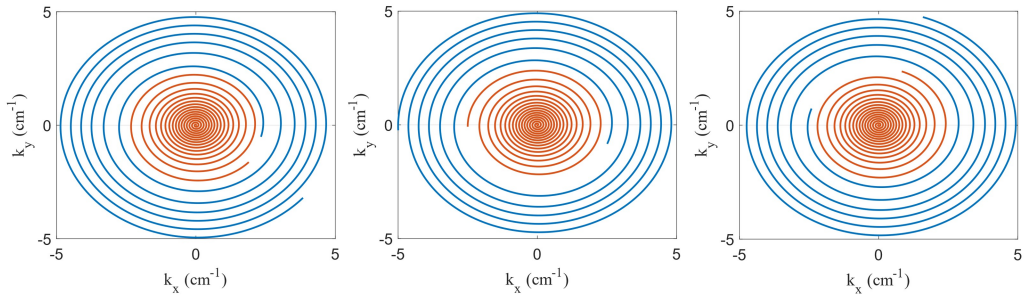


Figure 6.3: Schematic representation of the k-space trajectory of three-shot SPIK. The trajectory consists of optimally dense sampling of the central k-space (KK, displayed in red) and segmented interleaved acquisition of the peripheral k-space (KS, displayed in blue). Similar to a three-shot spiral imaging, every third spiral interleave samples the peripheral k-space at each scan.

6.2.2 Image Reconstruction

We introduced a blind dictionary learning reconstruction to leverage the correlations in the high spatial frequencies to alleviate the blurring artefacts during the recovery of the diffusion signal. In other words, the unique diffusion contrast information associated with each DW direction will be effectively retained while the correlated structural details of the images are learnt during the reconstruction, leading to improved overall quality and enhanced directional SNR. The proposed reconstruction is based on a blind compressed sensing (BCS) scheme [198], which was originally introduced for dynamic imaging [199], and later optimised for multi-parametric brain imaging [200].

The BCS model is built upon the assumption that there are a finite number of distinct tissue types with unique parametric values within the specimen being imaged that can be efficiently represented by linear combinations of basis functions contained within a dictionary. In Section 3.4.1, we briefly discussed the sparsifying transforms/dictionaries used in cDTI and mentioned that generally data-adaptive sparsifying dictionaries yield better sparsity representations and thus lower reconstruction errors than predefined transforms. BCS framework combines elements from both dictionary learning (DL) and low-rank (LR) methods, which made it a suitable candidate for our problem. It enables recovery of a high-dimensional signal from a small number of measurements, without any reliance on the prior knowledge of the sparsity basis or a training

dataset. BCS simultaneously estimates the dictionary and the coefficients directly from the undersampled measurements, as opposed to DL algorithms. This results in the robustness of the algorithm to subject-specific imperfections, including the inter-frame motion [200]. Additionally, the rows of the coefficient matrix $\mathbf{C}$ are constrained to be sparse in the BCS framework, as opposed to the LR representation. Therefore, the BCS model can be regarded as a locally low-rank representation [62]. However, instead of estimating the basis functions independently for each neighborhood, it estimates the overcomplete dictionary $\mathbf{\Gamma}$ from all voxels simultaneously, allowing to leverage correlations between voxels that are spatially distant from each other (nonlocal correlations).

6.2.3 Diffusion-weighted Signal Representation

As previously discussed, diffusion MRI involves obtaining a series of complex-valued diffusion-weighted images using different diffusion encoding gradient directions and/or weightings (q-points). We denote the spatio-angular signal as $\{s(\mathbf{x}, q)\}_{q=1}^{N}$, where $\mathbf{x}$ is the spatial location and q denotes the diffusion-encoding point. The DW dataset can be represented in a Casorati matrix form as:

$$\mathbf{S}_{M\times N} = \begin{bmatrix} s(\mathbf{x}_1, q_1) & \cdots & s(\mathbf{x}_1, q_N) \\ s(\mathbf{x}_2, q_1) & \cdots & s(\mathbf{x}_2, q_N) \\ \vdots & \ddots & \vdots \\ s(\mathbf{x}_M, q_1) & \cdots & s(\mathbf{x}_M, q_N) \end{bmatrix} \tag{6.1}$$

Here, M represents the number of voxels in the image ($M = m_x \times m_y \times m_z$), and N represents the number of diffusion-weighted images in the dataset.

The main goal is to recover the diffusion-weighted signal $s(\mathbf{x}, q_i) \in \mathbb{C}^{M\times 1}$ from undersampled MR measurements $\mathbf{d}_i \in \mathbb{C}^{D\times 1}$ ($D < M$). We model the measurement process as:

$$\mathbf{d}_i = \underbrace{(\Omega_i \circ \mathcal{NF} \circ \mathcal{E}_i)}_{\mathcal{A}_i}(\mathbf{S}) + n_i; \quad i = 1, ..., N \tag{6.2}$$

where $\mathcal{E}_i$ is an operator that extracts the ith column of $\mathbf{S}$, which corresponds to the image at q_i point, $\mathcal{NF} : \mathbb{C}^{M\times 1} \rightarrow \mathbb{C}^{N_k\times 1}$ is the non-uniform Fourier transform (NUFT), $\Omega_i : \mathbb{C}^{N_k\times 1} \rightarrow \mathbb{C}^{D\times 1}$ is the undersampling operator that extracts Fourier samples on the k-space trajectory, their product $\mathcal{A}_i$ constructs the sensing matrix corresponding to the ith DWI, and $n_i \in \mathbb{C}^{D\times 1}$ represents the additive measurement noise.

Building upon the BCS model, we represent the diffusion signal matrix $\mathbf{S}$ as the product of a sparse coefficient matrix $\mathbf{C} \in \mathbb{C}^{M \times R}$ and diffusion parametric basis functions in a finite dictionary $\mathbf{\Gamma} \in \mathbb{C}^{R \times N}$, i.e., $\mathbf{S} = \mathbf{C}\,\mathbf{\Gamma}$:

$$\begin{bmatrix} s(\mathbf{x}_1, q_1) & \cdots & s(\mathbf{x}_1, q_N) \\ s(\mathbf{x}_2, q_1) & \cdots & s(\mathbf{x}_2, q_N) \\ \vdots & \ddots & \vdots \\ s(\mathbf{x}_V, q_1) & \cdots & s(\mathbf{x}_V, q_N) \end{bmatrix} = \begin{bmatrix} c_1(\mathbf{x}_1) & \cdots & c_R(\mathbf{x}_1) \\ c_1(\mathbf{x}_2) & \cdots & c_R(\mathbf{x}_2) \\ \vdots & \ddots & \vdots \\ c_1(\mathbf{x}_V) & \cdots & c_R(\mathbf{x}_V) \end{bmatrix} \begin{bmatrix} \gamma_1(q_1) & \cdots & \gamma_1(q_N) \\ \gamma_2(q_1) & \cdots & \gamma_2(q_N) \\ \vdots & \ddots & \vdots \\ \gamma_R(q_1) & \cdots & \gamma_R(q_N) \end{bmatrix} \tag{6.3}$$

where R is the total number of basis functions in the dictionary. Here, we assume that the size of the dictionary R is greater than the rank of matrix $\mathbf{\Gamma}$, $(R > r)$ a condition that typically holds in most practical cases. With a larger number of basis functions in the dictionary, the BCS model offers a richer representation compared to the LR model. However, the sparsity assumption ensures that this richness does not translate into increased degrees of freedom. Moreover, the sparsity assumption aids in noise suppression and reduces blurring artefacts, resulting in sharper reconstructions, as will be shown in the results section.

6.2.4 Reconstruction Formulation

We formulate the recovery of the signal matrix $\mathbf{S}$ as a constrained optimisation problem by imposing sparsity constraints on both the dictionary $\mathbf{\Gamma}$ and the coefficients $\mathbf{C}$ in the respective transform basis $\mathcal{T}_1$ and $\mathcal{T}_2$:

$$\{\hat{\mathbf{C}}, \hat{\mathbf{\Gamma}}\} = \arg\min_{C,\Gamma} \left\{ \sum_{i=1}^{N} \|\mathcal{A}_i(\mathbf{C}\mathbf{\Gamma}) - \mathbf{d}_i\|_2^2 + \tau_1 \|\mathcal{T}_1(\mathbf{C})\|_{\ell_1} + \tau_2 \|\mathcal{T}_2(\mathbf{\Gamma})\|_{\ell_1} \right\} \tag{6.4}$$

The first term ensures the data consistency, and τ_1 and τ_2 are the regularisation parameters. The ℓ_1 norm promotes sparsity and is defined as the absolute sum of the matrix entries: $\|U\|_{\ell_1} = \sum_{i=1}^{M} \sum_{j=1}^{R} |u(i,j)|$. This constraint indicates that there are very few non-zero entries in $\mathbf{C}$ or equivalently only a few diffusion parametric basis functions are sufficient to model the diffusion profile at any given voxel. We refer to our proposed BCS formulation as sparse BCS (s-BCS).

The operators $\mathcal{T}_1$ and $\mathcal{T}_2$ can be chosen appropriately based on the specific diffusion MRI application. In this study, we consider the cases: (i) $\mathcal{T}_1 = \mathcal{I}$ and $\mathcal{T}_2 = \mathcal{I}$, (ii) $\mathcal{T}_1 = W$ (2D Wavelet operator) and $\mathcal{T}_2 = \mathcal{I}$, where $\mathcal{I}$ represents the identity matrix.

These choices were made to address scenarios involving breath-hold or ex vivo imaging and large inter-frame motion (e.g., free-breathing cDTI) or low SNR, respectively.

A key distinction of the proposed formulation from the original BCS setting [200] lies in the sparsity constraint imposed on $\mathbf{\Gamma}$ instead of the Frobenius norm constraint. The comparisons in Figure 6.8 demonstrate the performance improvement offered by introducing this additional sparsity constraint on the dictionary. The sparsity constraint effectively penalises basis functions that capture the streaking artefacts (noisy oscillations) and avoids learning of coefficient values that are arbitrarily small in magnitude. Another distinction lies in our consideration of sparsity of the coefficients in a transform domain (i.e., wavelet), which led to adaptation of a different solving algorithm instead of a majorise-minimise approximation for the ℓ_1 penalty on $\mathbf{C}$.

6.2.5 Optimisation Algorithm

In this subsection, we will outline the steps for solving the optimisation problem in Equation 6.4. We employed an Augmented Lagrangian (AL) optimisation scheme, specifically the alternating direction method of multipliers (ADMM), and a variable splitting strategy. The original optimisation problem is split into simpler sub-problems and the solution is obtained by iterating through these sub-problems.

We first drive an equivalent problem for Equation 6.4 by change of variables:

$$\begin{aligned} \arg\min_{C,\Gamma} \|\mathcal{A}(\mathbf{C\Gamma}) - \mathbf{d}\|_F^2 + \tau_1\|\mathbf{P}\|_{\ell_1} + \tau_2\|\mathbf{Q}\|_{\ell_1} \\ \textbf{s.t.}\ \ \mathcal{T}_1(\mathbf{C}) - \mathbf{P} = 0 \\ \textbf{s.t.}\ \ \mathcal{T}_2(\mathbf{\Gamma}) - \mathbf{Q} = 0 \end{aligned} \tag{6.5}$$

The augmented Lagrangian corresponding to constrained optimization problem 6.5 can be written as:

$$\begin{aligned} \mathcal{L}(\mathbf{C},\mathbf{\Gamma},\mathbf{P},\mathbf{Q}) =& \|\mathcal{A}(\mathbf{C\Gamma}) - \mathbf{d}\|_F^2 + \tau_1\|\mathbf{P}\|_1 + \tau_2\|\mathbf{Q}\|_1 \\ &+ \frac{\beta_1}{2}\|\mathcal{T}_1(\mathbf{C}) - \mathbf{P}\|_2^2 + \mathbf{\Lambda}_1^H(\mathcal{T}_1(\mathbf{C}) - \mathbf{P}) \\ &+ \frac{\beta_2}{2}\|\mathcal{T}_2(\mathbf{\Gamma}) - \mathbf{Q}\|_2^2 + \mathbf{\Lambda}_2^H(\mathcal{T}_2(\mathbf{\Gamma}) - \mathbf{Q}) \end{aligned} \tag{6.6}$$

where $\mathbf{P}$ and $\mathbf{Q}$ are auxiliary variables, $\mathbf{\Lambda}_1$ and $\mathbf{\Lambda}_2$ are the normalised Lagrangian multipliers, and β_1 and β_2 are the penalty parameters controlling the speed of convergence. The modified cost function should be minimised with respect to four variables $\mathbf{C}$, $\mathbf{\Gamma}$,

$\mathbf{P}$, and $\mathbf{Q}$. This formulation may seem more complex than the original problem in 6.4; however, the simplifications enable us to decouple the optimisation problem in Equation 6.6 into simpler subproblems. We use an alternating minimisation strategy to solve for a specific variable, assuming the other variables remain constant, and then systematically iterate through the subproblems until convergence. Each of the subproblems is relatively straightforward and can be efficiently solved, either through analytical methods or simple optimisation techniques as described below.

The subproblems for $\mathbf{P}$ and $\mathbf{Q}$ become:

$$\mathbf{P}^{k+1} = \min_{\mathbf{P}} \tau_1 \|\mathbf{P}\|_1 + \frac{\beta_1}{2}\|\mathbf{L}^k - \mathbf{P}\|_2^2 \tag{6.7}$$

$$\mathbf{Q}^{k+1} = \min_{\mathbf{Q}} \tau_2 \|\mathbf{Q}\|_1 + \frac{\beta_2}{2}\|\mathbf{Z}^k - \mathbf{Q}\|_2^2 \tag{6.8}$$

where $\mathbf{L}^k = \mathcal{T}_1(\mathbf{C}^k) + \mathbf{\Lambda}_1^k$ and $\mathbf{Z}^k = \mathcal{T}_2(\mathbf{\Gamma}^k) + \mathbf{\Lambda}_2^k$. Equations 6.7 and 6.8 can be analytically solved through soft-thresholding:

$$\mathbf{P}^{k+1} = \frac{\mathbf{L}^k}{|\mathbf{L}^k|}\Big(|\mathbf{L}^k| - \frac{\tau_1}{\beta_1}\Big)_+ \tag{6.9}$$

$$\mathbf{Q}^{k+1} = \frac{\mathbf{Z}^k}{|\mathbf{Z}^k|}\Big(|\mathbf{Z}^k| - \frac{\tau_2}{\beta_2}\Big)_+ \tag{6.10}$$

where $(.)_+$ represents the soft thresholding operator defined as $(u)_+ = \max\{0, u\}$, and k denotes the iteration number.

The subproblems for $\mathbf{C}$ and $\mathbf{\Gamma}$ can be written as follows:

$$\mathbf{C}^{k+1} = \min \|\mathcal{A}(\mathbf{C}\mathbf{\Gamma}^k) - \mathbf{d}\|_2^2 + \frac{\beta_1}{2}\|\mathcal{T}_1(\mathbf{C}) - \mathbf{P}^{k+1}\|_2^2 + \mathbf{\Lambda}_1^{k^H}(\mathcal{T}_1(\mathbf{C})) \tag{6.11}$$

$$\mathbf{\Gamma}^{k+1} = \min \|\mathcal{A}(\mathbf{C}^{k+1}\mathbf{\Gamma}) - \mathbf{d}\|_2^2 + \frac{\beta_2}{2}\|\mathcal{T}_2(\mathbf{\Gamma}) - \mathbf{Q}^{k+1}\|_2^2 + \mathbf{\Lambda}_2^{k^H}(\mathcal{T}_2(\mathbf{\Gamma})) \tag{6.12}$$

These equations are quadratic in nature, and therefore, can be effectively solved using a pre-conditioned conjugate gradient (PCG) algorithm. We used the steepest ascent method to update the Lagrangian multipliers at each iteration:

$$\mathbf{\Lambda}_1^{k+1} = \mathbf{\Lambda}_1^k + (\mathcal{T}_1(\mathbf{C}^{k+1}) - \mathbf{P}^{k+1}) \tag{6.13}$$

$$\mathbf{\Lambda}_2^{k+1} = \mathbf{\Lambda}_2^k + (\mathcal{T}_2(\mathbf{\Gamma}^{k+1}) - \mathbf{Q}^{k+1}) \tag{6.14}$$

The pseudo-code of the reconstruction algorithm is shown in 3. It is worth noting that by applying the ADMM algorithm, the convex optimisation problem converges to

an optimal solution, irrespective of the initialisation choice [201]. In our implementation, the algorithm initialises $\mathbf{\Lambda}_1^0$, $\mathbf{\Lambda}_2^0$ with a zero matrix and $\mathbf{\Gamma}^0$ with a random matrix. An initial estimation for the coefficient matrix $\mathbf{C}^0$ is then obtained by minimising only the data consistency term. To effectively perform soft-thresholding of the transform coefficients, we initialise $\frac{\tau_1}{\beta_1}$ and $\frac{\tau_2}{\beta_2}$ with the maximum value within the initial transform domains, resulting in small values of β_1, β_2. The algorithm iterates between Equations 6.9 and 6.12 in an inner loop while β_1 and β_2 are gradually increased in an outer loop. The continuation strategy on β_1 and β_2 improves the convergence and avoids falling into local minima by gradually eliminating the small coefficients in $\mathbf{C}$ while relearning the dictionary bases. The inner loop is terminated when the cost in Equation 6.4 stagnates to a threshold of 10^{-5}. We define convergence as the point where the constraint defined by:

$$\begin{aligned} \|\mathcal{T}_1(\mathbf{C}) - \mathbf{P}\|^2 &< 10^{-6} \\ \|\mathcal{T}_2(\mathbf{\Gamma}) - \mathbf{Q}\|^2 &< 10^{-8} \end{aligned} \tag{6.15}$$

are satisfied.

Algorithm 3 BCS Algorithm

Input: $\tau_1, \tau_2 > 0, k = 0$

Initialise: $\mathbf{\Gamma}^0, \mathbf{C}^0, \quad \beta_1, \beta_2 > 0$

for $k = 1$ **to** K **do**

while: Stopping criterion is not satisfied (Eq. 6.15)

Update $\mathbf{P}$: Shrinkage step $\leftarrow$ Eq. 6.9

Update $\mathbf{C}$: Quadratic subproblem $\leftarrow$ Eq. 6.11

Update $\mathbf{Q}$: Shrinkage step $\leftarrow$ Eq. 6.10

Update $\mathbf{\Gamma}$: Quadratic subproblem $\leftarrow$ Eq. 6.12

Update $\mathbf{\Lambda}_1$: Steepest ascend $\leftarrow$ Eq. 6.13

Update $\mathbf{\Lambda}_2$: Steepest ascend $\leftarrow$ Eq. 6.14

if $\frac{|\mathcal{C}^k - \mathcal{C}^{k-1}|}{|\mathcal{C}^k|} \leq 10^{-5}$

then $\begin{matrix} \beta_1 & = \beta_1 \times 5 \\ \beta_2 & = \beta_1 \times 5 \end{matrix}$: Continuation parameter update

end for

return $(\mathbf{C}, \mathbf{\Gamma})$

6.3 Experiments

6.3.1 Datasets

We evaluated the effectiveness of the proposed approach using the ex vivo human cDTI data explained in Section 5.3.1. To validate the clinical potential of the approach, we employed the MRXCAT-cDTI numerical phantom [202] to generate in vivo cardiac DTI data. The simulator utilises the extended Cardiac-Torso (XCAT) phantom to produce realistic anatomical masks that incorporate cardiac contraction and respiratory motion while accounting for variability in breathing states. In vivo cardiac DW images were realistically modelled in both breath-hold and free-breathing conditions using different diffusion encoding schemes. In the free-breathing simulations, 10 different motion states were randomly assigned to different DW directions. For spin echo simulations, the apparent diffusion coefficient (ADC) was assumed to be 1.45×10^{-3} mm^2/s, while it was set to 1×10^{-3} mm^2/s in STEAM simulations [107]. The T_1 and T_2 values of myocardium were taken as 1000 ms and 50 ms, respectively. Other parameters used for the simulation are summarised in Table 6.1. We considered a single coil that is placed on the chest and has the maximum sensitivity to the FOV containing the heart (30.25 $\times$ 10.75 cm^2) and a single-shot EPI readout[1]. We used one b_0 image and a b-value of 450 s/mm^2 for the diffusion weighting. We also considered different SNR values for the b_0 image to examine the performance of the proposed approach in both high and low SNR regimes.

Table 6.1: Parameters used in MRXCAT-CDTI simulations

DW Sequence	Respiratory State	TR/TE (ms)	DW direction	N_{slc}	NSA	SNR[2]
STEAM	Breath-hold	2000/35	9	3	1	50
M2-SE[3]	Free-breathing	1000/88	12	1	4	30
M2-SE	Free-breathing	1000/83	12	1	4	15

The Cartesian k-space data obtained from the simulations were on a 121 $\times$ 43 (frequency encodes $\times$ phase encodes) grid. The complex data were then rearranged to

[1] This was the only available option.
[2] SNR of the b_0 image
[3] Second-order motion-compensated (M2) spin-echo

cover an isotropic FOV of 12 × 12 cm^2 and regridded to the proposed SPIK sampling with an isotropic resolution of 1 mm. The resulting spiral data were then reconstructed using the proposed sparse BCS reconstruction. A non-rigid registration was applied to the reference and reconstructed data before computing the diffusion tensors [203]. Examples of this dataset are shown in Figures 6.4 and 6.5.

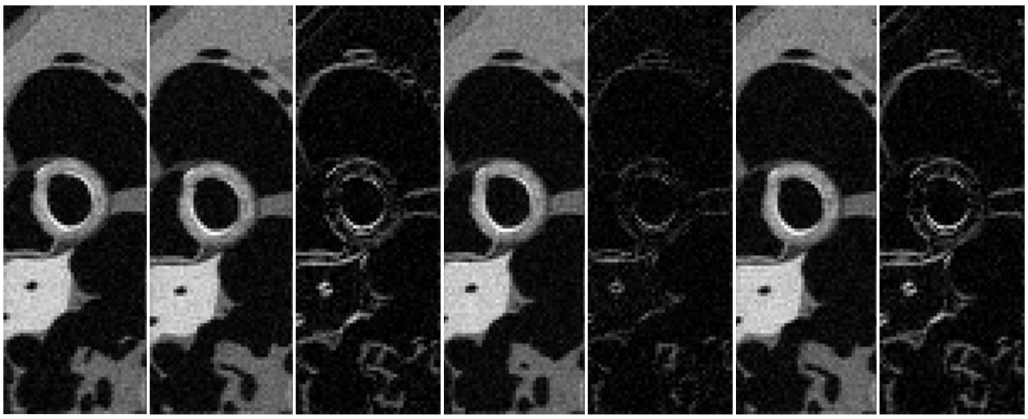


Figure 6.4: The MRXCAT DW images used to validate the proposed approach. One DW direction across four acquisition repetitions (a, b, d, f) simulated with distinct respiratory motion states is shown. The corresponding difference images (c, e, g) with respect to the first repetition (a) are included (with the same scaling as the DW images) for better visualisation of motion effects.

6.3.2 Sampling Schemes

To choose the optimal radius for the *keyhole* region, we studied a range of α values between 0.2 to 0.6. The performance of the proposed SPIK sampling is compared against the proposed multi-shot VDS trajectory by retrospectively undersampling the ex vivo cDTI dataset at varying acceleration factors, specified by the number of interleaves.

To demonstrate the utility of the proposed approach in accelerated 3D cDTI, we extended the SPIK sampling to three dimensions by stacking the trajectories along the k_z direction. Here, we discarded the k_z dependence of the effective FOV in the $k_x - k_y$ and utilised the same trajectories for all the k_z encoding steps. We referred to

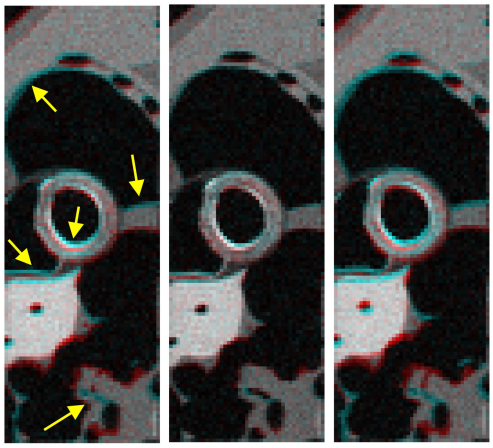


Figure 6.5: Representative free-breathing MRXCAT DW images of three different directions in one DTI examination. Arrows point to the areas with noticeable misalignment.

this k-space trajectory as the standard stack of SPIK. The trajectories were rotated by a small random angle in each k_z slice to increase the incoherence. We compared this sampling with a standard stack of multi-shot VDS sampling using the optimal sampling density function introduced in Chapter 4. We then applied an inverse FFT to the k_z direction and performed the reconstruction of each $k_x - k_y$ slice independently using the proposed sparse BCS reconstruction.

6.3.3 Comparison of Reconstruction Methods

We compared the proposed sparse BCS scheme with 1) the original BCS setting used in [200], where the reconstruction problem is constrained with an ℓ_1 norm on the coefficient matrix $\mathbf{C}$ and a Frobenius norm on the dictionary $\mathbf{\Gamma}$, 2) the joint low-rank and sparsity (LR + S) reconstruction [127], which integrates low-rankness and group wavelet sparsity constraints, and 3) the classical bidomain CS reconstruction.

The regularisation parameters in all the methods were selected to maximise the signal to error ratio between the reconstruction and the reference fully sampled data specified by $SER = 10 \log_{10} \left(\|\mathbf{S}_{rec}\|_2^2 / \|\mathbf{S}_{ref} - \mathbf{S}_{rec}\|_2^2 \right)$.

We introduced zero-mean Gaussian noise to the real and imaginary components of the raw k-space of the ex vivo data to reduce the SNR to approximately 15, resembling the SNR levels found in in vivo cases. The fully sampled k-space data of all stud-

ied datasets were retrospectively resampled using the previously described sampling patterns.

To determine the effect of the size of the dictionary on the performance of the BCS reconstruction, we varied the dictionary size R from 18 to 54. We consider the normalised reconstruction error (ζ), PSNR and SSIM metrics for comparing the performance of the methods in reconstructing the DWIs:

$$\zeta = \frac{\|\mathbf{S}_{rec} - \mathbf{S}_{ref}\|_2^2}{\|\mathbf{S}_{ref}\|_2^2} \tag{6.16}$$

6.3.4 Analysis Methods

We derived the diffusion parametric maps by fitting the mono-exponential model to the reconstructed images on a pixel by pixel basis using a weighted linear least-squares algorithm. We compared the mean square error (MSE in %) of the parameter maps obtained from different algorithms with those obtained from the fully sampled Cartesian data. For the analysis, we focused on the myocardium as the region of interest (ROI). Helix Angle Transmurality (HAT) measures the transmural continuum of HA, showing how HA changes from the endocardium to the epicardium. It is quantified as the slope of HA over transmural depth (TD) across the myocardium. We obtained this measurement by sampling HA along 25 transmural directions and performing linear regression.

6.4 Results

6.4.1 Influence of *Keyhole* Radius

Figure 6.6 shows the effect of the *keyhole* radius. The top row displays one DW image obtained with gridding (direct NUFT) and the proposed sparse BCS scheme using the proposed SPIK sampling with $\alpha = 0.2$, 0.4, and 0.6, respectively. It can be seen that when using direct NUFT, with smaller α values, no visible aliasing artefacts appear in the reconstruction, but the image is significantly blurry and the details could not be recovered precisely. With increasing α, the aliasing artefacts become visible in the reconstruction, yet the details are more distinguishable. Increasing α further results in the severe aliasing artefacts and degradation of the reconstructed images. However, the BCS reconstruction can greatly recover the original DW image and lead to visu-

ally similar image quality regardless of the α values. The bottom row compares the PSNR and normalised reconstruction error (ζ) scores as a function of the α for both reconstruction methods. It is observed that $\alpha = 0.4$ yielded the best reconstruction performance with the BCS method, and hence, was used for the all the comparisons in the following sections.

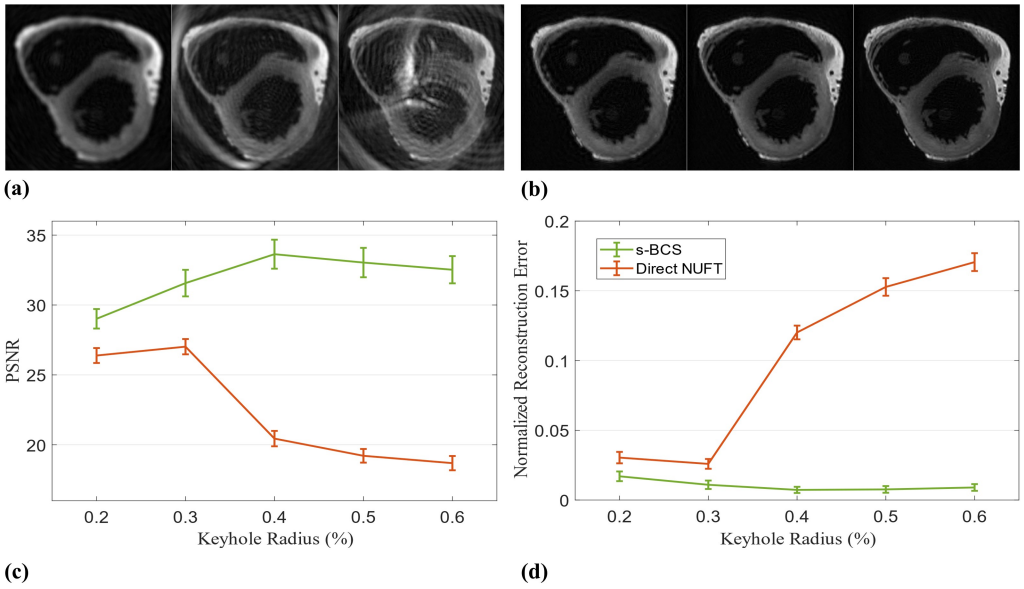


Figure 6.6: Influence of *keyhole* radius (α): A diffusion-weighted image reconstructed with (a) gridding and (b) sparse BCS using the proposed SPIK sampling with $\alpha = 0.2$, 0.4, and 0.6, respectively from left to right. The PSNR and normalised reconstruction error (mean $\pm$ SD) compared between the gridding and the BCS reconstructions as the function of the *keyhole* radius are plotted in the bottom row (c,d). As seen from the plots, $\alpha = 0.4$ yields better performance with the BCS reconstruction.

6.4.2 Algorithm Convergence

We examined the convergence of the optimisation algorithm, as illustrated in Figure 6.7. The reduction in the cost function, specified by Equation 6.4 is plotted against the number of iterations. At lower values of β_1 and β_2 parameters, the algorithm rapidly converges to a solution of 6.5. However, this yields a lower accuracy and falling to local minima due to the fact that the constraints in Equation 6.5 are not fully

satisfied. Increasing these parameters ensures that the constraints are satisfied but leads to slower convergence. Employing a continuation scheme, where β_1 and β_2 are progressively increasing from low values, significantly improves the convergence rate while maintaining accuracy.

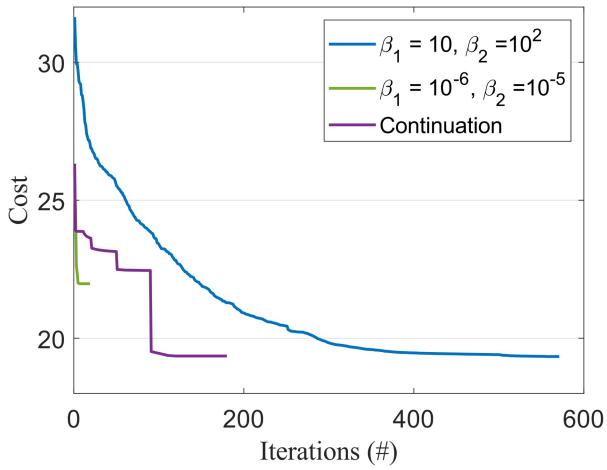


Figure 6.7: Convergence of the proposed continuation scheme as a function of β_1 and β_2. The evolution of the cost function specified by Equation 6.4 as a function of the iterations is shown. Note that the convergence of the algorithm is very slow if these parameters are chosen as high values, which is needed for the constraints in Equation 6.5 to be satisfied. In contrast, the algorithm converges very fast, when these parameters are set to low values. However, the solution of 6.5 is a poor approximation for the solution of 6.4. We observed that by properly selecting a continuation scheme, it is possible to significantly improve the convergence rate, while maintaining the accuracy.

6.4.3 Comparison of Reconstruction Schemes

Figure 6.8 compares the reconstruction performance of different approaches. Reconstructed images of one arbitrary DW direction from the ex vivo human cDTI data is obtained with different reconstruction schemes using the proposed SPIK sampling. The corresponding absolute error images and the SER values compared to the reference are also presented. The intensity profiles along the middle line are plotted in Figure 6.9. CS results in oversmoothing of image edges, reflected in larger differences in the myocardium boundaries in both the error maps and intensity profiles. The LR

+ S method was unable to completely suppress the residual streaking artefacts, leading to inaccuracies in the myocardium region and resulting in the largest errors in intensity profiles and qualitative error maps. In contrast, all BCS schemes outperform both methods, providing more accurate reconstructions, with the sparse BCS (case II) yielding the best performance.

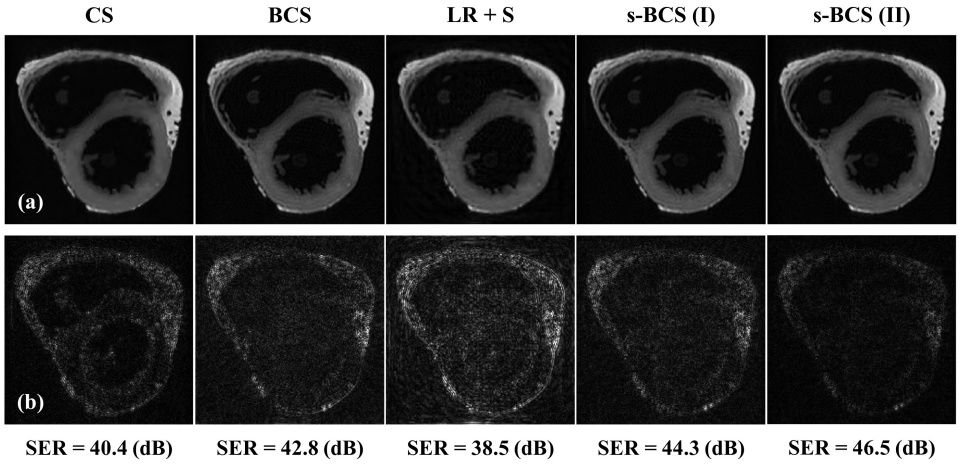


Figure 6.8: Performance Comparison of different reconstruction methods: (a) DW images obtained with different methods using the SPIK sampling, and (b) Absolute error maps relative to the reference, scaled by a factor of 5 ($|\mathbf{S}_{rec} - \mathbf{S}_{ref}| \times 5$). CS: classical compressed sensing; BCS: blind compressed sensing; LR + S: low-rank plus sparsity; s-BCS: proposed sparse BCS method with two implementations; SER: signal-to-error ratio.

SER plots comparing the reconstruction performance of various algorithms are displayed in Figure 6.10, highlighting the superior reconstruction quality achieved by the BCS schemes. The sparse BCS schemes provide higher SER values by an average of 4-6 dB compared to both LR + S and CS methods. We observed that the sparse BCS with ℓ_1-norm in the wavelet transform of the coefficients outperformed other BCS reconstruction schemes, providing the highest SER values.

The reconstructed FA and MD maps with their corresponding absolute error maps are shown in Figures 6.11 and 6.12, respectively, for different reconstruction schemes. The MSE scores are included in the Figures. The sparse BCS methods consistently outperformed other reconstructions in recovering the diffusion-related information and

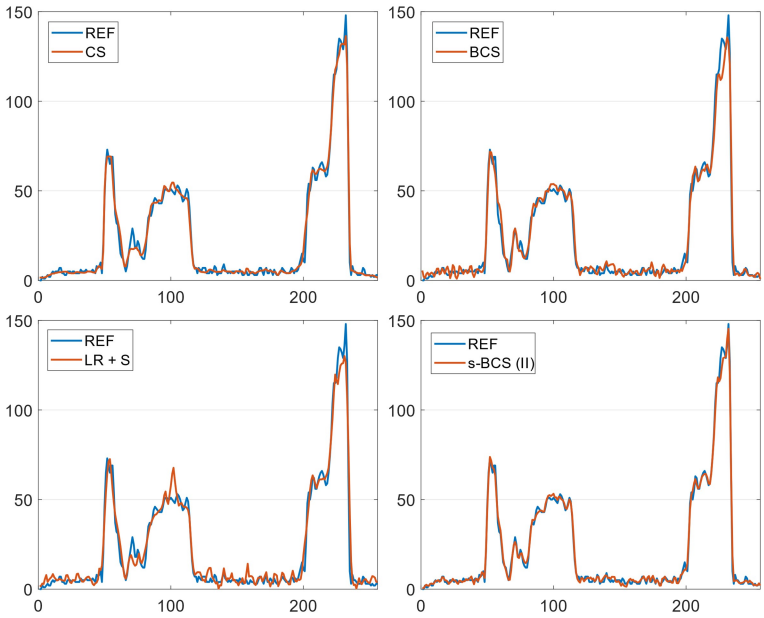


Figure 6.9: Intensity profiles corresponding to the middle line in Figure 6.8 for different reconstruction methods compared to the reference are shown.

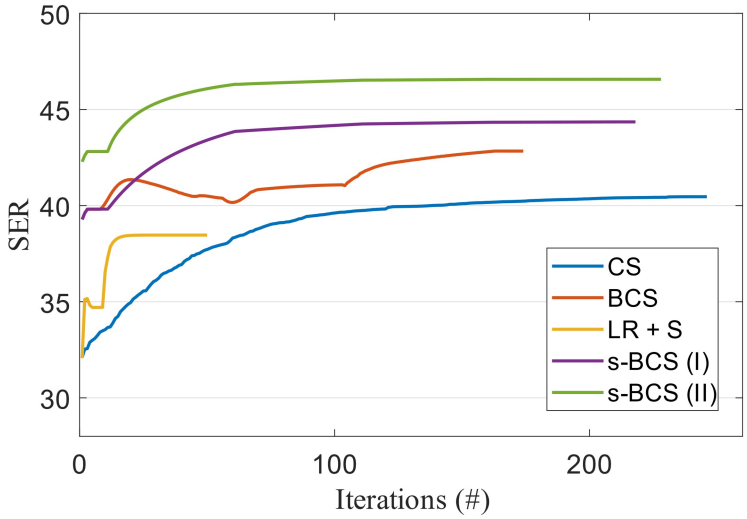


Figure 6.10: Signal-to-error (SER) scores as a function of number of iterations comparing the performance of different reconstruction methods.

resulted in lower MSE values.

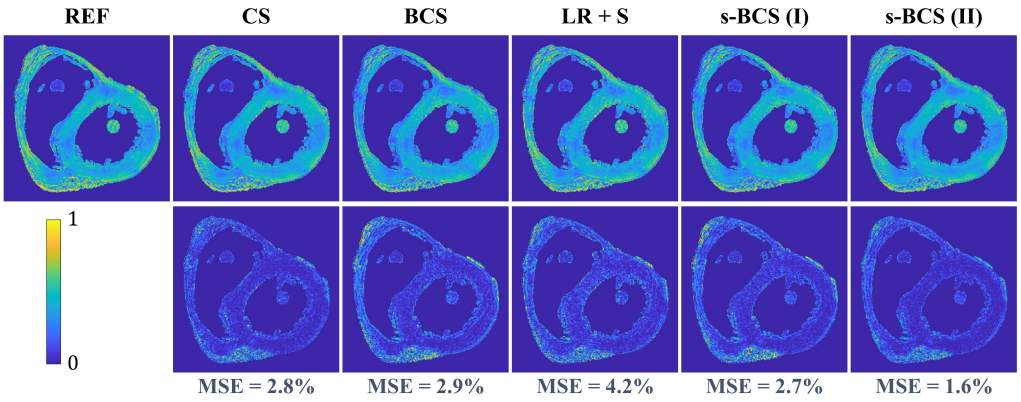


Figure 6.11: FA maps and the corresponding absolute error maps with respect to the reference images, scaled by a factor of 5 ($|\mathbf{I}_{rec} - \mathbf{I}_{ref}| \times 2$), obtained from different reconstruction methods.

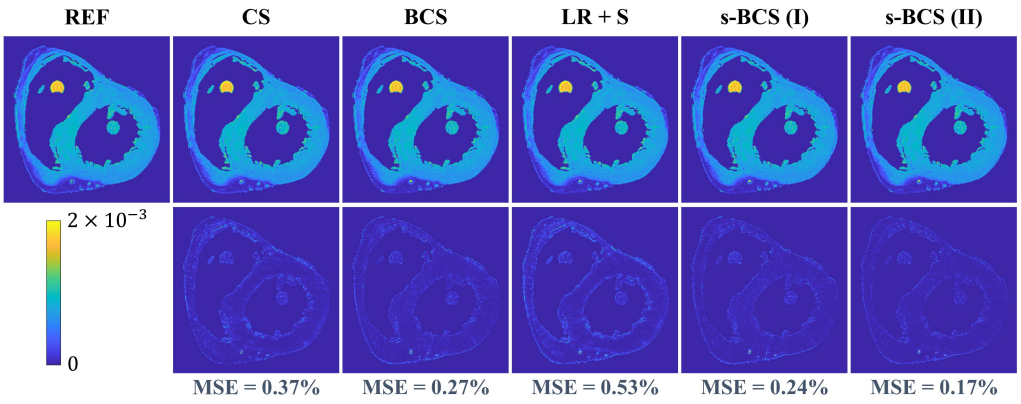


Figure 6.12: MD maps and the corresponding absolute error maps with respect to the reference images, scaled by a factor of 5 ($|\mathbf{I}_{rec} - \mathbf{I}_{ref}| \times 4$), obtained from different reconstruction methods.

Since the BCS reconstruction algorithm estimates the parametric bases and the spatial weights simultaneously from the undersampled data, the estimates are more accurate, as confirmed through results in these Figures. In contrast, the LR method only updates the spatial weights and the dictionary is estimated from the SVD of the initial reconstruction. We observed in Figure 6.6, the initial reconstruction contains

aliasing artefacts, particularly when higher acceleration factors are used, as is the case with our proposed SPIK sampling.

6.4.4 Comparison of Sampling Schemes

Figure 6.13 compares the reconstruction results obtained from the multi-shot spiral sampling with different number of interleaves N_s and the SPIK sampling. The standard stacks of multi-shot spirals with $N_s = 8$, 5, 3, and 2 were evaluated against the proposed SPIK sampling. One DW image, FA, MD, and HA maps with their corresponding error images relative to the Cartesian reference data are shown. As expected, with higher number of interleaves $N_s \geq 5$ in the optimised spiral sampling, the sparse BCS yielded an accurate reconstruction, while decreasing the number of interleaves resulted in significant aliasing artefacts in the reconstructions and the subsequent DTI parameter maps. However, the efficient sampling of the required information through the proposed SPIK sampling led to the improved performance compared to the three- and two-shot spiral imaging. Quantitative comparisons of the stack of SPIK sampling with 3D multi-shot spiral sampling is shown in Table 6.2.

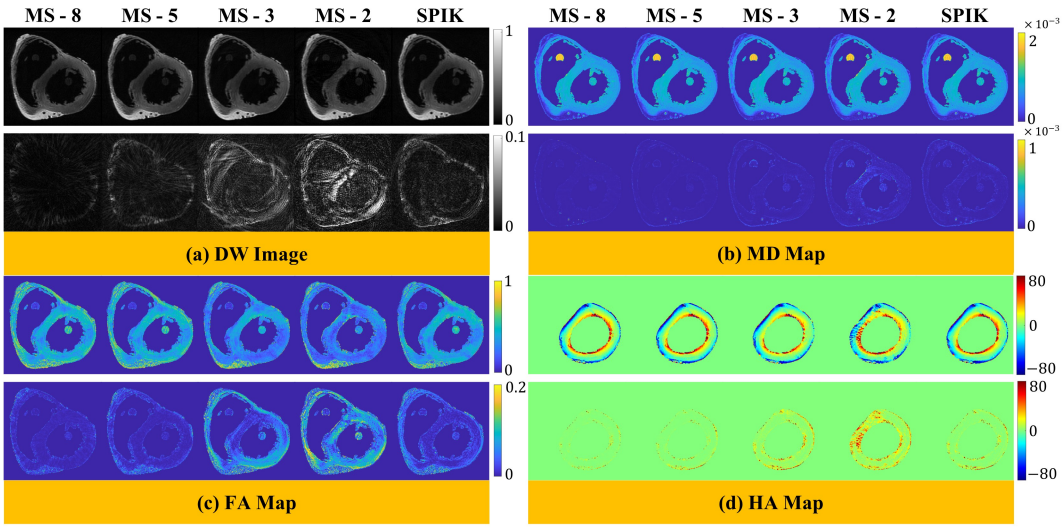


Figure 6.13: Comparison of the SPIK sampling and the optimised multi-shot spiral sampling: Reconstructed DW images, FA, MD (mm^2/s), and HA (°) maps, as well as the corresponding difference maps relative to the reference images, are shown.

Table 6.2: Quantitative comparison of the stack of SPIK sampling and the optimised multi-shot spiral sampling with different number of interleaves (N_s)

	MSE_{rec}	PSNR	RMSE_{FA}	RMSE_{MD}	HAT
$N_s = 8$	0.049	42.5	0.36	0.08	-0.98± 0.17
$N_s = 5$	0.064	41.1	0.86	0.09	-1.00± 0.18
$N_s = 3$	0.59	32.9	11.65	0.15	-0.90± 0.33
$N_s = 2$	1.19	30.3	20.15	0.6	-0.89± 0.45
$SPIK$	0.44	38.6	1.64	0.10	-0.95± 0.21

6.4.5 Dictionary Size

The relationship between the number of diffusion weighted signals in the dataset and the size of the dictionary or the number of dictionary atoms is evaluated in Figure 6.14. It is observed that when the number of basis functions in the dictionary is close to the number of diffusion signals, the estimation of parametric basis functions becomes poor, leading to degraded reconstruction quality and higher errors. Increasing the number of basis functions improves reconstruction accuracy, but at the cost of greater computational demands, especially when the number of DW directions is high. Our findings suggest that using approximately $2 \times N$ basis functions, where N is the number of DW directions in the dataset, provides a good compromise between reconstruction performance and computational efficiency.

6.4.6 Comparison on the MRXCAT cDTI phantom

Optimal regularisation parameters

We acknowledge that the approach of selecting the regularisation parameters based on fully sampled reference data is not practical in real-world applications where such reference data is unavailable. In such cases, heuristic approaches like the L-curve strategy [204] can be used. Here, the optimal regularisation parameters τ_1 and τ_2 were chosen using this method as shown in Figure 6.15. The values of 0.001 and 0.05 were respectively set for τ_1 and τ_2 to recover the MRXCAT cDTI data. However, a detailed discussion of more advanced approaches falls beyond the scope of our study.

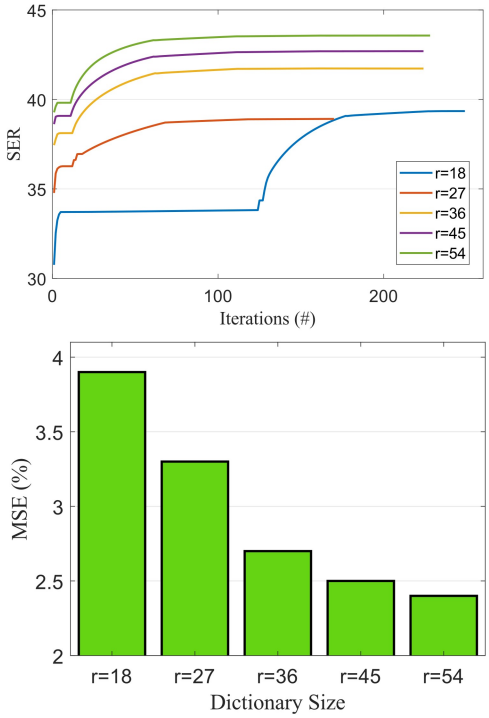


Figure 6.14: Evaluation of the performance of the proposed s-BCS method using varying numbers of dictionary atoms (r): (Top) Convergence is depicted as signal-to-error (SER) versus the number of iterations, and (Bottom) reconstruction accuracy is presented as the mean-square-error (MSE).

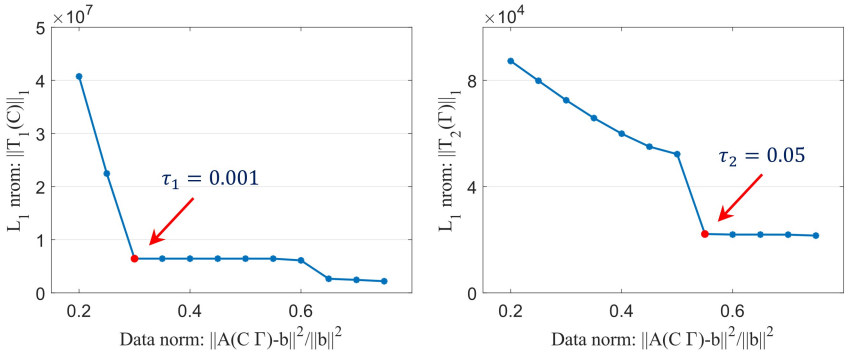


Figure 6.15: Choice of regularisation parameters τ_1 and τ_2. The parameters were optimised using the L-curve strategy. We vary τ_1 and τ_2 and plot the data consistency error against the ℓ_1 penalty. $\tau_1 = 0.001$ and $\tau_2 = 0.05$ were chosen as the regularising parameters for the MRXCAT cDTI data.

Breath-hold cDTI data

Representative three short-axis slices of one DW image and the corresponding diffusion parameter maps of the MRXCAT cDTI data obtained with the proposed approach are shown in Figure 6.16. Reference MRXCAT cDTI data were simulated using the STEAM sequence with one b_0 and 9 DW directions in the breath-hold condition. The reference Cartesian k-space data were resampled onto the SPIK sampling and then reconstructed using the sparse BCS method with ℓ_1 norm on $\mathbf{C}$. It can be observed that the proposed SPIK sampling could efficiently capture the required diffusion-related information, while the sparse BCS method, which relies on estimating the dictionary and the spatial coefficients directly from the data, could successfully recover the DW images with preserved details and reduced artefacts.

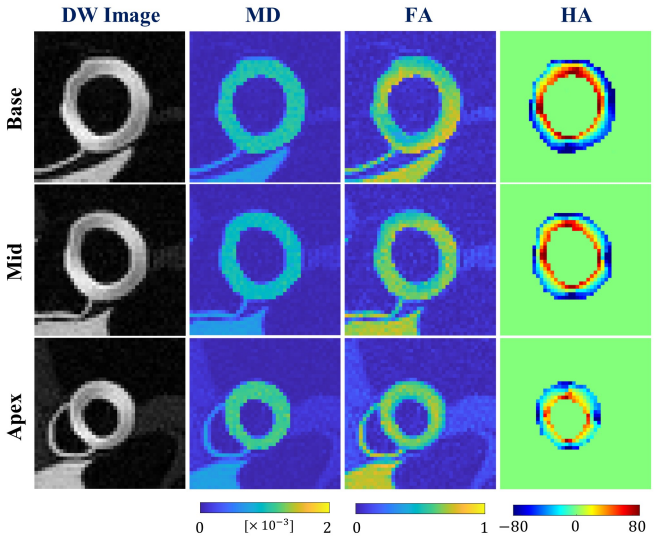


Figure 6.16: Reconstructed MRXCAT cDTI data simulated using the proposed sparse BCS (case I). The reference MRXCAT cDTI data were simulated with the STEAM sequence and along 9 DW directions in the breath-hold condition. Three base-to-apex short-axis slices are shown with corresponding DTI metrics.

Figure 6.17 shows one arbitrary DW image in three short-axis slices reconstructed with the sparse BCS scheme with ℓ_1 norm on the wavelet coefficients of the spatial weights $\mathbf{C}$ (case II). Absolute error images (5x-scaled) and the intensity profiles along the yellow line are also shown. To assess the accuracy in determining these profiles, we

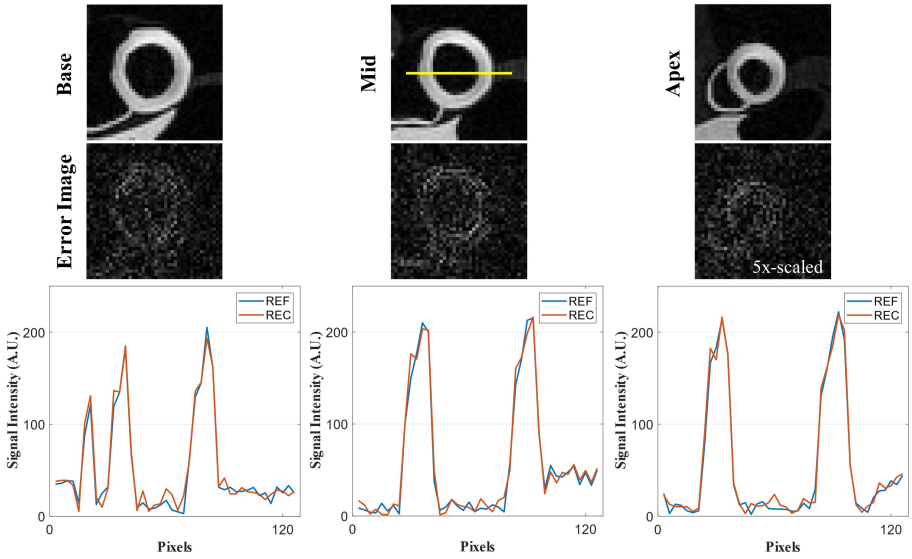


Figure 6.17: Reconstructed MRXCAT cDTI data using the proposed sparse BCS (case II) reconstruction. Three base-to-apex short-axis slices are shown with error images ($|\mathbf{S}_{rec} - \mathbf{S}_{ref}| \times 5$) and the intensity profiles along the yellow line.

also plotted the corresponding intensity profiles of the reference images. The close to accurate intensity profiles in the myocardium regions show the ability of the proposed approach to capture the diffusion information and to greatly recover the variations of the signal through reconstruction. This is also reflected in the noise-like residual images. Any inaccuracies here, could lead to false analysis in the subsequent diffusion tensor quantification.

Quantitative analysis of the proposed method in reconstructed DW images are reported in Table 6.3. Table 6.4 shows the RMSE of FA and MD maps in three slices and the global HAT compared with the reference parameter maps.

Free-breathing cDTI data

Figure 6.18 displays qualitative comparisons of the NUFT and the proposed sparse BCS reconstruction obtained from MRXCAT cDTI data in free-breathing condition. The b_0 image is simulated with a low SNR of 15. The NUFT reconstruction shows a high degree of blurring as the respiratory motion leads to the displacement of the k-space centre and thus is manifested as blurring artefacts in the reconstruction of spiral data.

Table 6.3: Quantitative analysis of the proposed approach using breath-hold MRXCAT cDTI data for three base-to-apex short axis slices. Mean values averaged over all DW images are reported.

	PSNR (dB)	DW_{RMSE}	SI
Base	29.1	0.043	0.9097
Mid	28.4	0.044	0.9037
Apex	29.5	0.041	0.9171

Table 6.4: Quantitative analysis of the proposed approach in DTI parameter maps using breath-hold MRXCAT cDTI data for three base-to-apex short axis slices

	FA_{MSE}	MD_{MSE} (mm^2/s)	HAT (HAT_{REF})
Base	0.97 %	0.43 %	-0.96 ± 0.33 (-1.05 ± 0.34)
Mid	1.39 %	0.64 %	-0.90 ± 0.18 (-0.92 ± 0.32)
Apex	0.95 %	0.31 %	-0.70 ± 0.13 (-0.77 ± 0.31)

However, since the BCS relies on data-specific features, the reconstruction exhibited significantly reduced blurring and improved image details. Because the reference data had a low SNR, we observed that the BCS scheme (case II) led to significantly more improved quality and suppressed noisy artefacts compared with the case I (the results are not included here). The PSNR and SSIM scores compared to the reference images are shown, highlighting the substantial quality improvement in the BCS method. The resulting FA and MD maps are also displayed. Reconstructed data were registered prior to tensor fitting. The RMSE of the FA and MD with respect to the Cartesian reference maps were 0.0327 and 2.5867×10^{-4} (mm^2/s), respectively, indicating the effectiveness of the proposed approach in free-breathing and even under low SNR regimes.

Figure 6.19 displays quantitative comparisons of the MRXCAT cDTI reconstructions for four simulations ($S_i;\ i = 1,..,4$) in free-breathing condition, comprising of different motion states. The mean ± standard deviation of the PSNR, RMSE, and

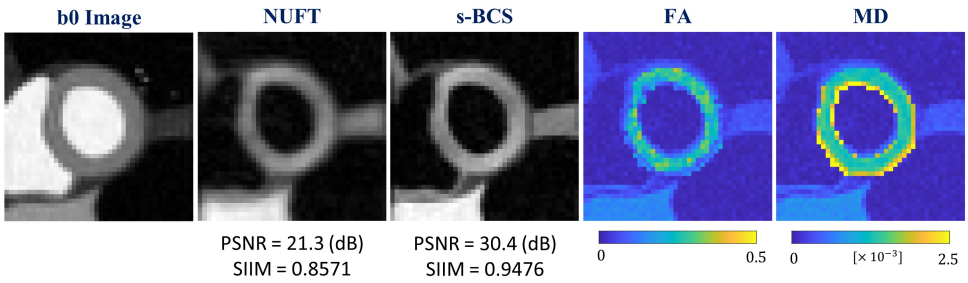


Figure 6.18: Reconstruction performance of the proposed approach in DTI parameters using noisy free-breathing MRXCAT cDTI data with one b_0 and 12 DW directions.

SSIM scores across all DW images are plotted. It can be observed that apart from the variations in the motion states of different datasets, the proposed sparse BCS reconstruction yielded a consistent quality. This is also reflected in the quality metrics obtained from the DTI parameter maps, as illustrated in Figure 6.20.

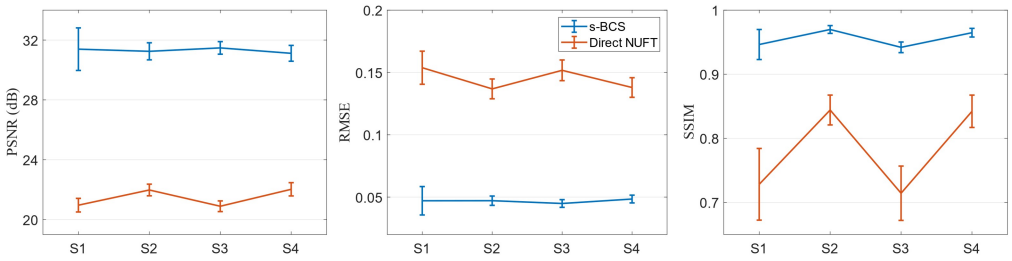


Figure 6.19: Reconstruction performance of the proposed approach against the direct NUFT reconstruction using four free-breathing ($S_i : i = 1 : 4$) MRXCAT cDTI data with one b_0 and 12 DW directions in different motion states. Mean ± SD across DW images are shown.

In Figure 6.21, we illustrate the SER of MRXCAT cDTI reconstructions plotted against the number of iterations for four simulations (S_i; $i = 1, .., 4$) in free-breathing condition, comprising of different motion states. The trend mirrors that of the real ex vivo comparisons. The proposed sparse BCS scheme consistently results in SER improvement in all four datasets.

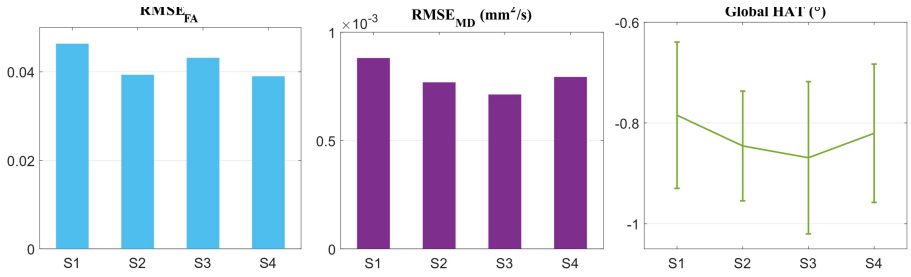


Figure 6.20: Reconstruction performance of the proposed approach in DTI parameters using four free-breathing ($S_i : i = 1 : 4$) MRXCAT cDTI data with one b_0 and 12 DW directions in different motion states. For the global HAT, mean ± SD across 25 transmural lines is plotted.

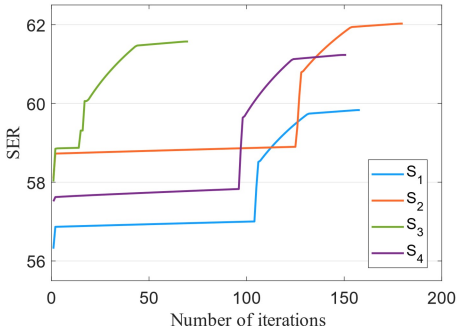


Figure 6.21: Reconstruction performance of the proposed approach in DTI parameters using noisy free-breathing MRXCAT cDTI data with one b_0 and 12 DW directions.

6.5 Discussion

In this chapter, we introduced a novel framework for accelerated high resolution cDTI. The proposed framework combines a new spiral sampling, referred to as SPIK, and a new reconstruction technique based on blind dictionary learning. The proposed model represents diffusion-weighted signals as sparse linear combinations of diffusion parametric basis functions obtained from a large dictionary. We formulated the reconstruction as a linear combination of a data consistency term and two transform sparsity constraints on the spatial coefficients and dictionary bases. We also developed an efficient algorithm to simultaneously estimate the sparse coefficient matrix and the dictionary from the data itself.

The proposed SPIK sampling supports high accelerations for cDTI while maintain-

ing sufficient coverage of both low spatial frequencies containing the diffusion contrast information and high spatial frequencies contributing to more details and increased spatial resolution. This sampling corresponds to an acquisition time of 30 ms for a single image slice with a high resolution of 1 mm. The significant consistency between the proposed SPIK sampling and the sparse BCS reconstruction resulted in well-preserved unique diffusion related information among the DW images and well-restoration of the the correlated boundary details. The overcomplete dictionary in the BCS formulation enabled accurate modelling of both specific and correlated image features by providing richer representation of data and through simultaneous learning of the dictionary and the coefficients.

The quantitative comparisons between different algorithms clearly demonstrate the ability of the proposed sparse BCS schemes to provide accurate reconstructions with minimal artefacts. Notably, the sparse BCS schemes consistently outperform the CS algorithm and the LR + S method. They showed about 4-6 dB SER improvement over the CS reconstruction and 6-8 dB SER improvement over the LR +S method. CS led to oversmoothing of the image details, especially in the edges (Figure 6.8), as it is unable to incorporate the correlations in the data, while the LR + S could not fully suppress the streaking artefacts resulting from the high undersampling rate in the proposed SPIK compared to the fully sampled data. This can be attributed to the fact that the LR method estimates the dictionary from the initial reconstruction, which is significantly blurry due to high undersampling and contains large streaking artefacts, and the dictionary bases are not updated through the reconstruction, thus leading to the learning of noisy spatial weights. However, in the BCS formulation, the dictionary and the spatial weights are simultaneously estimated from the data and are relearned in every iteration. Besides, the additional sparsity constraint on the dictionary bases and the coefficients avoids the unnecessary learning of small spatial coefficients, resulting in significantly improved reconstruction performance. Furthermore, the dictionary in BCS is much more extensive ($R \gg N$), and the parametric basis functions are not constrained to be orthogonal. The significantly larger number of basis functions in the BCS dictionary led to a notably better approximation of the diffusion-weighted signal, particularly for datasets with substantial inter-frame motion. A similar trend was observed for the reconstructed FA and MD maps. These visual observations are supported by MSE scores in Figures 6.11 and 6.12, with the corresponding maps obtained from

the Cartesian data taken as the reference.

We also demonstrated the superiority of the sparse BCS approach over the original BCS formulation, which relies on the sparsity of the coefficients and the unit Frobenius norm on the dictionary. We observed that the sparse BCS with an ℓ_1 norm on the wavelet coefficients yielded better reconstruction performance with the SPIK sampling, specifically when used with the low SNR free-breathing data.

Furthermore, we compared the performance of the multi-shot spiral acquisition with varying number of interleaves, as an indicator of the acceleration factor, with the proposed SPIK sampling. As expected, the reconstructions obtained from reduced number of interleaves ($N_s < 5$) exhibited noticeable aliasing artefacts, especially in the myocardium region of interest (see Figure 6.13), leading to a pronounced absolute error maps when compared with the reference. This also resulted in a high RMSE of ≥ 10 in the FA maps and distorted HA maps with reduced global HAT values. When increasing the number of interleaves $N_s \geq 5$, the reconstruction performance is significantly improved and the RMSE in the FA maps reduced to < 1 with improved mean PSNR score of ≥ 40. The SPIK sampling maintains a relatively high mean PSNR score of 38.6 in the reconstructed DW images and a relatively low RMSE of 1.64 in the FA map, which is comparable to the corresponding values obtained from the 5-shot spiral sampling. The lower the number of spiral interleaves, the faster the scan time. However, conventionally reducing the number of interleaves would result in significant image quality degradation, inaccurate diffusion tensor estimation and/or poorer spatial resolution. However, the SPIK sampling allows for efficient reduction of the scan time and the associated challenges in multi-shot spiral imaging, while maintaining the high spatial resolution and preserving the image quality and accurate diffusion tensor estimation. We observed that the size of the dictionary is closely related to the number of diffusion-weighted signals in the dataset. When a small number of DWIs is used, large dictionary sizes led to highly sparsified coefficients, which resulted in degraded reconstructions.

We performed the reconstruction of each image slice separately, however, future work involves the extension of the sparse BCS reconstruction to three dimensions. The parallelisation of the 3D algorithm over the number of slices using multiple cores, especially when dealing with large 3D matrices can help to accelerate the reconstruction. The method is easily extendable to multi-coil settings, allowing for further acceleration through parallel imaging techniques.

Similar to our findings with the ex vivo cDTI data, we observed significant performance improvement of the proposed sparse BCS method in reconstruction of the MRXCAT cDTI data. Results displayed in the three slices of the MRXCAT cDTI data (Figure 6.16) using the proposed SPIK sampling showed well-preserved image quality, especially in the myocardium LV. We found that both sparse BCS schemes provided comparable performance in reconstructing the ex vivo and breath-hold MRXCAT cDTI data, while the sparse BCS with an ℓ_1 wavelet norm on the coefficients led to improved reconstruction quality when dealing with free-breathing and low SNR MRXCAT cDTI data.

The proposed SPIK sampling and sparse BCS approach yielded similar and consistent image quality for simulations with varying respiratory motion states, as supported by the quantitative quality metrics measured on both the DW images (Figure 6.19) and the DTI parameter maps (Figure 6.20), using the Cartesian images as a reference. Capturing the most important part of the signal contrast in this technique through the *keyhole* spiral shot contributes to higher accuracy in the resulting diffusion maps. The high-resolution spiral shot allows for learning of the structural details, which are largely shared among the DWIs, instead of learning the random missing data. SPIK's continuous update of high spatial frequencies ensures that uncorrelated sections of the peripheral k-space are limited to a few number of scans. Future work will involve investigating efficient allocation of different diffusion directions to the sparse k-space of the SPIK sampling. As indicated in Chapter 4, an efficient sampling scheme needs to be optimised with respect to the individual DW directions. SPIK is a self-navigator trajectory. Since the navigators are collected simultaneous with the data, this will improve the SNR and their phase is consistent with the data. Further elaboration on the practical utility of the proposed SPIK sampling remains for future investigations.

The proposed scheme learns parametric basis functions directly from the data, without any assumptions about the spatial structure, allowing it to exploit the correlations in the data even in non-local regions. This property makes the proposed framework applicable to various diffusion imaging schemes, including the multi-b value DW imaging. Specifically, the proposed scheme holds significant promise for applications requiring a larger number of parametric images and more sophisticated models, particularly in scenarios where partial volume effects need to be addressed. Its ability to jointly estimate both the sparse coefficients and the subject-specific dictionary

from undersampled data makes it well-suited for these settings, potentially enabling the acquisition of more comprehensive and detailed parametric maps while addressing issues related to motion, inhomogeneity artefacts, and other deviations from ideal conditions. Moreover, with the enhanced spatial resolution offered by the SPIK sampling, it becomes more feasible to reliably obtain the 3D tractography of the whole myocardium. This capability opens up the potential for more robust investigations into the 3D myocardium microstructure.

The performance of the method could potentially be improved by utilising inpainting methods to estimate the non-acquired k-space samples from the acquired volumes before the reconstruction is performed.

6.6 Summary

This chapter proposed a novel sampling method, referred to as SPIK, combined with a data-adaptive reconstruction method for whole-heart cDTI with high spatial resolution and within the practical scan time range. The proposed sparse blind compressed sensing (s-BCS) approach utilised an overcomplete dictionary to model the DW images, allowing for sparser representations and enhanced reconstruction quality. The proposed method was retrospectively validated using both real ex vivo cDTI data and various realistic numerical phantom simulations, resembling in vivo cDTI setting. It is important to note that in this approach, we are simultaneously estimating both the sparse coefficients $\mathbf{C}$ and the subject-specific dictionary $\mathbf{\Gamma}$ directly from the undersampled k-space data acquired across the parametric dimension q. Therefore, the dictionary is tailored to each specific subject/sequence, ensuring that any deviations from the true parametric encoding, such as subject motion, field inhomogeneity, and chemical shift artefacts, are effectively learned by the basis functions in the dictionary.

Chapter 7

Conclusion And Future Work

Diffusion MRI (dMRI) is a valuable method for non-invasive quantification of microstructural information by probing water molecule displacements within biological tissues. Cardiac diffusion tensor imaging (cDTI) has gained popularity over the past two decades for its unique ability to characterise myocardial microstructure in both healthy and diseased states. The method holds significant clinical potential as it enables the identification of defects that lead to structural remodelling, providing valuable insights into various cardiac diseases. Despite many advantages, in vivo cDTI technique is currently not widely used in a clinical setting due to major technical limitations, including:

(i) Bulk motion artefacts,

(ii) Distortion/susceptibility artefacts,

(iii) Limited spatial resolution,

(iv) Limited spatial coverage,

(v) Long scan times ($\geq$ 20 minutes), and

(vi) Low signal-to-noise ratio.

This dissertation aimed at addressing some of these technical challenges with an ultimate aspiration to yield a clinically translatable in vivo cDTI technique. To address challenges (iii), (iv) and (v), we developed optimised spiral sampling patterns for cDTI within the framework of compressed sensing (CS), with the goal of achieving high spatial resolution and whole-heart coverage in comparable acquisition times to state-of-the-art 2D techniques. Whole-heart coverage would be advantageous for more accurate microstructural quantification of the myocardial syncytium and for eliminating misregistration errors in subsequent diffusion parameter maps. To fully harness the gain in efficient capturing of the diffusion-related information through the proposed sampling schemes, we introduced new reconstruction methods and efficient solving algorithms capable of utilising spatial and directional correlations within the acquired data.

7.1 Contributions

First, building upon the inherent advantages of spiral acquisitions, we developed and optimised a new parametrisation for the sampling density function well-adapted to the diffusion-weighted signal in the heart. The proposed multilevel parametrisation, as a function of the k-space radius (k_r) locations, allowed for more flexible design options and efficient spreading of samples along one spiral arm to maximise capturing the required information while covering a larger extent in k-space, thus enabling higher spa-

tial resolution imaging. An empirical study was conducted to provide qualitative and quantitative optimisation of the parameters. The analyses offer insights into designing application-specific undersampling patterns within a given setup of resolution and scan time and provide guidelines on how different parameters would affect the quality of reconstruction, aiding in the selection of appropriate parameter values based on the desired features. The new parametrisation enables creating variable density spiral (VDS) trajectories with high fidelity and efficiency, leading to significant reduction in the total acquisition time with minimal artefacts. The resulting incoherent diffused artefacts mainly appeared outside the myocardium region-of-interest (ROI), making the proposed VDS designs suitable for enhanced CS reconstruction performance and for achieving high undersampling factors. With this design in a multislice approach, a high in-plane resolution of 1 mm with 10 slices could be acquired in 30 seconds, instead of 20 minutes for the fully-sampled Cartesian reference, considering the same T_R of 1000 ms. The concept of a three-domain sampling density function is well-suited to the new findings on multi-level compressed sensing sampling patterns [164] and, therefore, can be applicable for optimising sampling for any other body organ.

Second, we further advanced the use of this optimised sampling density function by extending the proposed VDS design to three dimensions. Specifically, we incorporated the variable density in the k_z direction to create an efficient z-variable stack-of-spirals trajectory and in designing a fully 3D conical scheme. Our experiments with 3D cDTI data demonstrated the superior efficiency of the proposed z-variable VDS sampling pattern compared to the commonly used Self-navigated interleaved spiral (SNAILS) trajectory. Additionally, our method outperformed two widely used Cartesian undersampling patterns in previous accelerated cDTI studies. Particularly in high undersampling rates, the proposed VDS images exhibited fewer artefacts compared to other techniques, resulting in superior image quality in the reconstructed diffusion parameter maps. For instance, the average RMS errors in FA maps were less than 4% at undersampling factors as high as 20. To fully exploit the information acquired through the proposed sampling schemes, we employed a bidomain joint sparsity reconstruction method capable of leveraging correlations between DW images in both TV and wavelet domains. We employed an efficient fast solving algorithm, integrated an adaptive noise estimation method into the reconstruction process, and pre-calculated the Fourier coefficients on a dense grid to address the high computational cost of CS reconstructions

for high-dimensional DW data. Typically, such data has an image size of $N \times N \times N$ acquired at N_q DW directions, resulting in a large dataset ($N^3 \times N_q$). The computational challenge becomes particularly significant when dealing with non-Cartesian data, which necessitates non-uniform fast Fourier transform (NUFT) operations. The work presented in chapter 5 was delivered as an oral presentation at the ESMRMB 2021 Annual Meeting and is being prepared as a journal paper for submission to IEEE Transactions on Biomedical Engineering:

- S. Golshani, et al. "Towards accelerating cardiac DTI with undersampled 3D spiral acquisitions and a compressed sensing reconstruction based on joint sparsity," *Magn. Reson. Mater. Physics, Biol. Med.,* vol. 34, no. 1, pp. 61-62, 2021.
- S. Golshani, et al. "A Fast and Efficient Sampling/Reconstruction Framework for high-resolution 3D Cardiac Diffusion MRI", to be submitted to *IEEE TBE.*

Third, we introduced a novel sampling scheme, referred to as SPIK (*spiral with keyhole*), which involves sharing portions of the k-space between different DW images. SPIK sampling takes advantage of the fact that the central part of k-space determines the coarse appearance of an image, while the peripheral parts provide edge definition. This method incorporates two spiral shots with the proposed tailored VD design, one to acquire the central k-space of every DW image and one with reversed density to capture the correlated high spatial frequencies across all DW images. We developed a sparse blind dictionary method and an efficient algorithm to utilise spatio-directional correlations within DWIs, assuming that the directional evolution of the voxel-wise diffusion signal can be represented as a linear combination of properly chosen directional basis functions. The method eliminates the need for prior knowledge about the sparsifying transforms and thus can be readily used in diffusion MRI in other organs. The method showed to be promising in accelerating the acquisition time of cDTI with minimal degradation of image quality even under free-breathing condition using in vivo human heart phantoms. The work presented in chapter 6 has been prepared as journal paper for submission to Magnetic Resonance in Medicine:

- S. Golshani, et al. "Improved High-Resolution Cardiac DTI with a Novel Data-Adaptive Sampling and Reconstruction Strategy", to be submitted to *Magn. Reson. in Med..*

Spiral trajectories are known to be sensitive to confounding factors such as off-resonance effects. In addition, interleaved or multi-shot acquisitions, particularly in diffusion MRI, are susceptible to phase inconsistencies between the shots. However, in our experiments, we took these considerations into account either in the sampling design or in the reconstruction step. For instance, the considerations of the off-resonance effects were addressed during the sampling design by constraining the readout duration, while possible phase inconsistencies between shots arising from motion or other imperfections where addressed during the reconstruction. In practice, if those considerations were not sufficient, additional approaches can be effectively utilised to compensate for any residual artefacts/impurities.

To demonstrate the applicability of our proposed sampling techniques in different imaging protocols, we retrospectively evaluated their performance across a wide range of real DW images and considering both clinical and preclinical settings. The fidelity of the introduced approach was demonstrated using high-resolution ex vivo human and rat heart DTI datasets, as well as realistically simulated in vivo human cDTI images. Using the optimal gradient toolbox (OptGrad V2.0), we incorporated the gradient system specifications in the proposed designs, ensuring that the resulting trajectories and gradient waveforms will be implementable on actual MRI scanners. In the preclinical setting, where high-performance gradients with high amplitudes and greater slew rates are available, higher acceleration factors could potentially be achieved. The saved scan time can be incorporated in increasing the number of diffusion encoding directions to enable investigating more complex diffusion models than the diffusion tensor model. These investigations can provide context for more research in this area and facilitate the development of potentially more specific microstructural biomarkers, which has not been widely explored yet.

7.2 Limitations and Future Work

This study primarily focused on practical design considerations, given the vast flexibility in possible spiral trajectory designs, k, q sampling strategies, and reconstruction methods. Identifying the optimal parameters for an efficient spiral pulse sequence and reconstruction technique for high-resolution whole-heart cDTI is essential before conducting a formal clinical evaluation of the technique. While this work establishes the theoretical and conceptual foundation, its implementation and prospective acquisition

of in vivo cDTI, as well as further extensions to clinical applications and comparisons with established methods to demonstrate the efficiency of the proposed approach for improving spatial resolution and coverage, are subjects of our future work. The implementation of the proposed VD sampling schemes, which have taken into account the limitations of the gradient hardware, will be realisable with minor modifications on an MR scanner. However, it is important to be aware of effective techniques and their shortcomings that may be required to address additional practical challenges. As mentioned, spiral-like acquisitions require special care to account for off-resonance, gradient imperfections, and concomitant field effects, which may degrade image quality if left uncorrected. Even with considerations for gradient and slew rate limitations, the actual spiral trajectory realised can deviate from the intended one when the gradient waveform contains frequencies beyond the gradient system's bandwidth limit. The proposed method can also be readily combined with other acceleration techniques such as partial Fourier sampling, non-Cartesian parallel imaging, and reduced encoding techniques to further improve the scan time.

Additionally, potential improvements can be made to the proposed BCS reconstruction. Here, we considered a single dictionary to incorporate spatio-directional correlations within DWIs with atoms containing the diffusion parametric bases; however, the method can be easily extended to incorporate separate dictionaries for the angular and spatial dimensions, thus providing sparser representations. A preliminary work in this topic was presented in an ISMRM abstract (2024).

We have also demonstrated that the optimised sampling density parameters are dependent on the diffusion direction. Therefore, one aspect of the future work will be DW-dependent optimisation of the sampling density. It has been shown that for more robust estimation of the DTI parameters, acquiring more distinct diffusion directions is preferable to repeating a limited set of directions [49]. Hence, this optimisation can provide insights into the optimal selection of the diffusion encoding scheme (including the number of diffusion-weighting measurements and their directions) from a signal perspective. Simulation studies could be conducted to determine the optimal sampling profile for different diffusion-encoding directions. This would enhance the directional point spread function, thereby leading to improved CS reconstructions. With increasing interest in learning-based optimisation methods, joint optimisation of the reconstruction and sampling can be performed in cDTI applications, making it particularly

relevant within our proposed framework.

In a recent cardiac dMRI study [157], a diffusion preparation approach based on balanced steady-state free precession (SSFP) technique has been validated in in vivo cDTI, enabling multi-shot readouts. The method could reduce the acquisition time to approximately 1.5 minutes per slice. The primary advantage of this approach is the flexibility to independently optimise the diffusion encoding and the imaging readout. Therefore, an interesting future study would be to combine our 3D trajectories with this diffusion encoding scheme, which has the potential to significantly reduce the scan time and improve the spatial coverage. This can potentially provide better SNR and isotropic spatial resolution with an acceptable scan time.

In conclusion, our work demonstrates that the proposed VDS sampling schemes, when combined with dedicated CS reconstruction methods, hold significant promise for cDTI applications, allowing for the acquisition of high-quality DW images with whole-heart coverage (30 slices) and high in-plane resolution (1 mm) at heart rates of up to 100 beats per minute (bpm). Finally, a significant development of this work lies in its potential contribution to enhancing the high resolution characterisation of laminar myocardial structure and paving the way for improved diagnosis and understanding of global remodelling of myocytes within the left ventricle (LV) due to cardiac pathologies. This reduction in scan time may open up new applications for cardiac diffusion imaging, allowing for the characterisation of global cardiomyocyte orientations across the LV and potentially revealing subtle differences in complex microstructure in patients with conditions such as hypertrophic cardiomyopathy. Furthermore, whole-heart coverage with isotropic resolution datasets enabled by the proposed approach may benefit biomechanical mathematical models that rely on accurate myocyte orientations as inputs to finite element modeling.

Appendix A

Spiral Data Simulation (Simulation Principles)

In practice, the received MRI signal can be significantly influenced by various experimental inaccuracies. These include field inhomogeneities, object motion, inaccurate sequence timing, noise, and eddy currents. Identifying error sources solely from observed image artefacts is often challenging, particularly because different effects can lead to similar artefacts. Consequently, performing simulations of the MRI signal generation to replicate potential inaccuracies has proven highly beneficial. This approach enables systematic investigation of the response behaviour of acquisition and reconstruction techniques to a wide range of error sources. Moreover, simulations are valuable for understanding the imaging characteristics of different sampling schemes, which can be non-intuitive due to the involvement of Fourier transformation and rather complex for non-Cartesian patterns such as spiral.

Several strategies can be employed to simulate MRI experiments. The most general approach is based on utilising the Bloch equation (Equation 2.2), which describes how bulk magnetisation evolves over time in response to time-varying magnetic fields. This approach enables detailed analysis of all relevant mechanisms on a physical level, including the RF excitation process and macroscopic interference effects. In a realistic imaging scenario, all variables of the Bloch equation, especially the magnetic field strength when gradients are switched, vary locally. Therefore, it is necessary to compute an ensemble of independent magnetisation vectors at a sufficiently dense grid in the image space. The MRI signal is then obtained by summing over the transversal components of all vectors for each time point. However, this simulation procedure is computationally intensive due to the requirement for a high number of vectors to accurately capture realistic interference effects, such as signal cancellation from dephasing.

Hence, in this work, a different simulation approach was employed. It is based on the higher-level observation from Equation 2.14 that the MRI signal corresponds to the continuous Fourier transform of the object. Therefore, if the analytical Fourier transform of a numerical phantom is known, a simulated MRI signal can be obtained by computing the k-space position for each time point and evaluating the Fourier transform at these positions. The analytical Fourier transform of a rectangle is given by:

$$f_{rect}(k_x, k_y) = 2\pi \cdot s_x \cdot s_y \cdot sinc(s_x \cdot k_x) \cdot sinc(s_y \cdot k_y) \tag{A.1}$$

where s_x, s_y denote the size of the rectangle in the x- and y-direction. Further, the

Fourier transform of an ellipse is defined by:

$$f_{circ}(k_x, k_y) = \frac{a \cdot J_1(b \cdot \sqrt{(a/b \cdot k_x)^2 + k_y^2})}{\sqrt{(a/b \cdot k_x)^2 + k_y^2})} \tag{A.2}$$

where $J_1(x)$ is the first-order Bessel function of the first kind, and a, b denote the short and long axes of the ellipse. Due to the linearity of the Fourier transformation, we can derive an analytical Fourier transform of any combination of rectangles and ellipses by simply superimposing these functions, and a displacement of a single element can be realised with a linear phase modulation. Thus, arbitrary numerical phantoms can be constructed for the simulations. While this concept allows for fast computation of the signal, it has limitations. For instance, it cannot model the effects of RF pulses, thereby hindering the study of excitation or refocusing profiles of different pulse waveforms. Additionally, it provides a global signal, making it difficult to directly analyse the effect of local changes, such as field strength variations, as the signal formation process is not covered by the simulation procedure. However, this limitation can be partially overcome by leveraging the linearity of the Fourier transformation. For instance, deviations in the field strength cause a phase drift in the signal from the respective area. By adding a temporal phase drift to the Fourier transform of an object element, like a single circle, and combining it with the Fourier transforms of other object elements, the MRI signal corresponding to a magnetic field with a discrete jump at the location of the circle can be obtained. Although field strength deviations are typically smooth in practice, such analysis aids in understanding the resulting image artefacts qualitatively. Indeed, many effects observed in experimental settings, such as relaxation, saturation, or non-rigid object motion, can be analysed similarly using dedicated numerical phantoms. Hence, the simulation concept offers high flexibility and proves to be valuable during the development of new techniques presented in this work.

When a non-Cartesian k-space sampling pattern is used, it becomes necessary to interpolate the Fourier transform from the grid to the desired k-space positions, introducing some degree of inaccuracy into the simulation. While this effect can be mitigated by increasing the resolution of the source image to make the grid points denser, the approach must be used with caution to prevent interference with the artefacts under investigation.

Bibliography

[1] C. W. Tsao, A. W. Aday, Z. I. Almarzooq, A. Alonso, A. Z. Beaton, M. S. Bittencourt *et al.*, "Heart disease and stroke statistics - 2022 update: A report from the American Heart Association," *Circulation*, vol. 145, no. 8, pp. e153–e639, 2022.

[2] N. Townsend, D. Kazakiewicz, F. Lucy Wright, A. Timmis, R. Huculeci, A. Torbica *et al.*, "Epidemiology of cardiovascular disease in Europe," *Nature Reviews Cardiology*, vol. 19, no. 2, pp. 133–143, 2022.

[3] British Heart Foundation, *Heart & Circulatory Disease Statistics 2022*, 2022. [Online]. Available from: https://www.bhf.org.uk/what-we-do/our-research/heart-statistics/heart-statistics-publications/cardiovascular-disease-statistics-2022

[4] R. Luengo-Fernandez, M. Walli-Attaei, A. Gray, A. Torbica, A. P. Maggioni, R. Huculeci *et al.*, "Economic burden of cardiovascular diseases in the European Union: a population-based cost study," *European Heart Journal*, vol. 44, no. 45, pp. 4752–4767, 2023.

[5] The Texas Heart Institute, *Heart Information Center/Heart Anatomy*, 2023. [Online]. Available from: https://www.texasheart.org/heart-health/heart-information-center/topics/heart-anatomy/

[6] J. A. Shar, "Unraveling the Etiologies of Discrete Subaortic Stenosis: A Focus on Left Ventricular Outflow Tract Hemodynamics," Ph.D. dissertation, Wright State University, 2021.

[7] S.-J. Yoo and D. G. Perrin, *Radiology Key*, 2016. [Online]. Available from: https://radiologykey.com/normal-cardiovascular-anatomy-for-imaging/

[8] S. Göktepe, O. J. Abilez, K. K. Parker, and E. Kuhl, "A multiscale model for eccentric and concentric cardiac growth through sarcomerogenesis," *Journal of Theoretical Biology*, vol. 265, no. 3, pp. 433–442, 8 2010.

[9] S. A. Sheffield, *Cardiac muscle tissue: structure and function*, 2022. [Online]. Available from: https://www.getbodysmart.com/heart-anatomy/cardiac-muscle-tissue/

[10] I. J. LeGrice, B. H. Smaill, L. Z. Chai, S. G. Edgar, J. B. Gavin, and P. J. Hunter, "Laminar structure of the heart: ventricular myocyte arrangement and connective tissue architecture in the dog," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 269, no. 2, pp. H571–H582, 1995.

[11] T. Arts, K. D. Costa, J. W. Covell, and A. D. McCulloch, "Relating myocardial laminar architecture to shear strain and muscle fiber orientation," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 280, no. 5, pp. H2222–H2229, 2001.

[12] D. Trifunović-Zamaklar, I. Jovanović, J. Vratonjić, O. Petrović, I. Paunović, M. Tešić *et al.*, "The basic heart anatomy and physiology from the cardiologist's perspective: Toward a better understanding of left ventricular mechanics, systolic, and diastolic function," *Journal of Clinical Ultrasound*, vol. 50, no. 8, pp. 1026–1040, 2022.

[13] M. Takeuchi, T. Nishikage, H. Nakai, M. Kokumai, S. Otani, and R. M. Lang, "The assessment of left ventricular twist in anterior wall myocardial infarction using two-dimensional speckle tracking imaging," *Journal of the American Society of Echocardiography*, vol. 20, no. 1, pp. 36–44, 2007.

[14] E. Nagel, M. Stuber, B. Burkhard, S. Fischer, and M. Scheidegger, "Cardiac rotation and relaxation in patients with aortic valve stenosis," *European Heart Journal*, vol. 21, no. 7, pp. 582–589, 2000.

[15] J. Chen, W. Liu, H. Zhang, L. Lacy, X. Yang, S.-K. Song *et al.*, "Regional ventricular wall thickening reflects changes in cardiac fiber and sheet structure

during contraction: quantification with diffusion tensor MRI," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 289, no. 5, pp. H1898–H1907, 2005.

[16] K. A. Ammar, S. J. Jacobsen, D. W. Mahoney, J. A. Kors, M. M. Redfield, J. C. Burnett *et al.*, "Prevalence and prognostic significance of heart failure stages," *Circulation*, vol. 115, no. 12, pp. 1563–1570, 2007.

[17] J. L. Jefferies and J. A. Towbin, "Dilated cardiomyopathy," *The Lancet*, vol. 375, no. 9716, pp. 752–762, 2010.

[18] J. Ker, "A new phenotypic marker of hypertrophic cardiomyopathy," *The Lancet*, vol. 384, no. 9956, p. 1731, 2014.

[19] A. Das, C. Kelly, I. Teh, C. Nguyen, L. A. E. Brown, A. Chowdhary *et al.*, "Phenotyping hypertrophic cardiomyopathy using cardiac diffusion magnetic resonance imaging: the relationship between microvascular dysfunction and microstructural changes," *European Heart Journal - Cardiovascular Imaging*, vol. 23, no. 3, pp. 352–362, 2022.

[20] Society for Cardiovascular Angiography & Interventions, *Cardiomyopathy (Diseases of the Heart Muscle)*, 2022. [Online]. Available from: https://www.secondscount.org/condition/cardiomyopathy

[21] S. M. Factor, E. M. Okun, and E. S. Kirk, "The histological lateral border of acute canine myocardial infarction. A function of microcirculation." *Circulation Research*, vol. 48, no. 5, pp. 640–649, 1981.

[22] G. D. Buckberg, "Basic science review: The helix and the heart," *The Journal of Thoracic and Cardiovascular Surgery*, vol. 124, no. 5, pp. 863–883, 2002.

[23] D. D. Streeter, H. M. Spotnitz, D. P. Patel, J. Ross, and E. H. Sonnenblick, "Fibre orientation in the canine left ventricle during diastole and systole," *Circulation Research*, vol. 24, no. 3, pp. 339–347, 1969.

[24] J. D. Humphrey, *Cardiovascular Solid Mechanics*. New York: Springer, 2002.

[25] F. Torrent-Guasp, M. J. Kocica, A. F. Corno, M. Komeda, F. Carreras-Costa, A. Flotats *et al.*, "Towards new understanding of the heart structure and function," *European Journal of Cardio-Thoracic Surgery*, vol. 27, no. 2, pp. 191–201, 2005.

[26] G. Buckberg, A. Mahajan, S. Saleh, J. I. Hoffman, and C. Coghlan, "Structure and function relationships of the helical ventricular myocardial band," *The Journal of Thoracic and Cardiovascular Surgery*, vol. 136, no. 3, pp. 578–589, 2008.

[27] P. S. Jouk, Y. Usson, G. Michalowicz, and F. Parazza, "Mapping of the orientation of myocardial cells by means of polarized light and confocal scanning laser microscopy," *Microscopy Research and Technique*, vol. 30, no. 6, pp. 480–490, 1995.

[28] P. S. Jouk, A. Mourad, V. Milisic, G. Michalowicz, A. Raoult, D. Caillerie *et al.*, "Analysis of the fibre architecture of the heart by quantitative polarized light microscopy. Accuracy, limitations and contribution to the study of the fibre architecture of the ventricles during fetal and neonatal life," *European Journal of Cardio-Thoracic Surgery*, vol. 31, no. 5, pp. 915–921, 2007.

[29] R. M. Smith, A. Matiukas, C. W. Zemlin, and A. M. Pertsov, "Nondestructive optical determination of fibre organization in intact myocardial wall," *Microscopy Research and Technique*, vol. 71, no. 7, pp. 510–516, 2008.

[30] G. B. Sands, D. A. Gerneke, D. A. Hooks, C. R. Green, B. H. Smaill, and I. J. Legrice, "Automated imaging of extended tissue volumes using confocal microscopy," *Microscopy Research and Technique*, vol. 67, no. 5, pp. 227–239, 2005.

[31] R. L. Price, S. T. Haley, T. Bullard, J. Davis, T. K. Borg, and L. Terracio, "Confocal microscopy of cardiac myocytes." New York: Humana Press, 2014, pp. 185–199.

[32] D. E. Sosnovik and T. Geva, "Imaging the microstructure of the human fetal heart," *Circulation: Cardiovascular Imaging*, vol. 11, no. 10, 2018.

[33] P. Basser, J. Mattiello, and D. LeBihan, "MR diffusion tensor spectroscopy and imaging," *Biophysical Journal*, vol. 66, no. 1, pp. 259–267, 1994.

[34] P. Basser, J. Mattiello, and D. Lebihan, "Estimation of the effective self-diffusion tensor from the NMR spin echo," *Journal of Magnetic Resonance, Series B*, vol. 103, no. 3, pp. 247–254, 1994.

[35] S. Nielles-Vallespin, C. Mekkaoui, P. Gatehouse, T. G. Reese, J. Keegan, P. F. Ferreira *et al.*, "In vivo diffusion tensor MRI of the human heart: Reproducibility of breath-hold and navigator-based approaches," *Magnetic Resonance in Medicine*, vol. 70, no. 2, pp. 454–465, 2013.

[36] Z. Khalique, P. F. Ferreira, A. D. Scott, S. Nielles-Vallespin, D. N. Firmin, and D. J. Pennell, "Diffusion tensor cardiovascular magnetic resonance imaging: A clinical perspective," *JACC. Cardiovascular imaging*, vol. 13, no. 5, pp. 1235–1255, 2020.

[37] C. Mekkaoui, T. G. Reese, M. P. Jackowski, H. Bhat, and D. E. Sosnovik, "Diffusion MRI in the heart," *NMR in Biomedicine*, vol. 30, no. 3, 2017.

[38] E. W. Hsu, D. L. Buckley, J. D. Bui, S. J. Blackband, and J. R. Forder, "Two-component diffusion tensor MRI of isolated perfused hearts," *Magnetic Resonance in Medicine*, vol. 45, no. 6, pp. 1039–1045, 2001.

[39] D. Le Bihan, J. Mangin, C. Poupon, C. A. Clark, S. Pappata, N. Molko *et al.*, "Diffusion tensor imaging: Concepts and applications," *Journal of Magnetic Resonance Imaging*, vol. 13, no. 4, pp. 534–546, 2001.

[40] P. A. Helm, H. Tseng, L. Younes, E. R. McVeigh, and R. L. Winslow, "Ex vivo 3D diffusion tensor imaging and quantification of cardiac laminar structure," *Magnetic Resonance in Medicine*, vol. 54, no. 4, pp. 850–859, 2005.

[41] M. T. Wu, W. Y. I. Tseng, M. Y. M. Su, C. P. Liu, K. R. Chiou, V. J. Wedeen *et al.*, "Diffusion tensor magnetic resonance imaging mapping the fibre architecture remodelling in human myocardium after infarction," *Circulation*, vol. 114, no. 10, pp. 1036–1045, 2006.

[42] E. W. Hsu, A. L. Muzikant, S. A. Matulevicius, R. C. Penland, and C. S. Henriquez, "Magnetic resonance myocardial fiber-orientation mapping with direct histological correlation," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 274, no. 5, pp. H1627–H1634, 1998.

[43] D. E. Sosnovik, R. Wang, G. Dai, T. Wang, E. Aikawa, M. Novikov *et al.*, "Diffusion spectrum MRI tractography reveals the presence of a complex network of residual myofibers in infarcted myocardium," *Circulation: Cardiovascular Imaging*, vol. 2, no. 3, pp. 206–212, 2009.

[44] D. F. Scollan, A. Holmes, R. Winslow, and J. Forder, "Histological validation of myocardial microstructure obtained from diffusion tensor magnetic resonance imaging," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 275, no. 6, pp. H2308–H2318, 1998.

[45] L. Zhukov and A. Barr, "Heart-muscle fibre reconstruction from diffusion tensor MRI," In: *IEEE Visualization*. Seattle, WA, USA: IEEE, pp. 597–602.

[46] F. Carreras, M. Ballester, S. Pujadas, R. Leta, and G. Pons-Llado, "Morphological and functional evidences of the helical heart from non-invasive cardiac imaging," *European Journal of Cardio-Thoracic Surgery*, vol. 29, pp. S50–S55, 2006.

[47] J. C. Walker, J. M. Guccione, Y. Jiang, P. Zhang, A. W. Wallace, E. W. Hsu *et al.*, "Helical myofiber orientation after myocardial infarction and left ventricular surgical restoration in sheep," *The Journal of Thoracic and Cardiovascular Surgery*, vol. 129, no. 2, pp. 382–390, 2005.

[48] P. Helm, M. F. Beg, M. I. Miller, and R. L. Winslow, "Measuring and mapping cardiac Fibre and laminar architecture using diffusion tensor MR imaging," *Annals of the New York Academy of Sciences*, vol. 1047, no. 1, pp. 296–307, 2005.

[49] B. A. Landman, J. A. Farrell, C. K. Jones, S. A. Smith, J. L. Prince, and S. Mori, "Effects of diffusion weighting schemes on the reproducibility of DTI-derived fractional anisotropy, mean diffusivity, and principal eigenvector measurements at 1.5T," *NeuroImage*, vol. 36, no. 4, pp. 1123–1138, 2007.

[50] A. D. Scott, S. Nielles-Vallespin, P. F. Ferreira, L.-A. McGill, D. J. Pennell, and D. N. Firmin, "The effects of noise in cardiac diffusion tensor imaging and the benefits of averaging complex data," *NMR in Biomedicine*, vol. 29, no. 5, pp. 588–599, 2016.

[51] N. Tueni, J.-M. Allain, and M. Genet, "On the structural origin of the anisotropy in the myocardium: Multiscale modelling and analysis," *Journal of the Mechanical Behavior of Biomedical Materials*, vol. 138, p. 105600, 2023.

[52] J. P. Laissy, V. Gaxotte, E. Ironde-Laissy, I. Klein, A. Ribet, A. Bendriss *et al.*, "Cardiac diffusion-weighted MR imaging in recent, subacute, and chronic myocardial infarction: A pilot study," *Journal of Magnetic Resonance Imaging*, vol. 38, no. 6, pp. 1377–1387, 2013.

[53] G. J. Strijkers, A. Bouts, W. M. Blankesteijn, T. H. J. M. Peeters, A. Vilanova, M. C. van Prooijen *et al.*, "Diffusion tensor imaging of left ventricular remodeling in response to myocardial infarction in the mouse," *NMR in Biomedicine*, vol. 22, no. 2, pp. 182–190, 2009.

[54] C. Nguyen, Z. Fan, Y. Xie, J. Pang, P. Speier, X. Bi *et al.*, "In vivo diffusion-tensor MRI of the human heart on a 3 tesla clinical scanner: An optimized second order (M2) motion compensated diffusion-preparation approach," *Magnetic Resonance in Medicine*, vol. 76, no. 5, pp. 1354–1363, 2016.

[55] E. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006.

[56] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[57] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.

[58] C. Shannon, "Communication in the presence of noise," *Proceedings of the IRE*, vol. 37, no. 1, pp. 10–21, 1949.

[59] H. Nyquist, "Certain topics in telegraph transmission theory," *Transactions of the American Institute of Electrical Engineers*, vol. 47, no. 2, pp. 617–644, 1928.

[60] Y. Wu, Y.-J. Zhu, Q.-Y. Tang, C. Zou, W. Liu, R.-B. Dai *et al.*, "Accelerated MR diffusion tensor imaging using distributed compressed sensing." *Magnetic Resonance in Medicine*, vol. 71, no. 2, pp. 763–72, 2014.

[61] G. Adluru, E. Hsu, and E. V. R. Di Bella, "Constrained Reconstruction of Sparse Cardiac MR DTI Data," 2007, pp. 91–99.

[62] J. Huang, L. Wang, C. Chu, W. Liu, and Y. Zhu, "Accelerating cardiac diffusion tensor imaging combining local low-rank and 3D TV constraint," *Magnetic Resonance Materials in Physics, Biology and Medicine*, vol. 32, no. 4, pp. 407–422, 2019.

[63] I. Teh, D. McClymont, E. Carruth, J. Omens, A. McCulloch, and J. E. Schneider, "Improved compressed sensing and super-resolution of cardiac diffusion MRI with structure-guided total variation," *Magnetic Resonance in Medicine*, vol. 84, no. 4, pp. 1868–1880, 2020.

[64] Z.-P. Liang and P. C. Lauterbur, *Principles of magnetic resonance imaging: A signal processing perspective*. New York: IEEE Press, 2000.

[65] M. A. Bernstein, K. F. King, and X. J. Zhou, *Handbook of MRI pulse sequences*. Elsevier Inc., 2004.

[66] Radiology Cafe, *Introduction to MRI* , 2021. [Online]. Available from: https://www.radiologycafe.com/frcr-physics-notes/mr-imaging/introduction-to-mri/

[67] F. Bloch, "Nuclear induction," *Physical Review*, vol. 70, no. 7-8, pp. 460–474, 1946.

[68] Radiology Key, *Basic Principles of Cardiovascular Magnetic Resonance*, 2020. [Online]. Available from: https://radiologykey.com/basic-principles-of-cardiovascular-magnetic-resonance/

[69] P. C. Lauterbur, "Image formation by induced local interactions: Examples employing nuclear magnetic resonance," *Nature*, vol. 242, no. 5394, pp. 190–191, 1973.

[70] D. W. McRobbie, E. A. Moore, M. J. Graves, and M. R. Prince, *MRI from Picture to Proton*. New York: Cambridge University Press, 2006.

[71] D. B. Twieg, "The k-trajectory formulation of the NMR imaging process with applications in analysis and synthesis of imaging methods," *Medical Physics*, vol. 10, no. 5, pp. 610–621, 1983.

[72] P. C. Lauterbur, "Image formation by induced local interactions: Examples employing nuclear magnetic resonance," *Nature*, vol. 242, no. 5394, pp. 190–191, 1973.

[73] P. Mansfield, "Multi-planar image formation using NMR spin echoes," *Journal of Physics C: Solid State Physics*, vol. 10, no. 3, pp. L55–L58, 1977.

[74] C. Lazarus, "Compressed Sensing in MRI: optimization-based design of k-space filling curves for accelerated MRI," Ph.D. dissertation, Université Paris Saclay, 2018.

[75] M. A. Bernstein, K. F. King, and X. J. Zhou, *Handbook of MRI Pulse Sequences*. Burlington: Academic Press, 2004.

[76] J. D. O'Sullivan, "A fast sinc function gridding algorithm for Fourier inversion in computer tomography," *IEEE Transactions on Medical Imaging*, vol. 4, no. 4, pp. 200–207, 1985.

[77] J. Jackson, C. Meyer, D. Nishimura, and A. Macovski, "Selection of a convolution function for Fourier inversion using gridding," *IEEE Transactions on Medical Imaging*, vol. 10, no. 3, pp. 473–478, 1991.

[78] P. Beatty, D. Nishimura, and J. Pauly, "Rapid gridding reconstruction with a minimal oversampling ratio," *IEEE Transactions on Medical Imaging*, vol. 24, no. 6, pp. 799–808, 2005.

[79] D. K. Jones, PhD, Ed., *Diffusion MRI: Theory, Methods, and Applications* . Oxford University Press, 9 2012, vol. 1. ISBN 9780195369779

[80] A. A. Holmes, D. Scollan, and R. L. Winslow, "Direct histological validation of diffusion tensor MRI in formaldehyde-fixed myocardium," *Magnetic Resonance in Medicine*, vol. 44, no. 1, pp. 157–161, 2000.

[81] I. Teh, D. McClymont, M.-C. Zdora, H. J. Whittington, V. Davidoiu, J. Lee *et al.*, "Validation of diffusion tensor MRI measurements of cardiac microstructure with structure tensor synchrotron radiation imaging," *Journal of Cardiovascular Magnetic Resonance*, vol. 19, no. 1, p. 31, 2017.

[82] P. T. Callaghan, *Principles of Nuclear Magnetic Resonance Microscopy.* Oxford University PressOxford, 1991.

[83] N. Nissan, E. Furman-Haran, M. Feinberg-Shapiro, D. Grobgeld, E. Eyal, T. Zehavi *et al.*, "Tracking the Mammary Architectural Features and Detecting Breast Cancer with Magnetic Resonance Diffusion Tensor Imaging," *Journal of Visualized Experiments*, no. 94, 2014.

[84] E. L. Hahn, "Spin Echoes," *Physical Review*, vol. 80, no. 4, pp. 580–594, 1950.

[85] H. C. Torrey, "Bloch equations with diffusion terms," *Physical Review*, vol. 104, no. 3, pp. 563–565, 1956.

[86] S. Mori and P. B. Barker, "Diffusion magnetic resonance imaging: Its principle and applications," *The Anatomical Record*, vol. 257, no. 3, pp. 102–109, 1999.

[87] E. O. Stejskal and J. E. Tanner, "Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient," *The Journal of Chemical Physics*, vol. 42, no. 1, pp. 288–292, 1965.

[88] M. K. Tarzamni, A. Kiani Nazarlou, and J. Abdolmohammadi, "MRI imaging and pancreatic diseases," *Acta healthmedica*, vol. 1, no. 1, pp. 8–22, 2016.

[89] N. S. White, C. R. McDonald, N. Farid, J. Kuperman, D. Karow, N. M. Schenker-Ahmed *et al.*, "Diffusion-weighted imaging in cancer: Physical foundations and applications of restriction spectrum imaging," *Cancer Research*, vol. 74, no. 17, pp. 4638–4652, 2014.

[90] P. W. Kuchel, G. Pagès, K. Nagashima, S. Velan, V. Vijayaragavan, V. Nagarajan *et al.*, "Stejskal-tanner equation derived in full," *Concepts in Magnetic Resonance Part A*, vol. 40A, no. 5, pp. 205–214, 2012.

[91] D. Le Bihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis, and M. Laval-Jeantet, "MR imaging of intravoxel incoherent motions: Application to diffusion and perfusion in neurologic disorders," *Radiology*, vol. 161, no. 2, pp. 401–407, 1986.

[92] J. Mattiello, P. Basser, and D. Lebihan, "Analytical expressions for the b matrix in NMR diffusion imaging and spectroscopy," *Journal of Magnetic Resonance, Series A*, vol. 108, no. 2, pp. 131–141, 1994.

[93] S. H. Gilbert, A. P. Benson, P. Li, and A. V. Holden, "Regional localisation of left ventricular sheet structure: integration with current models of cardiac fibre, sheet and band structure," *European Journal of Cardio-Thoracic Surgery*, vol. 32, no. 2, pp. 231–249, 2007.

[94] T. D. Nguyen, O. E. Kadri, and R. S. Voronov, "An introductory overview of image-based computational modelling in personalized cardiovascular medicine," *Frontiers in Bioengineering and Biotechnology*, vol. 8, 2020.

[95] C. T. Nguyen, J. Dawkins, X. Bi, E. Marbán, and D. Li, "Diffusion tensor cardiac magnetic resonance reveals exosomes from cardiosphere-derived cells preserve myocardial fibre architecture after myocardial infarction," *JACC: Basic to Translational Science*, vol. 3, no. 1, pp. 97–109, 2018.

[96] C. Coghlan and J. I. E. Hoffman, "Leonardo da Vinci's flights of the mind must continue: Cardiac architecture and the fundamental relation of form and function revisited," *European journal of cardio-thoracic surgery : official journal of the European Association for Cardio-thoracic Surgery*, vol. 29 Suppl 1, pp. 4–17, 2006. [Online]. Available from: https://api.semanticscholar.org/CorpusID:26588075

[97] E. X. Wu, Y. Wu, J. M. Nicholls, J. Wang, S. Liao, S. Zhu *et al.*, "MR diffusion tensor imaging study of postinfarct myocardium structural remodeling in a porcine model," *Magnetic Resonance in Medicine*, vol. 58, no. 4, pp. 687–695, 2007.

[98] C. Mekkaoui, S. Huang, H. H. Chen, G. Dai, T. G. Reese, W. J. Kostis *et al.*, "Fiber architecture in remodeled myocardium revealed with a quantitative diffusion CMR tractography framework and histological validation," *Journal of Cardiovascular Magnetic Resonance*, vol. 14, no. 1, p. 70, 2012.

[99] H. Wei, M. Viallon, B. M. Delattre, V. M. Pai, H. Wen, H. Xue *et al.*, "In vivo cardiac diffusion tensor imaging in free-breathing conditions," *Journal of Cardiovascular Magnetic Resonance*, vol. 15, no. S1, p. P231, 2013.

[100] N. Toussaint, C. T. Stoeck, T. Schaeffter, S. Kozerke, M. Sermesant, and P. G. Batchelor, "In vivo human cardiac fibre architecture estimation using shape-based diffusion tensor processing," *Medical Image Analysis*, vol. 17, no. 8, pp. 1243–1255, 2013.

[101] P. Jezzard, A. S. Barnett, and C. Pierpaoli, "Characterisation of and correction for eddy current artifacts in echo planar diffusion imaging," *Magnetic Resonance in Medicine*, vol. 39, no. 5, pp. 801–812, 1998.

[102] T. Reese, O. Heid, R. Weisskoff, and V. Wedeen, "Reduction of eddy-current-induced distortion in diffusion MRI using a twice-refocused spin echo," *Magnetic Resonance in Medicine*, vol. 49, no. 1, pp. 177–182, 2003.

[103] C. Barmet, N. De Zanche, B. J. Wilm, and K. P. Pruessmann, "A transmit/receive system for magnetic field monitoring of in vivo MRI," *Magnetic Resonance in Medicine*, vol. 62, no. 1, pp. 269–276, 2009.

[104] J. Dou, T. G. Reese, W. I. Tseng, and V. J. Wedeen, "Cardiac diffusion MRI without motion effects," *Magnetic Resonance in Medicine*, vol. 48, no. 1, pp. 105–114, 2002.

[105] D. E. Sosnovik, R. Wang, G. Dai, T. G. Reese, and V. J. Wedeen, "Diffusion MR tractography of the heart," *Journal of Cardiovascular Magnetic Resonance*, vol. 11, no. 1, p. 47, 2009.

[106] D. Le Bihan, C. Poupon, A. Amadon, and F. Lethimonnier, "Artifacts and pitfalls in diffusion MRI," *Journal of Magnetic Resonance Imaging*, vol. 24, no. 3, pp. 478–488, 2006.

[107] C. von Deuster, C. T. Stoeck, M. Genet, D. Atkinson, and S. Kozerke, "Spin echo versus stimulated echo diffusion tensor imaging of the in vivo human heart," *Magnetic Resonance in Medicine*, vol. 76, no. 3, pp. 862–872, 2016.

[108] R. R. Edelman, J. Gaa, V. J. Wedeen, E. Loh, J. M. Hare, P. Prasad *et al.*, "In vivo measurement of water diffusion in the human heart," *Magnetic Resonance in Medicine*, vol. 32, no. 3, pp. 423–428, 1994.

[109] P. F. Ferreira, S. Nielles-Vallespin, A. D. Scott, R. de Silva, P. J. Kilner, D. B. Ennis *et al.*, "Evaluation of the impact of strain correction on the orientation

of cardiac diffusion tensors with in vivo and ex vivo porcine hearts," *Magnetic Resonance in Medicine*, vol. 79, no. 4, pp. 2205–2215, 2018.

[110] U. Gamper, P. Boesiger, and S. Kozerke, "Diffusion imaging of the in vivo heart using spin echoes-considerations on bulk motion sensitivity," *Magnetic Resonance in Medicine*, vol. 57, no. 2, pp. 331–337, 2007.

[111] J. G. Pipe and T. L. Chenevert, "A progressive gradient moment nulling design technique," *Magnetic Resonance in Medicine*, vol. 19, no. 1, pp. 175–179, 1991.

[112] C. T. Stoeck, C. von Deuster, M. Genet, D. Atkinson, and S. Kozerke, "Second-order motion-compensated spin echo diffusion tensor imaging of the human heart," *Magnetic Resonance in Medicine*, vol. 75, no. 4, pp. 1669–1676, 2016.

[113] C. L. Welsh, E. V. R. DiBella, and E. W. Hsu, "Higher-order motion-compensation for in vivo cardiac diffusion tensor imaging in rats," *IEEE Transactions on Medical Imaging*, vol. 34, no. 9, pp. 1843–1853, 2015.

[114] D. Noll, C. Meyer, J. Pauly, D. Nishimura, and A. Macovski, "A homogeneity correction method for magnetic resonance imaging with time-varying gradients," *IEEE Transactions on Medical Imaging*, vol. 10, no. 4, pp. 629–637, 1991.

[115] D. C. Noll, J. M. Pauly, C. H. Meyer, D. G. Nishimura, and A. Macovskj, "Deblurring for non-2D fourier transform magnetic resonance imaging," *Magnetic Resonance in Medicine*, vol. 25, no. 2, pp. 319–333, 1992.

[116] L.-C. Man, J. M. Pauly, and A. Macovski, "Multifrequency interpolation for fast off-resonance correction," *Magnetic Resonance in Medicine*, vol. 37, no. 5, pp. 785–792, 1997.

[117] A. Nylund, "Off-resonance correction for magnetic resonance imaging with spiral trajectories," Ph.D. dissertation, 2014.

[118] D. Atkinson, D. A. Porter, D. L. Hill, F. Calamante, and A. Connelly, "Sampling and reconstruction effects due to motion in diffusion-weighted interleaved echo planar imaging," *Magnetic Resonance in Medicine*, vol. 44, no. 1, pp. 101–109, 2000.

[119] S. J. Vannesjo, N. N. Graedel, L. Kasper, S. Gross, J. Busch, M. Haeberlin *et al.*, "Image reconstruction using a gradient impulse response model for trajectory prediction," *Magnetic Resonance in Medicine*, vol. 76, no. 1, pp. 45–58, 2016.

[120] L. Feng, T. Benkert, K. T. Block, D. K. Sodickson, R. Otazo, and H. Chandarana, "Compressed sensing for body MRI," *Journal of Magnetic Resonance Imaging*, vol. 45, no. 4, pp. 966–987, 2017.

[121] D. J. Larkman and R. G. Nunes, "Parallel magnetic resonance imaging," *Physics in Medicine and Biology*, vol. 52, no. 7, pp. R15–R55, 2007.

[122] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: sensitivity encoding for fast MRI." *Magnetic resonance in medicine*, vol. 42, no. 5, pp. 952–62, 1999.

[123] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang *et al.*, "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1202–1210, 2002.

[124] M. Lustig and J. M. Pauly, "SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space," *Magnetic Resonance in Medicine*, vol. 64, no. 2, pp. 457–471, 2010.

[125] K. G. Hollingsworth, "Reducing acquisition time in clinical MRI by data undersampling and compressed sensing reconstruction," *Physics in Medicine and Biology*, vol. 60, no. 21, pp. R297–R322, 2015.

[126] J. Huang, L. Wang, C. Chu, Y. Zhang, W. Liu, and Y. Zhu, "Cardiac diffusion tensor imaging based on compressed sensing using joint sparsity and low-rank approximation," *Technology and Health Care*, vol. 24, no. s2, pp. S593–S599, 2016.

[127] S. Ma, C. T. Nguyen, A. G. Christodoulou, D. Luthringer, J. Kobashigawa, S.-E. Lee *et al.*, "Accelerated cardiac diffusion tensor imaging using joint low-rank and sparsity constraints," *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 10, pp. 2219–2230, 2018.

[128] D. McClymont, I. Teh, H. J. Whittington, V. Grau, and J. E. Schneider, "Prospective acceleration of diffusion tensor imaging with compressed sensing using

adaptive dictionaries," *Magnetic Resonance in Medicine*, vol. 76, no. 1, pp. 248–258, 2016.

[129] D. Donoho and X. Huo, "Uncertainty principles and ideal atomic decomposition," *IEEE Transactions on Information Theory*, vol. 47, no. 7, pp. 2845–2862, 2001.

[130] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Transactions on Information Theory*, vol. 53, no. 12, pp. 4655–4666, 2007.

[131] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 33–61, 1998.

[132] M. Fornasier and H. Rauhut, "Iterative thresholding algorithms," *Applied and Computational Harmonic Analysis*, vol. 25, no. 2, pp. 187–208, 2008.

[133] Y. Nesterov, "Gradient methods for minimizing composite functions," *Mathematical Programming*, vol. 140, no. 1, pp. 125–161, 2013.

[134] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.

[135] Y. Nesterov, "A method for solving the convex programming problem with convergence rate $$O(1/k^2)$$," *Proceedings of the USSR Academy of Sciences*, vol. 269, pp. 543–547, 1983. [Online]. Available from: https://api.semanticscholar.org/CorpusID:145918791

[136] J. Huang, S. Zhang, H. Li, and D. Metaxas, "Composite splitting algorithms for convex optimization," *Computer Vision and Image Understanding*, vol. 115, no. 12, pp. 1610–1622, 2011.

[137] J. Huang, S. Zhang, and D. Metaxas, "Efficient MR image reconstruction for compressed MR imaging," *Medical Image Analysis*, vol. 15, no. 5, pp. 670–679, 2011.

[138] D. Liang, H. Wang, Y. Chang, and L. Ying, "Sensitivity encoding reconstruction with nonlocal total variation regularization," *Magnetic Resonance in Medicine*, vol. 65, no. 5, pp. 1384–1392, 2011.

[139] J. Huang and F. Yang, "Compressed magnetic resonance imaging based on wavelet sparsity and nonlocal total variation," In: *2012 9th IEEE International Symposium on Biomedical Imaging (ISBI)*. IEEE, 5 2012. doi: 10.1109/ISBI.2012.6235718. ISBN 978-1-4577-1858-8 pp. 968–971.

[140] F. Knoll, K. Bredies, T. Pock, and R. Stollberger, "Second order total generalized variation (TGV) for MRI," *Magnetic Resonance in Medicine*, vol. 65, no. 2, pp. 480–491, 2011.

[141] W. Guo, J. Qin, and W. Yin, "A New Detail-Preserving Regularization Scheme," *SIAM Journal on Imaging Sciences*, vol. 7, no. 2, pp. 1309–1334, 1 2014.

[142] D. Zhao, H. Du, Y. Han, and W. Mei, "Compressed sensing MR image reconstruction exploiting TGV and wavelet sparsity," *Computational and Mathematical Methods in Medicine*, vol. 2014, pp. 1–11, 2014.

[143] Lihan He and L. Carin, "Exploiting structure in wavelet-based Bayesian compressive sensing," *IEEE Transactions on Signal Processing*, vol. 57, no. 9, pp. 3488–3497, 2009.

[144] L. Jacques, L. Duval, C. Chaux, and G. Peyré, "A panorama on multiscale geometric representations, intertwining spatial, directional and frequency selectivity," *Signal Processing*, vol. 91, no. 12, pp. 2699–2730, 2011.

[145] X. Qu, W. Zhang, D. Guo, C. Cai, S. Cai, and Z. Chen, "Iterative thresholding compressed sensing MRI based on contourlet transform," *Inverse Problems in Science and Engineering*, vol. 18, no. 6, pp. 737–758, 2010.

[146] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde, "Model-based compressive sensing," *IEEE Transactions on Information Theory*, vol. 56, no. 4, pp. 1982–2001, 2010.

[147] M. Usman, C. Prieto, T. Schaeffter, and P. G. Batchelor, "k-t group sparse: A method for accelerating dynamic MRI," *Magnetic Resonance in Medicine*, vol. 66, no. 4, pp. 1163–1176, 2011.

[148] J. P. Haldar and Z.-P. Liang, "Spatiotemporal imaging with partially separable functions: A matrix recovery approach," In: *2010 IEEE International*

Symposium on Biomedical Imaging: From Nano to Macro. IEEE, 2010. doi: 10.1109/ISBI.2010.5490076. ISBN 978-1-4244-4125-9 pp. 716–719.

[149] Z.-p. Liang, "Spatiotemporal imaging with partially separable functions," In: *2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. IEEE, 2007. doi: 10.1109/ISBI.2007.357020. ISBN 1-4244-0671-4 pp. 988–991.

[150] M. Hong, Y. Yu, H. Wang, F. Liu, and S. Crozier, "Compressed sensing MRI with singular value decomposition-based sparsity basis," *Physics in Medicine and Biology*, vol. 56, no. 19, pp. 6311–6325, 2011.

[151] J. P. Haldar, D. Hernando, and Zhi-Pei Liang, "Compressed-sensing MRI With random encoding," *IEEE Transactions on Medical Imaging*, vol. 30, no. 4, pp. 893–903, 2011.

[152] N. Chauffert, P. Ciuciu, J. Kahn, and P. Weiss, "Variable Density Sampling with Continuous Trajectories," *SIAM Journal on Imaging Sciences*, vol. 7, no. 4, pp. 1962–1992, 2014.

[153] C. Boyer, N. Chauffert, P. Ciuciu, J. Kahn, and P. Weiss, "On the generation of sampling schemes for magnetic resonance imaging," *SIAM Journal on Imaging Sciences*, vol. 9, no. 4, pp. 2039–2072, 2016.

[154] J. Dou, W. I. Tseng, T. G. Reese, and V. J. Wedeen, "Combined diffusion and strain MRI reveals structure and function of human myocardial laminar sheets in vivo," *Magnetic Resonance in Medicine*, vol. 50, no. 1, pp. 107–113, 2003.

[155] A. Z. Lau, E. M. Tunnicliffe, R. Frost, P. J. Koopmans, D. J. Tyler, and M. D. Robson, "Accelerated human cardiac diffusion tensor imaging using simultaneous multislice imaging," *Magnetic Resonance in Medicine*, vol. 73, no. 3, pp. 995–1004, 2015.

[156] C. Mekkaoui, T. G. Reese, M. P. Jackowski, S. F. Cauley, K. Setsompop, H. Bhat *et al.*, "Diffusion tractography of the entire left ventricle by using free-breathing accelerated simultaneous multisection imaging," *Radiology*, vol. 282, no. 3, pp. 850–856, 2017.

[157] C. T. Nguyen, A. G. Christodoulou, J. Coll-Font, S. Ma, Y. Xie, T. G. Reese *et al.*, "Free-breathing diffusion tensor MRI of the whole left ventricle using second-order motion compensation and multitasking respiratory motion correction," *Magnetic Resonance in Medicine*, vol. 85, no. 5, pp. 2634–2648, 2021.

[158] I. Teh, D. McClymont, R. A. B. Burton, M. L. Maguire, H. J. Whittington, C. A. Lygate *et al.*, "Resolving fine cardiac structures in rats with high-resolution diffusion tensor imaging," *Scientific Reports*, vol. 6, no. 1, p. 30573, 2016.

[159] Y. Jiang, K. Pandya, O. Smithies, and E. W. Hsu, "Three-dimensional diffusion tensor microscopy of fixed mouse hearts," *Magnetic Resonance in Medicine*, vol. 52, no. 3, pp. 453–460, 2004.

[160] E. Candes and M. Wakin, "An introduction to compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, 2008.

[161] H. H. Wu, P. T. Gurney, B. S. Hu, D. G. Nishimura, and M. V. McConnell, "Free-breathing multiphase whole-heart coronary MR angiography using image-based navigators and three-dimensional cones imaging," *Magnetic Resonance in Medicine*, vol. 69, no. 4, pp. 1083–1093, 2013.

[162] J. H. Lee, B. A. Hargreaves, B. S. Hu, and D. G. Nishimura, "Fast 3D imaging using variable-density spiral trajectories with applications to limb perfusion," *Magnetic Resonance in Medicine*, vol. 50, no. 6, pp. 1276–1285, 2003.

[163] B. ADCOCK, A. C. HANSEN, C. POON, and B. ROMAN, "Breaking the coherence barrier: A new theory for compressed sensing," *Forum of Mathematics, Sigma*, vol. 5, p. e4, 2017.

[164] C. Boyer, J. Bigot, and P. Weiss, "Compressed sensing with structured sparsity and structured acquisition," *Applied and Computational Harmonic Analysis*, vol. 46, no. 2, pp. 312–350, 2019.

[165] G. Puy, P. Vandergheynst, and Y. Wiaux, "On variable density compressive sampling," *IEEE Signal Processing Letters*, vol. 18, no. 10, pp. 595–598, 2011.

[166] C. Lazarus, P. Weiss, A. Vignaud, and P. Ciuciu, "An empirical study of the maximum degree of undersampling in compressed sensing for T2*-weighted MRI," *Magnetic Resonance Imaging*, vol. 53, pp. 112–122, 2018.

[167] Y. Yang, C. M. Kramer, P. W. Shaw, C. H. Meyer, and M. Salerno, "First-pass myocardial perfusion imaging with whole-heart coverage using L1-SPIRiT accelerated variable density spiral trajectories," *Magnetic Resonance in Medicine*, vol. 76, no. 5, pp. 1375–1387, 2016.

[168] D. Kim, E. Adalsteinsson, and D. M. Spielman, "Simple analytic variable density spiral design," *Magnetic Resonance in Medicine*, vol. 50, no. 1, pp. 214–219, 2003.

[169] K. F. King, T. K. F. Foo, and C. R. Crawford, "Optimized gradient waveforms for spiral scanning," *Magnetic Resonance in Medicine*, vol. 34, no. 2, pp. 156–160, 1995.

[170] K. F. King, "Spiral scanning with anisotropic field of view," *Magnetic Resonance in Medicine*, vol. 39, no. 3, pp. 448–456, 1998.

[171] C. Liu, R. Bammer, D. Kim, and M. E. Moseley, "Self-navigated interleaved spiral (SNAILS): Application to high-resolution diffusion tensor imaging," *Magnetic Resonance in Medicine*, vol. 52, no. 6, pp. 1388–1396, 2004.

[172] H. E. Cline, X. Zong, and N. Gai, "Design of a logarithmic k-space spiral trajectory," *Magnetic Resonance in Medicine*, vol. 46, no. 6, pp. 1130–1135, 2001.

[173] H. Tan and C. H. Meyer, "Estimation of k-space trajectories in spiral MRI," *Magnetic Resonance in Medicine*, vol. 61, no. 6, pp. 1396–1404, 2009.

[174] M. Lustig, S.-J. Kim, and J. M. Pauly, "A fast method for designing time-optimal gradient waveforms for arbitrary k-space trajectories," *IEEE Transactions on Medical Imaging*, vol. 27, no. 6, pp. 866–873, 2008.

[175] Miki Lustig, *Time Optimal Gradient Design*. [Online]. Available from: https://people.eecs.berkeley.edu/~mlustig/Software.html

[176] J. Fessler, *Michigan Image Reconstruction Toolbox (MIRT)*. [Online]. Available from: https://web.eecs.umich.edu/~fessler/code/index.html

[177] M. Gorodezky, A. D. Scott, P. F. Ferreira, S. Nielles-Vallespin, D. J. Pennell, and D. N. Firmin, "Diffusion tensor cardiovascular magnetic resonance with a spiral trajectory: An in vivo comparison of echo planar and spiral stimulated

echo sequences," *Magnetic Resonance in Medicine*, vol. 80, no. 2, pp. 648–654, 2018.

[178] M. Gorodezky, P. F. Ferreira, S. Nielles-Vallespin, P. D. Gatehouse, D. J. Pennell, A. D. Scott *et al.*, "High resolution in-vivo DT-CMR using an interleaved variable density spiral STEAM sequence," *Magnetic Resonance in Medicine*, vol. 81, no. 3, pp. 1580–1594, 2019.

[179] J. Fessler and B. Sutton, "Nonuniform fast Fourier transforms using min-max interpolation," *IEEE Transactions on Signal Processing*, vol. 51, no. 2, pp. 560–574, 2003.

[180] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.

[181] M. Zanforlin, D. Munaretto, A. Zanella, and M. Zorzi, "SSIM-based video admission control and resource allocation algorithms," In: *2014 12th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*. IEEE, 5 2014. doi: 10.1109/WIOPT.2014.6850361. ISBN 978-3-901882-63-0 pp. 656–661.

[182] C. Lazarus, P. Weiss, N. Chauffert, F. Mauconduit, L. El Gueddari, C. Destrieux *et al.*, "SPARKLING: variable-density k-space filling curves for accelerated T2*-weighted MRI," *Magnetic Resonance in Medicine*, vol. 81, no. 6, pp. 3643–3661, 2019.

[183] P. Irarrazabal and D. G. Nishimura, "Fast three dimensional magnetic resonance imaging," *Magnetic Resonance in Medicine*, vol. 33, no. 5, pp. 656–662, 1995.

[184] G. Valvano, N. Martini, L. Landini, and M. F. Santarelli, "Variable density randomized stack of spirals (VDR-SoS) for compressive sensing MRI," *Magnetic Resonance in Medicine*, vol. 76, no. 1, pp. 59–69, 2016.

[185] S. Winkelmann, T. Schaeffter, T. Koehler, H. Eggers, and O. Doessel, "An optimal radial profile order based on the Golden Ratio for time-resolved MRI," *IEEE Transactions on Medical Imaging*, vol. 26, no. 1, pp. 68–76, 2007.

[186] T. Saito, Y. Takagaki, and T. Komatsu, "Three kinds of color total-variation semi-norms and its application to color-image denoising," In: *2011 18th IEEE International Conference on Image Processing.* IEEE, 9 2011. doi: 10.1109/ICIP.2011.6115716. ISBN 978-1-4577-1303-3 pp. 1457–1460.

[187] C. Liu, M. E. Moseley, and R. Bammer, "Simultaneous phase correction and SENSE reconstruction for navigated multi-shot DWI with non-Cartesian k-space sampling," *Magnetic Resonance in Medicine*, vol. 54, no. 6, pp. 1412–1422, 2005.

[188] T. Truong and A. Guidon, "High-resolution multishot spiral diffusion tensor imaging with inherent correction of motion-induced phase errors," *Magnetic Resonance in Medicine*, vol. 71, no. 2, pp. 790–796, 2014.

[189] J. Huang, C. Chen, and L. Axel, "Fast multi-contrast MRI reconstruction," *Magnetic Resonance Imaging*, vol. 32, no. 10, pp. 1344–1352, 2014.

[190] S. Chang, Bin Yu, and M. Vetterli, "Adaptive wavelet thresholding for image denoising and compression," *IEEE Transactions on Image Processing*, vol. 9, no. 9, pp. 1532–1546, 2000.

[191] Z. Hu, X. Ma, T. Truong, A. W. Song, and H. Guo, "Phase-updated regularized SENSE for navigator-free multishot diffusion imaging," *Magnetic Resonance in Medicine*, vol. 78, no. 1, pp. 172–181, 2017.

[192] G. H. Glover and S. Lai, "Self-navigated spiral fMRI: Interleaved versus single-shot," *Magnetic Resonance in Medicine*, vol. 39, no. 3, pp. 361–368, 1998.

[193] N. O. Addy, H. H. Wu, and D. G. Nishimura, "Simple method for MR gradient system characterization and k-space trajectory estimation," *Magnetic Resonance in Medicine*, vol. 68, no. 1, pp. 120–129, 2012.

[194] J. Y. Cheng, J. M. Santos, and J. M. Pauly, "Fast concomitant gradient field and field inhomogeneity correction for spiral cardiac imaging," *Magnetic Resonance in Medicine*, vol. 66, no. 2, pp. 390–401, 2011.

[195] J. D. Ianni and W. A. Grissom, "Trajectory Auto-Corrected image reconstruction," *Magnetic Resonance in Medicine*, vol. 76, no. 3, pp. 757–768, 2016.

[196] S. D. Yun and N. J. Shah, "Whole-brain high in-plane resolution fMRI using accelerated EPIK for enhanced characterisation of functional areas at 3T," *PLOS ONE*, vol. 12, no. 9, p. e0184759, 2017.

[197] S. D. Yun, R. Weidner, P. H. Weiss, and N. J. Shah, "Evaluating the utility of EPIK in a finger tapping fMRI experiment using BOLD detection and effective connectivity," *Scientific Reports*, vol. 9, no. 1, p. 10978, 2019.

[198] S. Gleichman and Y. C. Eldar, "Blind compressed sensing," In: *Lecture Notes in Computer Science*. Berlin: Springer, 2010, vol. 6365, pp. 386–393.

[199] S. G. Lingala and M. Jacob, "Blind Compressive Sensing Dynamic MRI," *IEEE Transactions on Medical Imaging*, vol. 32, no. 6, pp. 1132–1145, 6 2013.

[200] S. Bhave, S. G. Lingala, C. P. Johnson, V. A. Magnotta, and M. Jacob, "Accelerated whole-brain multi-parameter mapping using blind compressed sensing," *Magnetic Resonance in Medicine*, vol. 75, no. 3, pp. 1175–1186, 2016.

[201] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2010.

[202] R. J. H. van Gorkum, J. L. Weine, W. P. Segars, C. T. Stoeck, and S. Kozerke, "MRXCAT-CDTI: A numerical cardiac diffusion tensor imaging phantom," *CoRR*, vol. abs/2204.07966, 2022. [Online]. Available from: https://doi.org/10.48550/arXiv.2204.07966

[203] V. Vishnevskiy, T. Gass, G. Szekely, C. Tanner, and O. Goksel, "Isotropic total variation regularization of displacements in parametric image registration," *IEEE Transactions on Medical Imaging*, vol. 36, no. 2, pp. 385–395, 2017.

[204] P. C. Hansen and D. P. O'Leary, "The use of the L-curve in the regularization of discrete ill-posed problems," *SIAM Journal on Scientific Computing*, vol. 14, no. 6, pp. 1487–1503, 1993.